Fifth Biennial Report

June 1999 – August 2001

Contents

Part I

Overview – The Institute

Summary

When the last Biennial Report was written in Summer 1999, two new groups had just started their work at the institute: The Reactive and Hybrid Systems Group directed by Tom Henzinger and the Computer Graphics Group directed by Hans-Peter Seidel. While the build-up of the Computer Graphics Group proceeds as planned, Tom Henzinger and his wife have left Saarbrücken again at the end of 1999, due to private reasons.

During the subsequent search for a new director, an offer has been made to Thomas Lengauer, Director at the Institute for Algorithms and Scientific Computing at the Gesellschaft für Mathematik und Datenverarbeitung in Bonn. We are happy to report that

- Thomas Lengauer has accepted the offer and will join MPII on October 1, 2001. He directs the research unit on Computational Biology and Applied Algorithmics.

The other two groups for Algorithms and Complexity (Kurt Mehlhorn) and Programming Logics (Harald Ganzinger) continued with their successful work.

At present 34 research associates, 20 doctoral students and 12 postdocs are affiliated with the institute. The scientific staff is complemented by an administrative unit with 16 members (including secretaries), by a computing support unit (6 members of staff) and by our library (2 members of staff). The computing support unit currently operates a network of approximately 250 workstations.

In order to increase the number of independent research groups within MPII we also decided to install four *independent junior research groups*. These independent junior research groups are on the C3 level and provide an opportunity for outstanding young researchers to build up their own groups and conduct independent research in an intellectually stimulating environment for a period of up to five years.

- Offers for junior research groups have been made to Bruno Blanchet (CNRS, Paris), Fritz Eisenbrand (IASI, Rome, formerly at MPII), Marcus Magnor (Stanford University), and Matthias Rarey (GMD, Bonn), and we expect that the groups will start their work early next year.

Research Program

The institute is devoted to basic research in computer science, and in particular to the study of complex computer systems. Complexity in computer systems arises for various reasons:

A problem can be complex due to huge masses of data that have to be handled, sometimes in real time. For this sort of problem efficient algorithms and data structures as well as parallel processing are of great importance. Kurt Mehlhorn's group addresses this facet of complexity.

Complexity can mean logical complexity as we find it in large software systems, with many layers of abstraction, where applications from different problem domains interact with each other in often unpredictable ways. Here we need to apply methods based on mathematical logic in order to structure, reason about, and develop more systematically, such large systems. Harald Ganzinger's group addresses this facet of complexity.

Research in computational biology is aimed at understanding and modeling complex biological structures and processes. Typical examples are the extraction of biological knowledge from genomic data by protein structure prediction and protein family analysis, the interpretation of mRNA expression data, or the study of metabolic and regulatory networks. Tom Lengauer's group addresses this facet of complexity.

Computer systems are more and more used to realize and simulate parts of the real or an imaginary world. Such simulations require to model/acquire, to process, and to render complex objects. The goal of computer graphics is to turn abstract information into visual images and to allow the user to interact with complex objects and data in a natural and intuitive way. Hans-Peter Seidel's group addresses this facet of complexity.

In order to make progress along the lines above, our work is both theoretical and practical with a focus on

- first-class basic research on new algorithms,
- the integration of new algorithms into a system, and the evaluation of the system in practical applications.
- We also try to provide a stimulating environment for junior researchers that allows them to develop and build their own research programs and groups.

We do well in all three aspects. The institute has impact through publications, software and people. We did first-class research on a broad range of topics. Our work has appeared in top conferences and journals, our software has been used in a variety of applications, and researchers from MPII have spread out over other institutions.

Recruiting

During the last decade Saarbrücken has become a center of research in computer science that is visible to the whole world. We have recently taken several steps to exploit this fact and create additional opportunities for excellent students and junior researchers from abroad.

Members of the institute together with their colleagues at the CS department of Saarland University participate in the *Graduiertenkolleg (Graduate Research Center) Quality Guarantees for Computer Systems*. The graduate research center comprises a program of collaborative research and advanced studies on quality guarantees for computer systems with an emphasis on predictable running times, provable correctness, and guaranteed quality of service. The program offers several PhD fellowships.

In Summer 2000, MPII in cooperation with the CS department at Saarland University has been selected as an *International Max-Planck Research School for Computer Science* by the Max-Planck-Society. This is a graduate school in computer science (with courses partly taught in english) which offers both a Masters as well as a PhD program. Since its start in fall 2000 a total of 17 students have entered the school (4 from India, 3 from Croatia, 3 from Pakistan, 2 from Russia and 1 from each Lithuania, China, Corea, Japan and Iran).

In Fall 2000, MPII has been selected by the European Commission as a *Marie Curie Training Site for Strategies and Methods for Complex Computer Systems*. Within this program the institute offers fellowships to PhD students from the EU and associated countries. These fellowships are available for a period of 3 to 9 months and should also help to further strenghten the cooperation between MPII and european partner institutions. Three PhD students have visited the institute so far within this program, two more have been invited for the coming months.

In Summer 2001, the institute (in cooperation with the CS department at Saarland University) has been selected by the DAAD as an *International Quality Network*. The idea of this program is to support already existing cooperations with non german institutions. The foreign partners within this project are the Seoul National University in South Korea, the Khabarovsk State University in Russia and (specially for MPII) all Indian Institutes of Technology.

In 2000 and 2001, MPII has invited students from several institutions in India to spend their *summer internships* at our institute. Altogether 18 students have been involved in small projects that have been defined by members of the three research groups. Some of the students have meanwhile indicated that they may be interested in doing a PhD here at Saarbrücken.

Grants

The institute is involved in a number of projects related to research grants awarded by the European Union (EU), by the German-Israel Foundation (GIF), by the German Academic Exchange Service (DAAD), by the German Science Foundation (DFG) and by the German Ministry for Education and Research (BMBF), among others. Cooperation with industry has also substantially increased. Funding of all third-party projects in 2001 is about 1 051 TDM (it was 932 TDM in 2000). For the descriptions of these grants see sections III.14, IV.11 and V.16.

Results

In the parts to follow we describe in detail, for the Algorithms and Complexity group (Kurt Mehlhorn), Programming Logics Group (Harald Ganzinger) and the Computer Graphics Group (Hans-Peter Seidel) the research programs and results obtained in the period May 1999 through August 2001. We also briefly discuss the work of the new group. We have continued to be very successful in our research, as documented by our many scientific publications, including about 420 articles in journals, books or proceedings of major international conferences.

Many of the institute's results are, in addition, available to the public through computer programs such as the LEDA library of efficient algorithms, the SPASS theorem prover for first-order logic, and several graphics toolkits. Our software is widely used and has been validated in a variety of applications.

Teaching activities

The institute makes an effort to offer a variety of courses to computer science students of the Universität des Saarlandes. Courses taught during the period of this report are listed in Sections III.13, IV.9 and V.14; In the period of this report 21 doctoral dissertations and 6 habilitations have been successfully completed.

Professional activities

Members of the institute have been involved in the organization of 28 workshops and conferences. In 89 cases we have been invited to join the program committee of major international conferences, not counting program committee memberships for national and international workshops. Finally, we serve on the editorial board of 23 scientific journals.

Offers for Faculty Positions

The following members of the institute have accepted offers for faculty positions or equivalent research positions in the reported period: Susanne Albers (Universität Dortmund, Germany - now at Universität Freiburg), Christian Duncan (University of Miami), Rudolf Fleischer (University Waterloo, Canada - now at Hongkong University), Sherif Ghali (University of Alberta, Canada), Wolfgang Heidrich (University of British Columbia, Canada), Leif Kobbelt (RWTH Aachen, Germany), Hans-Peter Lenhof (Universität des Saarlandes, Germany), Petra Mutzel (TU Vienna,

Austria), Zeev Nutov (Open University of Israel, Ramat Aviv), Lorant Porkolab (Imperial College of Science, London), Jop Sibeyn (Umea University, Sweden), Cyril Soler (INRIA Rhone-Alpes, Grenoble, France), Steve Seiden (Louisiana State University), Roberto Solis-Oba (University of Western Ontario).

Part II

Overview – The Research Units

1 The Algorithms and Complexity Group

Vision

Kurt Mehlhorn's vision for AG1 is threefold:

- do first-class basic research in algorithms,
- build generally useful software libraries and implementations for specific application areas and contribute towards a theory of implementation,
- give junior researchers the possibility to work in a stimulating environment and to build their own research programs and groups.

We do well in all three aspects. The research group has impact through publications, software, and people. We did first-class research on a broad range of topics. Our work has appeared in top conferences and journals, our software libraries LEDA, CGAL, and BALL are widely used, and researchers from the group have spread out over other institutions[1]. Kurt Mehlhorn is the only permanent member of the group, other researchers stay in the group between one and seven years. More than half of the current group members joined the group since the last review; this influences our choice of research topics and our working style.

The efforts for items one and two relate approximately as 2 to 1. Many group members contribute to both items. Also, a considerable fraction of our theoretical work is strongly related to our implementation work and either leads to improved implementations or is motivated by shortcomings of our implementations.

In the sections that follow, I discuss group organization, our choice of research topics, software development, representative achievements, group development, cooperations, industrial take-up, and threats.

Organization

We are currently pursuing six research areas, each having its own coordinator. The subgroups and their coordinators are:

- Foundations, Coordinator Berthold Vöcking
- Computational Geometry, Coordinator Edgar Ramos
- Advanced Models of Computation, Coordinator Peter Sanders
- Assembly and Simulation, Coordinator Elmar Schömer
- Combinatorial Optimization, Coordinator Kurt Mehlhorn
- Software Libraries, Coordinator Kurt Mehlhorn. Starting October 1st, 2001 Lutz Kettner will coordinate the group. Lutz worked previously at FU Berlin, ETH Zürich and UNC and is one of the key developers of CGAL.

The coordinators coordinate research and teaching in their area and together with Kurt Mehlhorn form the steering committee of the group. The steering committee meets weekly to discuss and organize the work of the group. We use our noon seminar (twice weakly), mini courses, and group meetings to educate each other and to inform other group members about own work. Most group members contribute to more than one area.

[1] Mentioning only professorships in German speaking countries: Stefan Näher (Trier), Michael Kaufmann (Tübingen), Michiel Smid (Magdeburg), Torben Hagerup (Frankfurt), Susanne Albers (Freiburg), Petra Mutzel (Wien), Jop Sibeyn (now at Umeå and will move to Halle or Leipzig), and Hans-Peter Lenhof (Saarbrücken).

I encourage senior group members to build their own subgroups and to apply for their own research funding, which they can take with them once they leave the group. Peter Sanders, Elmar Schömer, and Berthold Vöcking have projects funded by the Deutsche Forschungs Gemeinschaft (German Research Foundation), see Section III.14.2 for additional information.

Research Topics

Our choice of research areas is long-term. We reconsider them, as senior researchers leave and as new opportunities arise. Hiring decisions on all levels (PhD students, Postdocs, Research Associates) are made on quality and fit into our research program. Within the research areas, we choose research topics according to international trends and the interests of our researchers. Over the past two years our research spectrum developed as follows:

Network algorithms, computational geometry, combinatorial optimization, scheduling and routing, exact geometric computation, models of computation, algorithm engineering, and software libraries retained their importance, our efforts in graph drawing[2], computational biology[3], and on-line and approximation algorithms[4] decreased. We were also successful to attract new research associates with new areas. Elmar Schömer brought in simulation and virtual reality, Peter Sanders strengthened algorithm engineering and external memory computation, Berthold Vöcking strengthened scheduling and routing, Lutz Kettner (starting October 1st 2001) will strengthen computational geometry and our software projects, and Susanne Schmitt brings computer algebra expertise into our exact computation project.

We recently started a new long term project to which about one-quarter of the group (Elmar Schömer, Lutz Kettner, Kurt Mehlhorn, Susan Hert, Edgar Ramos, Susanne Schmitt, Nicola Geismann, Christian Lennerz, Stefan Funke, and Thomas Warken) will contribute: Effective Computational Geometry (ECG). In the past computational geometry has concentrated on linear objects. The challenge of the next years is to extend algorithms, kernels, and software systems to curved objects. This will require theoretical progress as well as serious systems building. A successful effort will have to combine expertise in algorithms, algebra, numerical analysis, data structures, and software engineering. We also expect close collaboration with the members of the mesh processing group of AG4. We have formed the ECG group in 01, secured EU- and DFG-funding, ran an internal seminar series in the summer of 01, will teach (Elmar Schömer, Lutz Kettner, and Kurt Mehlhorn) a course in the winter term 01/02, and will run seminars and system projects in the winter term 01/02 and subsequent terms. See Sections III.10.3 and III.14.1.3 for further information.

Software Projects

A distinctive feature of our group is the combination of theoretical research and software development. LEDA (see Section III.10.2), CGAL (see Section III.10.1), and BALL (see Section III.9.4) have a long history in the group and achieved international recognition. The combined number of installations is several thousand. Smaller library efforts are SCIL (see Section III.5.2) and LEDA-SM (see Section III.6.4). We also develop software for specific applications, e.g., constraint programming (see Section III.5.3), assembly and simulation (see Section III.8), and bioinformatics (see Section III.9).

[2]Petra Mutzel's group moved to Vienna.

[3]Hans-Peter Lenhof's group moved to the CS department in Saarbrücken.

[4]Most members of the group moved to permanent positions at other institutions: Susanne Albers to Freiburg, Stefano Leonardi to Rome, Naveen Garg to IIT Dehli, Klaus Jansen to Kiel.

Some Achievements

Among our research achievements, the following are representative from the different subgroups:

- Scheduling Parallel Tasks at Twilight, see Section III.4.2. Hannah Bast received the Otto-Hahn Medaille for her work.
- Load Balancing, see Section III.6.2.
- Curve and Surface Reconstruction, see Sections III.5.4 and III.7.3.
- Quadric Surface Intersection, see Section III.8.5.
- The bioinformatics subgroup (Hans-Peter Lenhof) led a consortium that was awarded a five million DM grant from the DFG, see Sections III.9 and III.14.2.1.
- Stefan Näher and Kurt Mehlhorn finished the LEDA book (published by Cambridge University Press in November 99) and Tetsuo Asano (JAIST) completed an introductory book on LEDA (to be published in 2002). CGAL established itself in the computational geometry community.

Group Development

The composition of the group changed considerably over the past two years.

New additions to the group are: Peter Sanders moved from Postdoc to Research Associate and is now coordinating the research on advanced models of computation, Berthold Vöcking joined us in 2000 and is now coordinating our work in foundations; Elmar Schömer, Thomas Warken and Christian Lennerz form our new group on virtual reality and simulation (The group is directed by Elmar Schömer. He worked previously in the CS department); and Susanne Schmitt joined in 2000 (she brings computer algebra expertise into our exact computation project). The following postdocs joined in 2000: Christian Duncan, Panagiota Fatourou, Anil Kumar, and in 2001: Bela Csaba, Michiel Hagedoorn, Juha Kärkkäinen, Spyros Kontogiannis, Jae-Ha Lee.

Guido Schäfer, Naveen Srividasan, Rene Beyer, Thomas Warken, Christian Lennerz, Tobias Polzin and Rahul Ray are new PhD students.

The following researchers left the group: Jop Sibeyn to Umeå University (associate professorship), Susanne Albers to Dortmund University (associate professorship) and then to Freiburg University for a full professorship, Petra Mutzel to Technical University of Vienna for a full professorship (the graph drawing group – two PhD students and several master students – went with her); Petra received the ALCATEL-prize in 2000, Hans-Peter Lenhof to University of Saarland (Saarbrücken) for a full professorship in computational biology (the computational biology subgroup – three PhD-students – will go with him), Rudolf Fleischer to University of Waterloo and then to Hong-Kong University of Science and Technology, Stefan Schirra to industry (Think & Solve, consulting), and postdocs Roberto Solis Oba to University of Western Ontario (Canada), C.R. Subramanian to Indian Institute of Science, Bangalore (India), Panagiota Fatourou to University of Ioannina (Greece), Anil Kumar to Los Alamos National Laboratory (USA), Jae-Ha Lee to KAIST (South Korea), Christian Duncan to University of Miami (USA), Steve Seiden to Louisiana State University (USA), Lorant Porkolab to Imperial College of Science, London (UK), Zeev Nutov to Open University of Israel, Ramat Aviv (Israel).

Twelve PhD students finished in the past two years; we indicate their new affiliation in brackets: Hannah Bast (MPI), Knut Reinert (Celera Genomics), Volker Priebe (Allianz Leben), Mark Ziegelmann (Siemens Research), Stefan Funke (MPI), Piotr Krysta (MPI), Oliver Kohlbacher (Celera

Genomics), Ernst Althaus (International Computer Science Institute), Jordan Gergov (Cadence), Peter Müller (PricewaterhouseCoopers, consulting), Andreas Crauser (Algorithmic Solutions), and Thomas Schilz (Haufe Verlag).

Cooperations

We list some long-term cooperations (short-term ones for single research papers are not listed).

There is cooperation within the institute and with colleagues in the computer science (CS) and computational linguistics (CL) departments. Kurt Mehlhorn cooperates with Sherif Ghali from AG4, and Stefan Funke and Edgar Ramos' research on surface reconstruction has benefited from discussions with Leif Kobbelt and Mario Botsch from AG4. Work within the Effective Computational Geometry project has already triggered several discussions between the AG1 computational geometry and the AG4 mesh processing groups.

Elmar Schömer cooperates with Nicola Geismann from Prof. Raimund Seidel's group in CS; Kurt Mehlhorn and Sven Thiel cooperate with Denys Duchier and Joachim Niehren from Prof. Gert Smolka's group in CS, and with Alexander Koller from the CHORUS group in CL. Berthold Vöcking cooperates with Anja Feldmann's group in CS.

We are part of the EU-Projects ALCOM-FT (partners: Aarhus, Utrecht, Rome, Cologne, Paderborn, Patras, Barcelona, INRIA Roquencourt) and ECG (partners: Zürich, Groningen, INRIA Sophia-Antipolis, Berlin, Tel Aviv).

CGAL: We cooperate with Utrecht, INRIA Sophia-Antipolis, Berlin, Zürich, and Tel Aviv.

LEDA: We cooperate with Stefan Näher's group in Trier.

Constraint Programming: we cooperate with the computer linguists and constraint programmers in Saarbrücken and the constraint programmers at SICS.

SCIL: We cooperate with Michael Jünger and Mathias Elf in Cologne, Alexander Bockmayr in Nancy, and Thomas Kasper at SAP. Alexander and Thomas are former members of AG2.

Computational Biology: Hans-Peter Lenhof has built a strong research program in computational biology with strong links to our work in combinatorial optimization (mainly Integer Linear Programming). Hans-Peter Lenhof and his group have moved to the CS department. Of course, we are going to continue the cooperation and we look forward to the cooperation with Thomas Lengauer's group.

Virtual Reality and Simulation: Elmar Schömer and his group cooperates with the virtuality lab of Daimler-Chrysler.

Program Verification: we cooperate with David Basin in Freiburg, a former member of AG2.

Industrial Take-Up

Algorithmic Solutions (AS) is a spin-off of AG1. AS is marketing LEDA (Library of Efficient Algorithms), CGAL (computational geometry algorithms library), BALL (Biological Algorithms Library), AGD (Automatic Graph Drawing) under license agreements with the Max-Planck-Gesellschaft. In the case of LEDA, it is also responsible for the maintenance. Through its industrial contacts and through user feedback, AS brings interesting research problems into AG1. For example, Mark Zieglmann's PhD project on constrained shortest paths was suggested by a contact at DASA and Susan Hert's recent project on polygons with curved boundaries was suggested by a contact at Bubel Software.

LEDA continues to be a research and commercial success. It is now used at more than 1500 research groups worldwide and more than 300 companies licensed it. A spectacular use is at Celera Genomics. A significant fraction of the software used to sequence the human genome is based on

LEDA. CGAL is also widely used in academia and has found some commercial use, AGD has some commercial users, and BALL has a small but devoted group of academic users.

Algorithms developed by Elmar Schömer and his group have made their way into the virtual reality software at Daimler-Chrysler.

Threats

Kurt Mehlhorn feels that the group has a healthy balance of long-term and short-term research projects and of theoretical and applied work. The group has international standing through its publications, software, and people. Nevertheless, the quality of a research group with only temporary members is always precarious.

Kurt Mehlhorn was recently elected vice chairman of the Chemisch-Physikalisch-Technische (CPT) Sektion of the Max-Planck-Society. According to our rules, he will become chair from 2003 to 2006, and then be vice chair for another year. During the tenure as chair, he expects to spend about 40% of his time on CPT-affairs. Accordingly, he will have 60% for his group. He has already taken a number of measures in order to partially compensate: a new group organization with more responsibilities for the subgroup-coordinators and reduced involvement in other activities (he resigned from the editorial board of six journals, from the ESA steering committee, ...).

It is not unlikely that several (in the worst case, all) subgroup coordinators will leave the group in the next two years to take up permanent positions elsewhere. It is very important to find replacements of similar quality. Our pool of postdocs and recent PhD-graduates will be an important source; it contains a lot of talent.

2 The Programming Logics Group

The research unit "Programming Logics" applies methods of mathematical logic to a variety of problems in computer science. Computation is deduction, a principle that is taken literally in the area of Logic Programming. Formal specifications of software and hardware are formulae in logical systems. Program development, analysis, and verification is based on proving theorems about specifications and programs. Solving problems in specific application domains such as mathematical optimization or program analysis can be seen as deduction with respect to specific constraint structures such as the integers, reals or Herbrand terms.

Computation often means to simulate some model of the real world. While in logic programming and in program synthesis and verification one applies, to a large extent, the classical logics known from mathematics, simulation of the real world requires logics that allow one to treat incomplete and changing knowledge and to reason about beliefs, wishes, knowledge, and the like, of their agents. In that context non-classical logics have to be designed, investigated and applied.

Our work is both theoretical and practical in nature. A large fraction of it is essentially concerned with searching for new and better methods for finding proofs with the support of a computer. As the practical value of results in this area can often not be judged from the theory alone, we are engaged in various implementation projects in which we try to obtain experimental evidence of the practical potential of our results. Moreover, program analysis has become our main domain for applications. There we are trying to specifically develop and combine various deductive methods (model checking, constraint solving, constraint programming, theorem proving) to make analyses more powerful and more precise.

During the period covered by this report, a number of senior people (post-docs and research associates) have left the group to accept a permanent position at other institutions. Unfortunately we have not yet been able to fill all the vacant positions with candidates of comparable qualification.

Automated Theorem Proving for Predicate Logic (Coordinator: Harald Ganzinger)

Work in this area has remained to be both theoretical and practical. At present we are maintaining and further developing four automated theorem provers. Each of these provers is targeted towards specific applications. Hans de Nivelle's Bliksem system is being redesigned so that it can produce explicit proof objects that can easily be checked by a simple proof checker. This feature is a necessity when trying to integrate an automated (first-order) prover into an interactive (higher-order) proof checker. The Saturate system is being maintained by Harald Ganzinger and used for experiments in the area of automated complexity analysis. Together with Thomas Hillenbrand, his prover Waldmeister, the most powerful system for unit equations world-wide, also joined the group. It is presently being extended by new features that are theoretically well-understood, but difficult to implement efficiently. Last, not least, the SPASS system, despite the fact that its creator Christoph Weidenbach has joined industry, is being maintained and extended by more powerful redundancy elimination techniques. Of all our first-order provers, SPASS is the most dependable system for proof search, behaving decently in a number of different application domains.

On the theoretical side, interests have somewhat shifted. During the last two years work in deductive calculi, decidable fragments, and modal logics has not been as dominant as it used to be in previous years. It seems as if most of the natural problems in these subareas have been solved by now. New problems arise from combining methods and from applying them to particular domains. We have made some progress in combining different deductive methods. Uwe Waldmann has combined the Fourier-Motzkin (high-school) method of solving linear inequations into general equational reasoning for uninterpreted symbols. The resulting calculus is refutationally complete

and is expected to perform well on various first-order theories of ordered abelian groups. We have just started looking into the problem of combining black-box decision procedures for decidable theories with general equational reasoning where one specific problem is to extend known methods such as theory resolution to the equational case. Hans de Nivelle has tackled one of the problems to be solved for an integration of automated first-order provers into interactive proof checkers, namely that of generating detailed, yet small, proof objects that the proof checker is able to verify.

We are looking into ways of applying first-order reasoning methods to automated program analysis. Harald Ganzinger has identified deductive models that allow certain non-monotone reasoning processes in which deductive inference and elimination of redundancy are combined to be "executed" with guaranteed complexity bounds. Many type inference and program analysis algorithms can be specified in terms of such inference systems and the execution model yields complexity bounds often competitive with the best algorithms known.

Abstraction is a general method for (efficiently) computing approximations of functions that are computationally intractable. Abstraction should also help sometimes to prove theorems that otherwise would not be found due to the size of the search space. We have been looking at what one might call the relational abstraction of first-order formulae. Hereby one replaces functions by relations—their graphs—and abstraction consists in approximating these relations so that functions may become partial. In cases where representation or embedding theorems exist, the approximation is lossless. Viorica Sofronie-Stokkermans found a resolution-based procedure for deciding the positive theory of bounded distributive lattices exploiting Priestley representation. Harald Ganzinger has observed a strong relationship between applications of relational abstraction by Skolem, Evans, and Burris to PTIME uniform word problems and McAllester's analytic concept of local theories.

Logic Methods for Program Analysis (Coordinator: Andreas Podelski)

Work in this area has expanded considerably during the last two years. Program analysis is a difficult problem. Accordingly a large variety of methods, including checking finite models, abstract interpretation, type inference, constraint solving, and verification are being applied. One of the goals behind our work has been to clarify the relationship between the different approaches. Another goal is to find automatic ways of constructing useful abstractions of program execution for which one can check automatically whether certain kinds of properties are valid. When one has an abstraction at hand one would like to know how much information is lost by replacing the concrete semantics by the abstraction, or whether there are classes of properties for which the abstraction is precise. Moreover, there are many program calculi, type systems, and program logics specifically designed for modelling certain computation paradigms or for expressing specific kinds of program properties. For those systems one would like to know how difficult it is to verify or infer a property expressible in the logic against a formal model of the program defined in the calculus. And in order to be able to handle large programs and state spaces, program analysis algorithms should be efficient and, if possible, exploit the modular structure of programs. In practice often heuristics are applied to make program analysis or model checking succeed, and one would like to have a mathematical explanation of why they do so.

We have obtained a variety of results related to the various problems just described. Just to emphasize two of them, Andreas Podelski has shown that interprocedural reachability can be computed/decided in linear time for languages with dynamically generated threads provided communication is abstracted in a specific weak manner. Patrick Maier has obtained a very general result about soundness of assume-guarantee reasoning. Witold Charatonik and Jean-Marc Talbot have proved that model checking for a certain class of ambient logic formulas is PSPACE-complete for replication-free ambient calculus programs.

Constraint Solving (Coordinators: Fritz Eisenbrand, Andreas Podelski)

Under this headline we study both numerical constraints as they arise in mathematical optimization as well as constraints over symbolic structures (finite and infinite trees). Some of these logical formalisms are of general theoretical interest. Others have direct applications to program analysis.

The group on constraint programming for combinatorial optimization has focussed its interest to problems in integer programming and computer algebra. Fritz Eisenbrand has obtained two excellent results about integer programming and lattice basis reduction, respectively, in fixed dimension, exploiting Schönhage's techniques of fast computation of continued fractions. As a consequence, he has obtained faster algorithms for 2-variable integer programming and for computing shortest vectors in d-dimensional lattices. These algorithms are the theoretically fastest known today.

An unexpected result about monadic set constraints was obtained by Witold Charatonik and Andreas Podelski. As monadic terms are words, satisfiability of such constraints was believed to be decidable in PSPACE. However they could show the problem to be EXPTIME-hard. As a consequence, further approximations of the type of a program variable by sets of paths does not simplify type checking.

Logic and Uncertainty (Coordinators: Manfred Jaeger, Emil Weydert)

In the logic and uncertainty subgroup we have continued to pursue a broad range of subjects: from belief revision and qualitative decision theory to probabilistic reasoning and probability constraints.

Emil Weydert has designed a two-stage epistemic model with Spohn-type rankings and a corresponding minimal information revision strategy that eliminates some of the problems encountered with revision strategies based on cross-entropy minimization. Richard Booth has identified a connection between belief revision and default inference by showing, for a particular popular method of default entailment, how default entailment reveals itself as a form of belief revision.

One of the hopes behind Manfred Jaeger's relational Bayesian networks was that, due to the schematic nature of this first-order formalism, certain classes of inference problems can be solved more efficiently than with currently used inference methods. That might still be true in practice. In theory, however, unfortunately it is not the case, as Manfred Jaeger was able to show.

We have started to look into the combination of logic and probability as it arises in probabilistic verification, thus connecting research in this area to our research about deductive program analysis. Fairness conditions for nondeterministic systems can be interpreted as tests of randomness. Manfred Jaeger has proposed a concept of computable fairness and was able to show that for properties definable by deterministic Büchi automata this approximation of randomness is lossless.

3 The Computational Biology and Applied Algorithmics Group

The general research vision is to help advance areas inside or – as it happens – mostly outside computer science with computer science techniques. Methodically, Tom Lengauer's research has been concentrating on discrete modeling, algorithms, and optimization. Application areas have included (1) electronic circuit design – especially circuit layout – in the eighties and early nineties, (2) design and manufacturing in the textile, leather, and car industry, as well as (3) molecular biology, and medicinal chemistry since the early nineties and (4) organic and inorganic chemistry since about the mid nineties. While these areas seem disparate, at first glance, there actually is some underlying connection: In all cases the research concentrates on problems involving structural analyses in 2- and 3-space. However, since the quest is to make substantial contributions to the application, the actual research work increasingly transcends the mere aspect of analyzing spatial arrangements.

Currently about two thirds of the resources in Tom Lengauer's group are devoted to Computational Biology (3), the rest being equally divided between the fields (2) and (4). We now discuss these areas in a little more detail.

Computational Biology

Computational Biology is one of today's grandest challenges for applied computer scientists. The field is characterized by a strong demand for basic research in understanding and modeling biological structures and processes and, at the same time, high pressure to provide short- term practicable solutions for interpreting ever new kinds of very voluminous and sometimes quite noisy biological data. More so than in many other fields of applied computer science, the biological and the computer science part are inseparably intertwined. Modeling is the main challenge. The development of algorithms is intimately connected with the task of finding the right questions to ask.

Our vision in computational biology is to help provide bioinformatics tools that support biological innovation on the basis of the novel high-throughput experimental data in both molecular biology and chemistry. Currently the focus is on applications in pharmaceutics and medicine – providing new kinds of diagnoses and (drug) therapies. But the methods ought to be applicable also to agricultural and environmental and other biotechnological questions. We will approach these questions as much as we can establish and maintain cooperations with the respective biologists.

Application fields, in which we see strong demands for new solutions and where – in continuity with the past and present work of my group – we hope to be able to provide further helpful contributions include

- the extraction of biological knowledge from genomic data (e.g. completely sequenced genomes) such as by protein structure prediction and protein family analysis,
- the interpretation of mRNA expression data, e.g., based on mircoarrays, especially in the context of human diseases such as in the search for target proteins for drug design. Today microarray data still largely expose different cell states of a species for the eludication of the disease process and the development of drug therapy.
- Soon the interpretation of SNP data expressing genetic variations between individuals will strongly gain importance. Furthermore, proteomics data enabling the analysis of post- translational modifications of proteins will provide a more direct picture of the processes in a cell that involve proteins.

- As we increase our understanding of metabolic and regulatory networks, elucidating the correlations between molecular processes and phenotype will become an increasingly important aspect of computational biology. To this end, clinical data will enter the analysis in addition to the molecular data, to which the analysis is still largely confined today.
- In the area of searching for new drugs, the computer-based analysis of high-throughput screening data and the selection of sublibraries of potential ligands created with combinatorial chemistry methods are two challenges. However, the major challenge in this field is still to find suitable scoring functions.
- As time proceeds, the areas of computational biology and computational (pharmaceutical) chemistry will move closer and merge, eventually. Today, the areas are still largely sepa- rate. As the areas merge, questions of pharmacogenetics and aspects of side effects of drugs will become more accessible to computer analysis.

Besides biological and chemical insight, we need to advance computer science methods to help these applications to progress. On the methodical side, I expect to (continue to) provide contributions to

- advancing discrete optimization and statistical learning techniques for bioinformatics software. We have developed various combinatorial techniques, e.g., to optimize protein structures. We have employed decision trees, pattern matching, and support vector machines for the purpose of training and calibration. An important topic in this area is to analyze the statistical significance of predictions made by bioinformatics methods. Contributions like these can help a lot not only to make predictions more accurate but also to make bioinformatics tools easier to use.
- methods and systems for modeling and simulating interaction networks (regulatory and metabolic). We have used Petri nets as a basis for models of pathways and networks that can be queried and simulated. Extracting biological knowledge from computer- accessible literature is an important aspect of creating the databases. This area cries out for CS participation and has strong affinity with similar questions before the background of internet-based data mining.

Lastly, we see an especially strong demand in computational biology to get the newly developed methods out to the users. Therefore, we will continue to consider the positive feedback from a sizeable user community an important measure of the success of my work. In order to bring out the software to the users, we will continue to work on the establishment and maintenance of strong alliances for technology transfer.

Computational Chemistry

In the past, in Tom Lengauer's group has developed methods for the analysis of molecular structures. One project is to predict the crystal structures based on the knowledge of the structure of the molecule making up the crystal. This problem is akin to the molecular docking problem. Another project asked for the generation of models of amorphous clusters on the basis of the continuous random network model.

In this area, we are currently in a transition state. The degree to which we will continue work in this area depends on the degree to which I can build up cooperations with chemistry groups and

maintain synergy between the problems to be investigated here and the accumulated expertise on the modeling of molecular structures on the pharmaceutical side.

Possible future research problems include applying computational methods from biomolecular docking to questions in supramolecular chemistry.

Combinatorial optimization in engineering applications

This work is most directly derived from Tom Lengauer's activities in the 80s in the area of circuit layout. The expertise accumulated in this field has led to quite a steady inflow of requests from industrial sources. With the combinatorial groundwork provided back then, I have conducted quite applied projects, in this area. Projects are selected on the basis of their practical relevance and their methodical challenge. The research challenge is to employ robust combinatorial techniques for solving new and intricate optimization problems. E.g., the project on textile nesting has involved more complex shapes than we had considered before. A subsequent project on arrangement problems in the car industry has involved the additional challenge of dealing with three dimensions (as opposed to just two dimensions before). A recent project on cutting automobile body parts involves both grouping and two-dimensional arrangement. The typical project requires one scientist for a duration of about two years. As in the other fields above, cooperations with groups from the application area – here often from companies – are a prerequisite for successful work. These projects are almost always funded by external, often industrial, sources. Whether such projects can be continued in an MPG environment remains to be seen. We hope that this is possible.

4 The Computer Graphics Group

The computer graphics group has been newly established two years ago and is meanwhile up and running. Although the core of the group has moved jointly from Erlangen to Saarbrücken, this transition and the subsequent build-up of the group have not been trivial. The main reason for this is the fact that altogether 9 (!) researchers from the former Erlangen/ new Saarbrücken group have received offers for faculty positions within the last two years and have since moved on to establish their own groups [5].

Although we have recently been joined by K. Myszkowski (senior scientist (C3)), and J. Haber (who moved from postdoc to research associate), as a result of this, the current layer of senior researchers in the group is still somewhat thin, and several PhD students and postdocs had to take on additional responsibilities. As this report shows, there is obviously a lot of talent among this group of people, and we hope that several of them will stay with the group after finishing their PhD/postdoc period. Of course, the fact that there is now a well-established computer graphics group at Saarland University (Prof. P. Slusallek) also significantly contributes to the stability of our operation.

Vision

During the last decades computer graphics has established itself as a core technique within computer science and information technology. Computer systems are more and more used to realize and simulate parts of the real or an imaginary world. Such simulations require to model/acquire, to process and to render complex objects and scenes. The goal of computer graphics is to turn abstract information into visual images and to allow the user to interact with complex objects and data in a natural and intuitive way.

In contrast to the classic approach to computer graphics which takes as input a scene model – consisting of a set of light sources, a set of objects (specified by their shape and material properties), and a camera – and uses simulation to compute an image, we like to take the more integrated view of *3D Image Analysis and Synthesis* for our research, and consider the whole pipeline from data acquisition over processing to rendering in our work. In our opinion, this point of view is necessary in order to exploit the capabilities and perspectives of modern hardware, both on the input (sensors, scanners, digital photography, digital video) and output (graphics hardware) side. According to this point of view, one of the key research challenges then is the development of good models and modeling tools to efficiently handle the huge amount of data during the acquisition process and to facilitate further processing and rendering.

In order to make progress along the lines above our work is both theoretical and practical with a focus on first-class basic research on the development of novel algorithms, the integration of new algorithms into a system, and the evaluation of the system in practical applications. We also try to provide a stimulating environment for junior researchers that allows them to develop and build their own research programs and groups.

Research Topics and Structure of the Group

As mentioned above we consider the whole pipeline from data acquisition over processing to rendering in our work. Within this framework our choice of research areas is long-term. We reconsider

[5] T. Ertl (C4, Stuttgart), P. Slusallek (C4, Saarbrücken), G. Greiner (C4, Erlangen), L. Kobbelt (C4, Aachen), R. Westermann (C3, Aachen), W. Heidrich (Ass. Prof., Univ. British Columbia, Canada), C. Soler (Tenured Researcher, INRIA, France), S. Ghali (Ass. Prof., Univ. Alberta, Canda), A. Kolb (C3, FH Wedel)

them, as senior researchers leave and as new opportunities arise. Hiring decisions on all levels (PhD students, postdocs, research associates) are made on quality and fit into our research program. Our research is currently organized into the following eight research areas, each having its own small group of coordinators:

- Mesh Processing (C. Rössl, H.-P. Seidel)
- Free-Form Surfaces, Subdivision Surfaces, and Shape Analysis (J. Haber, H.-P. Seidel)
- Facial Modeling and Animation (J. Haber)
- Model Acquisition with Realistic Reflection Properties (M. Goesele, H. Lensch)
- Image-Based Rendering and Motion from Video (H. Schirmacher)
- Realistic Hardware-Supported Shading and Lighting (J. Kautz)
- Global Illumination (K. Myszkowski, A. Scheel)
- Perception Issues in Rendering and Animation (K. Myszkowski)

The coordinators coordinate the work in their areas and together with Hans-Peter Seidel form the AG4 steering committee. The steering committee meets on a weekly basis (Tuesday, 3 pm) and discusses all group related issues. In particular, it addresses topics such as recruiting, guests and seminars, teaching, project acquisition, mid-term and long-term strategic planning.

The whole group meets twice a week for the

- AG4 lab meeting (Thursday, 1 pm), where organizational issues are discussed and information is distributed by the members of the steering committee, and the
- AG4 graphics colloquium (Wednesday, 1pm), where people from within AG4 and the computer graphics group at Saarland University as well as visitors present their ongoing work to the group and to other interested people.

In addition we also have weekly *rendering group meetings* (technical issues in the areas of material acquisition, image-based rendering, and hardware rendering) and *geometric modeling meetings* (technical issues in the areas of mesh processing and free form surfaces).

Apart from these weekly meetings, there are several meetings and discussion groups that take place frequently, but not on a totally regular basis, such as paper discussion groups that discuss papers of special interest, especially immediately preceding major conference events; global illumination meetings where issues related with this research area are discussed, usually in cooperation with people from the graphics group at Saarland University; internship and practical course meetings where all people involved in internships or FoPras meet and discuss; and last but not least meetings dedicated to single projects, such as 3D scanning and material acquisition, photo studio, 3D visualization room, conference organization, and several others.

Some Achievements

We have been pursuing first-class research on a broad range of topics, and members of the group have actively published in the top conferences and journals (6 books, 112 articles, 17 book chapters) (see Section V.17 for details). We have participated in 62 program committees and have given 92 invited talks and tutorial presentations at major national and international events (see Section V.13

for details). Leif Kobbelt has received the Heinz-Maier-Leibnitz Award 2000. We have organized the workshop on Vision, Modeling and Visualization (VMV'00) here in Saarbrücken, and we will also organize both the ACM Symposium on Solid Modeling'02 and Eurographics'02 in Saarbrücken next year. In the following we briefly highlight some of our achievements in each of the eight research areas:

Mesh Processing We developed new algorithms for remeshing, for surface fairing, and for surface reconstruction. We extended the Marching Cubes Algorithm to feature sensitive surface extraction from volume data and successfully applied the extended algorithm to an NC milling simulation.

Free Form Surfaces, Subdivision Surfaces, and Shape Analysis We developed a new method for the smooth approximation and rendering of large scattered data sets, and studied interpolation operators for bivariate spline spaces. In subdivision we developed the new $\sqrt{(}3)$-Subdivision scheme and studied efficient evaluation and progressive transmission of subdivision surfaces. Finally, we studied the stability of the medial axis transformation using the hyperbolic Hausdorff distance.

Facial Modeling and Animation We have developed a system for photo-realistic facial modeling and animation which includes several tools that facilitate necessary tasks such as mesh processing, texture registration, and assembling of facial components. The resulting head model reflects the anatomical structure of the human head including skull, skin, and muscles. Semi-automatic generation of high-quality models from scan data for physics-based animation becomes possible with little effort.

Model Acquisition with Realistic Reflection Properties In order to conduct this research we have built a special photo studio for measurement purposes. We have developed a system for the acquisition of high-quality 3D models of real world objects representing the object's geometry as well as its material properties. The generated model allows for rendering under arbitrary viewing and lighting conditions and realistically reproduces the appearance of the original object.

Image-Based Rendering and Motion from Video Our research in image based rendering has focused on light fields. We have developed new algorithms for the efficient acquisition and rendering of light fields and use them to efficiently compute reflections and reflections. Research in motion from video focuses on the vision based capture of highly articulated body motion.

Realistic Hardware-Supported Shading and Lighting In our research we develop new algorithms for high quality shading and lighting using computer graphics hardware. In particular, we have developed algorithms for shadows, bump mapping, alternative material models, mirror reflections and glossy reflections off curved surfaces

Global Illumination The main goal of our research is to analyze problems and provide solutions towards making global illumination affordable in practical applications. In particular, we have improved final gathering in FEM radiosity computations, proposed a hybrid approach for bringing together interactivity and ray tracing, and studied the validation of global illumination solutions. We also investigated the computation of analytically correct shadows.

Perception Issues in Rendering and Animation In our research, we consider perceptual issues to improve the efficiency of global illumination computation, rendering, and image display. In particular, we developed perceptually guided solutions for high-quality rendering of animation sequences, an algorithm for perceptually guided corrective splatting, and a solution for tone mapping of interactive walkthroughs. We also developed a framework for evaluating and comparing the quality of various lossy image compression techniques.

Cooperations

There is local cooperation within the institute and with colleagues in the Saarland University CS department. Sherif Ghali has worked with Kurt Mehlhorn, and our work on surface reconstruction has benefited from discussions with Stefan Funke and Edgar Ramos. Hans-Peter Seidel has co-supervised several master theses with AG1 in this area. Work within the AG1 effective computational geometry project has triggered several discussions between the AG1 computational geometry and the AG4 mesh processing groups. The computer graphics group also participates in the Virtual Biolab project - a common proposal of MPII, Saarland University, and the German Research Center for Artificial Intelligence (DFKI), under the scientific leadership of the Bioinformatics group of AG1. Hans-Peter Seidel also participates in the Graduiertenkolleg (Graduate Research Center) Quality Guarantees for Computer Systems.

There have traditionally been very close ties with the computer graphics group at Saarland University, chaired by Philipp Slusallek, and the two groups collaborate in several joint research projects: Philipp Bekaert works with Philipp Slusallek and Ingo Wald towards interactive ray tracing. Hartmut Schirmacher and Ming Li collaborate on image-based viewing of dynamic scenes using an array of video cameras (the so-called "Lumi-Shelf"). Stefan Brabec, Christian Theobalt, Andreas Pomi and Philipp Slusallek work on fast approximate shadow generation, and the two groups collaborate in a joint project with Motorola on network-integrated multimedia middleware. The two groups also run a joint weekly graphics colloquium every Wednesday.

There is ongoing and close interaction with the graphics groups in Stuttgart (Prof. T. Ertl), Erlangen (Prof. Dr. G. Greiner), Aachen (Prof. L. Kobbelt, Prof. R. Westermann), Tübingen (Prof. W. Strasser) and Bonn (Prof. R. Klein, Prof. A. Weber). The joint workshop on Vision, Modeling and Visualization (joint with H. Niemann (Erlangen) and B. Girod (Stanford)) has meanwhile become a regular annual event. We also have close ties to the University of British Columbia (W. Heidrich) and to the University of Waterloo (M. McCool, S. Mann).

We are part of the EU-projects MINGLE (partners: Oslo, Tel Aviv, Munich, Haifa, Grenoble, Cambridge, Genova) and SIMULGEN (partners: Girona, Grenoble, Light Work). A further project (partners: Barcelona, Pisa) is currently in preparation. We also participate in a GIF-project (partners: Tel Aviv, Haifa, Tübingen, Bonn, Aachen) and in the DFG-Schwerpunktprogramm Verteilte Verarbeitung und Vermittlung digitaler Dokumente (partners: Braunschweig, Bonn, Tübingen, Stuttgart, Aachen, Leipzig, among others). Details on our projects are given in Section V.16.

We also have several cooperations with industry. In addition to the projects listed in Section V.16 we have established a link the the DaimlerChrysler Virtual Reality Competence Center in Ulm, and we have been working together with Mercedes-Benz on global illumination solutions for car interior design.

Group Development

The composition of the group has changed considerably over the past two years.

Karol Myszkowski who had previously been an associate professor at the University of Aizu, Japan, joined us as a senior scientist (C3) in 2000. Jörg Haber, who had originally joined us as a postdoc in 1999, has since moved from postdoc to research associate. Michael Goesele, Jan Kautz, Hendrik Lensch, Christian Rössl, Annette Scheel, and Hartmut Schirmacher all moved to research associate, although they are technically still PhD students. With the exception of Michael Goesele, who joined us from the University of Ulm, all had moved with the group from Erlangen to Saarbrücken.

The following postdocs joined us in 1999: Sherif Ghali, Cyril Soler, Frank Zeilfelder, 2000: Sung Woo Choi, Ulrich Schwanecke, Hitoshi Yamauchi, and in 2001: Philippe Bekaert, Vlastimil Havran, Ioannis Ivrissimtris. Our current PhD students are Stefan Brabec, Katja Daubert, Won-Ki Jeong, Kolja Kaehler, Ming Li, Takehiro Tawara, Christian Theobalt and Jens Vorsatz.

The following researchers left the group: Leif Kobbelt to RWTH Aachen for a full professorship (two PhD students, Stephan Bischoff and Mario Botsch, went with him), Leif received the Heinz-Maier-Leibnitz Award 2000; Wolfgang Heidrich to the University of British Columbia, Vancouver, Canada (Assistant Professor); Cyril Soler to INRIA Rhone-Alpes (Tenured Researcher); Sherif Ghali to the University of Alberta, Edmonton, Canada (Assistant Professor); Marc Stamminger received a Marie Curie Fellowship by the EU, and has moved to INRIA Antia-Sophipolis with a leave of absence; Robert Schneider to industry (Siemens UBMed); Ulrich Schwanecke to industry (DaimlerChrysler Virtual Reality Competence Center); Frank Zeilfelder back to the University of Mannheim.

Part III

The Algorithms and Complexity Group

1 Personnel

Director:
Prof. Dr. Kurt Mehlhorn

Research Associates:
Dr. Hannah Bast (since January 00)
Dr. Susan Hert (since January 99)
Dr. Edgar A. Ramos (since October 98)
Priv. Doz. Dr. Peter Sanders (since October 97)
Dr. Susanne Schmitt (since October 00)
Priv. Doz. Dr. Elmar Schömer (since January 00)
Michael Seel (since October 97)
Dr. Berthold Vöcking (since June 00)
Priv. Doz. Dr. Susanne Albers (until September 99; now Universität Freiburg)
Dr. Christoph Burnikel (until December 00; now ENCOM GmbH)
Priv. Doz. Dr. Rudolf Fleischer (until September 99; now Hong-Kong University of Science and Technology)
Priv. Doz. Dr. Hans-Peter Lenhof (until October 00; now Universität des Saarlandes)
Priv. Doz. Dr. Petra Mutzel (until September 99; now Technische Universität Wien)
Priv. Doz. Dr. Stefan Schirra (until September 00; now Think & Solve, consulting)
Priv. Doz. Dr. Jop Sibeyn (until December 00; now Umeå University, Sweden)

Postdocs:
Dr. Michiel Hagedoorn (since September 00)
Dr. Juha Kärkkäinen (since September 00)
Dr. Bela Csaba (since October 00)
Dr. Spyros Kontogiannis (since October 00)
Dr. Christian Duncan (until September 00; now University of Miami)
Dr. Panagiota Fatourou (until February 01; now University of Ioannina, Greece)
Dr. Xudong Fu (until September 99)
Dr. Anil Kumar (until August 01; now Los Alamos National Laboratory)
Dr. Jae-Ha Lee (until February 01; now KAIST, South Korea)
Dr. Z. Nutov (until August 99; now Open University of Israel, Ramat Aviv)
Dr. L. Porkolab (until August 99; now Imperial College of Science, London)
Dr. Steve Seiden (until August 99; now Louisiana State University, Baton Rouge)
Dr. Roberto Solis-Oba (until August 99; now University of Western Ontario, London – Canada)
Dr. C. R. Subramanian (until July 00; now Indian Inst. of Science, Bangalore)

Ph.D. students:
Rene Beier
Andreas Kerzmann
Christian Lennerz
Uli Meyer
Tobias Polzin
Rahul Ray
Guido Schäfer
Thomas Schäfer
Marite Sirava

Naveen Sivadasan
Sven Thiel
Thomas Warken
Ernst Althaus (finished April 01; now International Computer Science Institute, Berkeley)
Hannah Bast (finished February 00; now MPI)
Andreas Crauser (finished March 01; now Algorithmic Solutions)
Stefan Funke (finished August 01; now MPI)
Jordan Gergov (finished November 00; now Cadence)
Oliver Kohlbacher (finished January 01; now Celera Genomics)
Piotr Krysta (finished August 01; now MPI)
Peter Müller (finished November 99; now PricewaterhouseCoopers, consulting)
Volker Priebe (finished June 01; now Allianz Lebensversicherungs-AG)
Knut Reinert (finished August 99; now Celera Genomics)
Thomas Schilz (finished May 00; now Haufe Verlag)
Mark Ziegelmann (finished August 01; now Siemens Research)

Secretaries:
Ingrid Finkler
Petra Mayer
Vera Strasburger (until March 00)

2 Visitors

In the time period from June 1999 to June 2001, the following researchers visited our group:

Naveen Garg	21.06.99 - 23.07.99	IIT Delhi
	25.05.01 - 25.07.01	
Jeff Vitter	21.06.99 - 24.06.99	Duke University, Raleigh
James Abello	07.07.99 - 16.07.99	AT & T. Labs Research USA
Paolo Ferragina	12.07.99 - 18.07.99	University of Pisa
	01.08.00 - 06.08.00	
Stefano Leonardo	15.07.99 - 23.07.99	University of Rom
Guy Kortsarz	23.07.99 - 05.08.99	Open University of Israel, Tel-Aviv
Jonathan Schewchuk	01.08.99 - 08.08.99	University of California, Berkeley
Martin Fürer	01.08.99 - 07.08.99	Penn. State University
Larry Larmore	01.08.99 - 06.08.99	Rice University, Houston
Wolfgang Bein	02.08.99 - 06.08.99	University of New Mexico
John Kececioglou	04.08.99 - 13.08.99	Univ. of Georgia at Athens
Antonio Mesa	05.08.99 - 29.08.99	Universidad de la Habana
Martin Dietzfelbinger	06.09.99 - 30.09.99	University of Dortmund
Jörg Sack	06.01.00 - 29.02.00	Carleton University, Ottawa
Paul Spirakis	13.02.00 - 16.02.00	University of Patras
Stephen Kobourov	24.04.00 - 12.05.00	John Hopkins University, Baltimore
Michal Soch	03.05.00 - 27.05.00	Czech. Techn. University, Prague
Palash Sarkar	20.05.00 - 27.05.00	University of Waterloo
Marc de Kreveld	13.08.00 - 19.08.00	Utrecht University
Bogdan Chlebus	14.08.00 - 14.09.00	University of Warszawa
Marco Pellegrini	26.08.00 - 10.09.00	IMC-CNR Pisa
Roberto Solis-Oba	27.08.00 - 10.09.00	University of Western Ontario
Alberto Caprara	14.08.00 - 14.09.00	DEIS Univ. of Bologna
Pankaj K. Agarwal	13.08.00 - 23.08.00	Duke University, Raleigh
Anna Gambini	05.09.00 - 19.09.00	University Warszawa
Tamal K. Dey	18.09.00 - 24.09.00	Ohio State University, Columbus
Kavitha Telikepalli	28.09.00 - 15.12.00	School of Comp. Science, Humbai, India
Gemma Casa Garriga	09.10.00 - 15.10.00	Techn. University Barcelona
Hung-Sheung Poon	01.11.00 - 31.01.01	Hong-Kong Univ. of Science and Technology
	29.06.01 - 26.08.01	
Tetsuo Asano	11.12.00 - 17.12.00	JAIST, Japan
Antoine Vigneron	01.11.00 - 31.01.01	Hong-Kong Univ. of Science and Technology
Siavash Vahdati	01.03.01 - 30.04.01	University of Mannheim
Tobias Polzin	01.03.01 - 30.04.01	University of Mannheim
Chee Yap	15.03.01 - 16.03.01	New York University
Thomas Radzik	25.03.01 - 25.04.01	Kings College London
Sandeep Sen	15.05.01 - 31.06.01	IIT Delhi
Rohit Khandekar	25.05.01 - 25.07.01	IIT Delhi
Sajith Gopalan	01.05.01 - 31.07.01	ITT Delhi
Alan Frieze	31.05.01 - 01.06.01	Carnegie Mellon University, Pittsburgh
Artur Czumaj	31.05.01 - 01.06.01	New Jersey Institute of Technology

3 Group Organization

The group meets two to four times a week at 1.30 pm:

- On Monday and Wednesday (1.30 - 2.15) we have our noon seminar (sometimes also on Friday). It lasts about 45 minutes and is reserved for presentations of new results and ongoing research. We also ask our guests to give presentations in the noon seminar.
- On Tuesday and Thursday (1.30 - 3.00) we run the "Selected Topics in Algorithms" course. This course is reserved for two to four week intensive treatments of subjects of current interest.

On a monthly basis we run a "'Group-meeting" in which all the members of the group participate to discuss various topics regarding the group and to be informed about several other activities.

We are currently pursuing six research areas, each having its own coordinator. The subgroups and their coordinators are:

- Foundations, coordinator: Berthold Vöcking
- Computational Geometry, coordinator: Edgar Ramos
- Advanced Models of Computation, coordinator: Peter Sanders
- Assembly and Simulation, coordinator: Elmar Schömer
- Combinatorial Optimization, coordinator: Kurt Mehlhorn
- Software Libraries, coordinator Kurt Mehlhorn. Starting October 1st, 2001 Lutz Kettner will coordinate the group.

Most group members contribute to more than one area. The coordinators coordinate research and teaching in their area and together with Kurt Mehlhorn form the steering committee of the group which committee meets weekly to discuss and organize the work of the group.

4 Foundations – Coordinator: Berthold Vöcking

The goal that we address within in this research area is to understand the fundamental principles that govern the efficiency with which computational problems can be solved and, subject to these principles, to devise the most efficient algorithms possible. The topics of research that we cover are diverse and include, e.g., average case analysis for shortest-paths algorithms, scheduling under imperfect knowledge, approximation algorithms, and game theory applied to the Internet.

4.1 Analysis of Shortest-Paths Algorithms

Investigators: Ulrich Meyer, Volker Priebe, Guido Schäfer, Naveen Sivadasan

For arbitrary *undirected* networks with non-negative arc costs, it is known that single-source shortest-paths problems can be solved in linear time in the worst case [7]. Is this also true for *directed* networks? In [5], Meyer proves that on average, a similar result indeed holds. Problem instances are arbitrary directed network on n vertices and m arcs whose arc costs are randomly chosen according to the uniform distribution on $[0, 1]$, independently of each other. Meyer presents both label-setting and label-correcting algorithms that solve the single-source shortest-paths problem on such instances in time $O(n + m)$ with high probability. The algorithms do not use exact priority queues, but simple hierarchical bucket structures with adaptive splitting instead. The label-setting algorithm aims to split the current bucket until a single vertex remains in it, whereas the label-correcting algorithm adapts the width of the current bucket to the maximum degree of the vertices contained in it. Subsequently, Goldberg [2] suggested an alternative label-setting algorithm that achieves linear average-case time, too. Parallelizations of the label-correcting method are discussed in Section 6.) In [5, 4], Meyer proposes a general method of constructing networks with random arc costs on which many traditional algorithms are forced to run in non-linear time on average.

Meyer's result on the complexity of the single-source shortest-paths problem is of interest, for example, if networks have $m = o(n \log n)$ arcs. In his Ph. D. thesis [6], Priebe studies shortest-paths problems on randomly generated denser networks, $m = \Omega(n \log n)$. He proves that the all-pairs shortest-paths problem can be solved in time $O(n^2 \log n)$ with high probability with respect to various probability distributions on the set of inputs. For example, such a result holds if n-vertex networks are generated by two random experiments. Each arc in the directed graph is present with probability $p \geq C \cdot (\log n)/n$, independently of the presence of all other arcs; costs of arcs that were determined as present in this first experiment are then drawn at random in a second experiment, independently of each other, according to the uniform distribution on $[0, 1]$. The results also include the first theoretical analysis of the average behavior of shortest-paths algorithms with respect to the *vertex-potential model*, a family of probability distributions on complete networks with arbitrary real arc costs but without negative cycles; see also [1]. All algorithms mentioned exploit that with high probability, the single-source shortest-paths problem can be solved correctly by considering only a rather sparse subset of the arc set. Furthermore, in both models, a bound of $O(\log n)$ on the number of arcs on any shortest paths holds with high probability.

In [3], the authors present an algorithm that solves the all-pairs shortest-paths problem on a directed network with n vertices and m arcs in time $O(nm + n^2 \log n)$ in the worst case, where the arcs are assigned real, possibly negative costs. The algorithm is new in the following respect. It computes the distance $\mu(v, w)$ between each pair v, w of vertices even in the presence of negative cycles, where $\mu(v, w)$ is defined as the infimum of the costs of all directed paths from v to w. The algorithm thus outputs $\mu(v, w) = -\infty$ if there is a path from v to w that contains a negative cycle.

The algorithm is easy to implement, since it uses only standard algorithmic building blocks such as the computation of strongly connected components, Dijkstra's algorithm, or the Bellman–Ford algorithm.

References

- [1] C. Cooper, A. Frieze, K. Mehlhorn, and V. Priebe. Average-case complexity of shortest-paths problems in the vertex-potential model. *Random Structures Algorithms*, 16:33–46, 2000.

[2] A. V. Goldberg. A simple shortest path algorithm with linear average time. In F. Meyer auf der Heide, editor, *Algorithms — ESA '01*, number 2161 in Lecture Notes in Comput. Sci., pages 230–241. Springer-Verlag, Berlin, 2001.

- [3] K. Mehlhorn, V. Priebe, G. Schäfer, and N. Sivadasan. All-pairs shortest-paths computation in the presence of negative cycles. *Inform. Process. Lett.*, 2001. to appear.
- [4] U. Meyer. On the average-case performance of parallel and sequential single-source shortest-paths algorithms. submitted, 2001.
- [5] U. Meyer. Single-source shortest-paths on arbitrary directed graphs in linear average-case time. In *Proc. 12th Ann. Symp. on Discrete Algorithms*, pages 797–806. ACM–SIAM, 2001. Extended version: Technical Report MPI-I-2001-1-002, MPI Informatik, 2001.
- [6] V. Priebe. *Average-case complexity of shortest-paths problems.* PhD thesis, Universität des Saarlandes, Saarbrücken, 2001.

[7] M. Thorup. Undirected single-source shortest paths with positive integer weights in linear time. *Journal of the ACM*, 46:362–394, 1999.

4.2 Scheduling Parallel Tasks at Twilight

Investigators: Hannah Bast

Many real-world scheduling problems fit into the following abstract scenario: n tasks are to be scheduled on p processors such that the completion time of the last task, the so-called makespan of the schedule, becomes as small as possible. Each task can be assigned to any one of the processors, and the order in which the tasks are processed is arbitrary. Assignments are at runtime, in *chunks* of several tasks at a time, and for each such assignment an overhead is incurred, at least part of which is independent of the number of tasks in (= the size of) the chunk. Only vague information on the processing times of the tasks is available in advance, and the more tasks a processor is assigned at a time, the harder it is to predict, when this processor will have completed them. As is not hard to see the optimal tradeoff between the assignment overhead and the imbalance of processor loads is achieved by minimizing the *wasted time* of the schedule, which is the sum of all the overheads plus the idle times of those processors waiting for the last processor to finish.

In the past, a large number of heuristics for determining chunk sizes that minimize the wasted time have been developed and tested (e.g., [5], [3], [4]). We were the first to attempt a rigorous theoretical study, some results of which were already described in the last progress report (1997 - 1999). The fundamental results of our study are now published, unfortunately (and foolishly) in a single, somewhat overwhelming paper [1]. In short, [1] proposes a novel deterministic approach to modeling the vague-information issue, provides matching upper and lower bounds in the respective general model, derives upper bounds for randomly distributed processing times, and proves the corresponding almost matching lower bounds.

The work described in [2] puts emphasis on two more practical aspects, which played only a secondary role in [1]: the simplicity of the scheduling algorithm, and the constant factors in

the performance bounds. The motivation is that while most of the existing heuristics are quite elaborated, only the two simplest ones are actually in use: the *fixed-size* (FIX) heuristics, where the chunks size is kept fixed throughout the scheduling process, and what we call *geometric* (GEO) heuristics, where chunk sizes decrease at an exponential rate. The bottom line of [2] is that GEO can combine simplicity, efficiency, and robustness in the best conceivable manner, while, in comparison, the efficiency of FIX is always off by an order of magnitude. Another, maybe important, aspect of [2] is that it provides a fast and easy access to the results in [1], which, as a tribute to its generality, abstractness, and rigorousity, is certainly not a light paper.

References

- [1] H. Bast. On scheduling parallel tasks at twilight. *Theory of Computing Systems*, 33(5):489–563, 2000.
- [2] H. Bast. Scheduling at twilight the easy way. submitted, 2001.

[3] L. E. Flynn, S. F. Hummel, and E. Schonberg. Factoring: A method for scheduling parallel loops. *Comm. Assoc. Comput. Mach.*, 35:8:90–101, 1992.

[4] T. Hagerup. Allocating independent tasks to parallel processors: an experimental study. *J. Par. Dist. Comp.*, 47:185–197, 1997.

[5] C. P. Kruskal and A. Weiss. Allocating independent subtasks on parallel processors. *IEEE Trans. Software Eng.*, 11:1001–1016, 1985.

4.3 Approximation Algorithms for 2-Connectivity Problems

Investigators: Béla Csaba, Piotr Krysta, Anil Kumar

We study [8] some versions of the problem of finding the minimum size 2-connected subgraphs. The minimum 2-edge connected spanning subgraph problem (2-EC) is to find a 2-edge connected spanning subgraph with minimum number of edges. A graph is 2-edge connected (2-EC) if for any pair of vertices there are at least two edge disjoint paths between the vertices. The minimum 2-vertex connected spanning subgraph problem (2-VC) is defined analogously, but the paths must be vertex disjoint. These problems are NP-hard, even on cubic planar graphs, via a reduction from the Traveling Salesman Problem (TSP). They are also known to be MAX SNP-hard.

For the 2-EC problem Khuller and Vishkin [6] gave a $\frac{3}{2}$-approximation, which was improved by Cheriyan *et al.* [2] to $\frac{17}{12}$, and recently to $\frac{4}{3}$ by Vempala and Vetta [10]. For the 2-VC problem Khuller and Vishkin [6] gave an algorithm with approximation guarantee of $\frac{5}{3}$. The next known result for the 2-VC problems was a $\frac{3}{2}$-approximation algorithm, due to Garg *et al.* [4]. This was finally improved to $\frac{4}{3}$ by Vempala and Vetta [10].

We give better approximation algorithms for the 2-EC problem. We show a $(\frac{4}{3}-\epsilon)$-approximation algorithm for 2-EC on general graphs, where $\epsilon = \frac{1}{1344}$, improving on the results of Vempala and Vetta [10]. This result is a step towards establishing the best possible approximation ratio for this problem in the light of its MAX SNP-hardness.

We achieve better guarantees for special classes of graphs, on which 2-EC is still NP-hard. For planar graphs, we show for 2-EC a $\frac{93}{70}$-approximation, but in quasi-polynomial time. For 2-EC on graphs with maximum degree 3, we obtain a $\frac{21}{16}$-approximation using a simple *local search* heuristic. This implies an integrality gap of at most $\frac{21}{16}$ for the standard LP on such graphs.

Our results have connections to the well known metric $\frac{4}{3}$ TSP conjecture, due to Goemans [5], and a related conjecture for the 2-EC problem [2]. These conjectures say that the *integrality gaps* of the corresponding linear programming (LP) relaxations to these problems are at most $\frac{4}{3}$. Our $(\frac{4}{3}-\epsilon)$-approximation for 2-EC on general graphs proves that these conjectures are irrelevant to

the approximability of the 2-EC problem. The $\frac{21}{16}$-approximation shows a stronger version of the second conjecture (with $\frac{4}{3}$ replaced by $\frac{21}{16}$) for graphs with maximum degree 3.

We also consider in [7] generalizations, where we require 1 or 2 edge or vertex disjoint paths between any pair of the vertices. We call these problems $\{1,2\}$-EC and $\{1,2\}$-VC, respectively. We give $\frac{3}{2}$-approximations for both problems. This improves on straightforward 2-approximation, which can easily be deduced from the results of Nagamochi and Ibaraki [9]. Our algorithms are generalizations of the algorithms of Garg *et al.* [4] and of Khuller and Vishkin [6]. We also analyze the performance of the classical local optimization heuristics for these two problems.

In paper [3], we study the approximability of *dense* instances of the following problems: 2-EC and 2-VC, metric TSP with distances 1 and 2, maximum path packing, and the longest path (cycle) problems. The approximability of dense instances of these problems was left open in the paper of Arora *et al.* [1]. We resolve the approximability of all these problems by proving tight upper (via approximation algorithms) and lower bounds (via non-approximability).

References

[1] S. Arora, D.R. Karger, and M. Karpinski. Polynomial time approximation schemes for dense instances of np-hard problems. In *Proc. ACM Symposium on Theory of Computing (STOC)*, pages 284–293, 1995.

[2] J. Cheriyan, A. Sebő, and Z. Szigeti. An improved approximation algorithm for minimum size 2-edge connected spanning subgraphs. In *Proc. 6th International Integer Programming and Combinatorial Optimization Conference (IPCO)*, volume 1412 of *Lecture Notes in Computer Science*, pages 126–136. Springer Verlag, Berlin, 1998.

- [3] B. Csaba, M. Karpinski, and P. Krysta. Approximability of dense instances of minimum 2-connectivity, tsp and path problems. *Manuscript*, 2001. Submitted for publication.

[4] N. Garg, V. Santosh, and A. Singla. Improved approximation algorithms for biconnected subgraphs via better lower bounding techniques. In *Proc. 4th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 103–111, 1993.

[5] M.X. Goemans. Worst-case comparison of valid inequalities for the tsp. *Math. Programming*, 69:335–349, 1995.

[6] S. Khuller and U. Vishkin. Biconnectivity approximations and graph carvings. *Journal of the ACM*, 41:214–235, 1994.

- [7] P. Krysta. Approximating minimum size $\{1,2\}$-connected networks. *Technical Report MPI-I-2001-1-001, Max-Planck-Institut fuer Informatik, Saarbruecken, Germany*, 2001. Submitted to a journal.

- [8] P. Krysta and V.S. Anil Kumar. Approximation algorithms for minimum size 2-connectivity problems. In *Proc. 18th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 2010 of *Lecture Notes in Computer Science*, pages 431–442. Springer Verlag, Berlin, 2001.

[9] H. Nagamochi and T. Ibaraki. A linear-time algorithm for finding a sparse k-connected spanning subgraph of a k-connected graph. *Algorithmica*, 7:583–596, 1992.

[10] S. Vempala and A. Vetta. Factor 4/3 approximations for minimum 2-connected subgraphs. In *Proc. 3rd International Workshop APPROX*, volume 1913 of *Lecture Notes in Computer Science*, pages 262–273. Springer Verlag, Berlin, 2000.

4.4 Constrained Graph Exploration

Investigators: Christian Duncan, Anil Kumar

Researchers often model robot navigation as a problem of exploring an unknown graph, $G = (V, E)$. Starting at a fixed node s, the robot explores G by means of edge traversals. In many common

variants, the robot is able to recognize a previously visited node or edge. Therefore, it is able to construct a map of the explored portion of the graph. We say that a robot has successfully explored G once it has visited all vertices V and edges E in G. The performance is measured by the number of edge traversals made by the robot and ignores any "internal" computation done by the robot. Although much work has been done in the case of directed graphs [5, 1], we concentrate on exploring undirected graphs, that is all edges are undirected. Without any additional constraints, the classical search algorithms, breadth-first search (BFS) and depth-first search (DFS) are optimal within a constant factor, $\Theta(|E| + |V|)$ time.

One reasonable constraint, however, is that the robot cannot skip nodes. That is, it must only proceed from one node to another via edge traversals. This result makes breadth-first searches fail. However, depth-first searches are still successful. Panaite and Pelc [7] further minimize the constant dependency on E and prove an exploration time of $|E| + \Theta(|V|)$.

In [4], Betke, Rivest, and Singh introduce the piecemeal exploration problem. This variation forces the robot to periodically return to the start node, e.g. for refueling purposes. Under this constraint, the depth-first search algorithm also fails. In their paper, they present optimal $\Theta(|E| + |V|)$ results for very specific graphs on a grid. In [2, 3], Awerbuch *et al* extend the results to general graphs achieving a near-optimal bound of $O(|E| + |V| \log^2 |V|)$.

In [6], we prove an optimal linear time algorithm for the piecemeal exploration problem. In addition, we introduce the tethered robot problem. In this variation, the robot must not return for refueling but is attached to the start node by a fixed length tether. We show that this variation is as difficult as the refueling problem and prove a linear time algorithm for it as well.

We also extend the result to the *treasure hunting* problem. In this problem, the robot must find a treasure t in a (possibly infinite) graph G. The problem is similar to the shortest path problem (see, for example, [8]) where the robot must reach some given node t as quickly as possible. As in the exploration problem, the performance is measured by the number of edge traversals made and is compared to the length of the shortest path from s to t, $\delta(s, t)$. Unfortunately, for general graphs, we cannot bound the performance. As a simple example, let T be a complete binary tree and let the treasure lie on one of the leaves. An adversary can then place t on the last leaf node explored. The running time is then $|T|$ which is $2^{\delta(s,t)}$. Our algorithm finds t in time $O(|G'|)$, where G' is the subgraph induced by all vertices within distance $(1 + \alpha)\delta(s, t)$ from s for any constant $\alpha > 0$. We leave it open if a solution exists for α equal to 0.

References

[1] S. Albers and M. R. Henzinger. Exploring unknown environments. In *Proceedings of the Twenty-Ninth Annual Symposium on Theory of Computing (STOC)*, pages 416–425, 1997.

[2] B. Awerbuch, M. Betke, R. Rivest, and M. Singh. Piecemeal graph learning by a mobile robot. *Information and Computation*, 152(2):155–172, 1999.

[3] B. Awerbuch and S. Kobourov. Polylogarithmic-overhead piecemeal graph exploration. In *Proceedings of the 11th Annual Conference on Computational Learning Theory (COLT)*, pages 280–286, 1998.

[4] M. Betke, R. Rivest, and M. Singh. Piecemeal learning of an unknown environment. *Machine Learning*, 28:2-3:231–254, 1995.

[5] Deng and Papadimitriou. Exploring an unknown graph. *Journal of Graph Theory*, 32:265–297, 1999.

- [6] C. Duncan, S. Kobourov, and A. Kumar. Optimal constrained graph exploration. In *Proceedings of the 12th ACM-SIAM Symposium on Discrete Algorithms*, pages 807–814, 2001.

[7] P. Panaite and A. Pelc. Exploring unknown undirected graphs. *Journal of Algorithms*, 33:281–295, 1998.

[8] C. Papadimitriou and M. Yanakakis. Shortest paths without a map. *Theoretical Computer Science*, 84:127–150, 1991.

4.5 Algorithms, Games, and the Internet

Investigators: Piotr Krysta, Berthold Vöcking

The Internet is unique among all computer systems in that it is built, operated, and used by a multitude of diverse economic interests, in varying relationships of collaboration and competition with each other. This suggests that the mathematical tools and insights most appropriate for understanding the Internet may come from a fusion of algorithmic ideas with concepts and techniques from economics and game theory.

As a first step towards understanding the behavior of selfish players in a network environment, Koutsoupias and Papadimitriou introduced the *coordination ratio*, a game theoretic measure that aims to reflect the price of selfish routing in a network. In [1], we can show first tight bounds on this ratio. In particular, we show that the worst-case coordination ratio on m parallel links (of possibly different speeds) is

$$\Theta\left(\left(\frac{\log m}{\log\log m}\right)^2\right) \ .$$

Previously Mavronicolas and Spirakis [3] were only able to analyze very restrictive variants of selfish behavior. Our result is asymptotically tight and it entirely resolves the central open question posed in [2].

Our current work on this topic focuses on the effects of more realistic, convex latency functions as well as heavy-tailed traffic distributions on the coordination ratio. These aspects are of special interest as they are the major characteristics that were observed in numerous practical studies. In addition to these game theoretic aspects, we investigate approximation algorithms for the same realistic traffic models. Our goal is to compare the restrictions due to limited computational power with the restrictions due to the lack of coordination.

References

- [1] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria, submitted.

[2] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proc. of the 16th STACS*, pages 404–413, 1999.

[3] M. Mavronicolas and P. Spirakis. The price of selfish routing. In *Proc. of the 33rd ACM-STOC*, pages 510–519, 2001.

4.6 Edge-Coloring of Bipartite Graphs

Investigators: Stefan Schirra

Edge-coloring bipartite graphs is a classical and well studied problem in graph algorithms. The edge-coloring problem is to assign colors to the edges of a graph such that no vertex has two or more incident edges with the same color. Any bipartite (multi-) graph $G = (V, E)$ can be colored using D colors where D is the maximum degree of G. It is well known how to find a minimal edge-coloring for a bipartite graph in $O(|E| \log D)$ time, if D is a power of 2. Although already twenty years ago, algorithms have been published with running time $O(|E| \log D + T)$ for some disturbing

term $T = \Omega(|E| \log D)$, a general method to find a minimal edge-coloring in $O(|E| \log D)$ time was not known until recently. In [1] we present an algorithm that finds a minimal edge-coloring for bipartite (multi-) graphs within this time bound no matter what D looks like.

References

- [1] R. Cole, K. Ost, and S. Schirra. Edge-coloring bipartite multigraphs in $o(elogd)$ time. *Combinatorica*, 21:5–12, 2001.

4.7 Solving Recurrences: A Trick to Remember

Investigators: Hannah Bast

Solving a recurrence of the form $R(z) = R(f(z)) + g(z)$ yields to a sum of the form $g(z) + g(f(z)) + g(f(f((z))) + \cdots$. The "trick" referred to in the title is to approximate this sum by an integral of the function $x \mapsto g(x)/(x - f(x))$, as indicated in the following figure.

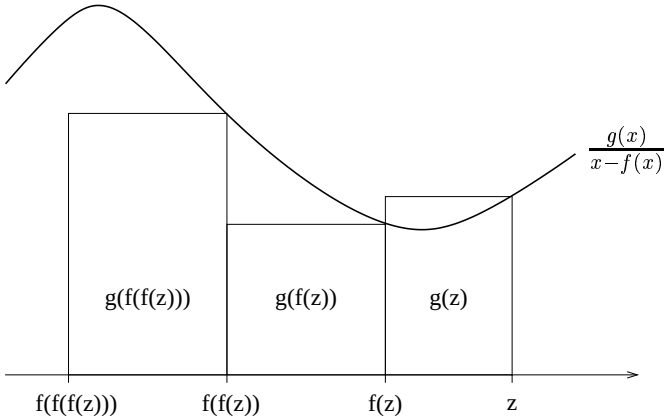

In [1], conditions on f and g are derived, so that

$$R(z) = R(z_0) + \Theta\left(\int_{z_0}^{z} \frac{g(x)}{x - f(x)}\, dx\right).$$

It is also shown how, in case these conditions are not satisfied from the outset, they can be established by a simple transformation of the recurrence.

References

- [1] H. Bast. Solving recurrences: A trick you should remember. Manuscript, 2001.

5 Combinatorial Optimization – Coordinator: Kurt Mehlhorn

We worked on specific problems (resource constrained shortest paths, curve and surface reconstruction, Steiner tree problem, protein docking and sequence alignment) and on generic methods (symbolic constraints for integer linear programming (SCIL) and efficient narrowing algorithms for constraint programming).

5.1 Resource Constrained Shortest Paths

Investigators: Kurt Mehlhorn, Mark Ziegelmann

The resource constrained shortest path problem (RCSP) is the problem of finding the shortest path between two nodes in a graph whenever the traversal of an edge consumes certain resources and the resources consumed along the path must lie within given limits. The problem may be illustrated as that of a traveler with a budget (a vector of resources) who has to reach a given destination as quickly as possible within the constraints imposed by its budget.
RCSP has numerous applications in operations research and mission planning.

Regarding the complexity, it can be shown that RCSP is NP-hard; however, there is an FPTAS for the problem [4].

Early work dealing with the exact solution of RCSP was done by Joksch [5] who gave a dynamic programming formulation. Other possibilities are to solve the corresponding ILP formulation or to simply adopt a path ranking procedure [2]. However, for large problem instances these approaches are often not efficient.

Handler and Zang [3] presented and algorithm for the single resource case based upon Lagrangean relaxation and subsequent path ranking to find the original optimum.

Beasley and Christofides [1] solved the multiple resources problem with a subgradient procedure and a subsequent tree search procedure also using problem reductions.

In [6], we consider RCSP for k resources and develop a method that solves the Lagrangean relaxation exactly. We associate each path with a point in $k+1$-dimensional space and try to find the facet of the lower hull that intersects the feasible region bottom most. This matches the result in [3] for the single resource case (however, we are the first to prove a tight bound for the running time) and improves on the approximate upper and lower bounds of [1] for multiple resources and thus provides a basis to solve multiple resource RCSP more efficiently. The duality gap can again be closed by subsequent path ranking to obtain a 2-step method for constrained shortest paths.

The 2-step approach can also be applied to a broader class of problems. We developed a software package CNOP that implements a generic 2-step method for constrained network optimization [7].

The PhD-thesis of Mark Ziegelmann [8] contains a summary of all these results.

References

[1] J. Beasley and N. Christofides. An Algorithm for the Resource Constrained Shortest Path Problem. *Networks*, 19:379–394, 1989.

[2] D. Eppstein. Finding the k shortest paths. *SIAM Journal on Computing*, 28(2):652–673, 1999.

[3] G. Handler and I. Zang. A Dual Algorithm for the Constrained Shortest Path Problem. *Networks*, 10:293–310, 1980.

[4] R. Hassin. Approximation Schemes for the Restricted Shortest Path Problem. *Math. Oper. Res.*, 17(1):36–42, 1992.

[5] H. Joksch. The Shortest Route Problem with Constraints. *Journal of Mathematical Analysis and Application*, 14:191–197, 1966.

- [6] K. Mehlhorn and M. Ziegelmann. Resource constrained shortest paths. In *8th Annual European Symposium on Algorithms (ESA), LNCS 1879*, pages 326–337, 2000.
- [7] K. Mehlhorn and M. Ziegelmann. CNOP - A package for constrained network optimization. In *3rd Workshop on Algorithm Engineering and Experiments (ALENEX), LNCS 2153*, pages 15–28, 2001.
- [8] M. Ziegelmann. *Constrained Shortest Paths and Related Problems.* PhD thesis, Universität des Saarlandes, 2001.

5.2 SCIL

Investigators: Ernst Althaus, Kurt Mehlhorn

Many combinatorial optimization problems are naturally formulated through constraints. Consider the Traveling Salesman Problem (TSP). It asks for the minimum cost Hamiltonian cycle in an undirected graph $G = (V, E)$ with edge costs $(c_e)_{e \in E}$. We want to select a subset S of the edges of G such that "S is a Hamiltonian cycle" and $\sum_{e \in S} c_e$ is minimized.

Our vision is that the paragraph above (written in some suitable language) suffices to obtain an efficient algorithm for the TSP. Efficiency is meant in a double sense. We want short development time (= efficient use of human resources) and runtime efficiency (= efficient use of computer resources). SCIL is our first step towards realizing this vision.

Integer linear programming (ILP) and, more specifically, the branch-and-cut-and-price paradigm (BCP-paradigm) for solving ILPs is one of the most effective approaches to the exact solution of hard combinatorial optimization problems. It has been successfully used to solve a wide range of problems, e.g. traveling salesman tours, max-cut problems, linear ordering problems, and general integer linear programs.

Although, the BCP-method works for every integer linear program, e.g., through the use of Gomory-cuts, it is most effective when problem specific cuts are used. The implementation of BCP-algorithms using problem specific cuts is non-trivial. Many components are needed: the maintenance of the branch-and-cut tree, the manipulation of linear programs, an LP-solver, separation and cut generation routines, pricing routines, and the communication between the components. Systems are available for some of these components: there are LP-solvers, ABACUS provides a framework for BCP-algorithms, and LEDA provides many of the data structures and algorithms needed for cut generation and pricing.

SCIL integrates the above and goes beyond. It offers high-level problem specification, a flexible and efficient BCP-solver, and a library of separation and pricing routines. We discuss the points in turn.

- SCIL offers a powerful language for specifying ILPs. More specifically, it offers high-level *constraints* such as the tour-constraint, the matching-constraint, the cutting-stock-constraint, the facility-location-constraint, ... A constraint defines a complex subset of the assignments for a set of variables. High-level constraints are also called symbolic constraints. We will call them simply constraints.

  Symbolic constraints have already been used very successfully in constraint programming. Bockmayr and Kasper [3, 5] have argued their usefulness for integer linear programming. SCIL offers them.

- SCIL offers a convenient method (*variable and inequality maps*) for connecting ILP-concepts, namely variables and inequalities, with problem oriented concepts such as graphs, sequences, and two-dimensional geometric objects. All LEDA-objects are available in SCIL. As a consequence, ILPs can be expressed at the application level and hence are more readable.
- SCIL generates BCP-solvers expressed in the ABACUS-framework and hence inherits the ABACUS-features: a powerful default BCP-solver and the flexibility to add separation and pricing routines and to formulate branching rules.
- SCIL has the concept of variable and inequality generators for generating variables or inequalities of a specific kind, e.g., degree equalities or set partition variables.
- SCIL is extendible. It provides the infra-structure for defining new constraints and for adding feasibility test, separation routines, and pricing routines.

SCIL has already been used for curve reconstruction [2], surface reconstruction from planar contours [4], docking with flexible side chains [1], and multiple sequence alignment.

References

- [1] E. Althaus, O. Kohlbacher, H.-P. Lenhof, and P. Müller. A combinatorial approach to protein docking with flexible side-chains. In Ron Shamir, Satoru Miyano, Sorin Istrail, Pavel Pevzner, and Michael Waterman, editors, *Proceedings of the 4th Annual International Conference on Computational Molecular Biology (RECOMB-00)*, pages 15–24. ACM Press, 2000. Full Version submitted to Journal of Computational Biology.
- [2] E. Althaus and K. Mehlhorn. TSP-based curve reconstruction in polynomial time. *SIAM Journal of Computing*, 31(1), 2001.

[3] A. Bockmayr and T. Kasper. Branch and infer: A unifying framework for integer and nite domain constraint programming. *INFORMS Journal on Computing*, 10(3):287–300, 1998.

- [4] C. Fink. Oberflächenrekonstruktion von planaren konturen. Master's thesis, Universität des Saarlandes, 2001.

[5] T. Kasper. *A Unifying Logical Framework for Integer Linear Programming and Finite Domain Constraint Programming.* PhD thesis, Fachbereich Informatik, Universität des Saarlandes, 1998.

5.3 Constraint Programming

Investigators: Kurt Mehlhorn, Sven Thiel

A constraint program $\mathcal{P}$ consists of a set $\mathcal{V}$ of variables and a set $\mathcal{C}$ of constraints. With every variable X we associate a finite domain $D(X)$ (usually a subset of the integers), which is the set of all values that may be assigned to X. An n-ary constraint C is defined by sequence $V(C) = [X_1, \ldots, X_n]$ which contains the variables that are constrained by C and an n-ary relation $R(C)$. An assignment a which is defined on a superset of $V(C)$ satisfies C iff $(a(X_1), \ldots, a(X_n)) \in R(C)$. For example, the relation for the constraint `alldifferent`$(X_1, \ldots, X_n)$ is the set of all n-tuples where all components are pairwise different. A solution s of $\mathcal{P}$ is an assignment defined on $\mathcal{V}$ that satisfies the domain constraint $s(X) \in D(X)\ \forall X \in \mathcal{V}$ and all constraints $C \in \mathcal{C}$.

When a constraint program is solved, the domains of the variables are reduced by propagators and by searching until a solution is found. Our research is mainly devoted to designing propagation algorithms for certain constraints. A propagator for a constraint C observes the domains of the variables in $V(C)$. It has to perform two tasks:

- It may detect *failure*, i.e. there is no assignment satisfying C and the domain constraint, or *entailment*, which means that all assignments satisfying the domain constraint also satisfy C.

- It may remove a value v from the domain of a variable X if there is no assignment a with $a(X) = v$ that satisfies C and the domain constraint. This process is called pruning or narrowing.

Combinatorial Constraints

We examined the `alldifferent` and the `sortedness` constraints. The latter constraint is a $2n$-ary constraint, its relation consists of all tuples $(d_1, \ldots, d_n, e_1, \ldots, e_n)$ such that the sequence $[e_1, \ldots, e_n]$ is obtained by sorting the values $d_1, \ldots, d_n$ in non-decreasing order. The `alldifferent` constraint is a classical constraint that appears in many applications and the `sortedness` constraint can be used in some scheduling applications [9].

In our studies we concentrated on the case where all domains are integer intervals $[a; b]$[6]. We were able to give propagation algorithms [8] which achieve bound-consistency, i.e. they narrow every domain to the smallest possible interval. Our algorithms, which are based on matching theory, improve upon previous results [10, 3] in some relevant cases.

Dominance Constraints

Dominance constraints are logical tree descriptions that are widely used in computational linguistics: they have been used for grammar formalisms [4], natural language semantics [5] and discourse analysis [6]. The problem is to assemble a tree out of given tree fragments obeying dominance requirements like "the fragment T_1 has to be contained in the subtree rooted at the leaf l of fragment T_2". It was shown in [7] that the solvability problem of dominance constraints is *NP*-complete. In our work [1] we describe a practically relevant subclass called *normal* dominance constraints for which solvability can be decided in polynomial time. We also give an efficient algorithm for enumerating all solutions.

Geometric Constraints

We considered the problem of placing the origins of a set of convex polygonal objects in the plane such that no two of them overlap. The shape of every object is fixed, it is described by a *shape polygon* P_{shape} with integer vertices. We require that one vertex v has the coordinates $(0, 0)$. The position of v is not fixed, it is given by two integer finite domain variables O_x and O_y, which are called *domain variables*. Further, we are allowed to specify a *domain polygon* P_{dom} (also with integer vertices) to constrain the placement of an object. The set of possible placements is the following $\{(x, y) + P_{shape} : x \in D(O_x) \wedge y \in D(O_y) \wedge (x, y) \in P_{dom}\}$.

Our propagation algorithm [2] is based on the notion of *forbidden regions*. A forbidden region for an object o_1 according to an object o_2 is a region R such that placing the origin of o_1 into R will cause overlapping of o_1 and o_2 no matter where the origin of o_2 is placed. We gave an efficient algorithm for computing forbidden regions and a pruning algorithm for the origin variables which is based on the sweepline paradigm.

[6] Here a, b are integers and $[a; b]$ is set of all integers i with $a \leq i \leq b$.

References

- [1] E. Althaus, D. Duchier, A. Koller, K. Mehlhorn, J. Niehren, and S. Thiel. An efficient algorithm for the configuration problem of dominance graphs. In *Proceedings of the 12th ACM-SIAM Symposium on Discrete Algoritms (SODA 01)*, pages 815–824, Washington DC, USA, January 2001. ACM-SIAM.
- [2] N. Beldiceanu, Q. Guo, and S. Thiel. Non-overlapping Constraints between Convex Polytopes. To appear in Proceedings of the 7th International Conference on Principles and Practice of Constraint Programming (CP-01).

[3] N. Bleuzen-Guernalec and A. Colmerauer. Optimal narrowing of a block of sortings in optimal time. *Constraints*, 5(1–2):85–118, 2000.

[4] D. Duchier and S. Thater. Parsing with tree descriptions: a constraint-based approach. In *Proceedings of NLULP-99*, Las Cruces, New Mexico, 1999.

[5] M. Egg, J. Niehren, P. Ruhrberg, and F. Xu. Constraints over Lambda-Structures in Semantic Underspecification. In *Proceedings COLING/ACL'98*, Montreal, 1998.

[6] C. Gardent and B. Webber. Describing discourse semantics. In *Proceedings of the 4th TAG+ Workshop*, Philadelphia, 1998.

[7] A. Koller, J. Niehren, and R. Treinen. Dominance constraints: Algorithms and complexity. In M. Moortgat, editor, *Proceedings of the Third Conference on Logical Aspects of Computational Linguistics (Dec. 1998)*, volume 2014 of *Lecture Note in Artificial Intelligence*, Heidelberg, 2001. Springer-Verlag.

- [8] K. Mehlhorn and S. Thiel. Faster algorithms for bound-consistency of the sortedness and the alldifferent constraint. In Rina Dechter, editor, *Proceedings of the 6th International Conference on Principles and Practice of Constraint Programming (CP-00)*, volume 1894 of *Lecture Notes in Computer Science*, pages 306–319, Singapore, September 2000. Springer.

[9] W.J. Older, G.M. Swinkels, and M.H. van Emden. Getting to the real problem: experience with BNR Prolog in OR. In *Proceedings of the 3rd International Conference on the Practical Application of Prolog (PAP-95)*, pages 465 – 478, April 1995.

[10] J.-F. Puget. A fast algorithm for the bound consistency of alldiff constraints. In *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI-98) and of the 10th Conference on Innovative Applications of Artificial Intelligence (IAAI-98)*, pages 359–366, Menlo Park, July 26–30 1998. AAAI Press.

5.4 Curve and Surface Reconstruction via LP and ILP

Investigators: Ernst Althaus, Christian Fink, Kurt Mehlhorn

An instance of the curve or surface reconstruction problem is a finite sample set V of an unknown collection of curves or surfaces. The task is to compute an approximation of the curve or surface.

In the case of curve reconstruction, we are interested in the polygonal reconstruction, i.e. we want to connect two points in V if they are adjacent on the curve. Giesen [5] showed recently that the Traveling Salesman tour of V solves the reconstruction problem for single closed curves under otherwise weak assumptions on the curve and the point set V; γ must be a single closed curve. We extended his result in several directions (see [2, 1])

- we weaken the assumptions on the sample,
- we show that the Traveling Salesman based reconstruction also works for single open curves (with and without specified endpoints) and for collections of closed curves,
- we give alternative proofs,

- we show that in the context of curve reconstruction, the Traveling Salesman tour can be constructed in polynomial time.

Furthermore we report on experiments with a number of recent curve reconstruction algorithms (see [3]).

In the case of surface reconstruction, we are interested in a set of triangles whose corners are sample points. This set of triangles should approximate the original surface. One distinguishes two cases of input points. Either the points are in general position, or they all lie in a few number of parallel slices. In the second case, we often assume that the reconstruction problem for the single slices are already solved and call the problem the surface reconstruction problem from planar contours.

We adapt the idea of minimizing the length in the case of curve reconstruction to minimizing the surface area in the case of surface reconstruction. We formulated this problem as integer linear program for the two cases described above. In the case of general position our we are not able to solve the instances arising in real problems. In the case of surface reconstruction from planar contours our algorithm can solve all instances in reasonable time (see [4]).

Our Computational Geometry subgroup has investigated other approaches to curve and surface reconstruction; see Section 7.3.

References

- [1] E. Althaus. *Curve Reconstruction and the Traveling Salesman Problem.* PhD thesis, Universität des Saarlandes, 2001.
- [2] E. Althaus and K. Mehlhorn. TSP-based curve reconstruction in polynomial time. *SIAM Journal of Computing*, 31(1), 2001.
- [3] E. Althaus, K. Mehlhorn, S. Näher, and S. Schirra. Experiments on curve reconstruction. In *Proc. 2nd Worksh. Algorithm Engineering and Experiments (ALENEX00)*, Lecture Notes in Computer Science. Springer-Verlag, Jan 2000.
- [4] C. Fink. Oberflächenrekonstruktion von planaren konturen. Master's thesis, Universität des Saarlandes, 2001.

[5] J. Giesen. Curve reconstruction, the TSP, and Menger's theorem on length. In *Proc. 15th Annu. ACM Sympos. Comput. Geom.*, pages 207–216, 1999.

5.5 Combinatorial Optimization for Bioinformatics

Investigators: Ernst Althaus, Oliver Kohlbacher, Hans-Peter Lenhof, Peter Müller, Knut Reinert

The introduction of combinatorial optimization methods to the field of computational biology has since then proven to be a useful tool for numerous applications such as RNA secondary structure alignment, computing the fit of three-dimensional structures and a general trace formulation. We investigated on two problems in computational biology.

Protein Docking

Rigid–body docking approaches are not sufficient to predict the structure of a protein complex from the unbound (native) structures of the two proteins. Accounting for side chain flexibility is an important step towards fully flexible protein docking. We describe an approach that allows conformational flexibility for the side chains while keeping the protein backbone rigid. Starting from candidates created by a rigid–docking algorithm, we demangle the side chains of the docking site,

thus creating reasonable approximations of the true complex structure. These structures are ranked with respect to the binding free energy. We present two new techniques for side chain demangling. Both approaches are based on a discrete representation of the side chain conformational space by the use of a rotamer library. This leads to a combinatorial optimization problem. For the solution of this problem we propose a fast heuristic approach and an exact, albeit slower, method that uses branch–&–cut techniques. As a test set we use the unbound structures of three proteases and the corresponding protein inhibitors. For each of the examples, the highest–ranking conformation produced was a good approximation of the true complex structure.

This work is published in [2, 1] and is submitted to the Journal of Computational Biology.

Multiple Sequence Alignment

Aligning DNA or protein sequences is certainly one of the dominant tools in computational molecular biology. The spectrum of methods ranges from extremely fast hashing-based methods over moderately expensive pairwise comparisons based on dynamic programming, to costly, exact multiple alignment formulations, which are either based on the natural extension of the dynamic programming paradigm, or on the application of combinatorial optimization techniques.

We have extended the formulation of the *gapped trace problem* proposed by Reinert [3] such that we can formulate a great variety of multiple sequence alignment problems, among them the weighted sum of pairs problem with arbitrary gap costs. To our knowledge this is the first algorithm that can deal with truly affine gap costs. Indeed our algorithm is independent of the choice of the gap cost function and can handle any function including convex gap costs which were proposed in several publications.

We will submit our work to RECOMB 2002.

References

- [1] E. Althaus, O. Kohlbacher, H.-P. Lenhof, and P. Müller. A branch-&-cut algorithm for the optimal solution of the side chain placement problem. Max-planck-institut für informatik : Forschungsbericht, Max-Planck-Institut für Informatik, 2000.
- [2] E. Althaus, O. Kohlbacher, H.-P. Lenhof, and P. Müller. A combinatorial approach to protein docking with flexible side-chains. In Ron Shamir, Satoru Miyano, Sorin Istrail, Pavel Pevzner, and Michael Waterman, editors, *Proceedings of the 4th Annual International Conference on Computational Molecular Biology (RECOMB-00)*, pages 15–24. ACM Press, 2000. Full Version submitted to Journal of Computational Biology.
- [3] K. Reinert. *A Polyhedral Approach to Sequence Alignment Problems.* PhD thesis, Universität des Saarlandes, 1999.

5.6 Steiner Tree Problems

Investigators: Tobias Polzin

The Steiner problem in networks is the problem of connecting a set of required vertices in a weighted graph at minimum cost. This is a classical NP-hard problem with many important applications in network design in general and VLSI design in particular (for an overview see [3]).

For combinatorial optimization problems like the Steiner problem which can naturally be formulated as integer programs, many approaches are based (explicitly or implicitly) on linear programming. For the Steiner problem there are many different linear relaxations and different ways

of solving (or approximating the solution of) a relaxation, but not much has been known about the (relative) quality of these relaxations and approximations.

In [6] we compared all classical and some new relaxations from a theoretical point of view with respect to their optimal values, building a hierarchy of relaxations. Furthermore, we presented a new relaxation which is the strongest linear relaxation of polynomial size for the Steiner problem in networks.

We also studied several algorithms for computing lower and upper bounds for the Steiner problem in networks using dual-ascent and primal-dual strategies. In [5] we showed that none of the known algorithms can both generate tight lower bounds empirically and guarantee their quality theoretically; and we presented a new algorithm which combines both features.

We achieved several algorithmical improvements for computing upper and lower bounds, and integrated them into a software package for solving the Steiner problem to optimality [7] that clearly outperformed all other published results (e.g. [1, 2, 4]) on a large set of benchmark instances (SteinLib).

References

[1] S. Chopra, E. R. Gorres, and M. R. Rao. Solving the Steiner tree problem on a graph using branch and cut. *ORSA Journal on Computing*, 4:320–335, 1992.

[2] C. W. Duin. *Steiner's Problem in Graphs.* PhD thesis, Amsterdam University, 1993.

[3] F. K. Hwang, D. S. Richards, and P. Winter. *The Steiner Tree Problem*, volume 53 of *Annals of Discrete Mathematics.* North-Holland, Amsterdam, 1992.

[4] T. Koch and A. Martin. Solving Steiner tree problems in graphs to optimality. *Networks*, 32:207–232, 1998.

- [5] T. Polzin and S. Vahdati. Primal-dual approaches to the steiner problem. In Klaus Jansen and Samir Khuller, editors, *Approximation Algorithms for Combinatorial Optimization*, volume 1913 of *Lecture Notes in Computer Science*, pages 214–225, Saarbrücken, Germany, 2000. Springer.
- [6] T. Polzin and S. Vahdati. A comparison of steiner tree relaxations. *Discrete Applied Mathematics*, 112:241–261, August 2001.
- [7] T. Polzin and S. Vahdati. Improved algorithms for the steiner problem in networks. *Discrete Applied Mathematics*, 112:263–300, August 2001.

6 Advanced Models of Computation – Coordinator: Peter Sanders

Basic algorithmic research usually assumes some variant of the von Neumann model of computation with a single processor and uniform memory. Driven by technological developments and new applications, the von Neumann model increasingly deviates from reality. Therefore much work in AG 1 is devoted to meeting these challenges by developing algorithms for more advanced models of computing.

We have decided to combine all the efforts concerning models in one research area in order to avoid artificial borders between parallel, distributed, or external computing, memory hierarchies, internet, or mobile computing. In particular, often one algorithmic approach helps in several models. To underline this effect, the subsequent more detailed descriptions follow guiding ideas and problem areas rather than particular models. Section 6.1 is concerned with communication in a network of computers. Currently, our main focus is on parallel computers and the internet. Another problem that is ubiquitous in parallel models is load balancing and scheduling. Section 6.2 discusses algorithms that schedule tasks on parallel computers and data accesses to parallel disks. Our studies on graph algorithms described in Section 6.3 exemplify how algorithmic ideas can be useful for both parallel and external implementation. Finally, Section 6.4 discusses software engineering issues in external computing.

6.1 Communication

We have produced several new results in the area of fundamental communication primitives. Such primitives have been intensively studied before because they form major building blocks for the efficient use of parallel computers. These results are remarkable since most important questions seemed to be largely closed for a decade. We describe the results going from low level primitives to more complex operations. A simple randomized packet routing algorithm for hypercube networks described in Section 6.1.1 closes a long standing factor two gap between the best known algorithm and the trivial lower bound [17]. A similar factor two improvement for the problem of broadcasting a long message from one processor to all others is described in Section 6.1.2. A natural generalization of broadcasting is the gossiping problem discussed in Section 6.1.3 where every processor has a message it wants to broadcast. Here we have developed powerful heuristics that can be applied to any specific network as well as a new algorithm for butterfly networks that breaks a long standing barrier for very large networks. A more complex yet very important setting is if every processor has different messages (of possibly different length) for each other processor. For one natural setting, Section 6.1.4 outlines an algorithm that beats an apparent lower bound and closes the gap to the real lower bound.

In the future we will work more on higher level problems that are important beyond classical parallel computing for example for internet settings. Section 6.1.5 describes an algorithm for keeping track of objects moving in a network.

6.1.1 Almost Optimal Permutation Routing on Hypercubes

Investigator: Berthold Vöcking

In [17], we investigate permutation routing on hypercube networks in the store-and-forward model. We introduce the first (on-line and off-line) algorithms routing any permutation on the d-dimensional hypercube in $d + o(d)$ steps. The best previously known results were $2d + o(d)$ for oblivious on-line routing (see, e.g., [16]) and $2d - 3$ for off-line routing (see, e.g., [13]) In particular, we present

- a randomized, oblivious on-line algorithm with routing time $d + O(d/\log d)$,
- a matching lower bound of $d + \Omega(d/\log d)$ for (randomized) oblivious on-line routing, and
- a deterministic, off-line algorithm with routing time $d + O(\sqrt{d \log d})$.

Previous algorithms lose a factor of two mainly because packets are first sent to intermediate destinations in order to resolve congestion. As a consequence, the maximum path length becomes $2d - o(d)$. Our algorithms use intermediate destinations as well, but we introduce a simple, elegant trick ensuring that the routing paths are not stretched too much. In fact, we achieve small congestion using paths of length at most d.

The main focus of our work, however, lies on the scheduling aspect. On the one hand, we investigate well-known and practical scheduling policies for on-line routing, namely *Farthest-to-Go* and *Nearest-to-Origin*. On the other hand, we present a new off-line scheduling scheme that is based on frugal colorings of multigraphs. This scheme might also be of interest for other sparse scheduling problems.

6.1.2 One-to-all Broadcasting

Investigators: Peter Sanders and Jop Sibeyn

In [9] we develop a new algorithm for broadcasting a long message from one processor to all other processors in a parallel computer. This "fractional tree" algorithm can be viewed as an interpolation between two well known simple algorithms: The sequential pipeline algorithm is almost unbeatable for very long messages whereas the the pipelined binary tree algorithm is up to a factor two slower for this case but much faster for small to medium message lengths. The new algorithm is good in all cases. For networks like meshes or clusters of shared memory multiprocessors which are not very densely connected, the new algorithm seems to be the best known method for the important case that each processor has only a single (possibly bidirectional) channel into the communication network.

6.1.3 All-to-all Broadcasting

Investigators: Rene Beier, Jop Sibeyn

In *gossiping*, each processor has one piece of information that should be communicated to all other processors. We study this problem for a network model where at any point in time, a processor can only exchange data with one of its neighbors in the network. In [3] and [2] we suggest a new heuristic for computing gossiping schedules for arbitrary networks. The algorithm refines previous approaches based on weighted matchings by introducing a new weighting scheme where not only the number of packets to be delivered but also the distance they need to travel is taken into account. Although this schemes was designed for working on arbitrary, irregularly structured networks, it even outperforms the best theoretical constructions for some very regular network structures (e.g butterflies, shuffle-exchange, DeBruijn and pancake graphs). The heuristic can also handle a model with communication costs proportional to the message lengths where only few theoretical results are known. As a contribution to this field we developed almost optimal constructions for gossiping on meshes and tori.

In [14] we gives new results for gossiping in butterfly and "regular" butterfly networks under the unit-cost model. For a butterfly of order k the best previously known gossiping algorithm require $2.5 \cdot k$ and $3 \cdot k$ communication rounds respectively. By new asymptotic methods we prove new bounds of $2.25 \cdot k + o(k)$ and $2.5 \cdot k + o(k)$ respectively.

6.1.4 Personalized All-to-all

Investigators: Peter Sanders, Roberto Solis-Oba

Many parallel algorithms are subdivided into phases where in each phase the processors first perform some local communication and then synchronize to exchange whatever information needs to be exchanged. For example, the popular BSP model of parallel computation [15] is based on this approach and offers only a single communication primitive known as h-relation or personalized all-to-all communication with nonuniform message sizes. In this primitive, each of P processors has $P-1$ messages with variable (possibly zero) length. In [10, 11] we study this problem on a very simple machine model that grasps important aspects of many practical machine architectures. We assume a fully connected network but allow a processor to participate only in a single data transfer (sending or receiving) at a time. Let h_i denote the total time it would take processor i to send out all its messages and to receive all the data destined for it. Let $h = \max_i h_i$ and call the corresponding communication problem an h-relation. It is easy to show that there are h-relations where it takes time at least $3h/2$ to exchange all messages if all messages are sent directly to their final destination.

Interestingly, this is not a lower bound. We show how the same problem can be solved in time close to $6h/5$ if we are willing to send some data over detours. There are h-relations where no better bound can be achieved. The idea of the algorithm is very simple. Using well known techniques based on edge coloring bipartite graphs that can be executed remarkably fast [5], we first decompose an h-relation into a set of permutations involving messages of identical lengths. Then we show how a permutation can be further decomposed. The main remaining problems are cyclic communications involving an odd number of processors. We pair such odd cycles A and B and use the resources of otherwise idle processors in A to speed up communication in B and vice versa.

Applications of this result beyond parallel computers are file transfers or data migration on wide area networks [4].

In cooperation with NEC Germany, we are currently developing simple practical algorithms for uniform message length on the increasingly important architecture where multiprocessor nodes are clustered to larger systems [12].

6.1.5 Data Tracking

Investigator: Berthold Vöcking

Consider an arbitrary distributed network in which large numbers of objects are continuously being created, replicated, and destroyed. A basic problem arising in such an environment is that of organizing a distributed directory service for locating object copies. In [8], we present a new data tracking scheme for locating nearby copies of objects in arbitrary distributed environments. The new scheme is based on a clever combination of ideas that were previously presented in [1, 6, 7].

Our tracking scheme supports efficient accesses to data objects while keeping the local memory overhead low. In particular, our tracking scheme achieves an expected polylog(n)-approximation in the cost of any access operation, for an arbitrary network. The memory overhead incurred by our scheme is $O(\text{polylog}(n))$ times the maximum number of objects stored at any node, with high probability. We also show that our tracking scheme adapts well to dynamic changes in the network.

References

[1] Y. Bartal. Probabilistic approximation of metric spaces and its algorithmic applications. In *37th IEEE Symposium on Foundations of Computer Science*, pages 184–193, 1996.

- [2] R. Beier. Eine heuristik fr das gossiping-problem. Master's thesis, Universität des Saarlandes, August 2000.
- [3] R. Beier and J. F. Sibeyn. A powerful heuristic for telephone gossiping. In *Proceedings of the 7th International Colloquium on Structural Information and Communication Complexity (SIROCCO-00)*, volume 7 of *Proceedings in Informatics*, pages 17–36, L'Aquila, Italy, August 2000. Carleton Scientific.

[4] E. G. Coffman, M. R. Garey, D. S. Johnson, and A. S. LaPaugh. Scheduling file transfers. *SIAM Journal on Computing*, 14(3):744–780, 1985.

- [5] R. Cole, K. Ost, and S. Schirra. Edge-coloring bipartite multigraphs in $O(E \log D)$ time. *Combinatorica*, 21(1):5–12, 2000.

[6] B. Maggs, F. Meyer auf der Heide, B. Vöcking, and M. Westermann. Exploiting locality for data management in systems of limited bandwidth. In *38th IEEE Symposium on Foundations of Computer Science*, pages 284–295, 1997.

[7] C. G. Plaxton, R. Rajaraman, and A. W. Richa. Accessing nearby copies of replicated objects in a distributed environment. In *9th ACM Symposium on Parallel Algorithms and Architectures*, pages 311–320, 1997.

- [8] R. Rajaraman, A. Richa, B. Vöcking, and G. Vuppuluri. A data tracking scheme for general networks. In *11th ACM Symposium on Parallel Algorithms and Architectures*, pages 247–254, 2001.
- [9] P. Sanders and J. Sibeyn. A bandwidth latency tradeoff for broadcast and reduction. In *6th Euro-Par*, number 1900 in LNCS, pages 918–926, 2000.
- [10] P. Sanders and R. Solis-Oba. How helpers hasten h-relations. In *8th European Symposium on Algorithms (ESA)*, number 1879 in LNCS, pages 392–402. Springer, 2000.
- [11] P. Sanders and R. Solis-Oba. How helpers hasten h-relations. *Journal of Algorithms*, 2001. to appear.
- [12] P. Sanders and J. L. Träff. The factor algorithm for all-to-all communication on clusters of SMP nodes. Manuscript, 2001.

[13] X. Shen, Q. Hu, and W. Liang. Realization of an arbitrary permutation on a hypercube. *Information Processing Letters*, 51:237–243, 1994.

- [14] J. F. Sibeyn. Faster gossiping on butterflies. In *28th International Colloquium on Automata, Languages and Programming*, volume 2076 of *Lecture Notes in Computer Science*, pages 785–796, Crete, Greece, 2001. Springer.

[15] L. Valiant. A bridging model for parallel computation. *Communications of the ACM*, 33(8), 1994.

[16] L. G. Valiant. *Handbook of Theoretical Computer Science , Volume A*, chapter General purpose parallel architectures, pages 943–971. Elsevier and MIT Press, 1990.

- [17] B. Vöcking. Almost optimal permutation routing on hypercubes. In *33rd ACM Symposium on the Theory of Computing*, pages 530–539, 2001.

6.2 Load Balancing

6.2.1 Load Balancing in the Presence of Overheads and Uncertainty

Investigator: Hannah Bast

If, in a dynamic load balancing scenario, the overhead associated with assigning work to a processor is nonnegligible and only vague information about the processing time of a chunk of work is available

in advance, the following tradeoff must be considered: the load balance tends to be better, when the work is assigned in many small chunks, while the overhead tends to be smaller, when the work is assigned in few large chunks. We have undertaken a deep theoretical study of this phenomenon, which is discussed in Section 4.2.

6.2.2 Threads

Investigator: Panagiota Fatourou

Dynamically growing multithreaded computations are nowadays quite common for parallel computers; therefore, multithreaded models of parallel computation have been proposed as a general approach to modeling dynamic, unstructured parallelism. The execution of a multithreaded computation begins at the main thread of the computation and unfolds dynamically as threads spawn other threads. Data or synchronization dependencies may exist between different threads of a multithreaded computation.

We have continued previous work of R. Blumofe and C. Leiserson [3] by presenting the first provably good, randomized, distributed, online, asynchronous, work-stealing scheduling algorithm for (the more general) *strict* multithreaded computations, that is, computations in which all dependencies from a thread go to the thread's ancestors [5]. We have analyzed the complexity of our algorithm and we have proven that it is very efficient in terms of all three performance parameters: *space*, *execution time* and *communication cost*.

Recent work on the same research area has been concentrated on the design of highly space-efficient scheduling algorithms. Blelloch, Gibbons, Matias, and Narlikar (e.g., [2, 9]) have conducted research in this direction. Most of the results for online versions of these algorithms, however, have been limited to computations in which threads can only synchronize with ancestor threads. Such computations do not include e.g., languages with futures or user-specified synchronization constraints. Moreover, almost all these algorithms result in high contention, while simultaneously they ignore the issue of locality.

We have designed [4] an online, work-stealing, space-efficient algorithm for the general class of programs with synchronization variables (such programs are produced by several parallel programming languages, like, e.g., Cool, ID, Sisal, Mul-T, OLDEN and Jade). The algorithm combines the *work-stealing* paradigm with the *depth-first* scheduling technique, resulting in high space efficiency and good time complexity. By automatically increasing the *granularity* of the work scheduled on each processor, our algorithm achieves good locality, low contention and low scheduling overhead, improving upon previous depth-first scheduling algorithms. Moreover, since it is provably efficient for the general class of multithreaded computations with write-once synchronization variables, it further improves upon most depth-first scheduling algorithms which have been proved to perform well only for the more restricted class of *nested parallel* computations.

6.2.3 Accessing Parallel Disks

Investigator: Peter Sanders, Berthold Vöcking

In a balanced high performance computing system build for processing large data sets there will be several hard disk for each processor because they offer the opportunity to store huge amounts of data at low cost. If we succeed to design algorithms that keep all the disks busy with coarse grained accesses, parallel disks also offer the opportunity for very high data bandwidths. However, until recently little was known about handling parallel disks algorithmically. When going from a single disk to parallel disks, the main additional difficulty is load balancing. In I/O bound applications,

we have to make sure that data is placed in such a way that all disks receive approximately uniform load.

Our approach is to use randomization, redundant storage, and sophisticated scheduling algorithms to achieve automatic load balancing that can be used in a wide spectrum of applications.

Let us start with the relatively simple problem where we have a sequence of L blocks each of which needs to be written to one of D disks. The application wants to logically write these blocks in the given order but are allowed to reorder the disk accesses using a buffer of size m blocks. It turns out that the obvious greedy algorithm that delays physical outputs as long as possible is optimal in this situation [6]. For random mapping of logical blocks to physical disks, this translates into an algorithm working in $(1 + O(D/m))L/D$ output steps if m buffer frames are available and $L = \omega(D \log D)$ [13, 6].

Using a newly discovered duality between writing to disks and prefetching, we get an optimal off-line prefetching algorithm with the same efficiency. One important application of these generic algorithms for write scheduling and prefetching is to combine them with external memory sorting algorithms like multi-way mergesort and multi-way quicksort. Using additional refinements for allocating data streams [15], we obtain the best currently known parallel disk sorting algorithms.

Very efficient read accesses to parallel disks can be achieved even without knowledge of future accesses if two randomly placed copies of each block are available. In this situation N blocks can be read from D disks using $\lceil N/D \rceil$ or $\lceil N/D \rceil + 1$ I/O steps with high probability if we use an optimal scheduling algorithm based on maximum flow computations [13, 11, 12]. This improves previous results achieving $O(N/D)$ [7] and $N/D + O(\log \log D)$ [1] using simpler but suboptimal scheduling algorithms. For example, $N = 3.84D$ blocks can usually be retrieved in 4 I/O steps leading to an efficiency very close to one. We also show how to implement the scheduling algorithms efficiently, how to modify the system so that disk failures can be tolerated and how to cope with communciation bottlenecks in the network connecting the disks with the processors.

The above analytical results have been obtained for batched access to the disks. For continuous streams of requests we have developed corresponding algorithms [10] that are analyzed using extensive simulations.

All the algorithms can be generalized to allocation strategies with lower redundancy or higher fault tolerance. The variant for the asynchronous case may be of some independent theoretical interest since it optimally solves a generalized form of the maximum cardinality bipartite matching problem in polynomial time.

Another theoretical problem that we first considered for accessing data streams on parallel disks, reappears for minimizing rotational delays of disks and for minimizing memory overhead for cache efficient memory allocation [8]. The performance of all these problems depends on the maximum overlap of of n randomly allocated circular arcs with average length of α circumferences of the circle. Recently, we were able to show that this overlap is bounded by $\alpha n + O(\sqrt{\alpha(1-\alpha)n \log(1/(\alpha(1-\alpha)))})$ with high probability [14].

References

[1] P. Berenbrink, A. Czumaj, A. Steger, and B. Vöcking. Balanced allocations: The heavily loaded case. In *32nd ACM Symposium on the Theory of Computing*, pages 745–754, 2000.

[2] G. E. Blelloch, P. B. Gibbons, Y. Matias, and G. J. Narlikar. Space-efficient scheduling of parallelism with synchronization variables. In *9th ACM Symposium on Parallel Algorithms and Architectures*, pages 12–23, 1997.

[3] R. D. Blumofe and C. E. Leiserson. Scheduling multithreaded computations by work stealing. *Journal of the ACM*, 46(5):720–748, 1999.

- [4] P. Fatourou. Low-contention depth-first scheduling of parallel computations with synchronization variables. In *13th ACM Symposium on Parallel Algorithms and Architectures*, pages 189–198, 2001.
- [5] P. Fatourou and P. Spirakis. Efficient scheduling of strict multithreaded computations. *Theory of Computing Systems*, 33(3):173–232, 2000.
- [6] D. A. Hutchinson, P. Sanders, and J. S. Vitter. Duality between prefetching and queued writing with applications to integrated caching and prefetching and to external sorting. In *9th European Symposium on Algorithms (ESA)*, LNCS. Springer, 2001. to appear.

[7] R. M. Karp, M. Luby, and F. Meyer auf der Heide. Efficient PRAM simulation on a distributed memory machine. In *24th ACM Sympposim on Theory the of Computing*, pages 318–326, 1992.

- [8] K. Mehlhorn and P. Sanders. Scanning multiple sequences via cache memory. submitted for publication, short version: ICALP '99, pages 655–664.

[9] Girija Narlikar. Scheduling threads for low space requirement and good locality. In *Eleventh Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA '99)*, pages 83–95, New York, June 1999. Association for Computing Machinery.

- [10] P. Sanders. Asynchronous scheduling of redundant disk array. In *12th ACM Symposium on Parallel Algorithms and Architectures*, pages 89–98, 2000.
- [11] P. Sanders. Reconciling simplicity and realism in parallel disk models. In *12th ACM-SIAM Symposium on Discrete Algorithms*, pages 67–76, Washington DC, 2001.
- [12] P. Sanders. Reconciling simplicity and realism in parallel disk models. *Parallel Computing*, 2001. Special Issue on Parallel Data Intensive Algorithms and Applications, to appear.
- [13] P. Sanders, S. Egner, and J. Korst. Fast concurrent access to parallel disks. In *11th ACM-SIAM Symposium on Discrete Algorithms*, pages 849–858, 2000.
- [14] P. Sanders and B. Vöcking. Random arc allocation and applications to disks, drums and DRAMs. submitted for publication, 2001.

[15] J. S. Vitter and D. A. Hutchinson. Distribution sort with randomized cycling. In *12th ACM-SIAM Symposium on Discrete Algorithms*, 2001.

6.3 Parallel and External Graph Algorithms

Investigators: Ulrich Meyer and Jop Sibeyn

The graph traversal routines breadth first search (BFS), depth first search (DFS), and single-source shortest-paths (SSSP) are fundamental ingredients of numerous combinatorial optimization algorithms. Unfortunately, neither I/O-efficient external memory versions nor fast and work-efficient parallel algorithms are known for any of these traversal routines on arbitrary sparse graphs: $\Omega(n)$ I/Os or phases are needed in the worst case. However, for many important graph classes we succeeded to come up with improved results:

In [3] we show how external graph algorithms can take advantage of a redundant graph representation and super-linear space. This yields the first $o(N)$ I/O BFS algorithm for arbitrary undirected graphs with constant maximum node degree. The result carries over to BFS, DFS and SSSP on undirected planar graphs with arbitrary node degrees. We further improved the performance for external DFS on planar graphs converting a work-inefficient parallel algorithm [9] to a rather efficient external algorithm [1]. In particular, this constitutes the first $o(N)$ I/O solution that works with linear space. In [2] we consider the influence of internal memory space restrictions

on the running time of graph traversal algorithms. We develop a space efficient bit state dictionary data structure that is useful in keeping track of already visited nodes. Ongoing research is carried out on semi-external DFS where the internal memory is sufficiently large to keep some constant size information for each node whereas the edges are stored externally. We have developed a heuristic that I/O-efficiently transforms a spanning tree into a DFS tree [8]. An implementation produces very promising results.

We intensified our studies on average-case efficient parallel SSSP algorithms on graphs with random edge weights and moderate maximum shortest path weight $\mathcal{L}$. In [7] we consider load-balancing and precomputing issues for a simple bucket-based SSSP approach. Unfortunately, the running time is linear in $\mathcal{L} \cdot d$ where d denotes the maximum node degree. Decoupling the performance from d and $\mathcal{L}$ was achieved by adaptive bucket-splitting [5], thus yielding the first sequential SSSP algorithm with linear average-case time for arbitrary directed graphs (see Section 4.1 for more details on this result). A basic parallelization of the bucket-splitting idea still requires $\Omega(d)$ phases. However, replacing buckets by a parallel heap structure with adaptive node selection [6] improves the performance for many graphs with unbalanced node degrees. For example, it results in the first $o(n^{1/4})$ time and $\mathcal{O}(n \log n + m)$ work parallel SSSP algorithm for graphs modeling the WWW. Recent work [4] features a split-free bucket structure and an improved node selection strategy which compared to [6] result in a logarithmic factor gain on both running time and work.

References

- [1] L. Arge, U. Meyer, and L. Toma N. Zeh. On external-memory planar depth first search. In *Proc. 7th Intern. Workshop on Algorithms and Data Structures (WADS 2001)*, LNCS. Springer, to appear August 2001.
- [2] S. Edelkamp and U. Meyer. Theory and practice of time-space trade-offs in memory limited search. In *Proc. Joint German/Austrian Conference on Artificial Intelligence (KI-2001)*, LNCS. Springer, to appear September 2001.
- [3] U. Meyer. External memory BFS on undirected graphs with bounded degree. In *Proc. 12th Ann. Symp. on Discrete Algorithms*, pages 87–88. ACM–SIAM, 2001.
- [4] U. Meyer. On the average-case performance of parallel and sequential single-source shortest-path algorithms. submitted, 2001.
- [5] U. Meyer. Single-source shortest-paths on arbitrary directed graphs in linear average-case time. In *Proc. 12th Ann. Symp. on Discrete Algorithms*, pages 797–806. ACM–SIAM, 2001. Extended version: Technical Report MPI-I-2001-1-002, MPI Informatik, 2001.
- [6] U. Meyer. Heaps are better than buckets: Parallel shortest paths on unbalanced graphs. In *Proc. Euro-Par 2001 Parallel Processing*, LNCS. Springer, to appear August 2001.
- [7] U. Meyer and P. Sanders. Parallel shortest path for arbitrary graphs. In *Proc. Euro-Par 2000 Parallel Processing*, volume 1900 of *LNCS*, pages 461–470. Springer, 2000.
- [8] U. Meyer and J.F. Sibeyn. Semi-external depth-first search. Manuscript, 2001.
- [9] Justin R. Smith. Parallel algorithms for depth-first searches. I. planar graphs. *SIAM Journal on Computing*, 15(3):814–830, 1986.

6.4 Software Aspects of External Computing

Investigator: Andreas Crauser

In 1997, we decided to develop a new library for secondary memory computation. The main idea was to make secondary memory computation publicly available in an easy way. The library, later

called *LEDA-SM* (LEDA for secondary memory), should be easy to use (even by non-experts) and easy to extend. LEDA-SM is a C++ library that is connected to the internal-memory library LEDA (Library of Efficient Data types and Algorithms) [3]. It therefore offers the possibility to use advanced and efficient internal memory data structures and algorithms in secondary memory algorithmic design and therefore saves important time in the development of new data structures and algorithms as recoding of efficient internal memory data structures or algorithms is not necessary. Details of the implementation and results can be found in [2].

We shortly review the main design features. The main design goal of the library was to follow the theoretical model of [5]. There each disk is a collection of a fixed number of blocks of fixed size; the blocks are indivisible and must be loaded into internal memory if computation should be performed on them. We therefore directly pay attention to the specific features of hard disks (sectors are the smallest accessible indivisible units) and file systems (fractions of logical disk blocks are the smallest accessible, indivisible units). The library is divided into two layers, *kernel layer* and *application layer*. The kernel layer is responsible for managing secondary memory and the access to it. The application layer consists of secondary memory data structures and algorithms. LEDA is used to implement the in-core part of the secondary memory algorithms and data structures and the kernel data structures of LEDA-SM. The kernel of LEDA-SM is divided into the abstract and the concrete kernel. The concrete kernel is responsible for performing I/Os, managing used and non-used disk blocks and managing users of disk blocks. There are several different implementations for each of these tasks. In particular, both file systems and raw devices are supported. The abstract kernel implements a user-friendly access interface to the concrete kernel. For managing blocks addresses and ownerships and for treating blocks as containers of objects.

The higher level of LEDA-SM implements data structures and algorithms using similar abstractions as LEDA. Besides simple data structure like stacks, queues, and arrays with different caching strategies there are more sophisticated data structures like priority queues, B-trees, buffer trees, and and suffix arrays. There are also algorithms for sorting, matrix arithmetics and some simple graph algorithms. To our best knowledge, no previous library [4, 1] implements such a wide spectrum of sophisticated algorithms in one consistent framework.

References

[1] A. Colvin and T. H. Cormen. ViC*: A Compiler for Virtual-Memory C*. Technical Report PCS-TR97-323, Dartmouth College, Computer Science, Hanover, NH, November 1997.

- [2] Andreas Crauser. *External Memory Algorithms and Data Structures in Theory and Practice.* PhD thesis, Universität des Saarlandes and MPII, 2001.
- [3] K. Mehlhorn and S. Näher. *The LEDA Platform of Combinatorial and Geometric Computing.* Cambridge University Press, 1999.

[4] D. E. Vengroff. *TPIE User Manual and Reference.* Duke University, 1995.

[5] J. S. Vitter and E. A. M. Shriver. Algorithms for parallel memory I: Two level memories. *Algorithmica*, 12(2–3):110–147, 1994.

7 Computational Geometry – Coordinator: Edgar A. Ramos

Our work in this area combines the theoretical investigation of fundamental and application oriented questions with the actual implementation of algorithms and its corresponding theoretical support. The implementation work is described in Section 10, in particular the continuing development of CGAL, and the further work on exact geometric computation. We describe below our theoretical work organized in the following themes: applications of geometric sampling, data structures for geometric queries, curve and surface reconstruction, motion planning, and theoretical support for implementations.

Recently, support for the new project "Effective Computational Geometry for Curves and Surfaces" has been granted by the European Union (for three years starting in May 2001). It is a join effort with other five European research groups and it aims to advance the handling of non linear objects in geometric algorithms, both in theory and practice. This requires interdisciplinary cooperation of computational geometry with computer algebra and numerical analysis, as well as with applied areas like computer aided design. In addition to our subgroup, also the AG1 subgroups "Assembly and Simulation" and "Software Libraries" are heavily involved in this project. See Sections 10.3 and 14.1.3 for more information.

7.1 Applications of Geometric Sampling

Investigator: Edgar A. Ramos

Randomization has been successfully applied to obtain simple and efficient algorithms for many geometric problems. On the other hand, for several problems the best deterministic algorithms mimic a randomized algorithm with each random sampling step replaced with a deterministic counterpart. We have continued research in applications of sampling, both randomized and deterministic.

Simple Polygon Triangulation in Linear Time

The problem of triangulating a simple polygon is a classic one in computational geometry and so the development of an efficient algorithm to solve it has attracted considerable interest. Starting with an $O(n \log n)$ time algorithm [8], where n is the size of the polygon, the time was improved first to $O(n \log\log n)$ [14], and then using randomization to $O(n \log^* n)$ [7, 13]; finally Chazelle [5] presented a deterministic algorithm with a linear running time. Unfortunately, Chazelle's algorithm is quite complex.

In [2], we have described a randomized algorithm for computing the trapezoidation of a simple polygon (formed by shooting a vertical ray through each vertex of the polygon, stopping each ray as soon as it hits another segment on the polygon). The expected running time of the algorithm is linear in the size of the polygon. It has been know for a while that a triangulation of a simple polygon can be obtained from its trapezoidation in linear time using a simple algorithm (in fact, most of the previous algorithms, including Chazelle's, actually described a trapezoidation algorithm). Our algorithm is considerably simpler than Chazelle's optimal deterministic algorithm and can be viewed as a combination of Chazelle's algorithm and the simple non-optimal randomized algorithms with running time $O(n \log^* n)$ mentioned above. An essential new idea is the use of random sampling on subchains of the initial polygonal chain, rather than on individual edges as is normally done.

3-D Diameter of a Point Set

We consider the problem of computing the *diameter* of a given set P of n points in $\Re^3$, that is, the maximum distance between any pair of points in P. A considerable amount of work on this problem has been performed: Yao [15] first showed that the problem can be solved in subquadratic time; he gave a deterministic algorithm with running time $O((n \log n)^{1.8})$; then, Clarkson and Shor [6] gave a simple randomized algorithm that runs in expected time $O(n \log n)$, which is optimal. Since then, it had been a challenge to match that time complexity with a deterministic algorithm. A sequence of works by several researchers made use of derandomization techniques and parametric search and got close to achieve this: $O(n^{4/3+\epsilon})$, $O(n^{1+\epsilon})$, $O(n \log^c n)$, $O(n \log^3 n)$ and finally $O(n \log^2 n)$ in [10] (afterwards, Bespamyatnikh [3] described an alternative algorithm also with running time $O(n \log^2 n)$ but avoiding derandomization and parametric search).

We have presented in [11] a new optimal randomized algorithm that can be easily derandomized using current standard techniques, while preserving the optimal running time. Specifically, the derandomization only makes use of the efficient construction of *epsilon-nets* by Matoušek. A new essential tool, combined with the use of random sampling, is the clustering of subproblems via planar-graph separators. parametric search is not necessary. The result is a relatively simple divide-and-conquer algorithm, modulo the epsilon-net computation.

Linear Programming Queries

We describe in [12] an approach for answering linear programming queries with respect to a set of n linear constraints in $\Re^d$, for a fixed dimension d. Solutions to this problem had been given before by Matoušek [9] using a multidimensional version of parametric search and by Chan [4] using randomization and Clarkson's approach to linear programming. These previous approaches use data structures for halfspace-range emptiness queries and reporting queries, respectively. Our approach is a generalization of Chan's. It also uses halfspace-range reporting data structures, Clarkson's approach to linear programming, and avoids parametric search. In contrast to Chan's, it makes possible to obtain deterministic constructions without considerable additional preprocessing overhead. The new solution is as good or improves the previous solutions in all the range of storage space: with $O(n^{\lfloor d/2 \rfloor} \log^{O(1)} n)$ storage space, it achieves query time $O(\log^{c \log d} n)$, where c is a small constant independent from d, in comparison to $O(\log^{d+1} n)$ for Matoušek's data structure and $O(n^\epsilon)$ for Chan's; with $O(n)$ storage space, it achieves, as Chan's data structure, query time $O(n^{1-1/\lfloor d/2 \rfloor} 2^{O(\log^* n)})$ after $O(n^{1+\epsilon})$ preprocessing, but without using randomization.

Arrangement of a Collection of Curves

Random sampling provides a natural approach for computing subcomplexes of arrangements of geometric objects via divide-and-conquer. Unfortunately, in general it does not lead to optimal algorithms. Most algorithms in the literature are plagued with extra n^ϵ or $\log^c n$ factors. By providing a better analysis, in [1] we show that in the case of the arrangement of n curve segments in the plane, the resulting running time is asymptotically optimal, namely $O(n \log n + k)$, where k is the number of pairwise intersections between the segments. We also introduce a new general approach, divide-and-conquer with partial clean-up, that adds certain globality to the plain divide-and-conquer approach. We apply this approach to the segments problem and obtain a divide-and-conquer algorithm that constructs a structure of optimal size $O(n+k)$, in contrast to $O(n \log \log n + k)$ for the first algorithm. These algorithms are easily parallelizable and can be made deterministic via derandomization techniques. In particular, the second algorithm can be implemented in the

EREW PRAM model, so that it uses optimal work and time $O(\log^{3/2} n)$, in contrast to $O(\log^2 n)$ previous algorithms. The divide-and-conquer with partial clean-up approach also applies to other problems like 3-d convex hulls, 2-d abstract Voronoi diagrams.

References

- [1] N. M. Amato, M. T. Goodrich, and E. A. Ramos. Computing the arrangement of curve segments: Divide-and-conquer algorithms via sampling. In *Proc. 11th Ann. ACM-SIAM Sympos. on Discrete Algorithms (SODA-00)*, pages 705–706, 2000.
- [2] N. M. Amato, M. T. Goodrich, and E. A. Ramos. Linear-time triangulation of a simple polygon made easier via randomization. In *Proc. 16th Ann. Sympos. on Comput. Geometry (SCG-00)*, pages 201–212, 2000.

[3] S. Bespamyatnikh. An efficient algorithm for the three-dimensional diameter problem. In *Proc. 9th Annu. ACM-SIAM Sympos. Discrete Algorithms (SODA 98)*, pages 137–146. SIAM, Jan. 1998.

[4] T. Chan. Fixed-dimensional linear programming queries made easy. In *Proc. 12th Annu. ACM Sympos. Comput. Geom.*, pages 284–290. ACM Press, June 1996.

[5] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete Comput. Geom.*, 6:485–524, 1991.

[6] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.

[7] K. L. Clarkson, R. E. Tarjan, and C. J. Van Wyk. A fast Las Vegas algorithm for triangulating a simple polygon. *Discrete Comput. Geom.*, 4:423–432, 1989.

[8] M. R. Garey, F. P. Preparata D. S. Johnson, and R. E. Tarjan. Triangulating a simple polygon. *Inform. Process. Lett.*, 7:175–179, 1978.

[9] J. Matoušek. Linear optimization queries. *J. Algorithms*, 14:432–448, 1993.

[10] E. A. Ramos. Constructing 1-d lower envelopes and applications. In *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, pages 57–66. ACM Press, June 1997.

- [11] E. A. Ramos. Deterministic algorithms for 3-d diameter and some 2-d lower envelopes. In *Proc. 16th Ann. Sympos. on Comput. Geometry (SCG-00)*, pages 290–299. ACM, 2000.
- [12] E. A. Ramos. Linear programming queries revisited. In *Proc. 16th Ann. ACM Sympos. on Comput. Geometry (SCG-00)*, pages 176–181, 2000.

[13] R. Seidel. A simple and fast incremental randomized algorithm for computing trapezoidal decompositions and for triangulating polygons. *Comput. Geom. Theory Appl.*, 1:51–64, 1991.

[14] R. E. Tarjan and C. J. Van Wyk. An $O(n \log\log n)$-time algorithm for triangulating a simple polygon. *SIAM J. Comput.*, 17:143–178, 1988.

[15] A. C. Yao. On constructing minimum spanning trees in k-dimensional space and related problems. *SIAM J. Comput.*, 11:721–736, 1982.

7.2 Data Structures for Geometric Search Queries

Proximity and Range Reporting Queries with k-d trees

Investigator: Christian Duncan

Spatial data structures are structures capable of answering queries about the relationship between objects in space, such as a set of points, and spatial queries, such as containment of objects inside a given region. Since the emergence of computational geometry, researchers have developed numerous spatial structures each typically specialized for certain objects and queries.

Our work concentrates on objects that are mainly sets of points lying in a d-dimensional space. The queries vary from determining the nearest (or farthest) point in the set to a given query point to reporting all points from the set contained in a given range, e.g. a circular region. In addition, we depart from the notion of calculating an exact solution allowing approximate solutions instead. For instance, rather than determining the point p^* in the set which is nearest to a query point q, we allow a solution p which is "almost" as close to q as p^*. More precisely, we guarantee that $\delta(p,q) \leq (1+\epsilon)\delta(p^*,q)$, where ϵ is some constant greater than 0 and $\delta(p,q)$ is the distance from p to q in some L^p metric (e.g. Euclidean). This relaxation of the problem greatly improves the worst-case analysis of the algorithm.

In [7], we focused on a popular variant of the well known k-d tree data structure [4, 10], proving the structure satisfies an important packing lemma. The consequence of this packing lemma is that standard priority search algorithms for performing approximate nearest-neighbor searching or range searching queries visit at most $O(\log^{d-1} n)$ nodes of the constructed tree T. Although researchers have been able to experimentally show that variants of the k-d tree exhibit polylogarithmic performance [5], surprisingly few theoretical results exist, e.g. see [12]. In recent years, researchers have successfully developed other more complex structures which have guaranteed logarithmic bounds [3, 2, 8, 9]. Our result, however, is the first one proving a worst-case polylogarithmic time bound for approximate geometric queries using the simple k-d tree data structure.

Nearest Neighbor Search in Metric Spaces

Investigators: Michiel Hagedoorn

Nearest-neighbor search is the problem of building a data structure on a set of data points such that for each query point, a closest data point can be found efficiently. This problem is also called the post-office problem and the closest-point problem. The main goal is to find an efficient tradeoff between the *storage space* of the data structure and the *query time* that is needed to find a nearest neighbor using the data structure. The time needed to build the data structure, the *preprocessing time*, is also important.

We mention the asymptotically most efficient data structures for (exact) nearest-neighbor search in d-dimensional Euclidean space. Clarkson [6] presents the randomized post-office (RPO) tree. For each choice of $\delta > 0$, using $O(n^{\lceil d/2 \rceil + \delta})$ expected preprocessing time, it is possible to build an RPO tree that has $2^{O(d)} \log n$ worst-case query time and $O(n^{\lceil d/2 \rceil + \delta})$ expected storage space. Meiser [13] presents a data structure for point location in hyperplane arrangements that allows for nearest-neighbor search in $O(d^5 \log n)$ worst-case query time, using $O(n^{2d+\delta})$ worst-case storage space, for each fixed choice of $\delta > 0$. Using the ray shooting algorithms, we can perform nearest-neighbor search with a time–space tradeoff of roughly $O(n/m^{1/\lceil d/2 \rceil})$ query time (ignoring a polylogarithmic factor) at the cost of $O(m)$ storage space [1].

The aim of our research [11] has been to find data structures that achieve polylogarithmic query time in combination with near-linear storage; each known algorithm for exact nearest-neighbor search achieves polylogarithmic query time only at the cost of $n^{\Omega(d)}$ storage, in the worst case. We achieve logarithmic query time in combination with linear storage space for a restricted class of data sets. This is achieved using a new data structure, which we call the DAS tree. This data structure is not restricted to the Euclidean setting: It works for any set of objects and any distance function (which does not have to be a metric). The DAS tree is a simple data structure: It is a binary tree in which each node stores (at most) two data points. Nearest neighbors are found by following a path from the root that leads to a leaf that stores a nearest neighbor. The path itself is determined

by comparing pairs of distances. We have proved asymptotic bounds on the performance of our data structure. Assume that over the data subsets of each given size, the average number of facets of the Voronoi cell containing any fixed query point is at most a constant. Then, in $O(n^2 \log n)$ expected time, we can build a DAS tree that has a $O(\log n)$ maximum worst-case query time and which uses $O(n)$ storage space.

References

[1] P. K. Agarwal and J. Erickson. Geometric range searching and its relatives. In B. Chazelle, J. E. Goodman, and R. Pollack, editors, *Advances in Discrete and Computational Geometry*, volume 223 of *Contemporary Mathematics*, pages 1–56. American Mathematical Society, 1999.

[2] S. Arya and D. M. Mount. Approximate range searching. In *Proc. 11th Annu. ACM Sympos. Comput. Geom.*, pages 172–181, 1995.

[3] S. Arya, D. M. Mount, N. S. Netanyahu, R. Silverman, and A. Wu. An optimal algorithm for approximate nearest neighbor searching. In *Proc. 5th ACM-SIAM Sympos. Discrete Algorithms*, pages 573–582, 1994.

[4] J. L. Bentley. Multidimensional binary search trees used for associative searching. *Commun. ACM*, 18(9):509–517, 1975.

[5] J. L. Bentley. K-d trees for semidynamic point sets. In *Proc. 6th Annu. ACM Sympos. Comput. Geom.*, pages 187–197, 1990.

[6] K. L. Clarkson. A randomized algorithm for closest-point queries. *SIAM J. Comput.*, 17(4):830–847, August 1988.

• [7] M. Dickerson, C. A. Duncan, and M. T. Goodrich. K-d trees are better when cut on the longest side. In *Proc. European Symposium on Algorithms, LNCS 1879*, pages 179–190. Springer-Verlag, 2000.

[8] C. A. Duncan, M. T. Goodrich, and S. G. Kobourov. Balanced aspect ratio trees and their use for drawing very large graphs. In *Proc. Graph Drawing '98, LNCS 1547*, pages 111–124. Springer-Verlag, 1998.

[9] C. A. Duncan, M. T. Goodrich, and S. G. Kobourov. Balanced aspect ratio trees: Combining the advantages of k-d trees and octrees. *Journal of Algorithms*, 38, 2001.

[10] J. H. Friedman, J. L. Bentley, and R. A. Finkel. An algorithm for finding best matches in logarithmic expected time. *ACM Trans. Math. Softw.*, 3:209–226, 1977.

• [11] M. Hagedoorn. A simple data structure for exact nearest-neighbor search. Submitted to SODA '02, 2001.

[12] D. T. Lee and C. K. Wong. Worst-case analysis for region and partial region searches in multidimensional binary search trees and balanced quad trees. *Acta Inform.*, 9:23–29, 1977.

[13] S. Meiser. Point location in arrangements of hyperplanes. *Information and Computation*, 106(2):286–303, October 1993.

7.3 Curve and Surface Reconstruction

Investigators: Stefan Funke, Edgar A. Ramos

The problem of reconstructing a shape from a finite set of points (sampling) on the shape has attracted much attention in the literature (computer vision and computer graphics) during the last twenty years. The most important problem, the reconstruction of a surface in 3-d space arises in a wide area of applications, mainly in reverse engineering. Many algorithms have been proposed that produce good approximations to the surface, but do not provide a guarantee for the correctness of

the returned solution. Recently, the reconstruction problem has been investigated in computational geometry. More specifically, the noise free case in which the reconstruction can be required to go through the sample points. This has resulted in algorithms that provably solve the reconstruction problem under certain restrictions for the surface and the sampling set: given a "good sampling", the algorithms output a "correct reconstruction."

During the last two years we have conducted research both on the problem of reconstructing surfaces in 3-d as well as on the simpler problem of reconstructing curves in the plane, using tools from computational geometry. This is described below.

Also in our group, Ernst Althaus and Kurt Mehlhorn have investigated a formulation of the curve reconstruction problem as a Traveling Salesman Problem (TSP) and used linear and integer linear programming methods to solve it. They have also extended this to reconstruction of surfaces. See Section 5.4.

Reconstructing a Collection of Curves with Corners and Endpoints

In the case where the shape is a *curve*, the correct solution can uniquely be defined as as the graph $G(S, \Gamma)$ with vertex set S and an edge between two vertices if and only if they are adjacent on some curve in Γ. The first algorithm which could provably reconstruct a collection of non-uniformly sampled closed smooth curves was the CRUST algorithm by Amenta, Bern and Eppstein [3]. Their definition of good sampling is based on the *local feature size*, which measures simultaneously the local curvature and the distance to other curve components. Subsequent papers extended this result to a larger class of curves, e.g. smooth curves with endpoints [8] and non-smooth curves with corners [11, 1].

In [9], we presented the most general algorithm so far which provably reconstructs collections of curves with corners *and* endpoints. For the purpose of analysis, we identify the corner areas and smooth areas. A modified sampling condition is required for the corner areas. When determining the neighbors of a particular sample, we can essentially use the same ideas already shown in [3] as long as we are in the smooth part of the curve. This makes use of the fact that the Voronoi cells are elongated on both sides of the normal direction. For a sample in a corner area, its Voronoi cell is also elongated in the normal direction, but possibly only to one side (the outside of the corner). Our algorithm uses this observation to determine the right connections in the corner area. The main idea is not to detect the corner by simple filtering, but rather start from the 'safe' smooth part of the curve and step by step explore the corner areas.

Smooth-Surface Reconstruction in Near-Linear Time

Recently, several algorithms for surface reconstruction with a correctness guarantee have been proposed. First, Amenta and Bern [2], extending the curve reconstruction work, gave appropriate definitions of good sampling and correct reconstruction and described a 3-d CRUST algorithm. A simplified version of this, the COCONE algorithm was described by Amenta *et al.* [4]. Subsequently, other algorithms have been proposed by Amenta *et al* [5] (POWER CRUST) and Boissonnat and Cazals [6] (interpolation based on natural coordinates). Given a "good sampling" P from a smooth and closed surface S, these algorithms output a "correct reconstruction" Σ: a piecewise linear interpolation of the samples that approximates S pointwise and in surface normal and is also topologically equivalent to it. They are all based on the computation of a (weighted) Voronoi diagram or its dual Delaunay (weighted) tetrahedrization for P, which can have size $\Theta(n^2)$ in the worst case. On the other hand, since the size of Σ is linear in n, a natural question is whether a surface reconstruction can be computed in near linear time. In [7], we described an implementation

of the CoCone algorithm which runs in time $O(n \log n)$ if the sampling is "locally uniform". Intuitively, a sampling is locally uniform if the density changes at most linearly with the distance (in other words, the density function is Lipschitz). Based on this work, in [10] we have devised a new algorithm that has a correctness guarantee and whose worst-case running time is $O(n \log n)$, for a general good sampling. As in some of the previous algorithms, the new algorithm outputs a triangulation that is a subcomplex of the 3-d Delaunay tetrahedrization; however, this is obtained by computing only the relevant parts of the 3-d Delaunay structure. The algorithm first estimates for each sample the surface normal and a parameter that is then used to "decimate" the original sampling. The resulting sampling is locally uniform and so a reconstruction based on it can be computed using the algorithm in [7]. In a last step, the decimated points are incorporated into the reconstruction. In order to achieve the faster running time, the algorithm uses approximations in the solution of several subtasks.

References

[1] E. Althaus and K. Mehlhorn. Traveling salesman-based curve reconstruction in polynomial time. *SIAM Journal on Computing*, 31(1):27–66, 2001.

[2] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discrete Comput. Geom.*, 22(4):481–504, 1999.

[3] N. Amenta, M. Bern, and D. Eppstein. The crust and the β-skeleton: Combinatorial curve reconstruction. *Graphical Models and Image Processing*, 60:125–135, 1998.

[4] N. Amenta, S. Choi, T. K. Dey, and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. In *Proc. 16th Ann. ACM Sympos. Comput. Geom.*, pages 213–222, 2000.

[5] N. Amenta, S. Choi, and R. K. Kolluri. The power crust, union of balls, and the medial axis transform. Submitted for publication, 2000.

[6] J. D. Boissonat and F. Cazals. Smooth surface reconstruction vial natural neighbor interpolation of distance functions. In *Proc. 16th Ann. ACM Sympos. Comput. Geom.*, 2000.

• [7] T. K. Dey, S. Funke, and E. A. Ramos. Surface reconstruction in almost linear time under locally uniform sampling. In *Proc. of the 17th Europ. Workshop on Computational Geometry (EWCG-01)*, pages 129–132, Berlin, Germany, 2001. Freie Universität Berlin.

[8] T: K. Dey, K. Mehlhorn, and E. A. Ramos. Curve reconstruction: Connecting dots with good reason. *Computational Geometry*, 15(4):229–244, 2000.

• [9] S. Funke and E. A. Ramos. Reconstructing a collection of curves with corners and endpoints. In *Proc. of the 12th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA-01)*, pages 344–353, January 2001.

• [10] S. Funke and E. A. Ramos. Smooth-surface reconstruction in near linear time. Submitted to SODA '02, 2001.

[11] J. Giesen. Curve reconstruction, the TSP, and Menger's theorem on length. In *Proc. 15th Ann. ACM Sympos. Comput. Geom.*, pages 207–216, 1999.

7.4 Motion Planning

Load Balancing via Polygon Partitioning

Investigators: Hannah Bast, Susan Hert

The first problem we consider arises in the context of terrain covering using a team of robots. There, the goal is to cause each point of the environment to be visited by at least one robot. The

algorithms that exist for single-robot terrain covering (*e.g.*, [1, 10]) can be used in this setting if the environment is appropriately divided among the robots. Since the time required by a robot to cover a certain region is roughly proportional to the area of the region, the geometric problem that arises here is how to partition a polygonal region into a certain number of connected pieces of given areas. Previous work on partitioning polygons into pieces with given areas has addressed the case of convex polygons only (*e.g.*, [5, 7]) or did not guarantee that the pieces would be connected (*e.g.*, [4]).

We have shown in [2] that producing an optimal area partition of a region with respect to the cut length is NP-hard in the number of pieces and that it is even hard to approximate to within any factor that is independent of the shape of the polygon. We have also presented two polynomial-time algorithms for producing non-optimal partitions. The first, presented in [2], is a simple $O(np)$-time sweep-line algorithm that partitions an n-gon into p pieces of given areas, where p is the number of robots. This algorithm can be seen as a static load balancing algorithm for the multiple-robot terrain covering problem. Static load balancing has the disadvantage that it can lead to an imbalance of work due to the need to rely on only an estimate of the amount of time required to do the work. The second algorithm, presented in [9] and summarized in [8], allows for the use of a combined static and dynamic load balancing technique to be used when assigning work to the robots. Such hybrid methods are known to produce better load balances [3]. To apply these techniques in a setting with p robots, one must first produce a partition of the polygonal region into $p+1$ pieces such that pieces $P_1, \ldots, P_p$ are each connected to piece P_{p+1} along an edge. The algorithm we present requires $O(n + q \log q + pq)$ time to divide a polygon with n vertices that has been partitioned into q convex pieces.

Path Coordination via Monitors

Investigators: Susan Hert

Collision avoidance in multiple-robot environment has been addressed most efficiently by planners that decouple the planning into two stages (*e.g.*, [6, 16]). First, the paths of the robots are computed independently of each other and then velocity profiles along the paths are computed to assure the robots do not collide. This second stage comprises what is know as the path coordination problem. In [9], we illustrate how the basic technique from parallel processing of using monitors to protect regions of shared memory where collisions may occur can be applied in the robotics setting. We describe a data structure in which the current locations as well as the next planned path segments of the robots are stored. Using this data structure, collision-free motion plans for a set of robots can be produced that have potentially shorter waiting times for the robots than those produced by purely off-line planners and can lead to better performance than purely sensor-based planners.

Visibility-based Path Planning

Investigators: Jae-Ha Lee

Both visibility and path planning questions have given rise to well-studied areas in computational geometry. The recent development of robots equipped with cameras requires new research at the interface between these two areas. In particular, visibility becomes a tool of robots for path planning. In the basic problem we have studied, the path planning goal is to detect mobile targets within a polygonal region. Previously, it was not known whether the feasibility of such a path planning is NP-hard or not and the best previous result was an exponential-time algorithm. We

first characterized the class of polygons in which no target can escape from the robot's camera forever and based on this characterization, presented an $O(n^2)$-time algorithm [15]. We also studied its variants in which the visibility of the robots is rather constrained [12, 14] and some optimization problems [13]. Another related researches include the path planning problem in which the output path should obey a curvature constraint [11].

References

[1] E. M. Arkin, S. P. Fekete, and J. S. B. Mitchell. Approximation algorithms for lawn mowing and milling. Technical report, Mathematisches Institut, Universität zu Köln, 1997.

- [2] H. Bast and S. Hert. The area partitioning problem. In David Bremner, editor, *Proceedings of the 12th Annual Canadian Conference on Computational Geometry (CCCG-00)*, pages 163–171, Fredericton, New Brunswick, Canada, August 2000. University of New Brunswick.

[3] H. Bast. On scheduling parallel tasks at twilight. *Theory of Computing Systems*, 33(5):489–563, November 2000.

[4] K. F. Böhringer, B. R. Donald, and D. Halperin. On the area bisectors of a polygon. *Discrete Computational Geometery*, 22:269–285, 1999.

[5] P. Bose, J. Czyzowicz, E. Kranakis, D. Krizanc, and D. Lessard. Near-optimal partitioning of rectangles and prisms. In *Proceedings 11th Canadian Conference on Computational Geometry*, pages 162–165, 1999.

[6] S. Buckley. Fast motion planning for multiple moving robots. In *Proceedings 1989 IEEE International Conference on Robotics and Au tomation*, pages 322–326, May 1989.

[7] R. Guàrdia and F. Hurtado. On the equipartitions of convex bodies and convex polygons. In *Proceedings 16th European Workshop on Computational Geometry*, pages 47–50, 2000.

- [8] S. Hert. Connected area partitioning. In *Proceedings of the 17th European Workshop on Computational Geometry*, pages 35–38, Berlin, Germany, March 2001. Freie Universitt Berlin.
- [9] S. Hert and B. Richards. Multiple robot motion planning = parallel processing + geometry. In Henrik Christensen and Greg Hager, editors, *Sensor Based Intelligent Robots*, Lecture Notes in Computer Science, pages 183–205, Schloss Dagstuhl, Wadern, Germany, November 2001. Springer.

[10] S. Hert, S. Tiwari, and V. Lumelsky. A terrain-covering algorithm for an AUV. *Journal of Autonomous Robots*, 3:91 – 119, 1996.

- [11] J.-H. Lee, O. Cheong, W.-C. Kwon, S. Y. Shin, and K.-Y. Chwa. Approximation of curvature-constrained shortest paths through a sequence of points. In *Proceedings of the 8th Annual European Symposium (ESA-00)*, pages 314–325, Germany, 2000.

[12] J.-H. Lee, S.-M. Park, and K.-Y. Chwa. Searching a polygonal room with one door by a 1-searcher. *International Journal of Computational Geometry and Applications*, 10(2):201–220, 2000.

- [13] J.-H. Lee, S.-M. Park, and K.-Y. Chwa. Optimization algorithms for sweeping a polygonal region by mobile guards. In *Proceedings of the 12th International Symposium on Algorithms and Computation*, 2001 (to appear).
- [14] S.-M. Park, J.-H. Lee, and K.-Y. Chwa. Characterization of rooms searchable by two guards. In *Proceedings of the 11th Conference on Algorithms and Computation*, pages 515–526, Taipei,Taiwan, 2000.
- [15] S.-M. Park, J.-H. Lee, and K.-Y. Chwa. Visibility-based pursuit-evasion in a polygonal region by a searcher. In *Proceedings of the 28th International Colloquium on Automata, Languages and Programming*, pages 456–468, Crete,Greece, 2001.

[16] C. W. Warren. Multiple robot path coordination using artificial potential fields. In *Proceedings 1990 International Conference on Robotics and Automat ion*, pages 500–505, May 1990.

7.5 Object Space Visibility

Investigators: Harald Keßler, Kurt Mehlhorn, and Sherif Ghali

To compute visibility in object space is to determine what is visible on the (real, Euclidean) projection plane. Unlike the rasterized output widely used in computer graphics, an object space output can be arbitrarily magnified by using only scale and translate operations.

In practice, Binary Space Partitions, or BSPs, appear to be the only method that is widely used for object space visibility. But BSP computation only provides the depth order of the features in a scene and not the set of visible features. BSP output also lacks adjacency information: Many applications require that the regions visible on the projection plane are stored as a planar map whereby the adjacent regions cross–reference each other.

Algorithms that compute the visibility as a planar map exist, but a wide gap exists between the theory and the practice of these algorithms. Indeed only a partial solution [1] is known for the simpler case of computing visibility without adjacency information. Our objectives are: 1– to study the practicality of the more efficient algorithms, 2– to implement a non–robust version of one algorithm, 3– to study the modifications necessary in the algorithm and its implementation to handle special inputs, and 4– to implement the modified robust algorithm. We are currently in step 2 of this work, having identified which algorithms are more amenable for conversion to a robust implementation.

References

[1] Jim Blinn. Fractional invisibility. *IEEE CG&A*, November 1988. Also available in *Jim Blinn's Corner, A Trip Down the Graphics Pipeline*, Chap. 10, Morgan Kaufman.

7.6 Theoretical Support for Implementations

A Separation Bound for Real Algebraic Expressions

Investigators: Christoph Burnikel, Stefan Funke, Kurt Mehlhorn, Stefan Schirra, Susanne Schmitt

Non-linear objects (circles, ellipses, ...) lead to expressions involving roots and hence an efficient method for computing signs of algebraic expressions is an essential basis for the robust implementation of geometric algorithms dealing with non-linear objects.

The separation bound approach is the most successful approach to sign computation; it is, for example, used in the number type `leda_real` [4] and the number type `Expr` of the `CORE` package [3]. A *separation bound* is an easily computable function sep mapping expressions into positive real numbers such that the value ξ of any non-zero expression E is lower bounded by $sep(E)$, i.e., either $\xi = 0$ or $|\xi| \geq sep(E)$. Separation bounds allow one to determine the sign of an expression by numerical computation. An error bound Δ is initialized to some positive value, say $\Delta = 1$, and an approximation $\widetilde{\xi}$ of ξ with $|\xi - \widetilde{\xi}| \leq \Delta$ is computed using approximate arithmetic, say floating point arithmetic with arbitrary-length mantissa. If $|\widetilde{\xi}| > \Delta$, the sign of ξ is equal to the sign of $\widetilde{\xi}$. Otherwise, $|\widetilde{\xi}| \leq \Delta$ and hence $|\xi| < 2\Delta$. If $2\Delta \leq sep(E)$, we have $\xi = 0$. If $2\Delta > sep(E)$, we halve Δ and repeat. The worst case complexity of the procedure just outlined is determined by the separation bound; $\log(1/sep(E))$ determines the maximal precision needed for the computation of $\widetilde{\xi}$ and we refer to $\log(1/sep(E))$ as the *bit bound*.

In [2] we prove a new separation bound for real algebraic expressions and compare it analytically and experimentally with previous bounds. The bound is used in the sign test of the number type `leda_real`.

Infimaximal Frames

Investigators: Michael Seel, Kurt Mehlhorn

Many geometric algorithms that are usually formulated for points and segments generalize easily to inputs also containing rays and lines. The sweep algorithm for segment intersection is a prototypical example. Implementations of such algorithms do, in general, not extend easily. For example, segment endpoints cause events in sweep line algorithms, but lines have no endpoints. We described a general technique, which we call infimaximal frames [6, 7], for extending implementations to inputs also containing rays and lines. The technique can also be used to extend implementations of planar subdivisions to subdivisions with many unbounded faces. We have used the technique successfully in generalizing a sweep algorithm designed for segments to rays and lines and also in an implementation of planar Nef polyhedra [5, 1].

References

[1] H. Bieri. Nef Polyhedra: A Brief Introduction. *Computing Suppl. Springer-Verlag*, 10:43–60, 1995.

- [2] C. Burnikel, S. Funke, K. Mehlhorn, S. Schirra, and S. Schmitt. A separation bound for real algebraic expressions. In Friedhelm Meyer auf der Heide, editor, *Proc. of the 9th Europ. Symp. on Algorithms (ESA-01)*, volume 2161 of *Lecture Notes in Computer Science*, pages 326–337, Aarhus, Denmark, August 2001. Springer.

[3] V. Karamcheti, C. Li, I. Pechtchanski, and C. Yap. A core library for robust numeric and geometric computation. In *Proc. of the 15th Intern. Ann. ACM Symp. on Computational Geometry (SCG-99)*, pages 351–359, Miami, USA, 1998. ACM Press.

[4] K. Mehlhorn and S. Näher. *LEDA: A Platform for Combinatorial and Geometric Computing*. Cambridge University Press, Cambridge, UK, 2000.

[5] W. Nef. *Beiträge zur Theorie der Polyeder mit Anwendungen in der Computer graphik*. Herbert Lang & Cie AG, Bern, 1978.

[6] M. Seel and K. Mehlhorn. Infimaximal Frames: a framework to make lines look like segments. Technical Report MPI-I-2000-1-005, Max-Planck-Institut für Informatik, Saarbrücken, 2000.

[7] M. Seel and K. Mehlhorn. Infimaximal Frames: a framework to make lines look like segments. In *Proceedings of the 17th European Workshop on Computational Geometry*, 2001.

8 Assembly and Simulation – Coordinator: Elmar Schömer

In automotive industry it is a costly and time consuming step to verify the design of a product with respect to manufacturability or assemblability. In addition to building physical prototypes engineers more and more use digital mock-ups because they provide the flexibility to explore alternative solutions more easily. In cooperation with the Virtual Reality Competence Center of DaimlerChrysler in Ulm we are investigating the possibility to perform assembly simulations and ergonomics studies in virtual reality environments. In a virtual assembly simulation the goal is to evaluate if and how the different component parts fit together and how much clearance space is there for easy removal and assembly of the parts. For this purpose the simulation software must be able to immediately interact with the user who wants to manipulate the motions of the parts by means of data gloves. The most important algorithmic problem arising in this context is real-time collision detection for moving rigid bodies. Most theoretical (e.g. [15]) and practical (e.g. [4]) work in this area has focused on polygonal models. In contrast to this we attack the collision detection problem for boundary models with *curved* surfaces. The main two reasons for choosing this approach are:

1. There is no need for converting the original CAD data to triangular meshes – a process, which can dramatically increase the complexity of the models and thus slows down the collision detection algorithms.

2. Coarse polygonal approximations of curved surfaces can inhibit the mountability of objects – a serious problem in assembly simulations.

Beyond the pure geometric problem of detecting collisions between moving objects we are also studying the physical reaction of colliding rigid bodies. In this area we have developed efficient and robust numerical methods for the calculation of contact [2] and friction forces [14], which are based on the solution of non-linear complementarity problems. Most of the algorithms have been implemented in our experimental simulation library for virtual reality an interactive applications, named SiLVIA.

8.1 Static Collision Detection for Objects with Curved Surfaces

Investigators: Elmar Schömer, Thomas Warken

A static collision detection algorithm determines whether two objects at given positions and orientations intersect. If we disregard the case that one object is completely enclosed by the other one, the algorithm must test whether the boundaries of the objects intersect. The static collision detection algorithm that we developed works as follows. Let O_1 and O_2 be two objects. In order to avoid describing special cases we assume that these objects are in general positions and orientations. For each pair f_1, f_2 of faces we perform the following test. First we check whether an edge of f_1 intersects f_2 and vice versa. If so we report a collision. Otherwise we know that each component of the intersection curve of the surfaces in which f_1 and f_2 are embedded lies completely inside or outside the faces. Therefore it suffices to compute one point on each such component and check whether it lies inside both f_1 and f_2. If all these points lie outside f_1 or f_2 we know that the faces do not intersect. Otherwise we report a collision.

This algorithm must perform the following computations:

- Compute the intersection between a curve and a surface.

- Compute one point on each component of the intersection curve of two surfaces.
- Check whether a point lies inside a face or on a curve segment.

We could adapt the results of [10] to show that for quadratic complexes, i.e. objects that are bounded by segments of quadrics and conic sections, these computations can be performed by finding the roots of polynomials of degree at most four. By taking the results of [8] into consideration we could show that this is also true for objects that are bounded by segments of natural quadrics, tori and conic sections. Further research in this direction will be on extending these classes of objects in natural ways. One such extension will be intersection curves of quadrics as parts of the boundary.

8.2 Dynamics Simulation for Objects with Curved Surfaces

Investigators: Elmar Schömer, Thomas Warken

As an aspect of the dynamic behavior of objects with curved surfaces we considered the rolling motion. This behavior occurs if the objects are in mutual contact and the sticking friction prevents them from sliding along each other. In [16] we developed the dynamics equations that describe the rolling motion of an arbitrarily shaped rigid object that is in a one- or two-point contact with an arbitrarily shaped surface. The starting point of this approach is a dynamics constraint which we call the *rolling condition*. If $\boldsymbol{v}$ and $\boldsymbol{\omega}$ are the linear and angular velocities of the object relative to the surface and $\boldsymbol{r}$ is the vector pointing from the center of mass of the object to any contact point then the rolling condition for that contact point is $\boldsymbol{v} + \boldsymbol{\omega} \times \boldsymbol{r} = \mathbf{0}$. This means that the relative contact velocity of each contact point is zero. This is just what *not sliding* means. Subtracting the rolling conditions for two different contact points yields $\boldsymbol{\omega} \times (\boldsymbol{r}_1 - \boldsymbol{r}_2) = \mathbf{0}$. This means that rolling implies that all contact points are collinear. Thus we can restrict ourselves to the cases of one- and two-point contacts. If we combine the rolling conditions with the Newton-Euler dynamics equations we obtain a system of first order ordinary differential equations of the form

$$\mathbf{A} \begin{bmatrix} \dot{\boldsymbol{v}} \\ \dot{\boldsymbol{\omega}} \end{bmatrix} = \boldsymbol{b}(\boldsymbol{v}, \boldsymbol{\omega}),$$

where the matrix $\mathbf{A}$ depends only on the geometry and $\boldsymbol{b}$ is a function of the linear and angular velocity. This system describes the rolling motion of the object on the surface. With this approach we could simulate the rolling motion of an oloid and a frustum on an inclined plane.

8.3 Distance Computation between Quadratic Complexes

Investigators: Christian Lennerz, Elmar Schömer

Given two quadratic complexes at a given point and orientation, the distance computation problem is to determine the Euclidean distance between them. It is obvious that this task can be accomplished by restricting oneself to the respective boundary representations. Since the faces are embedded on quadrics, the basic problem is to determine the distance between two patches of quadratic surfaces that are trimmed by quadratic curve segments.

One can consider the problem as a nonlinear, constrained optimization problem, but it can be solved much more efficiently by formulating it as a system of polynomial equations. Our formulation is based on the following observation: A pair of closest points between both patches is either minimizing the distance between the non-trimmed quadratic surfaces (surface-surface case) or it

consists of at least one point lying on a quadratic curve segment bounding its patch (curve-surface case). From the perpendicularity conditions of the surface-surface case we derived a system of non-linear polynomial equations in two variables. Though these polynomials have degree six in every variable, the total degree is also bounded by six. It is obvious that the curve-surface case can be handled by an analogous argument, leading to polynomial systems of lower complexity.

To solve these systems of bivariate polynomial equations we have applied different numerical and algebraic methods. We first examined the use of Interval Newton Methods since they guaranty global convergence to all roots as well as control over the numerical precision [12, 1, 13]. Our experiments have shown that though these techniques work very well on systems of low complexity like the curve-curve problem, they break down on the higher degree surface-surface formulation. To overcome this problem which is caused by over-estimating interval evaluations on polynomial equations, we have reformulated the problem as a system of non-polynomial equations.

Since the running times were still non-satisfying, we have also considered algebraic methods for solving such systems of bivariate polynomial equations. Generally these methods are based on elimination theory that provides techniques for reducing the bivariate root finding problem to the univariate case [3]. The elimination step transforms the bivariate system into a polynomial in one variable, the so-called Resultant polynomial. This polynomial has the property that its set of roots is just the projection of the bivariate solutions onto the non-eliminated variable. Repeating the elimination step by interchanging the roles of the variables gives the respective values for the other variable.

The Resultant polynomial is defined as the determinant of the so-called Sylvester Matrix, which consists of polynomials in the non-eliminated variable. To determine the power representation of the Resultant polynomial without computing this determinant symbolically, we evaluate the determinant at a sufficiently large set of sample points. Then the interpolating polynomial gives the coefficients of the Resultant polynomial which is solved by the Eigenvalue Method for univariate polynomials. When computing in double precision the method showed very promising running times but suffered from numerical problems. We already did many improvements with respect to numerical stability, but accuracy still heavily depends on an appropriate choice of sample points.

The main reason for the lack of numerical stability are the evaluations of the matrix determinant and the interpolation of the set of sample points. Both steps are needed to compute the power representation of the Resultant polynomial. To avoid this numerical bottleneck we exploited the fact that the roots of the Resultant polynomial form a subset of the eigenvalues of the so-called Generalized Companion Matrix [11]. Our experiments indicated that this approach shows sufficient numerical robustness whereas the sparsity of the matrix promises further improvements on the efficiency and accuracy of the eigenvalue computation.

8.4 Bounding Volume Hierarchies

Investigators: Christian Lennerz, Elmar Schömer

The elementary procedure of our algorithms for collision detection and distance computation is operating on a single pair of surface patches. They test whether both faces intersect and determine their Euclidean distance respectively. Since quadratic complexes may consist of a large number of surface patches, it is very inefficient to compare all faces of one object to all faces of the second object. By wrapping the patches in preferably small and simply shaped envelopes like spheres or boxes, many pairs of faces can be pruned away efficiently. Filtering is based on the fact that intersection of the bounding volumes is a necessary condition for the interior patches to collide. Analogously, the Euclidean distance between the envelopes gives a lower bound on the distance of

the enclosed faces.

The bounding volume hierarchy is a simple data structure that allows collision detection and distance calculation queries following divide and conquer strategies. It is represented as a binary tree of bounding volumes covering a hierarchical decomposition of the object. The root of the tree is associated with the full face set and every inner node corresponds to a proper subset of its parent's face set. Thereby both children form a partition of their parent's face set. We speak of a leaf if no further partitioning of the face set is possible. Every face set corresponding to a node is wrapped by a bounding volume of minimal size.

Collision detection and distance computation queries are parallel tree traversals of two bounding volume hierarchies representing the boundary of the respective quadratic complexes. Given a pair of nodes their subtrees can be cut off if the intersection test on the bounding volumes indicates that the enclosed face sets do not collide. In a distance computation query the child nodes can be pruned away if the lower bound of distance, given by the distance of the envelopes, is larger than the global upper bound. This upper bound is updated by the elementary distance computation procedure every time a pair of leaves is reached [9].

8.5 Quadric Surface Intersection

Investigators: Nicola Geismann, Elmar Schömer

In solid modeling it is important that the modeling kernel can reliably perform all basic operations such as calculating the intersection, the union or the difference of two solid objects. Commercial CAD systems often work with free-form surfaces, i. e. represented by Nurbs. Thus they depend on numerical approximation methods for calculating intersection curves, which must be determined when performing boolean operations. Especially in degenerate geometric situations these methods for curve tracing are error-prone because they are based on floating point arithmetic. In a strict sense they can not guarantee topologically correct results.

In mechanical solid modeling free-form surfaces only play an inferior role, since mechanical parts mainly consist of planar, spherical, cylindrical, conical or torical surface patches. Even for this restricted class of surfaces it is a challenging problem to develop an algorithm which can compute the correct topological structure of the boundary of an object, specified by a sequence of boolean operations on the given set of primitives.

In our approach to this problem [5] we assume that the primitives are quadric surfaces with rational coefficient matrices and that we want to compute their intersection. Thanks to the simple implicit representation of these surfaces we can calculate the projection of the pairwise intersection curves onto the xy-plane by means of the Sylvester resultant. The knowledge of self-intersection points and mutual intersection points in the 2-dimensional arrangement of the resulting algebraic curves is fundamental for inferring the topological structure of the intersection object in 3-space.

The key problem in constructing the arrangement of the projected curves consists in handling the intersection points, since these points do not possess rational coordinates in general. But their coordinates can be characterized as roots of univariate polynomials of degree at most 8 with integer coefficients. As in [7] we used root separation techniques and box hit counting in order to locate the critical points. In addition to that we introduced the concept of Jacobi curves which facilitate the localization of tangential intersection points. Self-intersection points could be described with the help of simple square root expressions.

We are currently working on a prototypical implementation of the algorithm sketched above. Alternatively we are investigating an approach which works with a suitable parameterization of the spatial intersection curves. (Unfortunately a rational parameterization does not exist.) This

latter approach makes use of interval arithmetic in order to identify critical degenerate geometric situations.

The dual to the intersection problem is the convex hull problem. We studied this problem for ellipsoids. A computer animation of the respective algorithm can be found in [6].

References

[1] G. Alefeld and J. Herzberger. *Introduction to interval computations.* Academic Press, 1983.

[2] M. Buck and E. Schömer. Interactive rigid body manipulation with obstacle contacts. *Journal of Visualization and Computer Animation*, 9:243–257, 1998.

[3] D. Cox, J. Little, and D. O'Shea. *Using algebraic geometry*, volume 185 of *Graduated texts in mathematics.* Springer, 1998.

[4] S. Ehmann and M. Lin. Accurate and fast proximity queries between polyhedra using surface decomposition. In *Computer Graphics Forum (Proc. Eurographics)*, 2001.

- [5] N. Geismann, M. Hemmer, and E. Schömer. Computing a 3-dimensional cell in an arrangement of quadrics: Exactly and actually! In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 264–273, 2001.
- [6] N. Geismann, M. Hemmer, and E. Schömer. The convex hull of ellipsoids. In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 321–322, 2001.

[7] J. Keyser, T. Culver, D. Manocha, and S. Krishnan. Mapc: A library for efficient and exact manipulation of algebraic points and curves. In *Proc. 15th Annu. ACM Sympos. Comput. Geom.*, pages 360–369, 1999.

[8] K.-J. Kim. *Torus and Simple Surface Intersection Based on a Configuration Space Approach.* Ph.D. thesis, Department of Computer Science and Engineering, POSTECH, February 1998.

- [9] C. Lennerz. Impulsbasierte Dynamiksimulation starrer Körper unter Verwendung von Hüllkörperhierarchien. Master's thesis, Fachbereich Informatik, Universität des Saarlandes, 2000.

[10] J. Levin. A parametric algorithm for drawing pictures of solid objects composed of quadric surfaces. *Commun. ACM*, 19(10):555–563, 1976.

[11] D. Manocha. Solving systems of polynomial equations. *IEEE Computer Graphics and Applications*, 1994.

[12] R. E. Moore. Methods and applications of interval analysis, 1979.

[13] A. Neumaier. Interval methods for systems of equations, 1990.

[14] J. Sauer and E. Schömer. A constraint-based approach to rigid body dynamics for virtual reality applications. In *ACM Symposium on Virtual Reality Software and Technology*, pages 153–161, 1998.

[15] E. Schömer and C. Thiel. Efficient collision detection for moving polyhedra. In *Proc. 11th Annu. ACM Sympos. Comput. Geom.*, pages 51–60, 1995.

- [16] T. Warken and E. Schömer. Rolling rigid objects. In *Proceedings of WSCG'2001*, volume 1, pages 57–62, 2001.

9 Bioinformatics – Coordinator: Hans-Peter Lenhof

In 1999 and 2000 the German national research council (DFG, Deutsche Forschungsgemeinschaft) initiated a Bioinformatics competition, the so-called "Intitiative Bioinformatics". Thirty-one German regions and cities entered the contest with concepts of how to build up centers of excellence in Bioinformatics. Saarbrücken is one of the five winners of the "Initiative Bioinformatics" and will thus receive financial support from the DFG. The project of Saarbrücken, the "Virtual Biolab", was awarded a five million DM grant from the DFG (for the first two years) and can potentially receive another five million DM for the following three years. The project "Virtual Biolab" is a common proposal of the Max Planck Institute for Computer Science (MPI), the University of Saarland, and the German Research Center for Artificial Intelligence GmbH (DFKI), under the scientific leadership of the Bioinformatics group of AG1. The Bioinformatics landscape in Saarbrücken has changed dramatically, and Bioinformatics is now a center of gravity at the University of Saarland. A center for Bioinformatics (center of excellence) was founded. The Bioinformatics group of AG1 moved to the university and Hans-Peter Lenhof became the first professor for Bioinformatics at the University of Saarland and one of the first Bioinformatics professors in Germany. A second professorship for Bioinformatics at the Saarland University will be installed soon. Thomas Lengauer, one of the world's leading researchers in the area of Bioinformatics, will become the director of the new working group at the MPI. Therefore, one of the strongest Bioinformatics groups in the world will soon join the MPI and Bioinformatics will play a much more important role at MPI. Oliver Kohlbacher, a former researcher of the Bioinformatics group of AG1, will lead the first independent young researcher group for Bioinformatics at the University of Saarland. Two other groups will be installed in 2002 and 2003. Two courses of study for Bioinformatics have been established at the University of Saarland, a Bachelor/Master course and a so-called "Diplom-Teilstudiengang". Both courses start for the first time in the winter term 2001/2002.

We continued our work on integer linear programming formulations and branch-and-cut algorithms for sequence alignment problems. A few years ago, we presented the first branch-and-cut algorithm for a special sequence alignment problem, the so-called Generalized Maximum Trace (GMT) problem. This branch-and-cut algorithm was able to solve large instances of GMT problems to optimality. In the GMT version of the multiple sequence alignment problem gaps were ignored and only a certain set of possible multiple sequence alignments were considered. In the meantime, we found an integer linear programming for the general multiple sequence alignment problem with gaps and arbitrary gap cost and developed and implemented a branch-and-cut algorithm for this general problem formulation. The new branch-and-cut algorithm is able to solve large problem instances that cannot be solved by dynamic programming. Up to now these instances could be solved only if the number of possible solutions (edges in the graph) had been reduced tremendously (by defining a somehow smaller problem).

The goal of protein docking research is the development of algorithms that enable the user to predict interactions between proteins and to compute the three-dimensional structure of the resulting protein complexes. Protein docking algorithms can be used to screen protein databases for possible inhibitors of a given virus enzyme. In the course of our DFG project "Protein-Protein-Docking", we developed and published new techniques for finding the optimal conformations of the amino acid side chains in the docking site. We presented two new algorithms for the side chain placement problem, a fast multi-greedy heuristic that is able to identify good approximations of the optimal conformations and a more sophisticated branch-and-cut algorithm that is able to identify

the optimal conformations, even for huge problem instances. Furthermore, we presented a new filtering technique for scoring potential protein complexes. The new filtering technique is based on 1H-NMR-spectra. It compares the experimental 1H-NMR-spectrum of the complex with calculated spectra of the candidates and sorts the candidates with respect to the area of the so-called difference spectra. For some of the docking examples, the new scoring function was able to separate the good approximations of the complex from the false positives.

The new scoring function and the side chain placement algorithm have been implemented within our C++ framework BALL (Biochemical Algorithms Library). BALL, which has been designed for rapid software prototyping, enabled us to implement and test each of the complex algorithms in less than two weeks. In the last two years, more functionality was added to BALL (e.g. Calculation of NMR-spectra). In 2000, Dr. Oliver Kohlbacher was awarded the "Heinz Billing Award for the Advancement of Scientific Computing" of the Max Planck Society for the development of BALL.

The goal of our BMBF project GELENA (GEne transfer systems based on LEctin-modified NAnoparticles) is the development of a new gene transfer method based on nanoparticles loaded with DNA vectors. The gene transfer method is intended to cure diseases caused by genetic defects such as cystic fibrosis. In close cooperation with the biotechnology company Across Barriers GmbH and the Institute for Biopharmacy and Pharmaceutical Technology of the University we are designing small peptide mimetics that have a high binding affinity to certain sugar molecules, so-called oligomers of N-acetyl-glucosamine (NAG). We are interested in these sugar molecules because they cover the surface of epithelial cells. We have developed a new model to predict the binding affinity of small sugar molecules and so-called lectins. The new approach is able to predict the binding affinity with high accuracy and produces an extremely good correlation between experimental and calculated affinities.

Our work in the area of approximate string matching concentrates on the development and improvement of filter algorithms. In the past we have combined a well–known filter technique, the q-gram filtration, with an indexing data structure, the suffix array, to build a powerful algorithm for high similarity approximate matching. Now we concentrate on improving the filter technique in order to increase the speed and the sensitivity of the complete algorithm. The use of word–neighborhoods allows for an increased sensitivity but also leads to a slightly slower filter algorithm. Another approach involving gapped q-grams is more flexible. It offers a wide range for improving either speed, sensitivity or a mixture of both, of course, with a tradeoff between them. We are currently working on creating more powerful filters that use sets of different gapped q-grams.

In 2000, we started a new project "Metabolic and Regulatory Pathways". The goal of the project is the development of new tools and algorithms (1) to handle biochemical reaction data, (2) to browse and to search in the large amount of available biochemical data in an easy and comfortable way, and (3) to visualize the complex results of the biochemical queries. First, we designed and implemented a class hierarchy that can be used to model all standard biochemical reactions and the so-called metabolic, regulatory and signaling pathways. The different pathways are represented by graphs where the nodes represent the biochemical compounds and the edges can be interpreted as biochemical reactions. Second, we implemented search algorithms to identify paths with some special properties. Third, in close cooperation with the group of Prof. Michael Kaufmann at the university of Tübingen, we developed new algorithms to visualize the pathways. We tested the data structures and algorithms with a first biological application, the identification

of new, previously unknown, metabolic pathways.

9.1 Multiple Sequence Alignment

Investigators: Ernst Althaus, Hans-Peter Lenhof, Knut Reinert

Let $S = \{S_1, S_2, \ldots, S_k\}$ be a set of k strings of length $n_1, \ldots, n_k$ over an alphabet Σ and let $\hat{\Sigma} = \Sigma \cup \{-\}$, where "$-$" (dash) is a symbol to represent "gaps" in strings. An *alignment* of S is a set $\hat{S} = \{\hat{S}_1, \hat{S}_2, \cdots, \hat{S}_k\}$ of strings over the alphabet $\hat{\Sigma}$ that satisfies the following two properties: (1) the strings in $\hat{S}$ all have the same length, and (2) ignoring any dashes, string $\hat{S}_i$ is identical to string S_i. An alignment in which each string $\hat{S}_i$ has length l can be interpreted as an array of k rows and l columns where row i corresponds to string $\hat{S}_i$. Two characters of distinct strings in S are said to be *aligned* under $\hat{S}$ if they are placed in the same column of the alignment array. A scoring function assigns to each alignment a score such that in as many situations as possible the score is in accordance with biology in the sense that alignments are assigned high scores if and only if they are *biologically meaningful*. Since sequence alignment is of central importance in the interpretation of protein and DNA sequences, a lot of research has been conducted in this area and dozens of algorithms for the different alignment problems have been published. Almost all published non-heuristic algorithms for these problems use dynamic programming. Since the time and space complexity of dynamic programming algorithms grows exponentially with the number of sequences, dynamic programming approaches can solve only rather small problem instances to optimality, even if elaborated bounding procedures are applied (see [4]).

In 1997, we published an integer linear programming formulation (ILP) for the so-called Maximum (Weight) Trace (MT) problem, an instance of the Multiple Sequence Alignment problem introduced by John Kececioglu [3], and we presented a branch-and-cut algorithm for this problem that was able to solve problem instances that cannot be handled by dynamic programming approaches [7]. In the MT problem we view the character positions of the k input strings in S as the vertex set V of a k-partite graph $G = (V, E)$ called the input *alignment graph*. The edge set E connects pairs of characters that one would like to have aligned. We call an edge in E *alignment edge* and say that an alignment edge is *realized* by an alignment if the endpoints of the edge are placed into the same column of the alignment array. The subset of E realized by an alignment $\hat{S}$ is called the *trace* of $\hat{S}$.

Based on this initial formulation we introduced the *Generalized Maximum Trace Problem* (GMT) [6, 2], in which we allow multiple edges between two vertices in the alignment graph G and in which we partition the edge set E into a set D of so-called *blocks*. In this way we were able to handle more generalized scoring schemes.

For both problems mentioned above we derived facet-defining inequalities for the corresponding polytopes and implemented branch-and-cut algorithms to solve them to optimality. Finally we proposed an ILP formulation for multiple sequence alignment with arbitrary gap costs. This formulation was since then extended and a branch-and-cut algorithm was implemented which at the time of this writing is about to be submitted to RECOMB 2002 [1]. In the following paragraph we will elaborate on the progress made for the *gapped trace problem*.

In order to express our problem as a graph we need the notion of a *mixed graph* which is a tuple $G = (V, E, A)$, where V is a set of vertices, E is a set of edges and A is a set of arcs.

The edge set E represents pairs of characters that one would like to have aligned in an alignment of the input strings. As in previous work we say that an edge $e = \{u, v\}$ is *realized* by an alignment if the endpoints of the edge are placed into the same column of the alignment array.

The arcs in $A = A_g \cup A_p$ represent positional constraints. Arcs in A_p run from each node to its

"right" neighbor (from v_j^i to v_{j+1}^i for $1 \leq i \leq k$ and $1 \leq j <\mid s^i \mid$). The arcs in A_g represent gaps in the alignment. Each substring of a string s^j can be aligned with gap characters in any other string s^i, or put it differently, no character in this substring of s^j is aligned with any character in s^i. Hence we introduce for each such substring an arc from the beginning of the substring in s^j to its end. Each of the edges in E and gap arcs in A is assigned a weight that corresponds to the benefit (or cost) of realizing the edge or arc. The goal is to identify the subgraph which has the highest weight given a suitable scoring scheme, and given certain constraint on the possible combinations of edges and arcs. It is relatively straightforward to translate this graph-theoretic problem into an ILP, for which we present an implementation in [1].

The solution of the *gapped trace problem* can thus formulate a great variety of multiple sequence alignment problems, among them the weighted sum of pairs problem with arbitrary gap costs. To our knowledge this is the first algorithm that can deal with truly affine gap costs ([5, 8] use what is called "quasi"-affine or natural gap costs). Indeed our algorithm is independent of the choice of the gap cost function and can handle any function including convex gap costs.

References

- [1] E. Althaus, A. Caprara, K. Reinert, and H.-P. Lenhof. Multiple sequence alignment with arbitrary gap costs – Computing an optimal solution using arbitrary gap costs –. In *Proceedings of the Sixth Annual International Conference on Computational Molecular Biology (RECOMB 02)*, 2002. sumitted for publication.
- [2] J. D. Kececioglu, H.-P. Lenhof, K. Mehlhorn, P. Mutzel, K. Reinert, and M. Vingron. A polyhedral approach to sequence alignment problems. *Discrete Applied Mathematics*, (104):143–185, 2000.

[3] J.D. Kececioglu. *Exact and approximation algorithms for DNA sequence reconstruction.* PhD thesis, University of Arizona, 1991.

- [4] M. Lermen and K. Reinert. The practical use of the $\mathcal{A}^*$ algorithm for exact multiple sequence alignment. *Journal of Computational Biology*, pages 655–671, 2000.

[5] D.J. Lipman, S.F. Altschul, and J.D. Kececioglu. A tool for multiple sequence alignment. *Proc. Natl. Acad. Sci. U.S.A.*, 86:4412–4415, 1989.

- [6] K. Reinert. *A polyhedral approach to sequence alignment problems.* PhD thesis, Universität des Saarlandes, 1999.

[7] K. Reinert, H.-P. Lenhof, P. Mutzel, K. Mehlhorn, and J. Kececioglu. A branch-and-cut algorithm for multiple sequence alignment. In *Proceedings of the First Annual International Conference on Computational Molecular Biology (RECOMB 97)*, pages 241–249, 1997.

- [8] K. Reinert, J. Stoye, and T. Will. An iterative methods for faster sum-of-pairs multiple sequence alignment. *BIOINFORMATICS*, 16(9):808–814, 2000.

9.2 Approximate String Matching in Genomic Databases

Investigators: Stefan Burkhardt, Juha Kärkkäinen

Large collections of DNA and protein sequences are nowadays a cornerstone of molecular biology. Many applications like clustering [10], or sequence assembly [15] rely on massive use of such databases.

There are well known programs for similarity search in sequence databases, e.g. BLAST [1] and FASTA [14]. These two and most other algorithms scan the whole database linearly in order to detect all approximate matches. Since the databases are growing rapidly, more sophisticated searching tools have to be developed to handle the computational challenges arising in new applications.

In the field of exact string matching, indexing data structures and efficient algorithms for building and preprocessing these data structures and for searching in them have been developed. Only few attempts have been made to adapt these indexing techniques to the similarity searches needed for the presented biological problems. Examples include the position tree by Martinez [12], work by Heumann [7] or the sublinear search algorithm suggested by Myers [13]. Another example is the IBM product FLASH [5]. However, most of these approaches require relatively large index structures.

In cooperation with the German Cancer Research Center in Heidelberg we created QUASAR, an algorithm designed for multiple high similarity searches in DNA databases. This algorithm is based on a filter technique introduced by Jokinen und Ukkonen [8]. Combined with a suffix array [11] this filter allows us to quickly locate all regions of the database that contain potential matches that are then examined more closely using a classical alignment algorithm. The algorithm together with the results of some "real world" experiments have been published in [2, 3]. QUASAR achieves a significant speedup over currently used search algorithms like BLAST and FASTA while requiring only a fairly small index.

The main restriction of QUASAR is that it can only guarantee successful filtration for fairly low error rates [6]. Reasonable runtimes can only be expected for error rates of up to 6–8 percent. Even though the algorithm will still find most of the hits with higher error rates, it still becomes a heuristic approach for these levels of similarity.

The reason for this lies in the nature of the filter criteriom, the q-gram Lemma. The main target of our work was to improve the filter sensitivity, i.e. to raise the allowed error percentage. A first improvement consisted in applying the idea of word-neighbourhoods to the q-gram Lemma. By using 1-neighbourhoods (with respect to the Levenshtein–distance) of substrings of the query one can increase the sensitivity of the filter phase to roughly 12 percent error rate while decreasing the speed of the search algorithm by a factor of about 2.

Further research on improving the filter led to the results on filtering with gapped q-grams [4]. In principle this result extends the approach of filtering by matching substrings of the query with the database. Instead of using substrings of a certain length one searches for so–called shapes. These shapes are basically strings that are only matched in certain positions, i.e. some positions are ignored when matched with the database. So far we have only applied them to approximate string matching that allows for character replacements (the Hamming–distance) but first experiments for the case of insertions and deletions (Levenshtein–distance) look promising. Gapped q-grams allow for higher error rates and/or filter speed and quality. The amount of improvement depends on the actual set of parameters used. Typically the amount of work for the filter step can be reduced by a factor of 50 to 100 for the same filter quality or the filter quality can be improved by two to three orders of magnitude when spending an equal amount of time for the filtration. A combination of improving filter speed and quality is also possible. The choice of the actual filter depends on the implementation of the filter and the verification phase.

One problem with gapped q-grams is the computation of the threshold value t, the minimum number of matching q-grams in a given region required for it to be a potential match. While it was trivial to do compute t for ungapped q-grams, this is not true for the gapped case. We came up with a generic method for computing t for arbitrary filters that may even be more complex than just using single gapped q-grams [9]. One could for example envision a filter that applies two different gapped q-grams simultaneously and uses 2-tuples of possible threshold values.

References

[1] S.F. Altschul, W. Gish, W. Miller, E.W. Myers, and D.J. Lipman. Basic local alignment search tool. *Journal of Molecular Biology*, 215:403–410, 1990.

- [2] S. Burkhardt, A. Crauser, P. Ferragina, H.-P. Lenhof, E. Rivals, and M. Vingron. q-gram based database searching using a suffix array (QUASAR). Technical Report MPI-I-98-024, 1998.
- [3] S. Burkhardt, A. Crauser, P. Ferragina, H.-P. Lenhof, E. Rivals, and M. Vingron. q-gram based database searching using a suffix array (QUASAR). In *Proc. of the Third Annual International Conference on Computational Molecular Biology (RECOMB 99)*, pages 77–83, 1999.
- [4] S. Burkhardt and J. Kärkkäinen. Better filtering with gapped q-grams. In *Proc. of the 12th Annual Symposium on Combinatorial Pattern Matching (CPM 2001)*, pages 73–85, 2001.

[5] A. Califano and I. Rigoutsos. FLASH: A fast look-up algorithm for string homology. In L. Hunter, D. Searls, and J. Shavlik, editors, *Proc. of the 1st International Conf. on Intelligent Systems for Molecular Biology*, pages 56–64, 1993.

[6] E.Sutinen and W. Szpankowski. On the collapse of q-gram filtration. In *Proceedings of the International Conference on FUN with Algorithms*, pages 178–193, 1998.

[7] K. Heumann. *Biologische Sequenzdatenanalyse großer Datensätze basierend auf Positionsbaumvarianten*. PhD thesis, 1997.

[8] P. Jokinen and E. Ukkonen. Two algorithms for approximate string matching in static texts. In *Proc. of the 16th Symposium on Mathematical Foundations of Computer Science*, volume 520 of *Lecture Notes in Computer Science*, pages 240–248, 1991.

- [9] J. Kärkkäinen. Computing the threshold for q-gram filters. July 2001.

[10] A. Krause and M. Vingron. A set-theoretic approach to database searching and clustering. *Bioinformatics*, 14:430–438, 1998.

[11] U. Manber and G.W. Myers. Suffix arrays: A new method for on-line string searches. In *Proc. of the first annual ACM-SIAM Symposium on Discrete Algorithms*, pages 319–327, 1990.

[12] H.M. Martinez. An efficient method for finding repeats in molecular sequences. *Nucleic Acids Research*, 11(13):4629–4634, 1983.

[13] E.W. Myers. A sublinear algorithm for approximate keyword searching. *Algorithmica*, 12(4/5):345–374, 1994.

[14] W.R. Pearson and D.J. Lipman. Improved tools for biological sequence comparison. *PNAS*, 85:2444–2448, 1988.

[15] J. Weber and G. Myers. Human whole genome shotgun sequencing. *Genome Research*, 7:401–409, 1997.

9.3 Protein Docking

Investigators: Oliver Kohlbacher, Hans-Peter Lenhof

9.3.1 Semi-flexible Docking

Docking of unbound protein structures using rigid body docking approaches suffers from a major drawback: side chains in the protein-protein interface tend to change their conformation upon binding. Hence, algorithms leaving those side chains rigid tend to produce physically impossible structures with overlapping side chains. Those conformations then cause major problems for the energetic evaluation of the tentative complex structures. We developed a new approach, which we call *semi-flexible docking*, including an additional side-chain demangling stage to our protein docking algorithm. In this stage, the side chains of the protein-protein interface of a tentative

complex structure produced by a rigid body docking algorithm are optimized with respect to their conformation to resolve overlaps. By discretizing conformational space of all side chains to a well defined set of rotamers (as first proposed by Ponder & Richards [6]), we can formulate this problem as a combinatorial optimization problem. The usual 50 – 60 side chains of a protein-protein interface give rise to 10^{50} – 10^{60} possible interface conformations. We proposed two algorithms, a fast greedy heuristic and a slower branch-and-cut algorithm solving the problem to optimality (with respect to the energy function employed). Both approaches were able to resolve the side-chain placement problem sufficiently well to allow a successful docking of unbound structures of three protease-inhibitor complexes [1, 2, 3].

9.3.2 NMR-based Protein Docking

One of the large unsolved problems in protein docking is the estimation of the free energy on binding for a tentative complex structure. We therefore explored the possibility of replacing the energy calculation, at least in part, by readily available experimental data. In collaboration with Peter Bayer (Max Planck Institute for Molecular Physiology), we developed a new approach to protein-protein docking based on Nuclear Magnetic Resonance Spectroscopy (NMR). NMR data contains a large amount of structural information. While the simplest form of a NMR spectrum, the one-dimensional ^{1}H-NMR spectrum, does not contain sufficient structural information to allow a complete structure elucidation of a protein, it usually contains enough information to allow a discrimination between true and false positive complex structures in docking. In contrast to methods for structure elucidation, we do not try to extract that structural information from the spectra directly, because this is not feasible from a single spectrum alone. Instead, we try to predict the spectrum of all tentative complex structures (based upon the chemical shift prediction for each proton) and compare these spectra to an experimentally determined ^{1}H-NMR spectrum of the complex. In an initial experiment, we predicted the complex structure of four protein-protein and protein-peptide complexes. In three of the four examples, we unambiguously ranked the true complex structure as the number one using an NMR-based scoring function, while the first approximation of the true complex structure was ranked second in the fourth example. Details of the method have been published in [5, 4, 3].

References

- [1] E. Althaus, O. Kohlbacher, H.-P. Lenhof, and P. Müller. A branch and cut algorithm for the optimal solution of the side–chain placement problem. Technical Report MPI-I-2000-1-001, Max-Planck-Institut für Informatik, Saarbrücken, 2000.
- [2] E. Althaus, O. Kohlbacher, H.-P. Lenhof, and P. Müller. A combinatorial approach to protein docking with flexible side–chains. *J. Comput. Biol.*, 2000. (submitted).
- [3] O. Kohlbacher. *New approaches to protein docking.* PhD thesis, Universität des Saarlandes, Saarbrücken, Germany, 2001.
- [4] O. Kohlbacher, A. Burchardt, A. Moll, A. Hildebrandt, P. Bayer, and H.-P. Lenhof. A NMR-spectra-based scoring function for protein docking. In Thomas Lengauer, David Sankoff, Sorin Istrail, Pavel Pevzner, and Michael Waterman, editors, *RECOMB 2001 – Proceedings of the Fifth Annual International Conference on Computational Molecular Biology*, pages 184–192. ACM press, 2001.
- [5] O. Kohlbacher, A. Burchardt, A. Moll, A. Hildebrandt, P. Bayer, and H.-P. Lenhof. Structure prediction of protein complexes by a NMR-based protein docking algorithm. *J. Biomol. NMR*, 20:15–21, 2001.

[6] J.W. Ponder and F.M. Richards. Tertiary templates for proteins – use of packing criteria in the enumeration of allowed sequences for different structural classes. *J. Mol. Biol.*, 193:775–791, 1987.

9.4 BALL: Biochemical ALgorithms Library

Investigator: Andreas Kerzmann, Oliver Kohlbacher, Hans-Peter Lenhof

Our Library BALL (Biochemical ALgorithms Library) has been released in May 2000 in a first public beta version and is available free of charge for research and teaching and (through Algorithmic Solutions Software GmbH) for commercial purposes as well. During the last two years, we spent lots of efforts on a complete overhaul of the kernel data structures, implemented copious functionality in the fields of solvation methods, force fields, and Nuclear Magnetic Resonance, and added bindings for the object-oriented script language Python [7]. Using Python, we are able to cut down turn-around times for the development of software prototypes (Python is an interpreted language and does not require the time-consuming compile and link phases of C++) and allow the user interactive manipulation of all his data structures. By employing an semi-automated translation of C++ headers to python classes, we ensure that the resulting Python code is easily translated to C++ and vice versa. We published several articles on the design of BALL, which was in 2000 awarded the "Heinz Billing Award for the Advancement of Scientific Computation" [1, 2, 5, 6, 3, 4].

References

- [1] N.P. Boghossian, O. Kohlbacher, and H.-P. Lenhof. BALL: Biochemical Algorithms Library. Technical Report MPI-I-99-1-002, Max-Planck-Institut für Informatik, Saarbrücken, 1999.
- [2] N.P. Boghossian, O. Kohlbacher, and H.-P. Lenhof. BALL: Biochemical Algorithms Library. In J.S. Vitter and C.D. Zaroliagis, editors, *Algorithm Engineering, 3rd International Workshop, WAE'99, Proceedings*, volume 1668 of *Lecture Notes in Computer Science (LNCS)*, pages 330–344. Springer, 1999.
- [3] O. Kohlbacher. BALL – a framework for rapid application development in molecular modeling. (in press).
- [4] O. Kohlbacher. *New approaches to protein docking.* PhD thesis, Universität des Saarlandes, Saarbrücken, Germany, 2001.
- [5] O. Kohlbacher and H.-P. Lenhof. Rapid software prototyping in computational molecular biology. In *Proceedings of the German Conference on Bioinformatics (GCB'99)*, 1999.
- [6] O. Kohlbacher and H.-P. Lenhof. BALL - rapid software prototyping in computational molecular biology. *Bioinformatics*, 2000.

[7] Guido van Rossum. Python version 2.1. `http://www.python.org`.

9.5 GELENA: Nonviral GEntransfer systems based on LEctin modified NAnoparticles

Investigators: Andreas Kerzmann, Oliver Kohlbacher, Hans-Peter Lenhof

Lectins are proteins which bind specifically to certain carbohydrate structures. Although their exact function in plants is still unknown, in animals the lectin-carbohydrate interaction mediates many cell-cell interactions (e. g. cell-cell-recognition and cell-adhesion). Since different tissues express various carbohydrates at differing levels of the same carbohydrate structure, lectins may be used to address a certain type of tissue specifically (drug targeting): lectin-drug-conjugates are taken up more readily by certains cells than the drug itself. Thus, lectins have been repeatedly proposed for applications like gene-transfer or even gene-therapy, enhanced drug-delivery, and drug targeting. *In-vitro* studies employing cell-culture models confirmed the usefulness of lectins in such experiments.

Unfortunately, lectins exhibited severe side-effects when used in animal studies. Designed lectins or lectin-like compounds ('lectinomimetics') possessing no such side-effects would be desirable.

Since the exact mechanism of the lectin-sugar interaction itself is still poorly understood, it is necessary to investigate the known structures and sequences of lectins and their complexes to identify amino acids necessary for specific sugar binding. A theoretical model to calculate binding free energies of sugars interacting with lectins has to be established, to compare the properties of carefully designed lectinomimetics with known lectins.

The model substance we chose for our studies is the wheat germ agglutinin, a N-acetglucosamin (GlcNAc) specific lectin which is remarkably resistant against chemical (pH changes, solvents) and physical influences (temperature, shear stress) and thus is of great interest from a pharmaceutical point of view.

In order to be able to design lectins according to our needs we analyzed the variance of the chosen lectins by comparing their sequences to homologous and functionally similar proteins. With the results we are able to identify mutable positions in the sequence which we can employ for changing the behavior and eventually the specificity of artificially designed lectins. As expected all those residue positions that are structurally important show little variance and therefore must not be changed. Others seem to be restricted to certain amino acid types, e. g. polar or aromatic ones.

Apart from developing a new theoretical model based on the Poisson-Boltzmann equation and incorporating structural information of the solvent, we developed a new energy function based on the AutoDock 3.0 energy function. With that we calculated the binding free energies of several GlcNAc derivatives with wheat germ agglutinin. These calculated binding free energies [3] correlate reasonably well with experimental data from literature [2, 1] and experiments carried out in the Department for Biopharmaceutics and Pharmaceutical Technology, Saarland University.

To our knowledge, similar correlations between experimental and calculated binding free energies have been reported only for single carbohydrate-lectin complexes, but not for a larger number of complexes. The correlations we obtained with the improved AutoDock 3.0 energy function are quite promising and will be verified with more experimental data in near future.

References

[1] G. Bains, R. T. Lee, et al. Microcalorimetric study of wheat germ agglutinin binding to N-acetylglucosamine and its oligomers. *Biochemistry*, 31(50):12624–12628, 1992.

[2] M. Monsigny and A. C. Roche. Sugar-lectin interactions: How does wheat-germ agglutinin bind sialoglyconjugates? *Eur. J. Biochem.*, 104:147–153, 1980.

- [3] D. Neumann, O. Kohlbacher, H.-P. Lenhof, and C.-M. Lehr. Lectin-sugar interaction: calculated versus experimental binding energies. In *Abstracts of papers presented at the 19th Interlec meeting, Fortaleza, Brazil*, page 126, 2001.

9.6 Metabolic and Regulatory Networks

Investigators: Hans-Peter Lenhof, Thomas Schäfer, Marite Sirava

The daily growing biochemical data databases (e.g. [2, 5]) can be used for further analysis of an organism's molecular biology and biochemistry. Bioinformatic is in charge of developing new tools and algorithms to handle the publicly available reaction sets and to search in this huge amount of data. This will – amongst other things – enable scientists to gain new insights in the biochemical world, to identify new target molecules that might help in developing new drugs and therapies, and to prevent or cure diseases. Biologists and biochemists will get tools to browse and to search the large amount of available biochemical data in an easy and comfortable way.

Our approach uses graphs to represent the metabolic and regulatory networks [7, 9] of organisms. Several other approaches where metabolic and regulatory pathways are represented by graphs or

Petri nets exist (e.g. [3, 4]). In order to develop a data model capable of storing all metabolic and regulatory pathways we developed a hierarchical class model in C++. We cooperate with the bioinformatics group at University of Manchester that developed a data model called GIMS (Genome Information Management System [8]). GIMS provides conceptual models describing genome sequence data and genome organization. We are about to compare the GIMS classes and the classes in our model in order to combine the two models.

We used the LEDA package [6] to implement our graph model, and we developed and implemented several algorithms to search certain reaction paths in the graph. To visualize the results of our algorithms in an easy and comfortable way we cooperate with the group of Prof. M. Kaufmann in Tuebingen that develops graph visualization tools.

As a first application we implemented a tool called Pathway "Generator". This tool can be used to identify unknown metabolic pathways in an organism's metabolism. It uses information about related organisms to build a set of all possible paths. Using a weighting function the algorithm ranks the generated pathways. At the moment, the weighting function is based on alignment scores [1] and phylogenetic distances. The next step will be to include more data by extending the weighting function (e.g. expression data and proteomic information) to improve our pathway reconstruction. Furthermore, we will extend our data model by including e.g. protein-protein interaction. Our next application is intended to find new target molecules for new drugs.

In order to test the "Pathway Generator" we defined several data sets from eubacteria and archaea. The program was able to identify the glycolysis pathway as well as the leucine biosynthesis pathway which could be proved with the help of known data from KEGG.

References

[1] S. F. Altschul, W. Gish, W. Miller, E. W. Myers, and D. J. Lipman. Basic local alignment search tool. *J. Mol. Biol.*, 215:403–410, 1990.

[2] A. Bairoch. The ENZYME database in 2000. *Nucleic Acids Research*, 28:304–305, 2000.

[3] T. Dandekar, S. Schuster, B. Snel, M. Huynen, and P. Bork. Pathway alignment: application to the comparative analysis of glycolytic enzymes. *Biochem J.*, 343:115–124, 1999.

[4] T. Gaasterland and E. Selkov. Reconstruction of metabolic networks using incomplete information. In *Proc Int Conf Intell Syst Mol Biol.*, volume 3, pages 127–135, 1990.

[5] M. Kanehisa and S. Goto. KEGG: Kyoto encyclopedia of genes and genomes. *Nucleic Acids Research*, 28:29–34, 2000.

[6] Kurt Mehlhorn and Stefan Näher. *LEDA : a platform for combinatorial and geometric computing.* Cambridge University Press, 1999.

[7] G. Michal, editor. *Biochemical pathways : Biochemie-Atlas.* Spektrum Akademischer Verlag, 1999.

[8] N. W. Paton, S. A. Khan, A. Hayes, F. Moussouni, A. Brass, K. Eilbeck, C. A. Goble, S. J. Hubbard, and S. G. Oliver. Conceptual modelling of genomic information. *Bioinformatics*, 16:548–557, 2000.

[9] G. N. Stephanopoulos, A. A. Aristidou, and J. Nielsen. *Metabolic engineering : principles and methodologies.* Academic Press, 1998.

10 Software Libraries – Coordinator: Kurt Mehlhorn

We continued our work on CGAL, LEDA, and BALL. ECG (computational geometry for curved objects) is a new library project. It is closely related to our work on exact geometric computation and on simulation and assembly. The work on network algorithms feeds into LEDA. The goal of the work on web-based testbeds is to ease evaluation of algorithmic software. Some of the effort on software libraries is described in other sections: LEDA-SM in Section 6.4, BALL in Section 9.4, SYLVIA in Section 8.

10.1 CGAL

We have continued development of CGAL, the Computational Geometry Algorithms Library [1] together with the other partners in the GALIA consortium (Section 14.1.2). The distinguishing features of the library are the careful and efficient treatment of robustness issues, the wide scope of the algorithms and data structures provided, and flexibility, extensibility, and ease of use. There have been three releases of the library since 1999 (2.1 in January 2000, 2.2 in October 2000, and 2.3 in August 2001) and efforts are underway to establish a company based around the library [2].

Generic Geometry Kernel

Investigators: Susan Hert, Michael Seel

Because computational geometry is used in many different application areas, one of the primary design goals of CGAL was flexibility. This flexibility is achieved following the generic programming paradigm and using templates as tools. In release 2.3 of the library, the new kernel design described in [3] has been realized. This design is completely compatible with the previous design [4], which supplied geometric objects that were each parameterized by a representation class, which was in turn parameterized by a number type. This design provided easy exchange of representation classes, was extensible, and provided limited adaptability of an existing representation class. The new design provides even more flexibility in that it allows geometric predicates as well as objects to be easily exchanged and adapted individually to users' needs.

Generic Higher-dimensional Geometry

Investigators: Kurt Mehlhorn, Stefan Schirra, Michael Seel

The first version of the package offering higher-dimensional Euclidean geometry was presented on the conference SoCG [5] and published in [6]. While CGAL evolved and matured, we transformed the kernel in correspondence with the lower dimensional kernels and integrated it into the current CGAL release [3].

Our geometric kernel consists of objects and primitives (predicates and constructions) that support the development of geometric algorithms in d-space. It consists of a generic and easy-to-use interface. Special care was taken to refine out the concepts that allow generic adaptation of the kernel from the original monolithic design, *e.g.*, the number types and their docking into the kernel functionality via a linear algebra module. To enhance usability the representation-based kernel families can at the same time be used as traits classes in the instantiation of application classes. The traits classes here encapsulate the geometric primitives that control the logical flow of algorithms. This is one general design pattern of CGAL starting with version 2.3.

Planar Nef Polyhedra

Investigators: Kurt Mehlhorn, Michael Seel

A planar Nef polyhedron is any set that can be obtained from open halfspaces by a finite number of set complement and set intersection operations. The set of Nef polyhedra is closed under the Boolean set operations. We implemented a data structure that realizes two-dimensional Nef polyhedra [7] and offers a large set of binary and unary set operations. The underlying set operations are realized by an efficient and complete algorithm for the overlay of two Nef polyhedra. The algorithm is efficient in the sense that its running time is bounded by the size of the inputs plus the size of the output times a logarithmic factor. The algorithm is complete in the sense that it can handle all inputs and requires no general position assumption. Our software module is part of CGAL and mainly based on three software components: the data structure is a selective plane map[7] whose vertices are embedded by extended points. The plane map is based on the generic CGAL halfedge data structure; the extended points are a generalization of affine points and ray based on infimaximal frames [8]. Binary operations are realized on top of a generic plane sweep module that calculates the overlay of a set of segments in the plane.

References

[1] CGAL. `http://www.cgal.org`, 2001.

[2] The Geometry Factory. `http://www.geometryfactory.com`, 2001.

- [3] S. Hert, M. Hoffmann, L. Kettner, S. Pion, and M. Seel. An adaptable and extensible geometry kernel. In *Proceedings of the 5th Workshop of Algorithmic Engineering*, 2001. to appear.

- [4] Andreas Fabri, Geert-Jan Giezeman, Lutz Kettner, Stefan Schirra, and Sven Schönherr. On the design of CGAL, the Computational Geometry Algorithms Library. *Software – Practice and Experience*, 30:1167–1202, 2000.

[5] K. Mehlhorn, M. Müller, S. Näher, S. Schirra, M. S eel, C. Uhrig, and J. Ziegler. A computational basis for higher-dimensional computational geometry a nd applications. In *Proceedings of the 13th International Annual Symposium on Computa tional Geometry (SCG-97)*, pages 254–263, Nice, France, 1997. Association of Computing Machinery (ACM), ACM Press.

[6] K. Mehlhorn, M. Müller, S. Näher, S. Schirra, M. Seel, C. Uhrig, and J. Ziegler. A computational basis for higher-dimensional computational geometry a nd applications. *Computational Geometry: Theory and Applications*, 10(4):289–304, July 1998.

- [7] M. Seel. Implementation of planar nef polyhedra. Technical Report MPI-I-2001-1-003, Max-Planck-Institut für Informatik, Saarbrücken, 2001. to be published.

- [8] M. Seel and K. Mehlhorn. Infimaximal Frames: a framework to make lines look like segments. Technical Report MPI-I-2000-1-005, Max-Planck-Institut für Informatik, Saarbrücken, 2000.

10.2 LEDA

Investigators: Kurt Mehlhorn, Michael Seel

We shortly describe the changes and developments concerning LEDA, the library of efficient data types and algorithms. LEDA is a collaboration between the research group of Kurt Mehlhorn at the MPII, the research group of Stefan Naeher at the university of Trier (previously Halle) and the

[7] a straight-line embedded bidirected graph enriched by faces of multiple boundary cycles.

external commercial partner Algorithmic Solutions Software GmbH. Major topics of our summary are documentation, distribution, and functionality.

The LEDA book of Kurt Mehlhorn and Stefan Näher appeared at the end of the year 1999. It treats the architecture, the implementation, and the use of the LEDA system. The presentation discusses the functionality of many data types and algorithms available in the system and elaborates on their implementation. The book is useful to different types of readers: readers with a general interest in combinatorial and geometric computing, casual users of LEDA, intensive users of LEDA, library designers and software engineers, and students taking an algorithms course.

In February 2001, Algorithmic Solutions Software GmbH (AS) became the sole distributor of LEDA. AS took over the responsibility for maintaining and developing LEDA. The research groups of Stefan Naeher and Kurt Mehlhorn partly migrated to new fields of research but still intent to disseminate the practical results via LEDA: as core functionality, extension packages or friend packages (depending on the scope of the research).

This change had several reasons. Maintenance of LEDA for a wide variety of platforms and compilers bound too much manpower in the research groups. Several key persons of the LEDA project left the research groups and moved to AS and thereby the expertise for maintaining and developing LEDA is now at AS. Also, AS will add functionality that will make LEDA even more useful, but is of limited scientific interest. One outcome of AS's support is the offer of a LEDA tutorial, that teaches introductory usage of LEDA for beginners.

We now come to the facts concerning the library's functionality. In the whole time period since 1999 LEDA evolved concerning platform and compiler support, new research results, and algorithmic revisions. There have been several releases (versions 3.8, 4.0, 4.1, and 4.2) including many functional updates. The current LEDA version 4.3 support all up-to-date C++ compilers, e.g., GNU gcc 3.0, MipsPro CC 7.3, SunPro CC 5.2, Microsoft VC++ 6.0, Borland bcc 5.5. One research topic of permanent attention have been the number types. There have been several improvements concerning, core assembler optimization, integral division [4, 2], and the integration of new separation bounds [3]. Of course, also the graphical interfaces like the window, graphwin and geowin packages underwent a general improvement.

Concerning graphs and graph algorithms there have been major extensions and revisions. Some new algorithms have been added, e.g., euler tours, transitive closure, transitive reduction. The graph package was extended. There's now a semi-dynamic graph representation. This representation needs only about half of the space of the fully-dynamic standard graph data structure and can significantly reduce the running time of graph algorithms (e.g., see some results concerning maxflow at http://www.informatik.uni-trier.de/ naeher/maxflow.html). Moreover several core modules have been revised, e.g. max flow and maximum weight matching in general and bipartite graphs and minimum spanning trees.

The geometric part of LEDA also contains new functionality. An excerpt of the topics are: planar constrained triangulations (of segments, polygons, embedded graphs), planar contour reconstruction (e.g., crust [1]), convex partitioning of polygons, minkowski sum and difference of polygons etc.

Finally general remarkable extension are: partially persistent dictionaries, tools for numerical analysis, a portable socket communication module, several updates in graph drawing.

An extensive list of revision can be found in the CHANGES file that is included in each LEDA package.

References

[1] N. Amenta, M. Bern, and D. Eppstein. The crust and the β-skeleton: Combinatorial curve reconst ruc tion. *Graphical Models and Image Processing*, 60:125–135, 1998.

[2] A. Borodin and R. Moenck. Fast modular transforms via division. In *Proceedings of the 13th Annual IEEE Symposium on Switching and Automata Theory*, pages 90–96, 1972.

[3] C. Burnikel, S. Funke, K. Mehlhorn, S. Schirra, and S. Schmitt. A separation bound for real algebraic expressions. In Friedhelm Meyer auf der Heide, editor, *Proceedings of the 9th European Symposium on Algorithms (ESA-01)*, volume 2161 of *Lecture Notes in Computer Science*, pages 326–337, Aarhus, Denmark, August 2001. Springer.

[4] C. Burnikel and J. Ziegler. Fast recursive division. Technical Report MPI-I-98-1-022, Max-Planck-Institut für Informatik, Saarbrücken, 1998.

10.3 Effective Computational Geometry for Curves and Surfaces

We started a new project on computational geometry for curved objects. The following two paragraphs are part of an announcement to Master and PhD students and outlines our goals.

"The algorithms group of the Max-Planck-Institute for Computer Science and Raimund Seidel's group in the Computer Science department has made important contributions to computational geometry over the past years. We invented algorithms, contributed decisively to the theoretical foundations of implementing geometric algorithms, and built systems (LEDA, CGAL, SILVIA) that make the algorithms and data structures of computational geometry available. A unique feature of our systems is their reliability. The systems are being used worldwide (for academic as well as commercial purposes). They are marketed by Algorithmic Solutions GmbH."

"Our current systems are mainly designed for linear objects (planar objects composed of points and line segments, three-dimensional objects that can be described by points, line segments, and triangles). Curved objects are approximated by linear objects. **Our goal for the next years is to handle curved objects directly.** We have won substantial inside and outside resources (MPG, DFG and EU) for the undertaking. The undertaking is a real challenge as it will force us to combine knowledge from many sources (computational geometry, combinatorial algorithms, numerical analysis, computer algebra, software engineering) and to make advances in all these areas. We also expect close collaboration with the members of the mesh processing group of AG4. We are in a unique position to attack the challenge, we know the available theory and its shortcomings, we have the best systems to build on, and we have prototypical implementations for some of the tasks to be tackled." See Section 8.5.

"Stefan Funke, Nicola Geismann, Susan Hert, Lutz Kettner, Christian Lennerz, Kurt Mehlhorn, Edgar Ramos, Susanne Schmitt, Elmar Schömer, and Thomas Warken (see `http://www.mpi-sb.mpg.de/units/ag1/people.html`) are committed to the project. The project has room for at least 20 students at all levels (programming projects, Diplomarbeit, or PhD thesis). In order to prepare you for the project we will offer the following courses, seminars and FOPRAs in the next two semesters (winter term 01/02 and summer term 02)."

See Section 14.1.3 for further information.

10.4 Exact Geometric Computation

Geometric algorithms are usually designed and proven to be correct in a computational model that assumes exact computation over the real numbers. Since no computer provides exact arithmetic on real numbers in hardware, programmers must find some substitution when implementing these

algorithms. Quite commonly, they resort to fast finite precision arithmetic due to its support by hard- and software as well as its convenient use. For some problems and restricted sets of input data, this approach works well, but in many implementations the effects of squeezing the infinite set of real numbers into the finite set of floating-point numbers can cause catastrophic errors in practice. The problem are geometric tests which affect the flow of control of the algorithm – so-called *geometric predicates*. Due to round-off errors, the results of different predicates might contradict each other leaving the algorithm in an undefined state. In the best case, the algorithm still produces some usable result, but most of the time, it produces completely inconsistent results, crashes or loops.

There are two approaches for solving this precision problem. The first is to change the model of computation: design *robust* algorithms that can deal with inconsistencies incurred by incorrect predicate decisions. Unfortunately, there is no generic way to derive such algorithms, and only for a small number of (easy) problems, such algorithms have been developed. The second approach is the *exact computation paradigm* [4]; it advocates to guarantee correctness of the implementation by ensuring that every single predicate is evaluated correctly. This is achieved by providing software-based exact number types and hence exact predicate implementations.

Exact arithmetic is not hard to implement for integer/rational arithmetic, but for arithmetic on the real numbers, it is not so clear how to guarantee exact results. Furthermore, the naive use of these exact number types and predicates has its cost, which is considerably higher then pure floating-point arithmetic; an overhead of a factor of 10-100 can be expected.

Many of our researchers have worked on these important problems over the last few years and produced quite a number of results.

Structural Filtering

Investigators: Stefan Funke, Kurt Mehlhorn

The evaluation of geometric predicates, such as the incircle or the side-of-circle predicate, amounts to the computation of the sign of an expression. One way to speed up exact sign determination is the use of so-called floating-point filters. The idea is first to compute a floating-point approximation together with an error bound for the maximal deviation from the true value. If the error bound is smaller than the absolute value of the approximation, approximation and exact value have the same sign and hence the sign of the approximation may be returned. Only if this fails, exact arithmetic has to be used.

In [3] we have developed the concept of *structural filtering* which is a generalization of floating-point filtering. Structural filtering can reduce the overhead compared to pure floating-point arithmetic further by allowing some predicates to err, without sacrificing the guarantee for an exact outcome. Structural filtering views the execution of an algorithm as a sequence of steps and applies filtering at the level of steps. A step can be anything between a simple predicate and the execution of the whole algorithm.

LOOK – a Lazy Object-Oriented Kernel for Geometric Computation

Investigators: Stefan Funke, Kurt Mehlhorn

In [2] we have developed the design of a geometric kernel called LOOK (Lazy Object Oriented Kernel for geometric computation), which is the first of its kind that makes use of filtering techniques not only on the predicate level but also on the level of geometric constructions. The main idea is to

represent geometric objects not by their coordinates, but by the geometric operation that produced them and compute coordinate representations only on demand in a lazy-fashion.

New LEDA number types

Investigators: Christoph Burnikel, Stefan Funke, Kurt Mehlhorn, Stefan Schirra, Susanne Schmitt

We have reimplemented all number types in LEDA, i.e., `integer`, `rational`, `real`, and `bigfloat` which allow exact computation on integer, rational, algebraic numbers and arbitrary precision calculations with floating-point numbers respectively. Compared with our old implementation we have gained in running time considerably. For example, the new LEDA `real` datatype is up to four times faster than the old implementation. With our new separation bound ([1], also see in the Geometry section) we are even up to 12 times faster. The `integer` datatype has also improved and can be regarded as one of the fastest datatype for exact integer computation available. We have also worked on the floating-point filter techniques used within the `real` datatype and geometric predicates, which also improves the running-time for many 'easy' instances.

References

- [1] C. Burnikel, S. Funke, K. Mehlhorn, S. Schirra, and S. Schmitt. A separation bound for real algebraic expressions. In Friedhelm Meyer auf der Heide, editor, *Proc. of the 9th Europ. Symp. on Algorithms (ESA-01)*, volume 2161 of *Lecture Notes in Computer Science*, pages 326–337, Aarhus, Denmark, August 2001. Springer.
- [2] S. Funke and K. Mehlhorn. LOOK - a lazy object-oriented kernel for geometric computation. In *Proc. of the 16th Ann. Symp. on Computational Geometry (SCG-00)*, pages 156–165, June 2000.
- [3] S. Funke, K. Mehlhorn, and S. Näher. Structural filtering: a paradigm for efficient and exact geometric programs. In *Abstracts for the 11th Canadian Conference on Computational Geometry (CCCG-99)*, pages 39–42, Vancouver, Canada, August 1999.

[4] C. K. Yap and T. Dubé. The exact computation paradigm. In D.-Z. Du and F. K. Hwang, editors, *Computing in Euclidean Geometry*, volume 4 of *Lecture Notes Series on Computing*, pages 452–492. World Scientific, Singapore, 2nd edition, 1995.

10.5 Network Algorithms

General Weighted Matching Algorithm

Investigators: Kurt Mehlhorn, Guido Schäfer

We implemented an $O(nm \log n)$ algorithm for the weighted matching problem in general graphs. The algorithm is a variant of the algorithm described by Galil, Micali and Gabow [3] and requires the use of concatenable priority queues. No previous implementation achieved a worst-case guarantee of $O(nm \log n)$.

The new code (which is contained in LEDA-4.2 or higher) is highly efficient in practice. In [5] we compared our new code to the experimentally fastest implementation (called Blossom IV) due to Cook and Rohe [2]. Blossom IV requires only very simple data structures. This experimental comparison showed that our matching algorithm is considerably faster on almost all instances.

We experimented with the following kinds of graph instances: Delaunay graphs, sparse random graphs, dense random graphs, (complete) geometric graphs, chains and an instance, which occurred in VLSI-Design. Our code is substantially faster (up to a factor of 1000) on all those instances except for the (complete) geometric instances. This is due to the fact that Blossom IV provides a

refined strategy, called price-and-repair, which leads to a significant improvement for this class of instances.

Bipartite Weighted Matching Algorithm

Investigators: Kurt Mehlhorn, Guido Schäfer

A new heuristic was implemented for the weighted matching algorithm in bipartite graphs. The necessary modifications are easy to implement. The algorithm still guarantees an asymptotic running time of $O(n(m + n \log n))$. Moreover, the heuristic can only help but never harm.

Experiments were performed on various random instances. The experimental findings indicate a running time improvement of up to a factor of 10.

In [6] we present a partial analysis for random graphs which gives some theoretical support to our believe that the heuristic helps in practice. The improved code is part of LEDA-4.3.

Min-Cost Flow Implementations

Investigators: Kurt Mehlhorn, Guido Schäfer, Andreas Bramer, Oliver Zlotowski

We closely inspected the following algorithms to solve the min-cost flow problem: the capacity scaling algorithm, the cost scaling algorithm (with heuristics) and the network-simplex algorithm (see e.g. [1] for more details).

The capacity scaling algorithm could be improved by some (small) constant factor using the same idea as described in [6]. For further improvement the algorithm seems to be too rigid.

The outcome of inspecting the cost scaling algorithm on the other hand looks very promising. An improvement by some magnitude is already achieved and we still believe that this can be improved further. The algorithm is highly flexible and many efficient heuristics are known (see e.g. [4]).

We finished a first implementation of the network-simplex algorithm. Preliminary experiments on random graphs look promising.

References

[1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network flows : theory, algorithms and applications.* Prentice Hall, 1993.

[2] W. Cook and A. Rohe. Computing minimum–weight perfect matchings. Technical Report 97863, Forschungsinstitut für Diskrete Mathematik, Universität Bonn, 1997.

[3] Z. Galil, S. Micali, and H. N. Gabow. An $O(EV \log V)$ algorithm for finding a maximal weighted matching in general graphs. *SIAM J. Computing*, 15:120–130, 1986.

[4] A. V. Goldberg. An efficient implementation of a scaling minimum-cost flow algorithm. *Journal of Algorithms*, 22(1):1–29, 1997.

- [5] K. Mehlhorn and G. Schäfer. Implementation of $O(nm \log n)$ weighted matchings in general graphs: The power of data structures. In *Workshop on Algorithm Engineering (WAE)*, Lecture Notes in Computer Science, to appear, 2000.

- [6] K. Mehlhorn and G. Schäfer. A heuristic for Dijkstra's algorithm with many targets and its use in weighted matching algorithms. In *Proceedings of the 9th Annual European Symposium on Algorithms (ESA 2001)*, Lecture Notes in Computer Science, to appear. Springer, 2001.

10.6 Experimental Algorithmics

Web-Based Testbed for Experimental Graph Algorithmics

Investigators: Kurt Mehlhorn, Guido Schäfer

A web-site (see `www.mpi-sb.mpg.de/~schaefer/MLLB/index.html`) functioning as a testbed for various graph algorithms has been set-up.

The web-site can be used to submit problem instances as well as problem generators for the following graph algorithms: all variants of matching algorithms (i.e. bipartite cardinality matching, bipartite weighted matching, general cardinality matching and general weighted matching) and a max-flow algorithm. Further graph algorithms, such as the min-cost flow algorithm, will be added in the near future.

The major motivation behind setting-up this web-site is to establish a set of both benchmark instances and problem generators for various graph algorithms. Moreover, the need for making graph algorithms more accessible to the public is achieved by an easy-to-use web-interface.

Submitted problem instances are run on a compute server of the MPII and the result is reported back to the sender via e-mail. The web-site thereby keeps track of the encountered worst-case instances for each algorithm. One may download these instances from the web-site.

Additionally, one may submit worst-case generators (as C++ source files). A submitted generator is accepted only if it follows a certain specification, which is given on the web-site, and if it passes a functionality test (which for the time being and probably also in the future cannot be done automatically). Each accepted problem generator is added to the web-site and can be downloaded.

11 Journal and Conference Activities

11.1 Editorial Positions

Kurt Mehlhorn has been an editor of
Algorithmica (since 1985),
Computational Geometry: Theory and Applications (since 1990),
Information and Computation (since 1985),
International Journal of Computational Geometry & Applications (since 1990),
Inernational Journal of Discrete and Computational Geometry (since 1988), and
SIAM Journal on Computing (since 1988).

Stefan Schirra has been
(sub-)editor of *Journal of Discrete Algorithms,*
co-editor of the *International Journal of Computational Geometry & Applications* special issue on the *16th Annual ACM Symposium on Computational Geometry, 2000.*

11.2 Conference and Workshop Activities

11.2.1 Membership in Program Committees

Panagiota Fatourou:
15th International Symposium on Distributed Computing, Lisbon, Portugal, October 2001
14th International Symposium on Distributed Computing, Toledo, Spain, October 2000
7th International Conference on High Performance Computing, Bangalore, India, December 2000

Hans-Peter Lenhof:
German Conference on Bioinformatics (GCB), Heidelberg, October 2000
German Conference on Bioinformatics (GCB), Braunschweig, October 2001
French Conference on Bioinformatics (JOBIM), Montpellier, May 2000
5th Annual International Conference on Computational Molecular Biology (RECOMB), Montreal, April 2001

Kurt Mehlhorn:
2nd International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 99), Berkeley, August 1999
5th Annual International Computing and Combinatorics Conference (COCOON 99), Tokio, July 1999
11th Annual International Symposium on Algorithms And Computation (ISAAC 00), Taipei, December 2000
33th Annual ACM Symposium on the Theory of Computing (STOC 01), Creta, July 2001
17th Annual ACM Symposium on Computational Geometry (SoCG 01), Medford, June 2001

Edgar A. Ramos:
17th Annual ACM Symposium on Computational Geometry (SoCG 01), Medford, June 2001

Peter Sanders:
28th International Colloquium on Automata, Languages and Programming (ICALP), Crete, July 2001

Stefan Schirra:

16th Annual ACM Symposium on Computational Geometry (SoCG 00), Hong-Kong, June 2000

5th Workshop on Algorithm Engineering (WAE), Aarhus, August 2001.

Jop Sibeyn:

EuroPar '99, Toulouse, August 1999.

11.2.2 Membership in Organizing Committees

Rudolf Fleischer:

Dagstuhl seminar "Experimental Algorithmics," Sep 11-15, 2000. More information in `http://www.dagstuhl.de/DATA/Reports/00371/`

Peter Sanders and Uli Meyer:

GI-Dagstuhl-Forschungsseminar "Algorithms for Memory Hierarchies." March 2002. More information in `http://www.mpi-sb.mpg.de/ sanders/gisem/`

Jop Sibeyn:

CONF 2000: The jointly organized conferences ESA 2000, WAE 2000 and APPROX 2000, which took place in Saarbrücken, Germany, from September 5 until September 8, 2000. More information in `http://www.mpi-sb.mpg.de/ conf2000`

Hannah Bast, Berthold Vöcking:

Summer School - 2nd Max-Planck Advanced Course on the Foundations of Computer Science (ADFOCS 2001), Max-Planck-Institut für Informatik, Saarbrücken, September 2001. More information in `http://www.mpi-sb.mpg.de/conferences/adfocs-01`

Anil Kumar, Panagiota Fatourou:

Summer School - 1st Max-Planck Advanced Course on the Foundations of Computer Science (ADFOCS 2000), Max-Planck-Institut für Informatik, Saarbrücken, September 2000. More information in `http://www.mpi-sb.mpg.de/conferences/adfocs00`

12 Dissertations, Habilitations, and Offers for Faculty Positions

12.1 Dissertations

Completed:

Knut Reinert: A Polyhedral Approach to Sequence Alignment Problems, August 1999.

Peter Müller: Parallel Melecular Dynamics Simulations for Synthetic Polymers, November 1999.

Hannah Bast: On the Tradeoff between Load Balance and Overhead in Dynamic Scheduling, February 2000.

Thomas Schilz: Effiziente Algorithmen fuer das verteilte Rechnen auf Workstation-Clustern, May 2000.

Jordan Gergov: Algorithms for Interval Coloring, Geometric Packing and Memory Optimization, November 2000.

Oliver Kohlbacher: New Approaches to Protein Docking, January 2001.

Andreas Crauser: LEDA-SM: External Memory Algorithms and Data Structures in Theory and Practice, March 2001.

Ernst Althaus: Curve Reconstruction and the Traveling Salesman Problem, April 2001.

Volker Priebe: Average-case Complexity of Shortest-Paths Problems, June 2001.

Mark Ziegelmann: Constrained Shortest Path Problems and Related Problems, August 2001.

Stefan Funke: Combinatorial Curve Reconstruction and the Efficient Exact Implementation of Geometric Algorithms, August 2001.

Piotr Krysta: Approximation Algorithms for Combinatorial Optimization Problems in Graph Coloring and Network Design, August 2001.

In preparation:

Michael Seel: Planar Nef Polyhedra and Generic Higher-Dimensional Geometry.

Uli Meyer: Design and Analysis of Shortest Path Algorithms: Sequential, Parallel and External.

Sven Thiel: Narrowing Algorithms for Constraint Programming.

Marite Sirava: Modeling and Analysis of Metabolic and Regulatory Pathways.

Andreas Kerzmann: Solvation: Enhancing the Energetic Evaluation of Protein-Docking.

Christian Lennerz: Effiziente Abstandsbestimmung für Objekte mit gekrümmten Oberflächen.

Thomas Warken: Dynamiksimulation für starre Körper.

Naveen Sivadasan: Randomized Algorithms on Networks.

Guido Schäfer: Matching and Network Algorithms.

Thomas Schäfer: Modeling and Analysis of Metabolic and Regulatory Pathways.

Tobias Polzin: Algorithmische Behandlung NP-schwerer kombinatorischer Optimierungsprobleme.

Rahul Ray: Surface Approximation and Reconstruction.

12.2 Habilitations

Completed:

Hans-Peter Lenhof, July 1999

Stefan Schirra, July 1999

Peter Sanders, December 2000

On-going:

Berthold Vöcking

12.3 Offers for Faculty Positions

Susanne Albers:
University of Dortmund, 1999
University of Freiburg, 2001

Rudolf Fleischer:
Hong-Kong University of Science and Technology, Hong-Kong, 2000

Panagiota Fatourou:
University of Ioannina, Greece, 2001

Klaus Jansen:
University of Kiel, 2000

Hans-Peter Lenhof:
University of Bielefeld, 2000
Ludwig Maximillian University Munich, 2000
University of Saarland, 2000

Petra Mutzel:
Technische Universität Wien, 1999

Peter Sanders:
University of Giesen, 2001

Jop Sibeyn:
University of Leipzig, 2001
University of Halle, 2001

Roberto Solis-Oba:
University of Western Ontario, 2000

13 Teaching Activities

13.1 Courses

The group contributes intensively to the curriculum of the Department of Computer Science at the Universität des Saarlandes. We teach core courses (like "Software Design Praktikum", "Grundlagen zu Datenstrukturen und Algorithmen", etc) and specialized courses. The details follow.

Winter Semester 1999/2000

LECTURES:
Bioinformatik (H.-P. Lenhof)
Software Design Praktikum (S. Schirra)
SEMINARS:
Parallele Algorithmen (J. Sibeyn)

Summer Semester 2000

LECTURES:
Grundlagen zu Datenstrukturen und Algorithmen (K. Mehlhorn, P. Sanders)
Software Design Praktikum (H.-P. Lenhof, J. Sibeyn)

Winter Semester 2000/2001

LECTURES:
Grundlagen zu Datenstrukturen und Algorithmen (H. Bast, E. Schömer)
Data Structures and Algorithms (K. Mehlhorn, P. Sanders)
Advanced Algorithms and Data Structures (P. Fatourou, U. Meyer, B. Vöcking)
Geometry of Mesh Generation (E. A. Ramos)

Summer Semester 2001

LECTURES:
Informatik für Hörer aller Fakultäten (H.-P. Lenhof)
Optimization (K. Mehlhorn, E. A. Ramos)
Grundlagen zu Datenstrukturen und Algorithmen (E. Schömer, S. Schmitt)
Randomized Algorithms (B. Vöcking)
SEMINARS:
Algorithmen für große Datenmengen (P. Sanders, U. Meyer)

Our group offers continually an advance course, called "Selected Topics in Algorithms", on various (advanced) topics in algorithms and complexity. This course is actually a sequence of mini-courses at a graduatelevel; each mini-course is taught by a set (usually singleton) of instructors which are group members and/or visitors. The advance course is mainly intented to our PhD students, but it is also attended by many members of the group.

13.2 Diploma Theses

During the last two years, the following diploma theses have been completed under guidance of members of our group.

Jung-Bae Son: *Berechnung konvexer Hüllen in erwarteter Linearzeit,* April 1999.

Ulrich Clanget: *Plattformunabhaengiges Toolkit fuer Grafische Benutzerschnittstellen (GUI),* May 1999.

Carsten Gutwenger: *Design und Implementierung einer Algorithmen-Bibliothek zum Zeichnen von Graphen,* May 1999.

Sven Thiel: *Persistente Suchbaeume,* May 1999.

Peter Hachenberger: *Dynamic shortest Path Algorithms for Planar and Outerplanar,* August 1999.

Boris Koldehofe: *Animation and Analysis of Distributed Algorithms,* August 1999.

Joainnis Potikoglou: *Parallele Algorithmen zur Lösung des Single-Source Shortest Path Problems,* October 1999.

Holger Sabo: *Ein schnelles Verfahren zur exakten Berechnung von Skalarprodukten,* November 1999.

Frank Hammerschmidt: *Konzeption und Entwicklung des streng strukturierten und erweiterbaren Datentyps html page fuer LEDA,* November 1999.

Jack Hazboun: *Molekuel-Docking-Reaktionen,* November 1999.

Thomas Buchheit: *Oberflaechenkonstruktion mithilfe von Voronoi-Diagrammen,* December 1999.

Oliver Ashoff: *Algorithmen der modularen Arithmetik,* January 2000.

Christian Behrens: *Solid Modeling durch Boolesche Operationen,* January 2000.

Sebastian Braun: *Entwicklung und Test von fuer statische Graphen optimierte Datentypen,* January 2000.

Klaus Brengel: *Externe Prioritaetswarteschlangen,* January, 2000.

Stefan Thome: *Erweiterung des hierarchischen Zeichenmodells in ArchEd,* January 2000.

Peter Dickert: *Berechnung von Einfach- und Zweifach-zusammenhangskomponenten auf einer PRAM,* February 2000.

Andreas Kerzmann: *Zwangsbasierte Dynamiksimulationen im R2 unter Verwendung von Kreiskanten,* February 2000.

Oliver Lambert: *Parallele und externe Implementierung eines Zusammenhangskomponentenalgorithmus,* February 2000.

Guido Schäfer: *Weighted Matchings in General Graphs,* June 2000.

Anne Stephan: *Implementierung eines heuristischen und eines exakten Ansatzes zur Berechnung minimaler Steinerbaeume in Netzwerken,* June 2000.

Gerhard Trippen: *Das Lokalisierungsproblem für mobile Roboter,* 2000.

Anastasios Semeloglou: Kuerzeste Wege in planare Graphen. Implementierung eines parallelen Algorithmus mithilfe von PAD, July 2000.

Dirk Luettke-Huettmann: Knickminimales Zeichnen 4-planarer Clustergraphen, July 2000.

Rene Beier: Eine Heuristik fuer das Gossiping-Problem, August 2000.

Markus Fries: Implementation of a Dynamic Convex Hall Algorithm in Arbitrary Dimension, November 2000.

Karsten Klein: Flussbasierte orthogonale Zeichenverfahren fuer Graphen mit variablen Knotengrössen, November 2000.

Igor Pouchkarev: ArchEd: Ein interaktives Zeichensystem fuer die Archaeologie, October 2000.

Michaela Wahl: Ein neuer Algorithmus fuer Online-Scheduling, October 2000.

Thomas Schanne: Visualisierung der Dialektgeographie und Topographie im deutschen Dialektraum, March 2001.

Georg Lambert: Ein Parametric Search Toolkit, June 2001.

Oliver Pabst: 'angela' Ein modularer Grapheneditor zum interaktiven Zeichnen, July 2001.

Ralph Schulte: Exaktes Lösen von linearen Programmen, July 2001.

Christan Fink: Oberflächenrekonstruktion von planaren Konturen, August 2001.

14 Grants and Cooperations

There is cooperation within the institute and with colleagues in the computer science and computational linguistics departments in the University of Saarland. Kurt Mehlhorn cooperates with Sherif Ghali from AG4, and Stefan Funke and Edgar Ramos' research on surface reconstruction has benefited from discussions with Leif Kobbelt from AG4. Work within the Effective Computational Geometry project has already triggered several discussions between the AG1 computational geometry and the AG4 mesh processing groups. Elmar Schömer cooperates with Nicola Geismann from Prof. Raimund Seidel's group in CS; Kurt Mehlhorn and Sven Thiel cooperate with Denys Duchier and Joachim Niehren from Prof. Gert Smolka's group in CS, and with Alexander Koller from the CHORUS group in CL. Berthold Vöcking cooperates with Prof. Anja Feldmann's group in CS.

Hans-Peter Lenhof and his Bioinformatics group cooperate in several projects: with the Theoretische Bioinformatik Gruppe of the Deutsches Krebsforschungszentrum (DKFZ) in Heidelberg (Dr. Martin Vingron), with the Max-Planck-Institut für Molekulare Physiologie in Dortmung (Dr. Axel Scheidig), with the Institut für Biopharmazie und Pharmazeutische Technologie der Universität des Saarlandes (Prof. Dr. Claus-Michael Lehr), with the Institut für Neue Materialien in Saarbrücken (Prof. Dr. Helmut Schmidt), with the Institut für Humangenetik der Universität des Saarlandes (Prof. Dr. Eckhart Meese), with the University of Tübingen (Prof. Dr. Michael Kaufmann), with the Bioinformatics group of the University of Manchester. and with the company Across Barriers GmbH in Saarbrücken (Dr. Ellen Haltner). Hans-Peter Lenhof has built a strong research program in computational biology with strong links to our work in combinatorial optimization (mainly Integer Linear Programming). He and his group have moved to the CS department. We are going to continue the cooperation and we look forward to the cooperation with Thomas Lengauer's group.

Some other cooperations:

LEDA: We cooperate with Stefan Näher's group in Trier.

Constraint Programming: We cooperate with the computer linguists and constraint programmers in Saarbrücken and the constraint programmers at SICS.

SCIL: We cooperate with Michael Jünger and Mathias Elz in Cologne, Alexander Bockmayr in Nancy, and Thomas Kasper at SAP. Alexander and Thomas are former members of AG2.

Program Verification: We cooperate with David Basin in Freiburg, a former member of AG2.

At an institutional level we are involved in several research projects with various sources of funds. The projects are funded by the European Union: GALIA, ALCOM-FT, ECG (Section 14.1); the German government through DFG (Section 14.2) and BMBF (Section 14.3), and by Industry (Section 14.4). The details are given in the rest of this section.

14.1 Projects funded by the European Union

14.1.1 ALCOM-FT

The ALCOM-FT project is a joint effort between ten of the leading groups in algorithms research in Europe. The aim of the project is to discover new algorithmic concepts, identify key algorithmic problems in important applications, and contribute to the accelerated transfer of advanced algorithmic techniques into commercial systems. ALCOM-FT works on five work packages. MPII is the work package leader in a work package on Experimental Algorithmics and contributes many results

on the work packages Massive Data Sets, Networks and Communication, and Generic Methods. All in all, 24 of our publications of the last year have been put in the ALCOM-FT report database.

The project takes place from June 2000 to June 2003. It is supported by the European Commision under the FET part of the Information Society Technologies programme of the Fifth Framework, as project number IST-1999-14186.

The partners and group leaders in the project are:

BRICS, Aarhus, Denmark *(Prof. E. Meineche Schmidt)*
Polytechnic University of Catalunya, Barcelona, Spain *(Prof. J. Diaz)*
University of Cologne, Germany *(Prof. M Jünger)*
INRIA Rocquencourt, France *(Prof. P. Flajolet)*
University of Paderborn, Germany *(Prof. B. Monien, Prof. F. Meyer auf der Heide)*
Computer Technology Institute, Patras, Greece *(Prof. P. Spirakis)*
University of Rome "La Sapienza", Italy *(Prof. G. Ausiello)*
Utrecht University, The Netherlands *(Prof. J. van Leeuwen)*
University of Warwick, United Kingdom *(Prof. M. Paterson)*

More information in `http://www.brics.dk/ALCOM-FT/`

14.1.2 GALIA

The goal of the ESPRIT Long Term Research project GALIA (Geometric ALgorithms for Industrial Applications) was to make the most important of solutions and methods developed in the field of computational geometry available to users in industry and academia in a software library. This software library is called CGAL (Computational Geometry Algorithms Library; `http:www.cgal.org`). The work on CGAL started in ESPRIT LTR project CGAL (project number 21957) and continued with ESPRIT LTR project GALIA (project number 28155). The Max-Planck-Institut für Informatik was prime contractor for GALIA.

The partners and group leaders of the GALIA project were:

Max-Planck-Institut für Informatik, Saarbrücken, Germany *(Prof. K. Mehlhorn, Dr. S. Schirra)*
ETH, Zürich, Switzerland, *(Prof. E. Welzl, Prof. P. Widmayer)*
Freie Universität Berlin, Germany *(Prof. H. Alt)*
INRIA Sophia-Antipolis, France *(Dr. J.-D. Boissonnat, Dr. M. Y vinec)*
Martin-Luther-Universität Halle, Germany *(Prof. S. Näher)*
Tel Aviv University, Israel *(Prof. D. Halperin)*
Utrecht University, The Netherlands *(Prof. M. Overmars, Dr. R. Veltkamp)*

Though the GALIA project ended in September 2000, the consortium continues to develop, promote, and distribute the library to an ever-increasing community of users.

14.1.3 ECG

The project "Effective Computational Geometry for Curves and Surfaces" is a join effort together with other five European research groups and it aims to advance the handling of non linear objects in geometric algorithms, both in theory and practice. This requires interdisciplinary cooperation of computational geometry with computer algebra and numerical analysis, as well as with software engineering and computer aided design. The project is split into four work areas: Geometric algorithms for curves and surfaces, algebraic issues, robustness issues, and approximation with

topological and geometric guarantees. Each work area will produce three types of deliverables: reports on theoretical advances, software prototypes and CGAL extension packages.

The project takes place between May 2001 and April 2004. It is supported by the European Commision under the FET part of the Information Society Technologies programme of the Fifth Framework, as project number IST-2000-26473. The EU grant supports a researcher and Ph.D. student for the three years, and in addition to this, the Institute is committing a workforce about twice as large spread between the subgroups "Computational Geometry", "Assembly and Simulation" and "Software Libraries".

The partners and group leaders of the ECG project are.

INRIA Sophia-Antipolis, France *(Dr. J.-D. Boissonnat, Dr. B. Mourrain, Dr. M. Teillaud)*
ETH, Zürich, Switzerland, *(B. Gärtner, Prof. E. Welzl, Prof. P. Widmayer)*
Freie Universität Berlin, Germany *(Prof. H. Alt, Prof. G. Rote)*
Rijksuniversiteit Groningen, The Netherlands *(Prof. G. Vegter)*
Max-Planck-Institut für Informatik, Saarbrücken, Germany *(Prof. K. Mehlhorn, Dr. E. Ramos)*
Tel Aviv University, Israel *(Prof. D. Halperin)*

INRIA is the coordinating site.

14.2 Projects funded by DFG

Four projects are funded by the German National Science Foundation (DFG = Deutsche Forschungsgemeinschaft).

14.2.1 Virtual Biolab

The project "Virtual Biolab" was one of the five winners of the "Initiative Bioinformatics" of the DFG and has been awarded a five million DM grant for the first two years and can potentially receive another five million DM for the following three years. The project is a common proposal of the Max Planck Institute for Computer Science (MPI), the University of Saarland, and the German Research Center for Artificial Intelligence GmbH (DFKI), under the scientific leadership of the Bioinformatics group of AG1. The objective is to construct "a virtual bio-laboratory, with which the number of laboratory experiments could be reduced to a minimum. The funds awarded by the DFG are to be used to set up a competency centre at which expertise from the life sciences and informatics is to be concentrated. The Max Planck Institute for Informatics is going to make a crucial contribution to this scheme. Interest focuses on the components of a virtual laboratory, which will be playing an important role in the development of new medicaments."

14.2.2 Management of Variable Data Streams in Large Networks

The project is part of the DFG-Schwerpunktprogramm Algorithmik großer und komplexer Netzwerke. The project is led by Berthold Vöcking, in cooperation with Anja Feldmann from the computer science department in the University of Saarland, and is for the period of August 2001 - July 2003. It deals with the development and analysis of algorithms for managing variable data streams in large networks, e.g., the Internet. The project is motivated by numerous empirical studies of the communication dynamics in the Internet showing that communication patterns typically are highly variable and cannot be modeled adequately by sharply concentrated probability distributions like, e.g., the Poisson or exponential distribution. We aim to analyze the impacts of

these observed dynamics on the efficiency of routing algorithms. In particular, we plan to devise algorithms for the mapping of data streams to Web servers and strategies for intra-domain routing.

14.2.3 Algorithm Engineering for Large Graphs and Memory Hierarchies

The project is also part of the DFG-Schwerpunktprogramm Algorithmik großer und komplexer Netzwerke. The project is led by Peter Sanders and is for the period August 2001 - July 2003. It deals with the study of fundamental graph algorithms like depth first search, breadth first search, shortest paths and minimum spanning trees for large inputs. One aspect are external algorithms and cache efficient algorithms for such problems that become important when the graphs do not fit into faster layers of the memory hierarchy. Another aspect are apparent gaps between theory and practice. We want to investigate how asymptotically superior theoretical approaches that are considered impractical can be made practical for large inputs.

14.2.4 Design and Implementation of Real-Time Algorithms for Simulating the Dynamics of Colliding Rigid Bodies

Planning and simulating assembly processes with the help of virtual prototypes requires intuitive and efficient means to interact with the objects of the virtual environment. Thereby it is important, that the physical behaviour of the virtual objects is as realistic as possible, i.e. collisions between objects must be detected and the acting forces and momentums must be determined. In this project we are developing new real-time algorithms for the simulation of colliding rigid bodies and we are testing these algorithms on the basis of assembly and disassembly studies in automotive industry. The project is led by Elmer Schömer and is for the period April 2000 - March 2002.

14.3 Project funded by BMBF

GELENA is an interdisciplinary project of the Institute for Biopharmacy and Pharmaceutical Technology of the University of Saarbrücken, the Institute for New Materials (INM) in Saarbrücken, the Institute for Human Genetics of the University of Saarbrücken, the biotechnology company Across Barriers, and our group. The project is supported by a two year grant (March 2000 – March 2002) from the German Ministry of Education, Science, and Technology (BMBF = Bundesministerium für Bildung und Forschung). The project goal is the development of a new non-viral gene transfer system. The gene transfer method is intended to cure diseases caused by genetic defects such as cystic fibrosis. See Section 9.5 for more information.

14.4 Cooperations with Industry

Algorithmic Solutions (AS) is a spin-off of AG1. AS is marketing LEDA (Library of Efficient Algorithms), CGAL (computational geometry algorithms library), BALL (Biological Algorithms Library), AGD (Automatic Graph Drawing) under license agreements with the Max-Planck-Gesellschaft. In the case of LEDA, it is also responsible for the maintenance. Through its industrial contacts and through user feedback, AS brings interesting research problems into AG1. For example, Mark Zieglmann's PhD project on constrained shortest paths was suggested by a contact at DASA and Susan Hert's recent project on polygons with curved boundaries was suggested by a contact at Bubel Software.

Elmar Schömer and his group cooperate with the virtuality lab of Daimler-Chrysler. Algorithms developed by them have made their way into the virtual reality software at Daimler-Chrysler.

Peter Sanders cooperates with NEC C&C Research Laboratories, Sankt Augustin, on communication in parallel computers, and with Philips Reserach, Eindhoven, on disk scheduling.

15 Publications

Books

[1] K. Mehlhorn and S. Näher. *LEDA: a platform for combinatorial and geometric computing.* Cambridge University Press, Cambridge, November 1999.

In Journals and Book Chapters

[1] S. Albers. Better bounds for online scheduling. *SIAM Journal on Computing*, 29(2):459–473, 1999.

[2] S. Albers, N. Garg, and S. Leonardi. Minimizing stall time in single and parallel disk systems. *Journal of the ACM*, 47(6):969–986, 2000.

[3] S. Albers and M. Mitzenmacher. Average-case analyses of first fit and random fit bin packing. *Random Structures & Algorithms*, 16(3):240 – 259, May 2000.

[4] E. Althaus and K. Mehlhorn. Traveling salesman-based curve reconstruction in polynomial time. *SIAM Journal on Computing*, 31(1):27–66, 2001.

[5] N. Amato, M. Goodrich, and E. A. Ramos. A randomized algorithm for triangulating a simple polygon in linear time. *Discrete and Computational Geometry*, 26(2):246–265, 2001.

[6] S. R. Arikati and K. Mehlhorn. A correctness certificate for the Stoer-Wagner min-cut algorithm. *Information Processing Letters*, 70:251–254, 1999.

[7] V. Auletta, Y. Dinitz, Z. Nutov, and D. Parente. A 2-approximation algorithm for finding an optimum 3-vertex-connected spanning subgraph. *Journal of Algorithms*, 32(1):21–30, 1999.

[8] H. Bast. On scheduling parallel tasks at twilight. *Theory of Computing Systems*, 33(5):489–563, November 2000.

[9] W. W. Bein, R. Fleischer, and L. L. Larmore. Limited bookmark randomized online algorithms for the paging problem. *Information Processing Letters*, 76(4-6):155–162, 2000.

[10] H. Brönnimann, C. Burnikel, and S. Pion. Interval arithmetic yields efficient dynamic filters for computational geometry. *Discrete Applied Mathematics*, 109(1/2):25–45, April 2001.

[11] H. Brönnimann, L. Kettner, S. Schirra, and R. Veltkamp. Applications of the generic programming paradigm in the design of CGAL. In M. Jazayeri, R. G. K. Loos, and D. R. Musser, editors, *Generic programming - International Seminar on Generic Programming (Dagstuhl – Germany, April 1998)*, volume 1766 of *Lecture Notes in Computer Science*, pages 206–216. Springer, Berlin, Germany, January 2000.

[12] C. Burnikel, R. Fleischer, K. Mehlhorn, and S. Schirra. A strong and easily computable separation bound for arithmetic expressions involving radicals. *Algorithmica*, 27(1):87–99, 2000.

[13] S. Chaudhuri, K. V. Subrahmanyam, F. Wagner, and C. D. Zaroliagis. Computing mimicking networks. *Algorithmica*, 26(1):31–49, 2000.

[14] S. Chaudhuri and C. D. Zaroliagis. Shortest paths in digraphs of small treewidth. Part i, Sequential algorithms. *Algorithmica*, 27(3/4):212–226, 2000.

[15] J. Cheriyan and K. Mehlhorn. An analysis of the highest-level selection rule in the preflow-push max-flow algorithm. *Information Processing Letters*, 69(5):239–242, 1999.

[16] R. Cole, K. Ost, and S. Schirra. Edge-coloring bipartite multigraphs in o(e log d) time. *Combinatorica*, 21(1):5–12, 2001.

[17] C. Cooper, A. Frieze, K. Mehlhorn, and V. Priebe. Average-case complexity of shortest-paths problems in the vertex-potential model. *Random Structures & Algorithms*, 16(1):33–46, 2000.

[18] A. Crauser and P. Ferragina. A theoretical and experimental study on the construction of suffix arrays in external memory. *Algorithmica*, page 32, 2001.

[19] A. Crauser, P. Ferragina, K. Mehlhorn, U. Meyer, and E. A. Ramos. Randomized external-memory algorithms for line segment intersection and other geometric problems. *International Journal of Computational Geometry and Applications*, 11(3):305–338, June 2001.

[20] T. K. Dey, K. Mehlhorn, and E. A. Ramos. Curve reconstruction: Connecting dots with good reason. *Computational Geometry*, 15(4):229–244, 2000.

[21] Y. Dinitz, N. Garg, and M. X. Goemans. On the single-source unsplittable flow problem. *Combinatorica*, 19(1):17–41, 1999.

[22] Y. Dinitz and Z. Nutov. A 3-approximation algorithm for finding optimum 4,5-vertex-connected spanning subgraphs. *Journal of Algorithms*, 32(1):21–30, 1999.

[23] H. N. Djidjev, G. E. Pantziou, and C. D. Zaroliagis. Improved algorithms for dynamic shortest paths. *Algorithmica*, 28(4):367–389, 2000.

[24] M. Eiglsperger, S. Fekete, and G. W. Klau. Orthogonal graph drawing. In M. Kaufmann and D. Wagner, editors, *Drawing Graphs: Methods and Models*, volume 2025 of *LNCS Tutorial*, chapter 6. Springer, Berlin, January 2001.

[25] A. Fabri, G.-J. Giezeman, L. Kettner, S. Schirra, and S. Schönherr. On the design of CGAL, a computational geometry algorithms library. *Software - Practice and Experience*, 30(11):1167–1202, September 2000.

[26] P. Fatourou and P. Spirakis. Efficient scheduling of strict multithreaded computations. *Theory of Computing Systems*, 33(3):173–232, May 2000.

[27] R. Fleischer. Decision trees: old and new results. *Information and computation*, 152:44–61, 1999.

[28] N. Garg, G. Konjevod, and R. Ravi. Approximation algorithm for the group Steiner tree problem. *Journal of Algorithms*, 37(1):66–84, 2000.

[29] T. Hagerup. Dynamic algorithms for graphs of bounded treewidth. *Algorithmica*, 27(3/4):292–315, 2000.

[30] T. Hagerup. Parallel preprocessing for path queries without concurrent reading. *Information and Computation*, 158(1):18–28, April 2000.

[31] S. Hert and V. Lumelsky. Motion planning in r^3 for multiple tethered robots. *IEEE Transactions on Robotics and Automation*, 15(4):623–639, August 1999.

[32] K. Jansen, R. Solis-Oba, and M. Sviridenko. Linear time approximation schemes for shop scheduling problems. In P. M. Pardalos, editor, *Approximation and complexity in numerical optimization*, volume 42 of *Nonconvex optimization and its applications*, pages 338–346. Kluwer Academic Publisher, Dordrecht, 2000.

[33] B. Jung, H.-P. Lenhof, P. Müller, and C. Rüb. Simulating synthetic polymer chains in parallel. *Future Generation Computer Systems*, 16(5):513–522, March 2000.

[34] J. D. Kececioglu, H.-P. Lenhof, K. Mehlhorn, P. Mutzel, K. Reinert, and M. Vingron. A polyhedral approach to sequence alignment problems. *Discrete Applied Mathematics*, 104(1/3):143–186, August 2000.

[35] L. Khachiyan and L. Porkolab. Integer optimization on convex semialgebraic sets. *Discrete & Computational Geometry*, 23(2):207–224, 2000.

[36] G. W. Klau and P. Mutzel. Automatic layout and labelling of state diagrams. In *Mathematics: Key Technology for the Future*. Springer, Berlin, January 2001.

[37] O. Kohlbacher and H.-P. Lenhof. BALL - Rapid software prototyping in computational molecular biology. *Bioinformatics*, 16(9):815–824, November 2000.

[38] S. Kontogiannis, G. Pantziou, P. Spirakis, and M. Yung. Robust parallel computations through randomization. *Theory of Computing Systems*, 33(5/6):427–464, November 2000.

[39] P. Krysta and L. Pacholski. The STO problem is NP-complete. *Journal of Symbolic Computation*, 27(2):207–219, February 1999.

[40] P. Krysta and R. Solis-Oba. Approximation algorithms for bounded facility location. *Journal of Combinatorial Optimization*, 5:233–247, 2001.

[41] V. S. A. Kumar. Coupling vs. conductance for the Jerrum-Sinclair chain. *Random Structures & Algorithms*, 18:1–17, January 2001.

[42] O. Lambert and J. F. Sibeyn. Parallel and external list ranking and connected components on a cluster of workstations. In S. Q. Zheng, editor, *Proceedings of the 11th International Conference on Parallel and Distributed Computing and Systems (PDCS-99)*, pages 454–460, Boston, November 1999.

[43] J.-H. Lee, S.-M. Park, and K.-Y. Chwa. Searching a polygonal room with one door by a 1-searcher. *International Journal of Computational Geometry and Applications*, 10(2):201–220, April 2000.

[44] J.-H. Lee, C.-S. Shin, and K.-Y. Chwa. Optimal embedding of multiple directed Hamiltonian rings into d-dimensional meshes. *Journal of Parallel and Distributed Computing*, 60(6):775–783, June 2000.

[45] H.-P. Lenhof, B. Morgenstern, and K. Reinert. An exact solution for the segment-to-segment multiple sequence alignment problem. *Bioinformatics*, 15(3):203–210, 1999.

[46] M. Lermen and K. Reinert. The practical use of the A* algorithm for exact multiple sequence alignment. *Journal of Computational Biology*, 7(5):665–671, 2000.

[47] B. M. Maggs and B. Vöcking. Improved routing and sorting on multibutterflies. *Algorithmica*, 28(4):438–464, December 2000.

[48] S. Mahajan, E. A. Ramos, and K. V. Subramaniam. Solving some discrepancy problems in NC. *Algorithmica*, 29(3):371–395, 2001.

[49] K. Mehlhorn, S. Näher, M. Seel, R. Seidel, T. Schilz, S. Schirra, and C. Uhrig. Checking geometric programs or verification of geometric structures. *Computational Geometry: Theory and Applications*, 12(1-2):85–104, 1999.

[50] K. Mehlhorn, V. Priebe, G. Schäfer, and N. Sivadasan. All-pairs shortest-paths computation in the presence of negative cycles. *Inform. Process. Lett.*, 2001. to appear.

[51] K. Mehlhorn and S. Schirra. Exact computation with LEDA-real - theory and geometric applications. In G. Alefeld, editor, *Symbolic-algebraic Methods and Verification Methods - Theory and Applications*, pages 200–224. Springer, Wien, January 2001.

[52] P. Mutzel and P. Eades. Graph drawing algorithms. In M. J. Atallah, editor, *Algorithms & Theory of Computation Handbook*, handbook 9, pages 9–1–9–26. CRC, Boca Raton, USA, January 1999.

[53] R. Niedermeier, K. Reinhard, and P. Sanders. Towards optimal locality in mesh-indexings. *Discrete Applied Mathematics*, 2000. to appear.

[54] Z. Nutov and M. Penn. On integrality, stability and composition of dicycle packings and covers. *Journal of Combinatorial Optimization*, 4(2):235 – 252, June 2000.

[55] T. Polzin and S. Vahdati. A comparison of steiner tree relaxations. *Discrete Applied Mathematics*, 112:241–261, August 2001.

[56] T. Polzin and S. Vahdati. Improved algorithms for the steiner problem in networks. *Discrete Applied Mathematics*, 112:263–300, August 2001.

[57] E. A. Ramos. An optimal deterministic algorithm for computing the diameter of a three-dimensional point set. *Discrete and Computational Geometry*, 26(2):233–244, 2001.

[58] U. Rathe, P. Sanders, and P. Knight. A case study in scalability: an adi method for the two-dimensional time-dependent Dirac equation. *Parallel Computing*, 25(5):525–534, 1999.

[59] K. Reinert, J. Stoye, and T. Will. An iterative methods for faster sum-of-pairs multiple sequence alignment. *BIOINFORMATICS*, 16(9):808–814, 2000.

[60] P. Sanders. Analysis of nearest neighbor load balancing algorithms for random loads. *Parallel Computing*, 25:1013–1033, 1999.

[61] P. Sanders. Reconciling simplicity and realism in parallel disk models. *Parallel Computing*, 2001. Special Issue on Parallel Data Intensive Algorithms and Applications, to appear.

[62] P. Sanders and R. Solis-Oba. How helpers hasten h-relations. *Journal of Algorithms*, 2001. to appear.

[63] P. Sanders, R. Vollmar, and T. Worsch. Feasible models of computation: Three-dimensionality and energy consumption. *Fundamenta Informaticae*, 2000. accepted.

[64] S. Schirra. Robustness and precision issues in geometric computation. In J. R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, chapter 14, pages 597–632. Elsevier, Amsterdam, The Netherlands, January 2000.

[65] S. S. Seiden. Randomized online scheduling with delivery times. *Journal of Combinatorial Optimization*, 3(4):399–416, December 1999.

[66] S. S. Seiden. Online randomized multiprocessor scheduling. *Algorithmica*, 28(2):173–216, 2000.

[67] J. F. Sibeyn. Solving fundamental problems on sparse-meshes. *IEEE Transactions on Parallel & Distributed Systems*, 11(12):1324–1332, 2000.

[68] J. F. Sibeyn, F. Guillaume, and T. Seidel. Practical parallel list ranking. *Journal of Parallel and Distributed Computing*, 56(2):156–180, February 1999.

[69] C. R. Subramanian. Minimum coloring k-colourable graphs in polynomial average time. *Journal of Algorithms*, 33:112–123, October 1999.

[70] C. R. Subramanian. Algorithms for colouring random k-colourable graphs. *Combinatorics, Probability and Computing*, 9(1):45–77, 2000.

[71] J. L. Träff and C. D. Zaroliagis. A simple parallel algorithm for the single-source shortest path problem on planar digraphs. *Journal of Parallel and Distributed Computing*, 60(9):1103–1124, 2000.

In Conference Proceedings

[1] S. Albers, S. Arora, and S. Khanna. Page replacement for general caching problems. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-99)*, pages 31–40, Baltimore, USA, 1999.

[2] S. Albers and K. Kursawe. Exploring unknown environments with obstacles. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-99)*, pages 842–843, Baltimore, USA, 1999.

[3] E. Althaus, D. Duchier, A. Koller, K. Mehlhorn, J. Niehren, and S. Thiel. An efficient algorithm for the configuration problem of dominance graphs. In *Proceedings of the 12th ACM-SIAM Symposium on Discrete Algoritms (SODA 01)*, pages 815–824, Washington DC, USA, January 2001.

[4] E. Althaus, O. Kohlbacher, H.-P. Lenhof, and P. Müller. A combinatorial approach to protein docking with flexible side-chains. In R. Shamir, S. Miyano, S. Istrail, P. Pezner, and M. Waterman, editors, *Proceedings of the 4th Annual International Conference on Computational Molecular Biology (RECOMB-00)*, pages 15–24, Tokyo, Japan, April 2000.

[5] E. Althaus and K. Mehlhorn. Tsp-based curve reconstruction in polynomial time. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-00)*, pages 686–695, San Francisco, USA, 2000.

[6] E. Althaus, K. Mehlhorn, S. Naeher, and S. Schirra. Experiments on curve reconstruction. In *Proceedings of the 2nd Workshop on Algorithm Engineering and Experiments (ALENEX-00) [Preliminary version]*, pages 103–114, Albuquerque, USA, January 2000.

[7] N. M. Amato, M. T. Goodrich, and E. A. Ramos. Linear-time triangulation of a simple polygon made easier via randomization. In *Proceedings of the 16th Annual ACM Symposium on Computational Geometry (SCG-00)*, pages 201–212, Clear Water Bay, Kowloon, Hong Kong, June 2000.

[8] L. Arge, U. Meyer, L. Toma, and N. Zeh. On external-memory planar depth first search. In F. Dehne, J.-R. Sack, and R. Tamassia, editors, *Proceedings of the 7th International Workshop on Algorithms and Data Structures (WADS 2001)*, volume 2125 of *Lecture Notes in Computer Science*, pages 471–482, Providence, Rhode Island, USA, August 2001. Springer.

[9] H. Bast and S. Hert. The area partitioning problem. In D. Bremner, editor, *Proceedings of the 12th Annual Canadian Conference on Computational Geometry (CCCG-00)*, pages 163–171, Fredericton, New Brunswick, Canada, August 2000. University of New Brunswick.

[10] R. Beier and J. F. Sibeyn. A powerful heuristic for telephone gossiping. In M. Flammini, E. Nardelli, G. Proietti, and P. Spirakis, editors, *Proceedings of the 7th International Colloquium on Structural Information and Communication Complexity (SIROCCO-00)*, volume 7 of *Proceedings in Informatics*, pages 17–36, L'Aquila, Italy, August 2000. Carleton Scientific.

[11] N. P. Boghossian, O. Kohlbacher, and H.-P. Lenhof. BALL: Biochemical algorithms library. In J. S. Vitter and C. D. Zaroliagis, editors, *Proceedings of the 3rd International Workshop on Algorithm Engineering (WAE-99)*, volume 1668 of *Lecture Notes in Computer Science*, pages 330–344, London, UK, 1999. Springer.

[12] F. J. Brandenburg, F. Schreiber, M. Jünger, J. Marks, and P. Mutzel. Graph drawing contest report. In J. Kratochvl, editor, *Proceedings of the 7th International Symposium on Graph Drawing (GD-99)*, volume 1731 of *Lecture Notes in Computer Science*, pages 400–409, Prag, Czech Republic, 1999. Springer.

[13] K. Brengel, A. Crauser, U. Meyer, and P. Ferragina. An experimental study of priority queues in external memory. In J. S. Vitter and C. D. Zaroliagis, editors, *Proceedings of the 3rd International Workshop on Algorithm Engineering (WAE-99)*, volume 1668 of *Lecture Notes in Computer Science*, pages 345–359, London, UK, July 1999. Springer.

[14] S. Burkhardt, A. Crauser, P. Ferragina, H.-P. Lenhof, E. Rivals, and M. Vingron. q-gram based database searching using a suffix array (quasar). In S. Istrail, P. Pevzner, and M. Waterman, editors, *Proceedings of the 3rd Annual International Conference on Computational Molecular Biology (RECOMB-99)*, pages 77–83, Lyon, France, April 1999.

[15] S. Burkhardt and J. Kärkkäinen. Better filtering with gapped q-grams. In A. Amir and G. Landau, editors, *Proceedings of the 12th Annual Symposium on Combinatorial Pattern Matching, CPM 2001*, volume 2089 of *Lecture Notes in Computer Science*, pages 73–85, Jerusalem, Israel, July 2001. Springer.

[16] C. Burnikel, R. Fleischer, K. Mehlhorn, and S. Schirra. Efficient exact geometric computation made easy. In *Proceedings of the 15th Annual ACM Symposium on Computational Geometry (SCG-99)*, pages 341–350, Miami Beach, USA, June 1999.

[17] C. Burnikel, S. Funke, K. Mehlhorn, S. Schirra, and S. Schmitt. A separation bound for real algebraic expressions. In F. Meyer auf der Heide, editor, *Proceedings of the 9th European Symposium on Algorithms (ESA-01)*, volume 2161 of *Lecture Notes in Computer Science*, pages 326–337, Aarhus, Denmark, August 2001. Springer.

[18] J. Comes and M. Ziegelmann. An easy to use implementation of linear perturbations within cgal. In J. S. Vitter and C. D. Zaroliagis, editors, *Proceedings of the 3rd International Workshop on Algorithm Engineering (WAE-99)*, volume 1668 of *Lecture Notes in Computer Science*, pages 169–182, London, UK, September 1999. Springer.

[19] A. Crauser and P. Ferragina. On constructing suffix arrays in external memory. In J. Nesetril, editor, *Proceedings of the 7th Annual European Symposium on Algorithms (ESA-99)*, volume 1643 of *Lecture Notes in Computer Science*, pages 224–235, Prague, Czech Republic, 1999. Springer.

[20] A. Crauser, P. Ferragina, K. Mehlhorn, U. Meyer, and E. A. Ramos. I/O-optimal computation of segment intersections. In J. M. Abello and J. S. Vitter, editors, *Proceedings of the DIMACS Workshop on External Algorithms and Visualization*, volume 50 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 131–138, New Brunswick, New Jersey, October 1999. AMS.

[21] A. Crauser and K. Mehlhorn. LEDA-SM: Extending LEDA to secondary memory. In J. S. Vitter and C. D. Zaroliagis, editors, *Proceedings of the 3rd International Workshop on Algorithm Engineering (WAE-99)*, volume 1668 of *Lecture Notes in Computer Science*, pages 228–242, London, UK, 1999. Springer.

[22] A. Crauser, U. Meyer, and K. Brengel. An experimental study of priority queues in external memory. *Journal of Experimental Algorithmics*, 5:24, 2001.

[23] T. Dey, S. Funke, and E. Ramos. Surface reconstruction in almost linear time under locally uniform sampling. In *Proceedings of the 17th European Workshop on Computational Geometry (EWCG-01)*, pages 129–132, Berlin, Germany, March 2001. Freie Universität Berlin.

[24] T. K. Dey, K. Mehlhorn, and E. A. Ramos. Curve reconstruction: Connecting dots with good reason. In *Proceedings of the 15th Annual ACM Symposium on Computational Geometry (SCG-99)*, pages 197–206, Miami Beach, Florida, 1999.

[25] M. Dickerson, C. A. Duncan, and M. T. Goodrich. K-d trees are better when cut on the longest side. In *Proc. 8th Annual European Symposium on Algorithms, LNCS 1879*, pages 179–190. Springer-Verlag, 2000.

[26] C. Duncan, S. Kobourov, and A. Kumar. Optimal constrained graph exploration. In *Proceedings of the 12th ACM-SIAM Symposium on Discrete Algorithms*, pages 807–814, 2001.

[27] S. Edelkamp and U. Meyer. Theory and practice of time-space trade-offs in memory limited search. In *Proceedings of the Joint German/Austrian Conference on Artificial Intelligence (KI-2001)*, volume 2xxx of *Lecture Notes in Computer Science*, Wien, Österreich, September 2001. Springer.

[28] R. Fleischer. Fun with implementing algorithms. In E. Lodi, L. Pagli, and N. Santoro, editors, *Proceedings of the International Conference on FUN with Algorithms (FUN-98)*, volume 4 of *Proceedings in Informatics*, pages 88–98, Elba, Italy, 1999. Carleton Scientific.

[29] R. Fleischer and M. Wahl. Online scheduling revisited. In M. Paterson, editor, *Algorithms - ESA 2000, Proceedings of the 8th Annual European Symposium (ESA-00)*, volume 1879 of *Lecture Notes in Computer Science*, pages 202–210, Saarbrücken, Germany, 2000. Springer.

[30] S. Funke and K. Mehlhorn. LOOK - a lazy object-oriented kernel for geometric computation. In *Proceedings of the 16th Annual ACM Symposium on Computational Geometry (SCG-00)*, pages 156–165, Hong Kong, China, June 2000.

[31] S. Funke, K. Mehlhorn, and S. Näher. Structural filtering: a paradigm for efficient and exact geometric programs. In *Abstracts for the 11th Canadian Conference on Computational Geometry (CCCG-99)*, pages 39–42, Vancouver, Canada, August 1999. University of British Columbia.

[32] S. Funke and E. A. Ramos. Reconstructing a collection of curves with corners and endpoints. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-12)*, pages 344–353, Washington DC, USA, January 2001.

[33] N. Geismann, M. Hemmer, and E. Schömer. Computing a 3-dimensional cell in an arrangement of quadrics: Exactly and actually! In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 264–273, 2001.

[34] N. Geismann, M. Hemmer, and E. Schömer. The convex hull of ellipsoids. In *Proc. 17th Annu. ACM Sympos. Comput. Geom.*, pages 321–322, 2001.

[35] M. Hagedoorn, M. Overmars, and R. C. Veltkamp. A new visibility partition for affine pattern matching. In G. Borgefors, I. Nystrom, and G. Sanniti di Baja, editors, *Discrete Geometry for Computer Imagery, Proceedings of the 9th International Conference (DGCI-00)*, volume 1953 of *Lecture Notes in Computer Science*, pages 358–370, Uppsala, Sweden, December 2000. Swedish Society for Automated Image Analysis, International Association for Pattern Recognition, Springer.

[36] S. Hert. Connected area partitioning. In *Proceedings of the 17th European Workshop on Computational Geometry*, pages 35–38, Berlin, Germany, March 2001. Freie Universität Berlin.

[37] S. Hert, M. Hoffmann, L. Kettner, S. Pion, and M. Seel. An adaptable and extensible geometry kernel. In *Proceedings of the 5th Workshop on Algorithm Engineering*, volume 2141 of *Lecture Notes in Computer Science*, Aarhus, Denmark, September 2001. Springer.

[38] S. Hert and B. Richards. Multiple robot motion planning = parallel processing + geometry. In H. Christensen and G. Hager, editors, *Sensor Based Intelligent Robots*, Lecture Notes in Computer Science, pages 183–205, Schloss Dagstuhl, Wadern, Germany, November 2001. Springer.

[39] B. Jung, H.-P. Lenhof, P. Müller, and C. Rüb. Simulating synthetic polymer chains in parallel. In P. Sloot, M. Bubak, A. Hoekstra, and B. Hertzberger, editors, *Proceedings of the 7th International Conference on High-Performance Computing and Networking Europe (HPCN Europe-99)*, volume 1593 of *Lecture Notes in Computer Science*, pages 13–22, Amsterdam, Netherlands, 1999. Springer.

[40] R. Karp, C. Schindelhauer, S. Shenker, and B. Vöcking. Randomized rumor spreading. In *Proceedings of the 41st Conference on Foundations of Computer Science (FOCS)*, pages 565–574, Redondo Beach, California, November 2001. IEEE Computer Society, IEEE Computer Society.

[41] G. W. Klau, K. Klein, and P. Mutzel. An experimental comparison of orthogonal compaction algorithms. In J. Marks, editor, *Proceedings of the 8th International Symposium on Graph Drawing (GD 2000)*, volume 1984 of *Lectur Notes in Computer Science*, pages 37–51, Colonial Williamsburg, VA, January 2001. Springer.

[42] G. W. Klau and P. Mutzel. Combining graph labeling and compaction. In J. Kratochvl, editor, *Proceedings of the 7th International Symposium on Graph Drawing (GD-99)*, volume 1731 of *Lecture Notes in Computer Science*, pages 27–37, Stirin Castle, Czech Republic, September 1999. Springer.

[43] G. W. Klau and P. Mutzel. Optimal compaction of orthogonal grid drawings. In G. Cornujols, R. E. Burkard, and G. J. Woeginger, editors, *Proceedings of the 7th International Conference on Integer Programming and Combinatorial Optimization (IPCO-99)*, volume 1610 of *Lecture Notes in Computer Science*, pages 304–319, Graz, Austria, June 1999. Springer.

[44] G. W. Klau and P. Mutzel. Optimal labelling of point features in the slider model. In D.-Z. Du, P. Eades, V. Estivill-Castro, X. Lin, and A. Sharma, editors, *Computing and Combinatorics, Proceedings of the 6th Annual International Conference (COCOON-00)*, volume 1858 of *Lecture Notes in Computer Science*, pages 340–350, Sydney, Australia, 2000. Springer.

[45] O. Kohlbacher and H.-P. Lenhof. Rapid software prototyping in computational molecular biology. In *Proceedings of the German Conference on Bioinformatics (GCB'99)*, pages 81–92, Hannover, Germany, 1999. GBF-Braunschweig / Dep. of Bioinformatics.

[46] A. Koller, K. Mehlhorn, and J. Niehren. A polynomial-time fragment of dominance constraints. In *Proceedings of the 38th Annual Meeting of the Association of Computational Linguistics (ACL '2000)*, Hong-Kong, 2001.

[47] P. Krysta and V. S. A. Kumar. Approximation algorithms for minimum size 2-connectivity problems. In A. Ferreira and H. Reichel, editors, *Proceedings of the 18th International Symposium on Theoretical Aspects of Computer Science (STACS-18)*, volume 2010 of *Lecture Notes in Computer Science*, pages 431–442, Dresden, Germany, 2001. Springer.

[48] P. Krysta and K. Lorys. Efficient approximation algorithms for the achromatic number. In J. Nesetril, editor, *Proceedings of the 7th Annual European Symposium on Algorithms (ESA-99)*, volume 1643 of *Lecture Notes in Computer Science*, pages 402–413, Prague, Czech Republic, July 1999. Springer.

[49] P. Krysta and R. Solis-Oba. Approximation algorithms for bounded facility location. In T. Asano, H. Imai, D. Lee, S. Nakano, and T. Tokuyama, editors, *Proceedings of the 5th Annual International Conference on Computing and Combinatorics (COCOON-99)*, volume 1627 of *Lecture Notes in Computer Science*, pages 241–250, Tokyo, Japan, July 1999. Springer.

[50] V. S. A. Kumar, S. Arya, and H. Ramesh. Hardness of set cover with intersection 1. In U. Montanari, J. D. P. Rolim, and E. Welzl, editors, *Automata, Languages and Programming, Proceedings of the 27th International Colloquium (ICALP-00)*, volume 1853 of *Lecture Notes in Computer Science*, pages 624–635, Geneva, Switzerland, 2000. EATCS, Springer.

[51] V. S. A. Kumar and H. Ramesh. Covering rectilinear polygons with axis-parallel rectangles. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC-99)*, pages 445–454, Atlanta, USA, May 1999. Association of Computing Machinery(ACM).

[52] V. S. A. Kumar and H. Ramesh. Markovian coupling vs. conductance for the Jerrum-Sinclair chain. In *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS-99)*, pages 241–251, New York, USA, 1999. IEEE, IEEE Computer Society.

[53] O. Lambert and J. F. Sibeyn. Parallel and external list ranking and connected components on a cluster of workstations. In S. Q. Zheng, editor, *Proceedings of the 11th International Conference on Parallel and Distributed Computing and Systems (PDCS-99)*, pages 454–460, Boston, November 1999.

[54] J.-H. Lee, O. Cheong, W.-C. Kwon, S.-Y. Shin, and K.-Y. Chwa. Approximation of curvature-constrained shortest paths through a sequence of points. In M. Paterson, editor, *Algorithms - ESA 2000, Proceedings of the 8th Annual European Symposium (ESA-00)*, volume 1879 of *Lecture Notes in Computer Science*, pages 314–325, Saarbrücken,Germany, September 2000. Springer.

[55] K. Mehlhorn. The engineering of some bipartite matching programs. In A. Aggarwal and C. P. Rangan, editors, *Algorithms and Computation, Proceedings of the 10. International Symposium (ISSAC-99)*, volume 1741 of *Lecture Notes in Computer Science*, pages 1–3, Chennai, India, 1999. Springer.

[56] K. Mehlhorn. Constraint programming and graph algorithms. In U. Montanari, J. D. P. Rolim, and E. Welzl, editors, *Automata, Languages and Programming, Proceedings of the 27th International Colloquium (ICALP-00)*, volume 1853 of *Lecture Notes in Computer Science*, pages 571–575, Geneva, Switzerland, 2000. Springer.

[57] K. Mehlhorn and S. Näher. From algorithms to working programs: On the use of program checking in LEDA. In L. Brim, J. Gruska, J. Zlatuska, and S. Näher, editors, *Proceedings of the 23rd International Symposium on Mathematical Foundations of Computer Science (MFCS-98)*, volume 1450 of *Lecture Notes in Computer Science*, pages 84–93, Brno, Czech Republic, August 1999. Springer.

[58] K. Mehlhorn and G. Schäfer. A heuristic for Dijkstra's algorithm with many targets and its use in weighted matching algorithms. In *Proceedings of the 9th Annual European Symposium on Algorithms (ESA 2001)*, Lecture Notes in Computer Science. Springer, 2001.

[59] K. Mehlhorn and G. Schäfer. Implementation of $o(nmlogn)$ weighted matchings: The power of data structures. In *Workshop on Algorithm Engineering (WAE)*, volume to appear of *Lecture Notes in Computer Science*, Saarbruecken, 2001. Springer.

[60] K. Mehlhorn and S. Schirra. Geometric computing with CGAL and LEDA. In P.-J. Laurent, P. Sablonnire, and L. L. Schumaker, editors, *Curve and surface design, Saint-Malo 1999*, pages 277–286, Saint-Malo, France, 2000. Vanderbilt University Press.

[61] K. Mehlhorn and S. Thiel. Faster algorithms for bound-consistency of the sortedness and the alldifferent constraint. In R. Dechter, editor, *Proceedings of the 6th International Conference on Principles and Practice of Constraint Programming (CP-00)*, volume 1894 of *Lecture Notes in Computer Science*, pages 306–319, Singapore, September 2000. Springer.

[62] K. Mehlhorn and M. Ziegelmann. Resource constrained shortest paths. In M. Paterson, editor, *Algorithms - ESA 2000, Proceedings of the 8th Annual European Symposium (ESA-00)*, volume 1879 of *Lecture Notes in Computer Science*, pages 326–337, Saarbrücken, Germany, September 2000. Springer.

[63] K. Mehlhorn and M. Ziegelmann. CNOP - a package for constrained network optimization. In A. Buchsbaum and J. Snoeyink, editors, *Proceedings of the 3rd Workshop on Algorithm Engineering and Experiments*, volume 2153 of *Lecture Notes in Computer Science*, pages 15–28, Washington, DC, USA, 2001. Springer.

[64] U. Meyer. External memory bfs on undirected graphs with bounded degree. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 87–88, Washington, DC, January 2001.

[65] U. Meyer. Heaps are better than buckets: Parallel shortest paths on unbalanced graphs. In *Proceedings of the 7th International Euro-Par Conference (Euro-Par 2001)*, Lecture Notes in Computer Science, Manchester, UK, August 2001. Springer.

[66] U. Meyer. Single-source shortest-paths on arbitrary directed graphs in linear average-case time. In *Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 797–806, Washington, DC, January 2001.

[67] U. Meyer and P. Sanders. Parallel shortest path for arbitrary graphs. In A. Bode, T. Ludwig, W. Karl, and R. Wismüller, editors, *Euro-Par 2000 Parallel Processing, Proceedings of the 6th International Euro-Par Conference (Euro-Par-00)*, volume 1900 of *Lecture Notes in Computer Science*, pages 461–470, Munich, Germany, August 2000. Springer.

[68] P. Mutzel and R. Weiskircher. Optimizing over all combinatorial embeddings of a planar graph. In G. Cornujols, R. E. Burkard, and G. J. Woeginger, editors, *Proceedings of the 7th International Conference on Integer Programming and Combinatorial Optimization (IPCO-99)*, volume 1610 of *Lecture Notes in Computer Science*, pages 361–376, Graz, Austria, June 1999. Springer.

[69] P. Mutzel and T. Ziegler. The constrained crossing minimization problem: a first approach. In P. Kall and H.-J. Lüthi, editors, *Operations Research Proceedings 1998*, pages 125–134, Zurich, Switzerland, 1999. Springer.

[70] P. Mutzel and T. Ziegler. The constrained crossing minimization problem. In J. Kratochvil, editor, *Graph Drawing, Proceedings of the 7th International Symposium (GD-99)*, volume 1731 of *Lecture Notes in Computer Science*, pages 175–185, Stirin Castle, Czech Republic, 2000. Springer.

[71] D. Neumann, O. Kohlbacher, E. Haltner, H.-P. Lenhof, and C.-M. Lehr. Modeling the sugar lectin interaction by computational chemistry relevant to drug design. In *Proceedings of the 3rd World Meeting on Pharmaceutics, Biopharmaceutics and Pharmaceutical Technology*, page 233, Berlin, Germany, 2000. Arbeitsgemeinschaft für Pharmazeutische Verfahrenstechnik e.V.

[72] S.-M. Park, J.-H. Lee, and K.-Y. Chwa. Characterization of rooms searchable by two guards. In D. Lee and S.-H. Teng, editors, *Proceedings of the 11th Conference on Algorithms and Computation*, pages 515–526, Taipei,Taiwan, 2000. Springer.

[73] S.-M. Park, J.-H. Lee, and K.-Y. Chwa. Visibility-based pursuit-evasion in a polygonal region by a searcher. In F. Orejas, P. G. Spirakis, and J. v. Leeuwen, editors, *Proceedings of the 28th International Colloquium on Automata, Languages and Programming*, pages 456–468, Crete,Greece, 2001. Springer.

[74] T. Polzin and S. Vahdati. Primal-dual approaches to the steiner problem. In K. Jansen and S. Khuller, editors, *Approximation Algorithms for Combinatorial Optimization*, volume 1913 of *Lecture Notes in Computer Science*, pages 214–225, Saarbrücken, Germany, 2000. Springer.

[75] R. Rajaraman, A. Richa, B. Vöcking, and G. Vuppuluri. A data tracking scheme for general networks. In *11th ACM Symposium on Parallel Algorithms and Architectures*, pages 247–254, 2001.

[76] E. A. Ramos. On range reporting, ray shooting and k-level construction. In *Proceedings of the 15th Annual ACM Symposium on Computational Geometry (SCG-99)*, pages 390–399, Miami Beach, Florida, 1999.

[77] E. A. Ramos. Deterministic algorithms for 3-d diameter and some 2-d lower envelopes. In *Proceedings of the 16th Annual ACM Symposium on Computational Geometry (SCG-00)*, pages 290–299, Clear Water Bay, Kowloon, Hong Kong, June 2000.

[78] E. A. Ramos. Linear programming queries revisited. In *Proc. 16th Ann. ACM Sympos. on Comput. Geometry (SCG-00)*, pages 176–181, 2000.

[79] E. A. Ramos, N. M. Amato, and M. T. Goodrich. Computing the arrangement of curve segments: Divide-and-conquer algorithms via sampling. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-00)*, pages 705–706, San Francisco, CA, USA, January 2000.

[80] P. Sanders. Accessing multiple sequences through set associative caches. In J. Wiedermann, P. van Emde Boas, and M. Nielsen, editors, *Proceedings of the 26th International Colloquium on Automata, Languages and Programming (ICALP-99)*, volume 1644 of *Lecture Notes in Computer Science*, pages 655–664, Prague, Czech Republic, 1999. EATCS, Springer.

[81] P. Sanders. Asynchronous random polling dynamic load balancing. In A. Aggarwal and C. P. Rangan, editors, *Proceedings of the 10th International Symposium on Algorithms and Computation (ISAAC-99)*, volume 1741 of *Lecture Notes in Computer Science*, pages 37–48, Chennai, India, December 1999. Springer.

[82] P. Sanders. Fast priority queues for cached memory. In M. T. Goodrich and C. C. McGeoch, editors, *Selected papers of the International Workshop on Algorithm Engineering and Experimentation (ALENEX-99)*, volume 1619 of *Lecture Notes in Computer Science*, pages 312–327, Baltimore, USA, January 1999. Springer.

[83] P. Sanders. Asynchronous scheduling of redundant disk arrays. In S.-H. Teng, editor, *Proceedings of the 12th Annual ACM Symposium on Parallel Algorithms and Architecture (SPAA-00)*, pages 89–98, Bar Harbor, Maine, 2000.

[84] P. Sanders, S. Egner, and J. Korst. Fast concurrent access to parallel disks. In *Proceedings of the 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-00)*, pages 849–858, San Francisco, USA, 2000.

[85] P. Sanders and J. F. Sibeyn. A bandwidth latency tradeoff for broadcast and reduction. In A. Bode, T. Ludwig, W. Karl, and R. Wismüller, editors, *Euro-Par 2000 Parallel Processing, Proceedings of the 6th International Euro-Par Conference (Euro-Par-00)*, volume 1900 of *Lecture Notes in Computer Science*, pages 918–926, Munich, Germany, 2000. Springer.

[86] P. Sanders and R. Solis-Oba. How helpers hasten h-relations. In M. Paterson, editor, *Algorithms - ESA 2000, Proceedings of the 8th Annual European Symposium (ESA-00)*, volume 1879 of *Lecture Notes in Computer Science*, pages 392–402, Saarbrücken, Germany, 2000. Springer.

[87] P. Sanders, T. Takkula, and D. Wedelin. High performance integer optimization for crew scheduling. In P. Sloot, M. Bubak, A. Hoekstra, and B. Hertzberger, editors, *Proceedings of the 7th International Conference on High-Performance Computing and Networking Europe (HPCN Europe-99)*, volume 1593 of *Lecture Notes in Computer Science*, pages 3–12, Amsterdam, NL, April 1999. Springer.

[88] S. Schirra. A case study on the cost of geometric computing. In M. T. Goodrich and C. C. McGeoch, editors, *Selected papers of the International Workshop on Algorithm Engineering and Experimentation (ALENEX-99)*, volume 1619 of *Lecture Notes in Computer Science*, pages 156–176, Baltimore, USA, 1999. Springer.

[89] J. F. Sibeyn. Better deterministic routing on meshes. In *Proceedings of the 13th International Parallel Processing Symposium, and 10th Symposium on Parallel and Distributed Processing (IPPS/SPDP-99)*, pages 420–425, St Juan, Puerto Rico, May 1999. IEEE Computer Society.

[90] J. F. Sibeyn. External selection. In C. Meinel and S. Tison, editors, *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS-99)*, volume 1563 of *Lecture Notes in Computer Science*, pages 291–301, Trier, March 1999. Springer.

[91] J. F. Sibeyn. Power and limitations of optical reconfigurable arrays. In C. Gavoille, J.-C. Bermond, and A. Raspaud, editors, *Proceedings of the 6th International Colloquium on Structural Information and Communication Complexity (SIROCCO 6)*, volume 5 of *Proceedings in Informatics*, pages 249–263, Lacanau-Ocan, France, July 1999. Carleton Scientific.

[92] J. F. Sibeyn. Row-major sorting on meshes. *SIAM Journal on Computing*, 28(3):847–863, June 1999.

[93] J. F. Sibeyn. Ultimate parallel list ranking? In P. Banerjee, V. K. Prasanna, and B. F. Sinha, editors, *Proceedings of the 6th International Conference on High Performance Computing (HiPC-99)*, volume 1745 of *Lecture Notes in Computer Science*, pages 197–201, Calcutta, December 1999. IEEE, Springer.

[94] R. Solis-Oba, K. Jansen, and M. Sviridenko. A linear time approximation scheme for the job shop scheduling problem. In D. Hochbaum, K. Jansen, J. D. P. Rolim, and A. Sinclair, editors, *Proceedings of the 3rd International Workshop on Randomization and Approximation Techniques in Computer Science, and 2nd International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (RANDOM-APPROX-99)*, volume 1671 of *Lecture Notes in Computer Science*, pages 177–188, Berkeley, U.S.A., August 1999. Springer.

[95] R. Solis-Oba, K. Jansen, and M. Sviridenko. Makespan minimization in job shops: a polynomial time approximation scheme. In *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC-99)*, pages 394–399, Atlanta, USA, 1999. Association of Computing Machinery, ACM.

[96] C. R. Subramanian. A generalization of Janson inequalities and its application to finding shortest paths. In *Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA-99)*, pages 795–804, Baltimore, USA, January 1999.

[97] B. Vöcking. Almost optimal permutation routing on hypercubes. In *33rd ACM Symposium on the Theory of Computing*, pages 530–539, 2001.

[98] T. Warken and E. Schömer. Rolling rigid objects. In *Proceedings of WSCG'2001*, volume 1, pages 57–62, 2001.

Others

[1] M. Bäsken, H. Brönnimann, F. Da, O. Devillers, E. Ester, A. Fabri, E. Flato, B. Gärtner, G.-J. Giezeman, D. Halperin, I. Hanniel, S. Har-Peled, T. Herrmann, S. Hert, S. Hirsch, M. Hoffmann, L. Kettner, O. Nechushtan, G. Neyer, D. Pasechnik, S. Pion, S. Schirra, S. Schönherr, M. Seel, M. Teillaud, R. Veltkamp, R. Wein, W. Wesselink, and M. Yvinec. *CGAL Reference and User Manuals*, 2.3 edition, August 2001.

[2] P. Sanders. Berechnungen mit grossen Datenmengen. In G. Plehn and S. Deutschmann, editors, *Jahrbuch der Max-Planck-Gesellschaft 1999*, pages 534–540. Vandenhoeck & Ruprecht, Göttingen, November 1999.

Part IV

The Programming Logics Group

1 Personnel

Director:
Prof. Dr. Harald Ganzinger

Senior research scientist:
Dr. Andreas Podelski

Researchers:
Dr. Witold Charatonik
Dr. Fritz Eisenbrand (–September 2001)
Dr. Manfred Jaeger
Dr. Hans de Nivelle
Dr. Viorica Sofronie-Stokkermans
Dr. Uwe Waldmann
Dr. Christoph Weidenbach (–December 1999)
Dr. Emil Weydert (–March 2000)

Post-doctoral fellows and long-term guests:
Dr. Richard Booth (–October 1999)
Dr. Georgio Delzanno (–November 1999)
Dr. Jean-Marc Talbot (–February 2001)
Dr. Margus Veanes (–July 1999)
Dr. Jinzhao Wu (–December 1999)

Ph.D. students:
Werner Backes (February 2000–)
Thomas Hillenbrand (May 2001–)
Jörn Hopf (–March 2001)
Patrick Maier
Christoph Meyer (–March 2000)
Supratik Mukhopadhyay (–June 2001)
Jürgen Stuber (–May 2000)
Miroslava Tzakova (–February 2000)

Secretaries:
Brigitta Hansen
Christine Kiesel

2 Visitors

1999

Hubert Baumeister	19.05.99-22.05.99	LMU München
Leo Bachmair	12.07.99-25.07.99	Suny Stony Brook
Christopher Lynch	20.06.99-27.06.99	Clarkson University
Lisa Fleischer	17.06.99-19.06.99	Columbia University
Bruce Spencer	01.08.99-21.08.99	University of Brunswick
Kahlil Hodgson	16.07.99-20.07.99	University of Canberra
Solofo Ramangalahy	19.07.99-30.07.99	
Sophie Tison	20.07.99-23.07.99	Univ. of Lille
Ruiming Li	19.08.99-21.08.99	
Dimitri Hendriks	03.10.99-08.10.99	University of Amsterdam
Moshe Vardi	30.10.99-02.10.99	Rice University
Ulrich Hustadt	02.11.99-09.11.99	Manchester Metropolitan University
Renate Schmidt	02.11.99-09.11.99	Manchester Metropolitan University
Ewa Orlowska	09.11.99-14.11.99	Inst. of Telecomm., Warschau
Li Wei	15.11.99-15.12.99	Beijing University

2000

Reinhard Kahle	02.03.00-03.03.00	Univ. Tuebingen
Enno Ohlebusch	14.03.00-15.03.00	Uni Bielefeld
David Basin	22.03.00-23.03.00	Freiburg University
Laurent Mauborgne	01.04.00-31.04.00	ENS Paris
Shilong MA	01.04.00-30.05.00	University Beijing
Fleming Nielson	01.04.00-30.06.00	University Aarhus
Jayati Ghoshal	01.07.00-31.08.00	University of California
Ian Pratt	16.05.00-26.05.00	University of Manchester
Christopher Lynch	24.5.00-30.05.00	Clarkson University
Abhinai Srivastava	11.05.00-31.07.00	ITT Delhi
Andy Gordon	13.06.00-16.06.00	Microsoft
Juan Heguiabehere	26.06.00-07.07.00	University of Amsterdam
Bertrand Jeannet	09.07.00-11.07.00	VERIMAG, Limoges
Bruce Spencer	06.08.00-27.08.00	University of Brunswick
Christopher Lynch	05.08.00-26.08.00	Clarkson University
Leo Bachmair	15.08.00-26.08.00	SUNY Stony Brook
Robert Nieuwenhuis	23.08.00-01.09.00	University Barcelona
Alex Leitsch	20.08.00-28.08.00	University Wien
Tobias Nipkow	20.08.00-25.08.00	University Muenchen
Andrei Voronkov	22.08.00-29.08.00	Manchester University
Leszek Pacholski	10.09.00-24.09.00	University of Wroclaw
Sergei Vorobyov	17.07.00-16.08.00	University of Uppsala
Kristian Kersting	07.09.00-08.09.00	Freiburg University

2001

Günter Rote	03.01.01-05.01.01	FU Berlin
Jürgen Dix	16.02.01	Koblenz University
Kristian Kersting	08.04.01-13.04.01	Freiburg University
R.K. Shyamasundar	30.04.01-31.07.01	TATA Institute Bombay
Anatol Slissenko	02.05.01-31.05.01	University Paris
Lav Rai	10.05.01-15.07.01	IIT
Ravi Krishna	10.05.01-15.07.01	IIT
Vinay Middha	10.05.01-15.07.01	IIT
C.R. Ramakrishnan	19.05.01-23.05.01	SUNY, StonyBrook
Christopher Lynch	25.05.01-28.06.01	Clarkson University
Stefan Schwoon	24.05.01-25.05.01	TU Muenchen
Mai Gehrke	17.06.01-19.06.01	New Mexico State University
John Slaney	25.06.01-27.06.01	Australian National University
Roberto Giacobazzi	27.06.01-30.06.01	Universitaet Verona
Sergei Vorobyov	02.07.01-31.07.01	University Uppsala
Henrik Bjorklund	03.07.01-31.07.01	University Uppsala
Viktor Petersson	03.07.01-31.07.01	University Uppsala
Stephane Demri	29.07.01-04.08.01	LSV-ENS
Jean-Marc Talbot	06.08.01-18.08.01	LIF Lille

3 First-Order Theorem Proving and Term Rewriting

In previous years, our theoretical research in the area of automated theorem proving for first-order logic concentrated on (i) the development of efficient deductive calculi, (ii) uniform methods for constructing decision procedures for fragments of first-order logic, (iii) the reconstruction of modal and other non-classical logics within fragments of first-order logic, (iv) combining algebraic and logic methods for varieties of groups, rings, fields, and lattices, and (v) analyzing the complexity of logics and formal theories.

During the period covered by this report, a few more results related to (i)–(iii) were obtained. It appears, however, as if substantial further progress in this area would require new methods resulting either from new ways of looking at the problem, or from the specialization of the known methods to particular application domains. We have started to concentrate on two different kinds of applications. We are now also working on (vi) the general problem of integrating an automated prover for first-order logic into interactive provers or proof checkers usually based on higher-order logic, and (vii) on efficient deductive methods suited for the automatic analysis of large programs.

With regard to (vi) the current situation is that the existing higher-order provers or proof checkers are equipped with more or less naive methods for automating the proof of "easy" subgoals, or employ standard decision procedures for "classical" theories such as Presburger arithmetic or congruence closure. So far, none of the more sophisticated first-order provers have been integrated successfully in systems such as Coq, HOL, Isabelle, Elf, or PVS. Attempts have been made by several people, but the outcome was generally disappointing.

There are many reasons why the situation is like this. First-order provers only speak first-order logic, and often do not have good ways of dealing with types. First-order provers normally do not perform well even on trivial problems if standard structures (numbers, lists, arrays, sets) are involved. Also, they normally do not produce proof objects acceptable to the higher-order proof checker. First-order provers are difficult to control. Their inference systems and search strategies are highly parameterized (atom and term orderings, selection functions, redundancy elimination strategies), but only for clever problem-specific settings of these parameters will they perform acceptably well.

We do not have the capacity of dealing with all these problems simultaneously. We have started to look at the problem of producing acceptable and small proof objects, and we have continued to investigate our long-term topic (iv), where the goal is to obtain efficient automated reasoning for standard theories, such as Abelian groups, rings, or lattices, in the presence of additional free function symbols.

Regarding (vii), the specialization of efficient first-order methods to analyzing large programs, we have combined methods of constraint programming, model checking and abstract interpretations (cf. Section 4), and we have looked into execution models for logic programs which allow certain inference systems to be "executed" with guaranteed complexity bounds, often competitive with the best algorithms known (cf. Section 3.2).

During the period covered by this report, substantial effort was devoted to summarizing previous research and presenting it in a form accessible to a larger community. Current and former members of the group contributed 6 chapters [1, 8, 2, 7, 4, 9] to the Handbook of Automated Reasoning which was finally published in June 2001. A tutorial about theorem proving and constraints appeared in [5]. Another overview paper is [6] where the fundamental properties of symmetric Heyting algebras of order n are summarized. The paper [3] is an overview article about translation-based resolution methods for modal logics.

References

- [1] L. Bachmair and H. Ganzinger. A theory of resolution. In J. A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 2, pages 19–99. Elsevier, Amsterdam, the Netherlands, 2001.
- [2] A. Bockmayr and V. Weispfenning. Solving numerical constraints. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 12, pages 751–842. Elsevier, Amsterdam, the Netherlands, January 2001.
- [3] H. de Nivelle. An overview of resolution decision procedures. In M. Faller, S. Kaufmann, and M. Pauly, editors, *Formalizing the Dynamics of Information*, volume 91 of *CSLI Publications*, pages 115–130, Stanford University, Palo Alto, USA, 2000. Center for the Study of Language and Information.
- [4] C. G. Fermüller, A. Leitsch, U. Hustadt, and T. Tammet. Resolution decision procedures. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 25, pages 1793–1849. Elsevier, Amsterdam, the Netherlands, January 2001.
- [5] H. Ganzinger and R. Nieuwenhuis. Constraints and theorem proving. In H. Comon, C. Marché, and R. Treinen, editors, *Contraints in Computational Logics, International Summer School (CCL-99)*, volume 2002 of *Lecture Notes in Computer Science*, pages 159–201, Gif-sur-Yvette, France, 2001. Springer.
- [6] L. Iturrioz and V. Sofronie-Stokkermans. SHn-algebras (Symmetric Heyting algebras of order n). In L. Iturrioz, E. Orlowska, and E. Turunen, editors, *COST Action 15 (Many-Valued Logics for Computer Science Applications) ATLAS of Many-Valued Structures*, number 75 in Department of Information Technology. Mathematics Report, pages 1–11. Tampere University of Technology (Mathematics), Tampere, January 2000.
- [7] A. Nonnengart, H. J. Ohlbach, and D. M. Gabbay. Encoding two-valued non-classical logics in classic logic. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 2, chapter 21, pages 1403–1486. Elsevier, Amsterdam, the Netherlands, January 2001.
- [8] A. Nonnengart and C. Weidenbach. Computing small clause normal forms. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 6, pages 335–367. Elsevier, Amsterdam, the Netherlands, January 2001.
- [9] C. Weidenbach. Combining superposition, sorts and splitting. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 27, pages 1965–2013. Elsevier, Amsterdam, the Netherlands, January 2001.

3.1 Decision Procedures

The Two-Variable Fragment with Equality Investigator: Hans de Nivelle

The 2-variable fragment with equality is one of the well-known decidable fragments of first-order logic, for which no saturation-based decision procedure had been found in the past. (In fact, for few fragments with equality only do we know of resolution- or superposition-based decisions procedures, as opposed to naive semantic methods. One of the major exceptions is a superposition decision procedure for the guarded fragment with equality reported in [4].) A decision procedure for the two-variable fragment without equality was known before [1]. In addition, [5] gives a very elegant, but unfortunately impractical, semantic decision procedure for the two-variable fragment with equality, exploiting the finite-model property of this fragment.

The practical motivation for considering this fragment arises from applications in natural language processing: One can give a simple, but useful semantics to natural languages such as English by restricting the ways in which anaphora are resolved. Referring words like 'he/she/it/himself' are allowed to refer only to the nearest noun-phrase (in the parse tree). In addition only two-place

predicates are allowed. The subset of English obtained in this way, corresponds exactly to the 2-variable fragment with equality, see [6].

We have found a decision procedure, based on resolution, for the two-variable fragment with equality in [3]. It provides a practical method for deciding provability and satisfiability, which can be easily integrated into existing theorem provers. It is a two-level procedure. First, the clause set is saturated under a restricted resolution rule. After that, negative equalities can be replaced by an arbitrary irreflexive relation. Positive equalities can be either ignored, or explicitly axiomatized. The resulting clause set is in the 2-variable fragment without equality. to which our previous resolution method can be applied.

Resolution in Modal and Description Logics Investigator: Hans de Nivelle

For theorem proving in non-standard logics, there exist in principle two approaches, direct theorem proving, and translation-based theorem proving.

Translation-based methods have the obvious advantage that one does not need to design a new theorem proving strategy for each non-classical logic. The translations are usually straightforward. On the translated formula, standard theorem provers can be used. On the other hand the complexity can be increased by the translation. One might for example translate a decidable modal logic into an undecidable logic, or a modal logic of PSPACE-complexity into a fragment of EXPTIME-complexity. Identifying the proper decidable fragment, and the right decision procedure may be as hard as designing a decision procedure in the original logic itself. Direct theorem proving has the advantage that the proofs are in the logic itself. Non-standard logics are often designed for the purpose of being able to naturally express certain notions, and it may be important to have natural representations of proofs as well.

One can conclude that it is useful to pursue both lines of research. The non-classical logics can be used for finding decision procedures for subsets of classical logic, and the decision procedures can be used for designing decision procedures in non-classical logics.

In [2], a direct resolution decision procedure for Description Logic is presented. Description logics are very restricted logics, mainly used for describing databases. The decision procedure is tuned towards producing natural proofs. In order to achieve this goal, Skolemization is avoided.

Paper [1] is an overview article of translation-based resolution methods for modal logics. It is explained how resolution can be made a decision procedure for many modal and description logics. The paper focuses on the modal logic $K_m(\cap, \cup, {}^{-1}, \neg)$, which contains many modal and description logics. A number of decidable subsets of clausal first order logic are defined, together with the decision procedures and the translation schemes. The clause class DL^* resulting from clausifying formulas in $K_m(\cap, \cup, {}^{-1}, \neg)$, can be decided by ordered resolution. The guarded fragment can be decided by ordered resolution with selection. Some other fragments are discussed, corresponding to natural fragments of $K_m(\cap, \cup, {}^{-1}, \neg)$.

References

- [1] H. de Nivelle. An overview of resolution decision procedures. In M. Faller, S. Kaufmann, and M. Pauly, editors, *Formalizing the Dynamics of Information*, volume 91 of *CSLI Publications*, pages 115–130, Stanford University, Palo Alto, USA, 2000. Center for the Study of Language and Information.
- [2] H. de Nivelle, C. Areces, and M. de Rijke. Prefixed resolution: A resolution method for modal and description logics. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction (CADE-16)*, volume 1632 of *Lecture Notes in Artificial Intelligence*, pages 187–201, Trento, Italy, July 1999. Springer.

- [3] H. de Nivelle and I. Pratt-Hartmann. A resolution-based decision procedure for the two-variable fragment with equality. In R. Gore, A. Leitsch, and T. Nipkow, editors, *Proceedings of the 1st International Joint Conference on Automated Reasoning (IJCAR-2001)*, volume 2083 of *Lecture Notes in Artificial Intelligence*, pages 211–225, Siena, Italy, July 2001. Springer.
- [4] H. Ganzinger and H. de Nivelle. A superposition decision procedure for the guarded fragment with equality. In G. Longo, editor, *Proceedings of the 14th Annual IEEE Symposium on Logic in Computer Science (LICS-99)*, pages 295–303, Trento, Italy, 1999. IEEE Computer Society, IEEE Computer Society.

[5] E. Grädel, P. G. Kolaitis, and M. Vardi. On the decision problem for two-variable first order logic. *The Bulletin of Symbolic Logic*, 3(1):53–68, 1997.

[6] I. Pratt-Hartmann. On the semantic complexity of some fragments of english. Technical Report UMCS–00–5–1, University of Manchester, Department of Computer Science, 2000.

3.2 Meta-Complexity Theorems

Investigator: Harald Ganzinger

Meta-complexity theorems for programming languages are an active area of research. One line of work, inspired by ideas of Bellantoni and Cook [3], and, more recently, by Hofmann [9], aims at identifying suitable type systems so that well-typed programs run in polynomial time. As the operational semantics of these (functional) programs is based on standard reduction mechanisms, no non-trivial algorithmic paradigms are provided internally, but have to be explicitly programmed. Hence, these languages are, at present, difficult to use. In fact they do not appear to be designed for human use, but rather as intermediate languages for intelligent compilers that do not yet exist.

Our directions of research are different. We are interested in logic programs rather than functional programs. Also, we provide quite sophisticated, partly automated methods of analyzing program complexity that, up to now, cannot be expressed by local typing rules. Finally, our intention is to design deductive machinery that employs fundamental algorithmic paradigms (hashing, dynamic programming, union-find, priority queues) internally. Externally a programmer is only faced with corresponding abstract and familiar logic concepts (structured term data with non-linear variable occurrences, subformula property, equality, rule priorities). We want our logic programs to be competitive with the best known algorithms. As we expect deductive program analysis to be a main area of application for our techniques, efficiency is crucial for being able to handle large programs.

Relating fundamental algorithmic concepts to fundamental concepts in logic and theorem proving is another more general goal of this research.

New Results about Order Locality One wants trivial theorems to be proved automatically and efficiently. A useful definition is to regard a theorem as trivial if it can be formulated within a logic or theory that is decidable in polynomial time. Accepting this definition of triviality, we have the problems of (i) recognizing when a set of axioms presents a theory in PTIME, and (ii) efficiently deciding — with guaranteed complexity — whether or not a formula is valid in some such theory.

In previous research we have shown that any time complexity class can be represented by the class of sets of clauses saturated up to redundancy with respect to ordered resolution. Certain characteristic properties of the ordering determine which complexity class is captured [1]. The decision procedures considered in that context are based on the concept of order locality, a generalization of McAllester's concept of locality [10]. [Order] locality, in turn, is a generalization of the concept of an analytic inference system.

In our previous paper, the concept of saturation up to redundancy was somewhat weak in that it was only sufficient, but not necessary, for a set of clauses to present an order-local theory. In the extended version [2] of our previous paper, we have now given a refined definition of saturation precisely characterizing order locality. Capturing non-trivial complexity classes, any such characterization must, however, be undecidable. In [2] we have also shown that our previous incomplete, but practically useful, criterion for order locality is decidable for a large class of orderings.

Meta-Complexity Theorems for Bottom-up Logic Programs In our attempts at applying efficient deductive methods to program analysis, we have reconsidered some of the ideas behind order-locality. Basically two problems had been left open. One problem was that the complexity bounds obtained from the deductive model behind order-local theories were not always competitive with more specialized known algorithms. In particular for congruence closure, when presented by the usual axioms of a congruence relation, we would only get a cubic-time decision procedure, compared with the much better bound of $O(n \log n)$ for the Downey/Sethi/Tarjan procedure. Hence, a specific problem left open was how to extend these methods to equational reasoning, in other words, to efficiently deciding the uniform word problem for suitable classes of presentations of varieties and quasi-varieties.

In parallel to our own attempts, McAllester in [11] studied the application of these and related ideas to inference systems for program analysis (pointer analysis, subtype analysis, interprocedural reachability, and many more). He showed that any theory presented by "bottom-up" Horn clauses can be decided in time proportional to the number of prefix firings of the theory on a given query clause. Given a set of Horn clauses P and a query clause $\Gamma \rightarrow A$, deciding $P \models \Gamma \rightarrow A$ means to decide whether A is contained in $I_{P \cup \Gamma}$, the least Herbrand model of $P \cup \Gamma$. A prefix firing of a theory P on Γ is a substitution instance of a prefix of the body of a clause in P that is contained in $I_{P \cup \Gamma}$. Prefix firings represent partially successful attempts at verifying the antecedent of an inference rule.

Proving this meta-complexity theorems requires to define a deductive machinery involving efficient shared term representations and dynamic programming. McAllester applied this theorem to a number of program analysis problems and in many cases was able to reconstruct the best known bounds. In each case the specification of an analysis by a set of bottom-up Horn clauses was natural and was easily proved correct — in sharp contrast to the original algorithmic formulations which are often tedious and impenetrable.

The weakness of McAllester's approach was again related to not properly dealing with equality. In many forms of program analysis one is faced with computing reachability facts with respect to a quasi-ordering. Being able to deal efficiently with the congruence generated by such an ordering (that is, with the strongly connected components of the reachability graph) is what makes an analysis usable in practice on large programs [13].

In [7] we presented a partial solution to this problem. We proved a meta-complexity theorem for an extended concept of bottom-up logic programs. In these programs one is also able to specify deletion of redundant information. Redundancy is one of the major sources of inefficiency in deduction processes as conclusions derived from redundant facts are redundant themselves and the resources needed to derive them are wasted.

We have worked with rules that have marked atoms in their antecedents — those atoms that are to be deleted after the rule has fired. As in the presence of deletion saturation of a database of facts under a set of rules is non-deterministic, priorities can be assigned to rules to restrict that non-determinism. We have shown that McAllester's previous meta-complexity theorem can be generalized to this case. Again one gets an upper bound proportional to the number of possible

prefix firings of the rules. Hence deletion and priorities are for free. We also showed how to declaratively implement a congruence closure module for conditional equations in this language with guaranteed complexity in $O(\min(n^2, m \log n))$, with n the number of different terms and m the number of different (condition) equations. In that way one internalizes those aspects of equality that do not involve logical variables (and equation solving). Based on the congruence closure module, Henglein's subtype analysis [8] can easily be specified by an inference system with a quadratic number of prefix firings.

Future work should refine the concept of priorities in an instance-based manner, allowing different instances of a rule to have different priorities. That would give one direct means of formalizing algorithms that would normally have to be defined via priority queues. For instance, minimal spanning trees can be computed by a two-rule program on top of union-find if rules referring to edges in graphs can be processed in an order related to their associated costs.

Polynomial-Time Uniform Word Problems In [6] we compare three approaches to polynomial time decidability for uniform word problems for quasi-varieties. Two of the approaches, by Evans [5] and Burris [4], respectively, are semantical, referring to certain embeddability and axiomatizability properties of partial algebras. One may consider these criteria as formalizations of a method dating back to Skolem [12]. What makes words problems undecidable in general are the congruence axioms $x = y \rightarrow f(x) = f(y)$ for function symbols. In a purely relational (Datalog) world, the congruence axioms $x = y, r_f(x, z) \rightarrow r_f(y, z)$, as well as all the other axioms for an equivalence relation, remain within the (polynomially decidable) Datalog world. Skolem's method was based on this observation and he gave examples (in particular lattice theory) where the replacement of functions by partial functions (unique relations) does not affect [un-]satisfiability. [Un-]satisfiability is preserved precisely when certain embeddability criteria for finite partial algebras can be established.

An at first sight quite different approach is proof-theoretic in nature, inspired by McAllester's concept of local inference [10]. We defined two closely related notions of locality for equational Horn theories and proved that both the criteria by Evans and Burris lie in between these two concepts. In particular, our new notion of weak locality subsumes both Evans' and Burris' method. Hence we showed that the relational approach by Skolem, Evans, and Burris, and the subterm approach of McAllester for PTIME-decidability are essentially equivalent.

It might be possible to exploit the relation between the semantic and proof-theoretic concepts for mutually transferring further techniques. In particular, certain (yet to be established) amalgamation properties for free partial algebras would induce combination results for restricted kinds of unions of local theories.

References

[1] D. Basin and H. Ganzinger. Complexity analysis based on ordered resolution. In *Proc. 11th IEEE Symposium on Logic in Computer Science*, pages 456–465. IEEE Computer Society Press, 1996.

• [2] D. A. Basin and H. Ganzinger. Automated complexity analysis based on ordered resolution. *Journal of the ACM*, 48(1):70–109, 2001.

[3] S. Bellantoni and S. Cook. A new recursion-theoretic characterization of the polytime function. *Computational Complexity*, 2:97–110, 1992.

[4] S. Burris. Polynomial time uniform word problems. *Mathematical Logic Quarterly*, 41(2):173–182, 1995.

[5] T. Evans. Embeddability and the word problem. *J. London Math. Soc.*, 28:76–80, 1953.

- [6] H. Ganzinger. Relating semantic and proof-theoretic concepts for polynomial time decidability of uniform word problems. In D. A. Williams, editor, *Proceedings of the 16th IEEE Symposium on Logic in Computer Science (LICS-01)*, pages 81–90, Boston, USA, 2001. IEEE Computer Society.
- [7] H. Ganzinger and D. McAllester. A new meta-complexity theorem for bottom-up logic programs. In R. Goré, A. Leitsch, and T. Nipkow, editors, *Proceedings of the 1st International Joint Conference on Automated Reasoning (IJCAR-2001)*, volume 2083 of *Lecture Notes in Artificial Intelligence*, pages 514–528, Siena, Italy, 2001. Springer.

[8] F. Henglein. Breaking through the n^3 barrier: Faster object type inference. *Theory and Practice of Object Systems*, 5(1):57–72, 1999.

[9] M. Hofmann. Linear types and non size-increasing polynomial time computation. In *Proceedings of the 14th Symposium on Logic in Computer Science*, pages 464–473, 1999.

[10] D. McAllester. Automatic recognition of tractability in inference relations. *J. Association for Computing Machinery*, 40(2):284–303, 1993.

[11] D. McAllester. On the complexity analysis of static analyses. In A. Cortesi and R. Filé, editors, *Static Analysis — 6th International Symposium, SAS'99*, LNCS 1694, pages 312–329, Venice, Italy, September 1999. Springer-Verlag.

[12] T. Skolem. Logisch-kombinatorische Untersuchungen über die Erfüllbarkeit und Beweisbarkeit mathematischer Sätze nebst einem Theorem über dichte Mengen. *Skrifter utgitt av Videnskapsselskapet i Kristiania*, I(4):1–36, 1920. Also in: Th. Skolem, Selected Works in Logic (Jens Erik Fendstad, ed.), Scand. Univ. Books, Universitetsforlaget, Oslo, 1970, pp. 103–136.

[13] Z. Su, M. Fähndrich, and A. Aiken. Projection merging: reducing redundancies in inclusion constraint graphs. In *Proceedings of the 27th ACM SIGPLAN-SIGACT symposium on Principles of programming languags*, pages 81–95, 2000.

3.3 Combination of Algebraic and Logic Methods

Investigators: Harald Ganzinger, Viorica Sofronie-Stokkermans, Uwe Waldmann

To be usable in practice, a theorem prover must be able to deal with problems that involve both standard algebraic theories and uninterpreted function and predicate symbols that are specific for a particular domain. Many standard theories are defined by axioms like associativity, commutativity, or transitivity that allow a huge number of inferences. Handling these axioms naïvely leads to an explosion of the search space. It is therefore necessary to incorporate algebraic theories into the inference rules of the calculus in a more sophisticated way.

We have continued to work on combining algebraic and logic methods in automated theorem proving for varieties of groups, rings, fields, and lattices. We are pursuing two methods. One method is to refine superposition and chaining inference systems to specific theories. Equational reasoning, or more generally, reasoning about transitive relations, in the presence of the associativity and commutativity axioms is known to be difficult – theoretically as well as practically. Using AC-unification and extended clauses the worst inefficiencies of a naïve approach can be avoided, but still the need for extended clauses leads to numerous variable overlaps, one of the most prolific types of inferences in resolution or superposition style calculi. The hope was that more specialized calculi would avoid many, if not all, of these implicit variable overlaps. Inequality reasoning as it is required for reasoning about orderings, in particular ordered groups and lattices, suffers from related problems. There, even overlaps below variables have to be considered.

Our second method employs representation theorems with the intention of encoding, at the logical level, properties of algebraic structures in a way such that logical deduction with implication and equivalence (which automated provers can do well) directly simulates reasoning with inequations and equations.

Abelian Groups, Rings, and Fields In previous work we have developed a superposition calculus for theorem proving modulo divisible torsion-free Abelian groups (DTAGs) [28] and a chaining calculus for dense total orderings without endpoints [5]. In order to avoid highly inefficient variable overlaps during the saturation, these calculi consist of two parts: The first part is an inference system that operates on clauses in which every variable is shielded (that is, it occurs at least once below a free function symbol). Shielded variables cannot be maximal terms in a clause, so by the ordering restrictions of the inference rules, overlaps with shielded variables are excluded. The inference system is then combined with a variable elimination algorithm that transforms clauses with unshielded variables into logically equivalent clauses without unshielded variables.

The vast majority of applications of divisible torsion-free Abelian groups (and in particular of the rationals or reals) requires also an ordering. It is therefore natural to investigate the combination of the theories of divisible torsion-free Abelian groups and dense total orderings without endpoints. In [31, 30] we have presented a superposition and chaining calculus for totally ordered divisible Abelian groups (TODAGs). Again there is a base calculus for clauses that contain only shielded variables and an algorithm to eliminate unshielded variables. While the base calculus is a fairly obvious combination of its two ancestor calculi, chaining and DTAG-superposition, proving its completeness turns out to be a tedious task: The model constructions for both chaining and DTAG-superposition are variants of the model construction for standard superposition, but they extend the latter in different directions. It is possible to adapt the completeness proof for DTAG-superposition, but the technical details are rather involved.

Like the DTAG-superposition calculus, the TODAG-superposition calculus is guaranteed to terminate if all free function symbols are the result of skolemization. (The proof is virtually the same as in [29].) It is therefore a decision procedure for the elementary theory of totally ordered divisible Abelian groups.

We are currently working on an implementation of DTAG-superposition within SPASS, which should then be extended to TODAG-superposition.

One of the drawbacks of both DTAG-superposition and TODAG-superposition is the need for term abstraction. Is it possible to extend the paramodulation calculus modulo Abelian groups by Godoy and Nieuwenhuis (which avoids abstraction) in such a way that it becomes usable for either divisible torsion-free Abelian groups or totally ordered divisible Abelian groups? If so, is it also possible to combine it with variable elimination? (Note that the combination method that we have used so far would not work if syntactic cancellation were mandatory as in Godoy and Nieuwenhuis's calculus [14].)

There are further algebraic structures for which variable elimination is possible. Can one design similar theorem proving calculi for some of these structures? How do the base calculi look like and how can they be combined with the variable elimination algorithms?

Theorem Proving in Varieties of Distributive Lattices with Operators In [22], we proposed a method based on representation theorems for automated theorem proving in the universal theory of certain varieties of distributive lattices with operators.

In [26] we extended our previous results from [22] to a wider class of operators, encompassing generalizations of the modal operators for necessity and possibility from modal logic and generalizations of weakened negations and implications in non-classical logics. This also allowed us to consider (generalizations of) residuation conditions.

More specifically, our method is based on the fact that every bounded distributive lattice is isomorphic to a lattice of sets, a result also known as the Priestley representation theorem for distributive lattices [20]. This representation theorem was extended by Goldblatt [15] to include

so-called join- and meet-hemimorphisms on lattices, operators that generalize the possibility and necessity operators in modal logic, and by us in [23, 25] to operators which are hemimorphisms in some arguments and "anti"-hemimorphisms in other arguments. The representation theorems allow us in many cases to give structure-preserving translations to clause form; the satisfiability of the set of clauses we obtain can for instance be checked by resolution.

The main advantage of the method we propose is that it avoids the explicit use of the full algebraic structure of (distributive) lattices. Instead, we use sets endowed with a reflexive and transitive relation and with additional functions and relations that correspond to the operators in the lattices in a standard way. The lattice operators are thus encoded into logical conjunction and disjunction, which can be handled by any automated theorem prover for first-order logic. Decidability and complexity results can be obtained by using ordered resolution or ordered chaining [5]. The results in [22] extend in a natural way to the more general classes of operators mentioned above. In addition we showed that, by using our method, the uniform word problem for the class of all distributive lattices with operators that satisfy a family of (generalized) residuation conditions can be decided in exponential time.

One direction of future work is to see for which other classes of lattices or Heyting algebras with operators the ideas above lead to decision procedures. We would also like to obtain a better understanding of the link between the decidability and complexity of the theory of (bounded distributive) lattices, or Boolean algebras and the theory obtained by also considering operators on such algebras that are monotone or commute with part of the lattice structure.

Applications to Automated Theorem Proving in Non-Classical Logics The embedding into classical logic mentioned above can be used for automated theorem proving in many classes of non-classical propositional logics. Some such logics are analyzed in [27]. There we show that similar ideas can also be applied for logics that are sound and complete with respect to classes of lattices with operators (not necessarily distributive), or even semilattices with operators. In particular, by using representation theorems for semilattices and lattices we obtain Kripke-style semantics for several classes of substructural logics.

We would like to understand better up to which extent such representation theorems can be helpful for efficient automated theorem proving in more general classes of algebras. In some specific cases there might exist more general ways of proving decidability and tractability of the uniform word problems. This could be achieved for instance by using various criteria for polynomial decidability of uniform word problems (semantical, referring to certain embeddability and axiomatizability properties [21, 8, 6], proof-theoretic, using McAllester's concept of local inference [19], or using the relationship between these two types [9]).

Chaining Calculi for Finitely-Valued Logics In [11] we show that general saturation-based techniques for first-order theories of transitive relations, in particular congruences and partial or total orderings — the chaining calculi introduced by Bachmair and Ganzinger [4, 5] —, can easily be specialized to deal with many-valued logics. With this method the previous results on automated theorem proving for many-valued logics can be greatly improved. Apart from reconstructing known completeness results for existing methods, including many-valued resolution [3], regular hyper-resolution [16], and annotated resolution [17, 18], the inference systems which we obtain are much more restricted, in particular by ordering constraints and selection functions. The specialization of the general chaining inference systems is very direct and does not involve any sophisticated encodings.

The general concept of redundancy for clauses and inferences that is compatible with the chaining calculi allows us to equip the inference systems with strong techniques for simplification and for the elimination of global redundancies. Without simplification any theorem proving method, regardless of how much restricted the inferences are that are employed for proof search, will be hopelessly inefficient. Knowing what simplification techniques are admissible is, therefore, extremely important in practice. In particular, all our inference systems are compatible with the eager rewriting of subformulas by "smaller" equivalent formulas. The method allows us to use existing efficient implementations of chaining techniques including SPASS [32] and Saturate [10].

Unification for Bounded Distributive Lattices In [24] we give a resolution-based procedure for deciding unifiability in the variety D_{01} of bounded distributive lattices. From an algebraic point of view, unification can be seen as solving (systems of) equations in the initial or free algebra of an equational theory. Apart from its theoretical interest, unification is used e.g. in resolution-based theorem proving and in term rewriting to deal with certain equational axioms (such as associativity and commutativity). Unification problems in semilattice- and lattice-based structures are becoming of increasing interest in computer science (cf. e.g. the results of Baader and Narendran on unification of concept terms in description logics [1]; similar possible applications in set constraints may also be of interest).

In [12], Gerhard and Petrich give a criterion for unifiability (with free constants) of two terms in the theory of distributive lattices. Ghilardi [13] showed that the equational class D of distributive lattices has unification type 0, i.e. there exist D-unification problems with no minimal complete set of unifiers. We are not aware of any other results on unification for distributive lattices, e.g. concerning its complexity.

In [24] we show that the use of the Priestley representation for D_{01} allows us to give an algorithm based on resolution with constrained clauses for the unification problem in this class. In a first step (structure-preserving translation to clause form), testing the satisfiability of a unification problem S is reduced to the problem of checking the satisfiability of a set Φ_S of clauses. The structure-preserving translation to clause form uses the description of the free bounded distributive lattice over C as the lattice of all upwards-sets of $(P(C), \subseteq)$, which is a consequence of the Priestley representation theorem. Expressing Φ_S as a set of constrained clauses further simplifies the representation of the problem. Next, ordered resolution with selection for (constrained) clauses is used for testing the satisfiability of Φ_S. We show that similar ideas can be used for unification with linear constant restrictions.

The main advantage of our approach is that it makes it much easier to treat the unification problem for bounded distributive lattices, by using results in resolution theory. As a byproduct, using Prop. 5.6 in [2], our results also show that resolution can be used for deciding the positive theory of the variety of bounded distributive lattices.

It seems that many of the results in this paper can be extended without major difficulties to other varieties in which the free algebras have a description similar to the free bounded distributive lattices. This is the case for many subvarieties of the variety of Ockham algebras (bounded distributive lattices with a lattice antimorphism), such as, e.g., the variety of De Morgan algebras. We think that the ideas might be even more general, and applicable to classes of algebras for which representation theorems in terms of homomorphisms into a finite algebra exist (or e.g. natural dualities [7]).

References

[1] F. Baader and P. Narendran. Unification of concept terms in description logics. In H. Prade, editor, *Proceedings of ECAI'98*, pages 331–335. Wiley, 1998.

[2] F. Baader and K. Schulz. Unification in the union of disjoint equational theories: Combining decision procedures. *J. Symbolic Computation*, 21:211–243, 1996.

[3] M. Baaz and C. Fermüller. Resolution-based theorem proving for many-valued logics. *J. of Symbolic Computation*, 19:353–391, 1995.

[4] L. Bachmair and H. Ganzinger. On restrictions of ordered paramodulation with simplification. In *Proceedings of the 10th Conference on Automated Deduction (CADE-10)*, LNCS 449, pages 427–441. Springer Verlag, 1990.

[5] L. Bachmair and H. Ganzinger. Ordered chaining calculi for first-order theories of transitive relations. *Journal of the ACM*, 45(6):1007–1049, 1998.

[6] S. Burris. Polynomial time uniform word problems. *Mathematical Logic Quarterly*, 41(2):173–182, 1995.

[7] D. Clark and B. Davey. *Natural dualities for the working algebraist*, volume 57 of *Cambridge studies in advanced mathematics*. Cambridge University Press, 1st ed. edition, 1998.

[8] T. Evans. The word problem for abstract algebras. *J. London Math. Soc.*, 26:64–71, 1951.

• [9] H. Ganzinger. Relating semantic and proof-theoretic concepts for polynomial time decidability of uniform word problems. In D. A. Williams, editor, *Proceedings of the 16th IEEE Symposium on Logic in Computer Science (LICS-01)*, pages 81–90, Boston, USA, 2001. IEEE Computer Society.

[10] H. Ganzinger, R. Nieuwenhuis, and P. Nivela. The Saturate system. Available on the World-Wide Web under URL http://www.mpi-sb.mpg.de/SATURATE/Saturate.html, 1994.

• [11] H. Ganzinger and V. Sofronie-Stokkermans. Chaining techniques for automated theorem proving in many-valued logics. In *Proceedings of the 30th IEEE International Symposium on Multiple-Valued Logic (ISMVL-00)*, pages 337–344, Portland, Oregon, 2000. IEEE Computer Society Technical Committee on Multiple-Valued Logic; Oregon Center for Advanced Technology Education, IEEE Computer Society.

[12] J. Gerhard and M. Petrich. Unification in free distributive lattices. *Theoretical Computer Science*, 126(2):237–257, 1994.

[13] S. Ghilardi. Unification through projectivity. *J. Logic and Computation*, 7(6):733–752, 1997.

[14] G. Godoy and R. Nieuwenhuis. Paramodulation with built-in abelian groups. In *Fifteenth Annual IEEE Symposium on Logic in Computer Science*, pages 413–424, Santa Barbara, CA, USA, June 26–28, 2000. IEEE Computer Society, Los Alamitos, CA, USA.

[15] R. Goldblatt. Varieties of complex algebras. *Annals of Pure and Applied Logic*, 44(3):153–301, 1989.

[16] R. Hähnle. Exploiting data dependencies in many-valued logics. *J. of Appl. Non-Classical Logics*, 6(1):49–69, 1996.

[17] M. Kifer and M. Lozinskii. A logic for reasoning with inconsistency. *J. of Automated Reasoning*, 9:179–215, 1992.

[18] J. Lu, N. Murray, and E. Rosenthal. A framework for reasoning in multiple-valued logics. *J. of Automated Reasoning*, 21(1):39–67, 1998.

[19] D. McAllester. Automatic recognition of tractability in inference relations. *J. Association for Computing Machinery*, 40(2):284–303, 1993.

[20] H. Priestley. Representation of distributive lattices by means of ordered Stone spaces. *Bull. London Math Soc.*, 2:186–190, 1970.

[21] T. Skolem. Logisch-kombinatorische Untersuchungen über die Erfüllbarkeit und Beweisbarkeit mathematischer Sätze nebst einem Theorem über dichte Mengen. *Skrifter utgitt av Videnskapsselskapet i Kristiania*, I(4):1–36, 1920. Also in: Th. Skolem, Selected Works in Logic (Jens Erik Fendstad, ed.), Scand. Univ. Books, Universitetsforlaget, Oslo, 1970, pp. 103–136.

- [22] V. Sofronie-Stokkermans. On the universal theory of varieties of distributive lattices with operators: Some decidability and complexity results. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction (CADE-16)*, volume 1632 of *Lecture Notes in Artificial Intelligence*, pages 157–171, Trento, Italy, 1999. Springer.
- [23] V. Sofronie-Stokkermans. Duality and canonical extensions of bounded distributive lattices with operators and applications to the semantics of non-classical logics. Part I. *Studia Logica*, 64(1):93–132, February 2000.
- [24] V. Sofronie-Stokkermans. On unification for bounded distributive lattices. In D. McAllester, editor, *Proceedings of the 17th International Conference on Automated Deduction (CADE-17)*, volume 1831 of *Lecture Notes in Artificial Intelligence*, pages 465–481, Pittsburgh, Pennsylvania, USA, 2000. Springer.

[25] V. Sofronie-Stokkermans. A note on representation theorems for distributive lattices with operators. Manuscript, 2001.

- [26] V. Sofronie-Stokkermans. On the universal theory of varieties of distributive lattices with operators: Some decidability and complexity results. Research Report MPI-I-2001-2-005, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, 2001.
- [27] V. Sofronie-Stokkermans. Representation theorems and the semantics of (semi)lattice-based logics. In *Proceedings of the 31st IEEE International Symposium on Multiple-Valued Logics*, pages 125–134, Warsaw, Poland, 2001. IEEE Computer Society, Institute of Computer Science Polish Academy of Sciences, IEEE Computer Society.

[28] U. Waldmann. Superposition for divisible torsion-free abelian groups. In C. Kirchner and H. Kirchner, editors, *Proceedings of the 15th International Conference on Automated Deduction (CADE-98)*, volume 1421 of *Lecture Notes in Artificial Intelligence*, pages 144–159, Lindau, Germany, 1998. Springer.

- [29] U. Waldmann. Cancellative superposition decides the theory of divisible torsion-free abelian groups. In H. Ganzinger, D. McAllester, and A. Voronkov, editors, *Proceedings of the 6th International Conference on Logic for Programming and Automated Reasoning (LPAR-99)*, volume 1705 of *Lecture Notes in Artificial Intelligence*, pages 131–147, Tbilisi, Georgia, 1999. Springer. Max-Planck-Institut für Informatik, Saarbrücken.
- [30] U. Waldmann. Superposition and chaining for totally ordered divisible abelian groups. Research Report MPI-I-2001-2-001, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, April 2001.
- [31] U. Waldmann. Superposition and chaining for totally ordered divisible abelian groups (extended abstract). In R. Goré, A. Leitsch, and T. Nipkow, editors, *Proceedings of the 1st International Joint Conference on Automated Reasoning (IJCAR-2001)*, volume 2083 of *Lecture Notes in Artificial Intelligence*, pages 226–241, Siena, Italy, 2001. Springer.

[32] C. Weidenbach, B. Afshordel, U. Brahm, C. Cohrs, T. Engel, E. Keen, C. Theobalt, and D. Topic. System description: SPASS version 1.0.0. In H. Ganzinger, editor, *Proceedings of the 16th Conference on Automated Deduction (CADE-16)*, LNAI 1632, pages 314–318. Springer Verlag, 1999.

3.4 Integration of Automated First-Order Provers into Higher-Order Provers

Proof Generation for Resolution Investigator: Hans de Nivelle

One of the obstacles for the integration of first order provers into interactive systems is the problem of reliability. Interactive theorem provers usually implement a simple but powerful logical calculus.

There is a small part of the code in the interactive system that does the actual manipulation of the formulae and the proofs. All other parts of the interactive prover, proof editors and tactical proof searchers do not manipulate the proofs directly, but do this by making calls to the trusted code. In this way only a small part of the interactive theorem prover needs to be really trusted. Users of such systems are not willing to give up this principle, and to trust a first order theorem prover. First order theorem provers implement complicated calculi. To make things worse, they contain extremely complicated data structures for maximal technical efficiency. Trusting such a system for safety critical applications would not be a sensible idea.

Because of this, it is necessary to be able to generate proofs from a resolution prover in some simple calculus. We have chosen type theory for proof representation. The rules of type theory are simple, it is formally defined what a proof is, and how it is represented. A proof checker for type theory can be easily implemented. Yet on the other side, the calculus is very powerful. It implements full higher order logic, lemmas and definitions.

We have studied the problem of how to efficiently generate type theory proofs from resolution proofs. A resolution proof for a first order formula consists of a clausal normal form transformation, followed by a resolution refutation of the clausal normal form. In [2], it was studied how proofs can be generated from the resolution refutation on the clause level. The paper presents methods that yield type theory proofs in linear time from resolution proofs. In order to obtain linear time, it is necessary to use a non-standard representation of clauses. The clause $A_1 \vee \cdots \vee A_n$ with free variables $\overline{x}$ has to be represented as $\forall x(\neg A_1 \to \cdots \to \neg A_n \to \bot)$.

We have also essentially solved the problem of efficiently generating short proofs for various optimized clausal normal form transformations [1]. The clausal normal form transformation can be seen as a (generalized) confluent rewrite system. One can apply replacement rules in any convenient order, until one reaches a normal form.

In order to prove that the original formula logically implies the normal form, one needs a justification for each rule, and rules for lifting justifications into each possible context. For example, if one replaces A by B in context $F[\ ]$, one needs to transform the proof of $A \to B$ into a proof of $F[A] \to F[B]$. The main obstacle is the so called context problem: The complexity of the proof of $F[A] \to F[B]$ is quadratic in $F[\ ]$, where it is only linear in A and B. Because of this, it becomes important in which order the replacements are applied. One should try to combine consecutive rule applications in the same context. This would make the algorithm more complicated. However, in [1], it is shown that if one applies appropriate proof normalizations on the constructed proofs, then the result is independent of the order in which replacements were made. Combining [2] and [1], the complete resolution process is now covered.

The type theory proof generator on the clause level has been implemented in an experimental way, as part of the theorem prover Bliksem `http://www.mpi-sb.mpg.de/~bliksem` . After the implementation was completed, it became clear that one cannot add a proof generator to a theorem prover afterwards. It has to be integrated into the design from the beginning. Because of this, a full rewrite is necessary, which is planned for next year.

Implementation of the clausal normal form transformation with proof generation has started.

References

- [1] H. de Nivelle. Translation of resolution proofs into higher order natural deduction using type theory. Research Report MPI-I-2001-2-004, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, August 2001.
- [2] H. de Nivelle, M. Bezem, and D. Hendriks. Automated proof construction in type theory using resolution. In D. McAllester, editor, *Proceedings of the 17th International Conference on Automated Deduction*

(CADE-17), volume 1831 of *Lecture Notes in Artificial Intelligence*, pages 148–163, Carnegie Mellon University, Pittsburgh, PA, USA, June 2000. Springer.

3.5 Implementation Techniques for First-Order Provers

Investigator: Harald Ganzinger

Indexing data structures have a crucial impact on the performance of automated theorem provers. Examples are discrimination trees, which are like tries where terms are seen as strings and common prefixes are shared, and substitution trees, where terms keep their tree structure and all common contexts can be shared. In [1] we describe a new indexing data structure, called context trees, where, by means of a limited kind of context variables, also common subterms can be shared, even if they occur below different function symbols. The paper also provides some initial evidence for the practical value of context trees. We showed how to implement context trees based on Curry terms and on an extension of substitution trees with equality constraints, where one also does not distinguish between internal and external variables. Experiments with matching benchmarks show that our preliminary implementation is already competitive with tightly coded current state-of-the-art implementations of the other main techniques. In particular space consumption of context trees is significantly less than for other index structures.

References

- [1] H. Ganzinger, R. Nieuwenhuis, and P. Nivela. Context trees. In R. Nieuwenhuis, P. Nivela, R. Goré, A. Leitsch, and T. Nipkow, editors, *Proceedings of the 1st International Joint Conference on Automated Reasoning (IJCAR-2001)*, volume 2083 of *Lecture Notes in Artificial Intelligence*, pages 242–256, Siena, Italy, 2001. Springer.

4 Logical Methods for Program Analysis

There is a whole spectrum of verification problems according to the level of abstraction in the problem formulation. The state of the art is roughly that high-level properties (related to the system design, such as safety in a mutual-exclusion protocol) can be checked for high-level system descriptions, and low-level properties (related to the implementation of the design, e.g. through arrays) can be checked for programs at the source level (written in, say, C or Java). We want to gradually increase the 'doable' quotient between the system level and the property level. This is because we envisage future programming languages that will allow (force) the programmer to express more and more high-level properties in the source code and have the compiler check them.

The level of abstraction is inherently connected with the formalism for specifying the system and the property, respectively. Examples of formalisms for expressing complex control dependencies are the ambient calculus or other process algebras; formalisms for expressing complex data dependencies are timed or hybrid automata, machine models with fifo queues etc.. At a high level of abstraction, one abstracts away (exactly) the irrelevant details and captures the essence of complex computational phenomena (such as communication between concurrent processes) in a concise way. Temporal logic or other modal logics are widely used formalisms to express complex dependencies in high-level properties describing the execution behavior. A low level of abstraction is used to account for implementation details; here, verification problems today target only invariance properties (regarding operations on arrays, pointers or other low-level data structures) that can be specified by types, although formalisms to specify restricted cases of temporal properties of programs start to appear.

We need to work on both ends in order to "increase the doable quotient", i.e. on methods for 'high-level' verification problems (formulated in formalisms such as the ambient calculus and corresponding modal logics), and on automated abstraction methods that, in a sense, map a low-level description to a high-level one. For example, if a method maps a program to a high-level system description in the form of a finite-state model, then the verification problem is turned into one to which existing methods can be applied (a method may also leave that abstract model implicit). Today, we know how to map interprocedurally recursive programs into a process algebra that is amenable to verification algorithms, but adding parallel threads makes this an open problem. We know abstraction methods that map infinite data domains into finite ones, essentially by partitioning, but we do not yet know how to automatically find the predicates that induce a 'good' partitioning. One possible outcome of this research is that we will know the right mix of 'ingredients' of the abstraction (of the control or the data dependencies) that are computed resp. provided by the programmer. For example, if an inductive invariant can be composed from a finite set of atomic formulas (provided e.g. by the programmer) then an automated tool can synthesize an inductive invariant.

4.1 Set-Based Analysis

Investigators: Witold Charatonik, Andreas Podelski, Jean-Marc Talbot

Set-based analysis consists of two steps: the extraction of a set constraint from the syntax of a given program and the resolution of the constraint, i.e. the computation of a particular solution. A least or greatest solution is computed when the semantics of the initial program is defined by means of a least resp. greatest fixed point computation, depending on the execution property to be analyzed.

In [3] we consider the 'alternation-free' fragment of the Horn μ-calculus. The Horn μ-calculus [1]

formalizes the idea of analyzing programs with respect to the least and greatest fixpoint semantics; it may be viewed as a logic programming language allowing arbitrary nesting of least and greatest fixed points. We give new, optimal algorithms for our fragment that is based on tree-automata techniques. Our original question whether these techniques can be pushed beyond that fragment to the full Horn μ-calculus remains as an open problem.

In [2] we showed that the restriction to unary signatures in set-based analysis does not yield (at least theoretically) any speed-up in the analysis. This is because the satisfiability of set constraints that are both definite and unary is a DEXPTIME-complete problem (that is, it has the same complexity as in the non-unary case). The result gives a negative answer to a question raised in 1991 by Frühwirth, Shapiro, Vardi and Yardeni and is quite surprising: due to its connections with the theory of word automata, the problem was expected to be PSPACE-complete. This is interesting also because many implementors of set-based analysis believed that introducing yet another degree of approximation by using unary signatures (which amounts to using the simpler domain of path-closed sets of trees) changes the algorithmic complexity of the analysis problem. Our result refutes that belief.

References

[1] W. Charatonik, D. McAllester, D. Niwinski, A. Podelski, and I. Walukiewicz. The horn mu-calculus. In V. Pratt, editor, *Proceedings of the 13th Annual IEEE Symposium on Logic in Computer Science (LICS-98)*, pages 58–69, Indianapolis, Indiana, June 1998. IEEE Technical Committee on Mathematical Foundations of Computing, IEEE Computer Society Press.

- [2] W. Charatonik, A. Podelski, and J.-M. Talbot. Paths vs. trees in set-based program analysis. In *Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL-00)*, pages 330–337, Boston, Massachusetts, USA, January 2000. ACM.

- [3] J.-M. Talbot. On the alternation-free Horn mu-calculus. In M. Parigot and A. Voronkov, editors, *Proceedings of the 7th International Conference on Logic for Programming and Automated Reasoning (LPAR-2000)*, volume 1955 of *Lecture Notes in Artificial Intelligence*, pages 418–435, La Reunion, France, 2000. Springer.

4.2 Mobile Ambients

Investigators: Witold Charatonik, Supratik Mukhopadhyay, Jean-Marc Talbot

With the development of the internet, mobile computation receives more and more scientific attention. One tries to capture its essential features in a way similar to that of λ-calculus capturing functional computation. The ambient calculus of Cardelli and Gordon is widely accepted as a good model for mobile computation over the world-wide web. In this calculus, an ambient is a named cluster of running processes and nested sub-ambients. Each computation state has a spatial structure, the tree induced by the nesting of ambients. Mobility is abstractly represented by re-arrangement of this tree: an ambient may move inside or outside other ambients.

Process calculi are already recognized in specification and verification of concurrent systems, especially communication protocols; our motivation is to generalize this work to mobility protocols or, more generally, to programs that contain mobility instructions. The starting point is to classify particular cases of the general verification problem according to the very first complexity criterion, namely decidability. Then the hope is that the methods obtained for the easy decidable cases together with an abstraction mechanism from programs on the source level to ambient calculus will give rise to verification algorithms for mobile programs on the source level.

To reason about distributed and mobile computations programmed in the ambient calculus, Cardelli and Gordon introduced a modal logic that apart from standard temporal modalities for describing the evolution of processes includes novel spatial modalities for describing the tree structure of ambient processes. They also proved that for a fragment of the calculus and logic (no name restriction, no replication, no composition adjunct) the model-checking problem is decidable. In [1, 2, 3] we settle the complexity bounds of the model checking problem for the ambient calculus against the ambient logic. We show that if either the calculus contains replication or the logic contains composition adjunct, the problem is undecidable, which answers a question raised earlier by Cardelli and Gordon. In the case of replication-free calculus and logic without composition adjunct we prove that the problem is PSPACE-complete, even if both calculus and logic contain the respective constructs for restricted names. For the complexity upper-bound, we devise a new representation of processes that remains of polynomial size during process execution; this allows us to keep the model checking procedure in polynomial space. Moreover, we prove PSPACE-hardness of the problem for several quite simple fragments of the calculus and the logic; this suggests that there are no interesting fragments with polynomial-time model checking algorithms.

Our current work includes an extension of the decidable fragment of the ambient calculus with a restricted form of recursion (in a way different from replication) that allows more flexibility in modeling computation without increasing the complexity of the model-checking problem. In particular, in contrast to Cardelli and Gordon, we are able to model check processes with infinite behavior.

References

- [1] W. Charatonik, S. Dal-Zilio, A. D. Gordon, S. Mukhopadhyay, and J.-M. Talbot. The complexity of model checking mobile ambients. In F. Honsell and M. Miculan, editors, *Foundations of Software Science and Computation Structures. Proceedings of the 4th International Conference (FOSSACS-01). Held as Part of the Joint European Conferences on Theory and Practice of Software (ETAPS-01)*, volume 2030 of *Lecture Notes in Computer Science*, pages 152–167, Genova, Italy, April 2001. Springer.
- [2] W. Charatonik, S. Dal-Zilio, A. D. Gordon, S. Mukhopadhyay, and J.-M. Talbot. The complexity of model checking mobile ambients. Technical Report MSR-TR-2001-03, Microsoft Research, Microsoft Corporation, Redmond, USA, May 2001.
- [3] W. Charatonik and J.-M. Talbot. The decidability of model checking mobile ambients. In *Computer Science Logic, 15th Annual Conference of the EACSL*, Lecture Notes in Computer Science, Paris, France, September 2001. European Association for Computer Science Logic (EACSL), Springer. To appear.

4.3 Software Model Checking

Investigator: Andreas Podelski

The goal of any automated program verification method is to statically compute (possibly partial) answers to questions about the runtime behavior of a program. Such questions may relate to data structures at specific program points as in data flow analysis ("Does an expression always have a constant value?") or to more general properties of program execution sequences (viz. a temporal property, such as termination) as in model checking.

In [3], we express the program property to be analyzed as a specific solution of an appropriate constraint. Computing the property amounts to solving the constraint, i.e. transforming it into a solved form (which directly exhibits the specific solution); the transformation is a sequence of inference steps adding more and more consequences under a set of inference rules that is chosen according to the data domain of the program. Fixpoint iteration as in model checking for finite

systems or for timed or hybrid automata is one special instance in this schema; model checking algorithms for programs with stacks, queues or trees are other instances. We thus obtain a framework in which we can uniformly explain the different techniques from theorem proving, program analysis and model checking (deduction for reasoning over different datatypes, fixpoint iteration for the computation of abstract program semantics, and tests of temporal runtime properties). The common framework is a possible starting point for developing model checking algorithms for programs with mixed data domains, and for developing 'logic' abstraction mechanisms (i.e. with automated reasoning using the translation of the program as a logical formula).

A fundamental problem arises in software verification when threads can be generated dynamically, each of them calling recursive procedures that again spawn threads etc.. In [2] we take two simplifying assumptions: finite data domains and a weak form of synchronization, which occurs in abstractions of programs, and give algorithms to compute the sets of reachable (or backward reachable) configurations in linear time in the program size. This is a potential starting point for addressing the general problem (without the simplifying assumptions). We already know how to abstract data domains to finite domains, but it is an open problem how a complex synchronization mechanism between threads can be abstracted to a simpler one (without giving away the precision needed to verify the given property of the program).

The question is how a 'good' abstraction of a program into one with finite data domains can be computed efficiently. This question must be preceded by the formulation of the problem in algorithmic terms. The ad-hoc idea of abstracting states according to their evaluation under a finite number of state predicates ("predicate abstraction") leads to unrealistic implementations of the abstraction procedure. In [1] we give a precise definition of a 'better' abstraction for which a realistic implementation exists. We trade the gain of efficiency with a loss of precision that may arise in general but does not arise in the case of deterministic programs. This is important since each single instruction in a program is usually deterministic. The computation of the abstraction is very costly. Once the goal of this computation is established, and only after such a goal is formulated as an algorithmic problem, it is possible to propose and evaluate optimizations. Our work enables this to happen.

References

- [1] T. Ball, A. Podelski, and S. K. Rajamani. Boolean and cartesian abstraction for model checking C programs. In T. Margaria and W. Yi, editors, *Proceedings of the 7th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2001)*, volume 2031 of *Lecture Notes in Computer Science*, pages 268–283, Genova, Italy, April 2001. Springer.
- [2] J. Esparza and A. Podelski. Efficient algorithms for pre* and post* on interprocedural parallel flow graphs. In *Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL-00)*, pages 1–11, Boston, Massachusetts, USA, January 2000. ACM, ACM.
- [3] A. Podelski. Model checking as constraint solving. In J. Palsberg, editor, *Proceedings of the 7th International Symposium on Static Analysis (SAS-00)*, volume 1824 of *Lecture Notes in Computer Science*, pages 221–237, Santa Barbara, California, USA, 2000. Springer.

4.4 Infinite-State Systems

Investigators: Giorgio Delzanno, Supratik Mukhopadhyay, Andreas Podelski

In [4] we explain the algorithm underlying the Uppaal model checking tool for timed systems in terms of the top-down query evaluation of Constraint Data Bases. This new connection has lead us to a direct implementation of the algorithm using the constraint facilities of the Sicstus Prolog

system. Using the formal basis of the algorithm, we could extend its applicability to a temporal logic enriched with a restricted class of (unbounded) liveness properties. The open problem is to determine how far this extension can be pushed, and thus determine the boundary for the applicability of local model checking algorithms (which were shown to be better in practice than the global ones).

The model of communicating finite-state machines is used to specify protocols for distributed systems connected with FIFO channels. In [1] we address the problem of automatically validating such protocols. Our method is based on the symbolic representation of sets of channel contents by first-order terms, and on the correspondence between the ground and the non-ground semantics of logic programs. Our practical experiments are promising, but we still do not have a systematic or formally founded way to design and analyze verification methods for systems with queues.

Forward analysis for timed automata is possibly non-terminating. In [3] we show that it does terminate for all automata where the (real-valued) clock variables are used to measure the progress of time; i.e., they are reset before they are read. This result explains why forward analysis succeeds in practical experiments (which is again why it is employed in tools such as Uppaal, although alternatives in the form of terminating algorithms exist). As a potential application of our sufficient termination condition, its test could be part of model checking experiments (e.g. a specification could be 'tuned' to meet the condition). The test would be practical only if it could be applied to the components instead of the parallel composed timed automaton as a whole. Such compositional tests are a topic of future research.

The bisimilarity of timed systems can be computed on the finite region graph which, however, may be prohibitively large. In [6] we characterize timed bisimilarity in terms of models of constraint databases; this allows us to present a local algorithm for deciding timed bisimilarity that is symbolic and thus avoids the need for constructing the region graph.

Forward analysis for timed systems is a simple fixpoint iteration. In [5], we use our constraint-based setup to implement widening operators that accelerate the fixpoint iteration in practical examples, as we could find out in our experiments. These widening operators are novel in that they reason on the program text; we hope to apply them to other settings than timed automata.

What are the BDD's for infinite data domains? In [2] we develop data structures to represent collections of upwards closed sets over numerical domains and evaluate them in reachability algorithms for well-structured infinite-state systems like unbounded Petri Nets, Vector Addition Systems, Lossy Petri Nets, and Broadcast Protocols. The experiments show their potential usefulness. They also indicate that one may gain a lot by developing specialized data structures for very specific specialized application areas.

References

- [1] P. Argon, G. Delzanno, S. Mukhopadhyay, and A. Podelski. Model checking for communication protocols. In *Proceedings of the 28th Annual Conference on Current Trends in Theory and Practice of Informatics*, Lecture Notes in Computer Science, Piestany, Slovak Republic, 2001. Springer. To appear.

- [2] G. Delzanno and J.-F. Raskin. Symbolic representation of upward-closed sets. In S. Graf and M. I. Schwartzbach, editors, *Proceedings of the 6th Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS-00); Held as Part of the European Joint Conferences on the Theory and Practice of Software (ETAPS-00)*, volume 1785 of *Lecture Notes in Computer Science*, pages 426–440, Berlin, Germany, 2000. Springer.

- [3] S. Mukhopadhyay and A. Podelski. Beyond region graphs: Symbolic forward analysis of timed automata. In R. Ramanujam and V. Raman, editors, *Proceedings of the 19th Conference on Foundations*

of Software Technology and Theoretical Computer Science (FST&TCS-99), volume 1738 of *Lecture Notes in Computer Science*, pages 232–244, Chennai, India, December 1999. IARCS, Springer.

- [4] S. Mukhopadhyay and A. Podelski. Model checking for timed logic processes. In J. W. Lloyd, V. Dahl, U. Furbach, M. Kerber, K.-K. Lau, C. Palamidessi, L. M. Pereira, Y. Sagiv, and P. J. Stuckey, editors, *Proceedings of the 1st International Conference on Computational Logic*, volume 1861 of *Lecture Notes in Computer Science*, pages 598–612, Imperial College, London, UK, 2000. Springer.

- [5] S. Mukhopadhyay and A. Podelski. Accurate widenings and boundedness properties of timed systems. In A. Zamulin, M. Broy, and D. Bjorner, editors, *Perspectives of System Informatics: 4th International Andrei Ershov Memorial Conference*, Lecture Notes in Computer Science, Novosibirsk, Russia, 2001. Springer. To appear.

- [6] S. Mukhopadhyay and A. Podelski. Constraint database models characterizing timed bisimilarity. In I. Ramakrishnan, editor, *Proceedings of the 3rd International Symposium on Practical Aspects of Declarative Languages*, volume 1990 of *Lecture Notes in Computer Science*, pages 245–258, Las Vegas, USA, March 2001. Compulog Americas, Springer.

4.5 Modular Verification

One of the main problems that arise in the verification of systems composed of many parallel processes is the state explosion problem: the state space can grow exponentially with the number of processes. Symbolic representations of states and symbolic model checking have greatly increased the size of the systems that can be verified. However, many realistic systems are still too large to be handled. It is therefore important to find techniques that can be used to further extend the size of the systems that can be verified. One possibility is to exploit the structure of the systems and check properties in a modular way. The basic idea is to verify the properties of the individual components (possibly using additional assumptions about properties of an environment that represents the rest of the system), infer that these properties also hold in the system obtained by combining the individual components, and then use them to deduce additional properties of the system.

There are several problems that arise in modular verification. One problem is that not all properties of systems are preserved by composition: for instance deadlocks might occur when combining deadlock free systems. Thus, it is important to find classes of properties that are preserved by composition. Such a class is identified in [3].

Another problem is how to resolve circularity. When verifying a property of some component, one usually has to specify assumptions about the component's environment (i.e. the rest of the system). Then, in order to infer that the combined system satisfies the property, we may have to discharge circular assumptions. In [2], a proof rule is presented to deal with this situation.

Sheaf Models for Interacting Systems Investigator: Viorica Sofronie-Stokkermans

The goal of [3] is to give a syntactical characterization of a class of properties that are preserved when combining systems and thus can be checked in a modular way.

We showed that many notions important when expressing properties of systems (such as states, parallel actions, transitions, behavior, and time) can be modeled as sheaves over a suitable topological space. Results from geometric logic were then used to show that those properties of systems that can be expressed by Cartesian axioms are preserved after combining the systems. A cartesian axiom relative to a theory $\mathbb{T}$ is a formula of the form $(\forall x)(\phi(x) \Rightarrow \psi(x))$ where ϕ and ψ are formulae built up from atomic formulae using only the connective $\wedge$ and the quantifier $\exists$ over

$\mathbb{T}$-provably unique variables. It follows that, under additional assumptions about the environment when verifying the individual components, Cartesian properties can be checked in a modular way.

Also recent papers on model checking use decomposition of systems to avoid the state explosion problem. It was for instance shown that formulae in universal computation tree logics ($ACTL^*$) can be checked in a modular way [1]. Both in formulae in $ACTL^*$ and in Cartesian axioms restrictions are imposed in the use of existential quantification and negation. We would like to gain a better understanding about the possible links between the results presented in our paper and other results on modular model checking such as those mentioned above.

Assume-Guarantee Reasoning Investigator: Patrick Maier

In modular verification, one often wants to infer that a system consisting of subsystems S_1 and S_2 satisfies the properties P_1 and P_2 from 'modular' premises stating that S_1 satisfies P_1 assuming its environment satisfies P_2, while S_2 satisfies P_2 assuming its environment satisfies P_1. Due to the circular nature of the assumptions, such reasoning is generally unsound. Despite this fact, soundness was established in various modeling formalisms, e.g. in TLA by Abadi and Lamport and in Reactive Modules by Alur and Henzinger. However, all published soundness proofs obscured their core argument, a circularity breaking induction, by technical details of the particular formalism, thus making it hard to find the essence of circular reasoning and transfer it to other formalisms.

In [2], we have presented a circular assume-guarantee proof rule in a very abstract setting, where systems and properties are modeled as subsets of a poset. The rule has a mathematically concise side condition which expresses exactly the requirements needed to establish soundness; thus the circularity breaking induction becomes clear and simple. Yet, the setting is so general that from our rule we can derive assume-guarantee rules in a variety of more complex modeling formalisms. All we need to do is to instantiate the abstract setting and check the side condition; in particular we need not redo the notorious circularity breaking argument again. We have actually provided such instantiations for various formalisms (including Moore and Mealy machines) which model systems and properties as (possibly infinite) transition graphs and are based on trace containment. Instantiations based on tree containment resp. simulation as well as instantiations to process-algebra-like formalisms will be goals in future research.

References

[1] E. Clarke, O. Grumberg, and D. Long. Model checking. In *Nato ASI Series F*, volume 152, New York, Heidelberg, Berlin, 1996. Springer-Verlag.

- [2] P. Maier. A set-theoretic framework for assume-guarantee reasoning. In F. Orejas, P. G. Spirakis, and J. van Leeuwen, editors, *Proceedings of the 28th International Colloquium on Automata, Languages and Programming (ICALP 2001)*, volume 2076 of *Lecture Notes in Computer Science*, pages 821–834, Crete, Greece, July 2001. Springer.

- [3] V. Sofronie-Stokkermans and K. Stokkermans. Modeling interaction by sheaves and geometric logic. In G. Ciobanu and G. Paun, editors, *Proceedings of the 12th International Symposium Fundamentals of Computation Theory (FCT-99)*, volume 1684 of *Lecture Notes in Computer Science*, pages 512–523, Iasi, Romania, 1999. Springer.

5 Constraint Solving

Constraint satisfaction problems usually arise from logic modeling of certain application domains (e.g. performance analysis, computational linguistics), or as well-defined subproblems of logic deduction in general theorem proving. We speak of numerical constraints if the underlying domain is a set of numbers (such constraints occur in a natural way e.g. in integer programming and computer algebra); otherwise we speak of symbolic constraints.

5.1 Integer Programming and Computer Algebra

Members of our group started to conduct research in the area of integer programming and computer algebra mainly due to the immediate connection of these two fields within the area of constraint-logic programming. In the meantime, we have contributed substantial results which have been recognized in the corresponding communities.

5.1.1 Cutting Plane Algorithms

Investigator: Friedrich Eisenbrand

Integer Programming is concerned with the optimization of a linear function over the integer points in a polyhedron P. Many industrial optimization problems such as crew scheduling, warehouse location and planning can be modeled as an integer optimization problem. Among the most successful methods for solving integer optimization problems is the cutting plane method. A Gomory-Chvátal cutting plane for P is an inequality $c^T x \leq \lfloor \delta \rfloor$, where c is an integral vector and $c^T x \leq \delta$ is valid for P. The addition of a cutting plane to the system of inequalities defining P results in a better approximation of the integer hull. The intersection of a polyhedron with all its Gomory-Chvátal cutting planes is called the elementary closure P' of P. If P is rational, then P' is a rational polyhedron again.

To describe P', one can restrict the cutting planes $c^T x \leq \lfloor \delta \rfloor$ to those corresponding to a totally dual integral (TDI) system defining P. The number of inequalities of a minimal TDI-system for a polyhedron P is exponential in the size of P, even in fixed dimension. Integer programming in fixed dimension can be done in polynomial time. From this viewpoint the latter upper bound is undesirable. We proved that the elementary closure can be described with a polynomial number of inequalities in fixed dimension and we provided a polynomial algorithm (in varying dimension) for finding cutting planes from this description [3, 1, 2]. Based on this result, we developed a polynomial algorithm in varying dimension for computing Gomory-Chvátal cutting planes of simplicial cones. This algorithm was meanwhile implemented [4] and performs well, especially on problems which cannot be successfully attacked by combinatorial methods.

References

- [1] A. Bockmayr and F. Eisenbrand. Cutting planes and the elementary closure in fixed dimension. *Mathematics of Operations Research*, 26(2):304–312, May 2001.
- [2] F. Eisenbrand. *Gomory-Chvátal Cutting planes and the Elementary Closure of Polyhedra.* Phd thesis, Universität des Saarlandes, July 2000.
- [3] F. Eisenbrand and A. Bockmayr. Combining logic and optimization in cutting plane theory. In H. Kirchner and C. Ringeissen, editors, *Proceedings of the Workshop on Frontiers of Combining Systems (FROCOS-2000)*, volume 1794 of *Lecture Notes in Artificial Intelligence*, pages 1–17, Nancy, France, March 2000. Springer.

- [4] I. Grenner. Die Erzeugung von Schnittebenen mit maximalem Verletzungsgrad und deren Einsatz im Branch-and-Cut Verfahren. Masters thesis, Universität des Saarlandes, 2001.

5.1.2 Fast Integer Programming in Small Dimensions

Investigator: Friedrich Eisenbrand

It is well known that some basic number theoretic problems, such as the greatest common divisor or best approximations of rational numbers can be formulated as integer linear programs in two variables. Thus it is not surprising that current polynomial methods for integer programming in fixed dimension use lattice reduction methods, related to the reduction which is part of the classical Euclidean algorithm for integers, or the computation of the continued fraction expansion of a rational number. Therefore, integer programming in fixed dimension has a strong flavor of algorithmic number theory, and the running times of the algorithms also depend on the binary encoding length of the input.

The classical Euclidean algorithm for computing the greatest common divisor (GCD) of two s-bit integers requires $\Theta(s)$ arithmetic operations and $\Theta(s^2)$ bit operations in the worst case. For example, when it is applied to two consecutive Fibonacci numbers, it generates all the predecessors in the Fibonacci sequence. Schönhage's algorithm [2] improves this complexity to $O(M(s)\log s)$ bit operations, where $M(s)$ is the bit complexity of s-bit integer multiplication. Thus the greatest common divisor of two integers can be computed with a close to linear number of bit operations, if one uses the fastest methods for integer multiplication. The speedup technique by Schönhage was not yet incorporated into methods for two variable integer programming. The best known algorithms for the integer programming problem in two dimensions use $\Theta(s)$ arithmetic operations and $\Omega(s^2)$ bit operations when the number of constraints is fixed.

We have shown [1] that integer programming in two variables with a fixed number of constraints is not harder than greatest common divisor computation. As one allows an arbitrary number of constraints, the nature of the problem also becomes combinatorial. For this general case we presented an algorithm which requires $O(\log m)$ gcd-like computations, where m is the number of constraints. Our algorithm for 2-variable integer programming is the fastest which is known.

References

- [1] F. Eisenbrand and G. Rote. Fast 2-variable integer programming. In B. Gerards and K. Aardal, editors, *Proceedings of the 8th Conference on Integer and Combinatorial Optimization (IPCO-01)*, volume 2081 of *Lecture Notes in Computer Science*, pages 78–89, Utrecht, Netherlands, June 2001. Springer.

  [2] A. Schönhage. Schnelle Berechnung von Kettenbruchentwicklungen. (Speedy computation of expansions of continued fractions). *Acta Informatica*, 1:139–144, 1971.

5.1.3 Lattice Basis Reduction

Investigators: Friedrich Eisenbrand, Werner Backes

Lattice basis reduction and the related shortest vector problem are a classical area of the algorithmic geometry of numbers. Very recently this area has had a great impact on computer science in particular in cryptography and in complexity theory with Ajtai's proof of a longstanding conjecture of NP-hardness of the shortest vector problem for randomized reductions. Central to all algorithms in this area is the notion of lattice basis reduction, a concept invented by Hermite. A lattice $\Lambda \subseteq Q^d$ is a set of the form $\Lambda = \Lambda(A) = \{Ax \mid x \in Z^k\}$, where $A \in Q^{d\times k}$ is a rational matrix of full column rank. The matrix A is a basis of the lattice Λ and its columns are the basis vectors. The lattice

determinant of Λ is the number $\det \Lambda = |\det(A)|$ and the lattice basis $A = (x_1, \ldots, x_k)$ is reduced if

$$\prod_{i=1}^{k} \|x_i\| \leq \gamma \ \det \Lambda \tag{5.1}$$

for some constant γ. The Lattice Reduction Problem is the problem of computing a reduced basis for a given lattice. We delivered both theoretical, as well as practical results in this area.

Theoretical Contributions In fixed dimension, a shortest vector can be easily extracted from a reduced basis of Λ. Schönhage has shown that a shortest vector of a two dimensional lattice can be computed with $O(M(n) \log n)$ bit operations, if n denotes the binary input encoding length of the lattice and $M(n)$ denotes the bit-complexity of n-bit integer multiplication. We have shown [3] that a shortest vector of a 2-dimensional integral lattice with respect to the ℓ_∞-norm can be computed with a constant number of extended-gcd computations, one common-convergent computation and a constant number of arithmetic operations. It follows that in two dimensions, a fast basis-reduction algorithm can be solely based on Schönhage's classical algorithm on the fast computation of continued fractions and the reduction algorithm of Gauß.

The known algorithms for shortest vectors are based on the famous LLL-reduction algorithm. In fixed dimension, this algorithm requires $O(M(n)\, n)$ bit-operations, where n is the binary input encoding length of the lattice. We provided a lattice basis reduction algorithm [4] for fixed dimension d, which requires $O(M(n) \log^{d-1} n)$ bit-operations. This is a generalization of Schönhages classical result on the fast computation of continued fraction expansions. Our algorithm is the fastest lattice basis reduction algorithm in fixed dimension. The running time improvement has immediate impact on the running time of algorithms for integer programming in fixed dimension.

Practical Contributions The LLL algorithm was the first polynomial time algorithm guaranteed to compute lattice bases of relatively short vectors. Since then, lattice basis reduction methods have been further developed. Schnorr and Euchner introduced a variant of the LLL algorithm, based on floating point arithmetic that performed well in practice. The ability to reduce large instances of lattice bases is particularly important in the area of cryptography and cryptanalysis. Especially since more and more cryptographic protocols are nowadays based on lattice problems.

Despite the big step towards applicability of Schnorr and Euchner, lattice basis reduction methods are still impractical for large problem instances or do not perform very well in some cases. We focused on the stability of the algorithms, the running time and the quality of the reduction. Extensive test have been made to check known heuristics and to develop new ones. To take the Schnorr-Euchner algorithm to further limits, we used different data types for the approximation of the lattice basis and introduced a new technique called dynamic approximation which performs well for a special type of lattice bases that appear often in the context of cryptography. We used iterative and modular heuristics to increase the performance of the reduction algorithm. Modular methods make sense for lattice bases where the lattice determinant is known and reasonably small. For better quality of reduction, we tested the deep insertion variant ([1]), tried out different scalar-products and tested the combination of lattice basis reduction and sorting of the lattice basis vectors. We were able to improve the quality of the reduced basis dramatically [1, 2].

References

- [1] W. Backes and S. Wetzel. New results on lattice basis reduction in practice. In W. Bosma, editor, *Proceedings of the 4th International Algorithmic Number Theory Symposium (ANTS-IV)*, volume 1838 of *Lecture Notes in Computer Science*, pages 135–152, Leiden, Netherlands, 2000. Springer.
- [2] W. Backes and S. Wetzel. Lattice basis reduction with dynamic approximation. In S. Näher and D. Wagner, editors, *Proceedings of the 4th Workshop On Algorithm Engineering (WAE-2000)*, volume 1982 of *Lecture Notes in Computer Science*, pages 63–73, Saarbrücken, Deutschland, 2001. Max-Planck-Institut für Informatik.
- [3] F. Eisenbrand. Short vectors of planar integral lattices via continued fractions. *Information Processing Letters*, 79:121–126, 2001.
- [4] F. Eisenbrand and G. Rote. Fast reduction of ternary quadratic forms. In J. Silverman, editor, *Proceedings of the 1st Conference on Lattices and Cryptography (CALC-2001)*, volume 2146 of *Lecture Notes in Computer Science*, pages 99–111, Brown University, Providence, USA, 2001. Springer.

5.1.4 Comprehensive Gröbner Bases

Investigator: Patrick Maier

Let $K[\mathbf{U};\mathbf{X}]$ be the polynomial ring in the indeterminates $\mathbf{U} = U_1, \ldots, U_m$ (parameters) and $\mathbf{X} = X_1, \ldots, X_n$ (main indeterminates) over the field K. To solve parametric problems in algebraic geometry, Weispfenning introduced comprehensive Gröbner bases $G \subseteq K[\mathbf{U};\mathbf{X}]$ such that $\sigma(G) \subseteq K'[\mathbf{X}]$ is a Gröbner basis for every specialization $\sigma : K[\mathbf{U}] \to K'$ of the parameters $\mathbf{U}$ into any extension field K' of K. His algorithm to compute such a comprehensive Gröbner basis from a set of generating polynomials was known to perform at most exponentially more operations (reductions, formation of S-polynomials) than Buchberger's algorithm for computing a non-parametric Gröbner basis. In collaboration with Manfred Göbel, we have shown that this exponential gap arises already when restricting to just two generating polynomials in only one main indeterminate [1]. In this case, no S-polynomials are computed, hence the exponential complexity is really introduced by reduction in the presence of parameters.

Considering the ring $K[\mathbf{X}]^G$ of polynomials invariant under a finite matrix group G, [1] also deals with the problem of finding so-called SAGBI bases [2] which enjoy reduction properties similar to Gröbner bases. We have found out that considering the conjugates of G can lead to finite SAGBI bases even if there are only infinite SAGBI bases for $K[\mathbf{X}]^G$. In this context, the computation of comprehensive SAGBI bases might be used as a systematic search procedure for appropriate conjugates and their bases.

References

- [1] M. Göbel and P. Maier. Three remarks on comprehensive Gröbner and SAGBI bases. In V. G. Ganzha, E. W. Mayr, and E. V. Vorozhtsov, editors, *Proceedings of the 3rd Workshop on Computer Algebra in Scientific Computing (CASC-2000)*, pages 191–202, Samarkand, Uzbekistan, 2000. Springer.

[2] L. Robbiano and M. Sweedler. Subalgebra bases. In W. Bruns and A. Simis, editors, *Workshop on Commutative Algebra*, Lect. Notes Math. 1430, pages 61–87. Springer, 1988.

5.2 Symbolic Constraints

5.2.1 Set Constraints

Investigators: Witold Charatonik, Andreas Podelski, Jean-Marc Talbot

Set constraints denote relations between sets of trees. Syntactically, they are conjunctions of inclusions between expressions built over variables, constructors (constants and function symbols from a given alphabet) and a choice of set operators that defines the specific class of set constraints.

The satisfiability problem for the general class of set constraints is NEXPTIME-complete. Due to the applications in set-based program analysis and type inference, many authors investigated restrictions of this general class hoping to find subclasses with better algorithmic behaviour; our contribution to this area, including [1, 4, 5], was discussed in the previous report. In particular, it was known that the restriction to either definite set constraints or unary set constraints (that is, constraints with function symbols of arities at most one) reduces the complexity of the satisfiability problem from to NEXPTIME-complete to DEXPTIME-complete. It was conjectured that due to the connections with the theory of word automata both these restrictions together should further decrease the complexity to PSPACE. Surprisingly, in [2] we showed that it is not the case: the satisfiability problem for unary definite set constraints remains DEXPTIME-complete. The result has also consequences in program analysis that we discuss in Section 4.

The satisfiability problem for set constraints can be seen as the satisfiability problem for the existential part of the first-order theory of set constraints. In this theory, the atomic formulas are inclusions between set expressions. General formulas are built by adding usual connectives (conjunction, disjunction, negation, implication) and quantifiers. By a reduction to the satisfiability for set constraints, the existential part of the theory is decidable; by duality the universal part is decidable, too. It is not difficult to show that the full first-order theory of set constraints is undecidable. In 1996 Seynhaeve, Tison and Tommasi started a kind of competition with the aim to find the simplest undecidable fragment of the theory. They first proved that the $\exists^*\forall^*\exists^*$-fragment is undecidable, and then, together with Treinen in 1997 they improved the result by showing that already one alternation of quantifiers leads to undecidability; the latter result was further improved by Charatonik in 1998 by restricting the constraints to be atomic (that is, set constraints without any Boolean set operators). In [3] we were able to find the simplest known undecidable fragment. We showed that already three variables with the quantifier prefix $\exists\forall\forall$ give an undecidable fragment of the theory of atomic set constraints.

References

- [1] W. Charatonik and A. Podelski. Set constraints with intersection. *Information and Computation*, 2001. To appear.
- [2] W. Charatonik, A. Podelski, and J.-M. Talbot. Paths vs. trees in set-based program analysis. In *Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL-00)*, pages 330–337, Boston, Massachusetts, USA, January 2000. ACM.
- [3] J.-M. Talbot. The $\exists\forall^2$ fragment of the first-order theory of set constraints is π_1^0-hard. *Information Processing Letters*, 74(1/2):27–33, 2000.
- [4] J.-M. Talbot, P. Devienne, and S. Tison. Generalized definite set constraints. *Constraints*, 5(1/2):161–202, January 2000.

- [5] J.-M. Talbot, J. Niehren, and M. Müller. Entailment of atomic set constraints is PSPACE-complete. In G. Longo, editor, *Proceedings of the 14th Annual IEEE Symposium on Logic in Computer Science (LICS-99)*, pages 285–294, Trento, Italy, 1999. IEEE Computer Society, IEEE Computer Society.

5.2.2 Second-Order Unification

Investigator: Margus Veanes

Second-order unification generalizes first-order unification by allowing variables to occur also in the position of function symbols. For example, the second-order terms $G(f(c,c))$ and $f(f(c,c),G(c))$ (with the second-order variable G) have a unifier that maps G to a λ-term $\lambda x.f(x,x)$. Application of that unifier to either term yields the term $f(f(c,c),f(c,c))$. Second-order unification and restricted forms thereof (such as context unification) play a fundamental role in several areas.

Second-order unification is undecidable in general. In 1991 Farmer improved the general undecidability result by proving it for very restricted second-order term languages. Using Farmer's technique and exploiting the connection between second-order unification and simultaneous rigid E-unification that we discussed in the previous biennial report [1], we were able to find even simpler second-order term languages with undecidable second-order unification problem [2]. Formally, we proved that there is an integer n, such that second-order unification is undecidable in all non-monadic second-order term languages with at least n first-order variables, and even if the arguments of all second-order variables are ground terms of size bounded by n and the total number of variable occurrences is at most n.

References

- [1] J. Levy and M. Veanes. On the undecidability of second-order unification. *Information and Computation*, 159:125–150, 2000.
- [2] M. Veanes. Farmer's theorem revisited. *Information Processing Letters*, 74(1/2):47–53, April 2000.

6 Logic and Uncertainty

In the logic and uncertainty subgroup we have continued to pursue a broad range of subjects: from belief revision for conditional evidence and qualitative decision theory to probabilistic reasoning with Bayesian networks and probability constraints. The common underlying theme is how to formalize and automate the human capability of processing uncertain and sometimes contradictory information so as to reach defeasible conclusions on the basis of which decisions may be taken.

With regard to more qualitative models for these reasoning processes we were especially concerned with the further exploitation of the ranking construction paradigm (sections 6.1.1,6.1.2). In the area of quantitative models we pursue probabilistic approaches, the focus being on the relational Bayesian network framework (section 6.2.1). Apart from research rooted in artificial intelligence, we have also begun to look into issues that arise from the combination of logic and probability in probabilistic verification, thus bringing our research closer to other core interests represented in the Programming Logics group (section 6.2.3).

6.1 Semi-Qualitative Approaches

6.1.1 Belief Revision

Investigator: Emil Weydert

Belief revision in the front of conflicting evidence is a major task for cognitive agents. The usual strategy is to exploit epistemic preferences – like probability measures or plausibility rankings – to determine a suitable revised set of beliefs. But iterated revision forces us to update these epistemic priorities as well. Propositional revision may then be seen as a special instance of preference revision.

For instance, within Spohn's prominent revision framework for surprise rankings [1] and our quasi-probabilistic generalizations thereof, revising a ranking R with a proposition A corresponds to revising R with the ranking constraint $R(\neg A) \geq 1$. We have generalized Spohn's update strategy to finite sets of conditional ranking constraints $R(\neg A'_i|A_i) \geq a_i$ using a minimal ranking construction strategy [2, 4]. The resulting proposal is the most powerful around and exhibits some nice properties. Because the ranking constraints can also be used to model default knowledge (A normally implies B may be translated by $R(\neg B|A) \geq 1)$), this also gives us a powerful revision strategy for default conditional evidence.

If we characterize epistemic states by probability measures and incoming evidence by linear probability constraints, a traditional, well-motivated update strategy has been cross-entropy minimization, i.e. relative information minimization. However, this approach is neither general enough – e.g. to model epistemic states by convex sets of distributions – nor does it always produce intuitively appropriate results – e.g. if we consider successive revision steps. We have therefore proposed a new sophisticated two-stage epistemic model with Spohn-type rankings on top of probability distributions and a corresponding – experimental – minimal information revision strategy [3].

A major result of this work has been that probabilistic belief revision is much more complicated than common wisdom assumes. In this context, we have also presented some rough ideas on how to bridge the gap between standard belief revision and scientific epistemology [5].

References

[1] W. Spohn. Ordinal conditional functions : A dynamic theory of epistemic states. In W. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change and Statistics*. D. Reidel, Dordrecht, Netherlands, 1988.

- [2] E. Weydert. JZBR - iterated belief change for conditional ranking constraints. In H. Rott, C. Albert, G. Brewka, and C. Witteveen, editors, *Proceedings of the 4th Dutch-German Workshop on Nonmonotonic Reasoning Techniques and their Applications (DGNMR-99)*, pages 57–66, Amsterdam, The Netherlands, 1999. ILLC Scientific Publications.
- [3] E. Weydert. How to revise ranked probabilities. In W. Horn, editor, *Proceedings of the 14th European Conference on Artificial Intelligence (ECAI-00)*, volume 54 of *Frontiers in Artificial Intelligence and Applications*, pages 38–42, Berlin, Germany, August 2000. IOS Press.
- [4] E. Weydert. Rankings in flux. In J. Cunningham and D. Gabbay, editors, *Proceedings of the 3rd International Conference on Formal and Applied Practical Reasoning (FAPR-00)*, pages 37–48, London, United Kingdom, September 2000. Imperial College Department of Computing.
- [5] E. Weydert. Thoughts on evolving science and belief. In *Proceedings of the Workshop "Scientific Reasoning in AI and Philosophy of Science" at the 14th European Conference on Artificial Intelligence (ECAI-00)*, pages 62–69, Berlin, Germany, 2000. ECAI.

6.1.2 Default Reasoning

Investigators: Emil Weydert, Richard Booth

System JZ is a powerful and well-motivated default inference notion based on the construction of a canonical ranking model for each finite consistent default knowledge base [5]. However, the construction procedure is a bit cumbersome and we might prefer a preferential semantic interpretation. To achieve this, we have tried to encode the minimal surprise and minimal construction philosophy within a preference order over Spohn-type rankings, the JLX-order. It turns out that each finite consistent default base has exactly one JLX-minimal model, which is usually different from the JZ-model. This is not surprising, because we can show that there is no preferential semantic reconstruction of system JZ. However, if we restrict ourselves to ranking models which are justifiably constructible w.r.t. the default base (a non-redundancy condition), the JLX-model coincides with the JZ-model. Because both approaches exhibit comparable default reasoning features, JLX-entailment represents a real alternative to system JZ [7, 6].

The construction methodology can be seen as an application of belief revision, more precisely of Spohn's approach, to default inference. A different connection between belief revision and default inference is explored in In [2, 1] we have shown how one popular method of default entailment, namely lexicographic closure [3] can be formulated in terms of a certain method of belief revision due to Nayak [4]. Our investigations touch on a couple of interesting side-issues, such as the question of how to define an iterable belief revision operator which takes sets of sentences as inputs rather than the usual single sentences, and of how we may apply traditional principles of belief revision theory in the context of default reasoning.

References

- [1] R. Booth. The lexicographic closure as a revision process. *Journal of Applied Non-Classical Logics*. To appear.
- [2] R. Booth. The lexicographic closure as a revision process. In C. Baral and M. Truszczynski, editors, *Proceedings of the 8th International Workshop on Non-Monotonic Reasoning (NMR 2000)*, 2000. Extended version to appear in "Journal of Applied Non-Classical Logics".

[3] D. Lehmann. Another perspective on default reasoning. *Annals of Mathematics and Artificial Intelligence*, 15:61–82, 1995.

[4] A. C. Nayak. Iterated belief change based on epistemic entrenchment. *Erkenntnis*, 41:353–390, 1994.

[5] E. Weydert. System JZ : How to build a canonical ranking model of a default knowledge base. In A. G. Cohn, L. Schubert, and S. C. Shapiro, editors, *Proceedings of the 6th International Conference on Principles of Knowledge Representation and Reasoning (KR-98)*, pages 190–201, Trento, Italy, 1998. Morgan Kaufmann.

- [6] E. Weydert. Rankings we prefer: a minimal construction semantics for default reasoning. In *Proceedings of the 6th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU-01)*, Toulouse, France, September 2001. Springer. To appear.

- [7] E. Weydert. Rankings we prefer: a rational minimal construction semantics for default reasoning. In *Proceedings of the 5th Dutch-German Workshop on Nonmonotonic Reasoning Techniques and their Applications (DGNMR-01)*, pages 1–12, Potsdam, Germany, April 2001. Institut für Informatik, Universität Potsdam.

6.1.3 Qualitative Decision Theory

Investigator: Emil Weydert

Qualitative decision theory is concerned with reasoning about preferences and expectations, e.g. desires and defaults, and exploits these utility and plausibility considerations for planning. This approach is particularly promising for the modeling of autonomous agents. In collaboration with Amsterdam and Toulouse, we have continued our investigation of relevant issues in this area.

We have proposed a nonmonotonic logic of desires with a utilitarian semantics, where desires are understood as constraints on utility losses and/or utility gains. Each set of desires – nonmonotonically – induces a set of "distinguished" utility functions, obtained by weighing and adding up the utility losses/gains of the individual desires in some suitable way, which then determines a preference relation [2].

To do justice to the wide variety of desires, we have introduced a more expressive framework where the meaning of desires is fine-tuned by the use of three different parameters (strength, lifting, polarity). In particular, they are relevant for the interaction of the desires when constructing the overall preference relation [4, 3].

We have also started to investigate genuine conflicts between desires in the context of a generalization of Boutilier's logic of qualitative decision [1]. Our preliminary conclusion is that there are two (possibly overlapping) types of conflicts, those based on utopian worlds respectively hidden uncertainty [2].

References

[1] C. Boutilier. Toward a logic for qualitative decision theory. In P. T. Jon Doyle, Erik Sandewall, editor, *Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning*, pages 75–86, Bonn, FRG, May 1994. Morgan Kaufmann.

- [2] J. Lang, L. W. N. van der Torre, and E. Weydert. Two kinds of conflicts between desires (and how to resolve them). In *Proceedings of the AAAI Symposium on Game and Decision Theoretic Agents*, pages 1–5, Stanford, USA, March 2001. AAAI.

- [3] J. Lang, L. W. N. van der Torre, and E. Weydert. Utilitarian desires. *Autonomous Agents and Multi-Agent systems*, 2001. To appear.

- [4] E. Weydert and L. van der Torre. Risk parameters for utilitarian desires. In J. Bell, editor, *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-99), Workshop on Practial Reasoning and Rationality*, pages 48–54, Stockholm, Sweden, 1999. IJCAI.

6.2 Probabilistic Reasoning

6.2.1 Relational Bayesian Networks

Investigator: Manfred Jaeger

Relational Bayesian networks, originally proposed in [1], remained the focus of attention in our work on integrating first-order style representations with the Bayesian network paradigm. Relational Bayesian networks represent classes of probability distributions on finite relational structures: for a given relational vocabulary S and each $n \in \mathbb{N}$ a probability distribution P_n on the set of all S-structures with a domain of size n is defined. The basic inference problem one wants to solve has two inputs: a domainsize n, and some ground S-literals. To be computed is the probability according to P_n of the conjunction of the given literals (i.e. the constants appearing in the literals are assumed to be taken from a domain of n elements).

The method usually proposed for the solution of this inference problem relies on the construction of a standard Bayesian network with one node for each ground atom constructible with S and constants for each of the n domain-elements. There thus exists an apparent mismatch between the high-level representation language and the proposed inference procedures, which involve a complete propositionalization of the original first-order representation. However, in [2] we have shown that by potential inference methods that directly operate on first-order representations no efficiency gains in terms of worst-case complexity can be expected. The advantage of such methods over inference on the propositional level could only lie in a lower complexity in terms of the size of the domain considered in a particular inference problem. What we have shown in [2] is that this complexity is not polynomial (assuming NETIME$\neq$ETIME). Nondeterministic polynomial complexity, on the other hand, can already be achieved by the standard inference methods. Motivated by this result, an optimized "propositional" inference procedure for relational Bayesian networks was developed [3]. The optimizations here introduced are aimed at constructing for the solution of particular queries standard Bayesian networks that are as small and as sparsely connected as possible, thus minimizing the complexity of the subsequent computations in the standard networks.

References

[1] M. Jaeger. Relational bayesian networks. In D. Geiger and P. P. Shenoy, editors, *Proceedings of the 13th Conference of Uncertainty in Artificial Intelligence (UAI-13)*, pages 266–273, Providence, USA, 1997. Morgan Kaufmann.

- [2] M. Jaeger. On the complexity of inference about probabilistic relational models. *Artificial Intelligence*, 117:297–308, 2000.
- [3] M. Jaeger. Complex probabilistic modeling with recursive relational bayesian networks. *Annals of Mathematics and Artificial Intelligence*, 2001. To appear.

6.2.2 Probability Constraints

Investigator: Manfred Jaeger

A linear probability constraint is a linear inequality of the form $a_1p_1 + \ldots + a_np_n \leq a_{n+1}$, where the a_i are real numbers and the p_i are variables that range over the probabilities of the states in an n-element state space. Partial information about a particular probability distribution very often is obtained in the form of probability constraints. Given a set of such constraints, two basic probabilistic inference problems arise: the problem of probabilistically valid inference, which consists of finding the tightest bounds strictly implied by the constraints for the probability of

some event of interest, and the problem of measure selection, which consists of selecting a unique probability distribution on the state space as the best guess for the true distribution described by the constraints.

There exists a fairly extensive amount of literature on how to solve the first inference problem by means of inference rules for the deduction of new constraints from given ones (example: from $P(A) \geq r$ and $P(B) \geq s$ deduce $P(A \cap B) \geq \max\{0, r+s-1\}$). In particular, many different such inference rules have been introduced and individually proved to be correct. In [1] we have shown that the process of deriving suitable inference rules itself can be automated: we have developed an algorithm that for given input constraints (e.g. $P(A) \geq r$, $P(B) \geq s$) and a specified target probability (e.g. $P(A \cap B)$) automatically derives a parametric expression (here: $\max\{0, r+s-1\}$) that defines the optimal lower or upper bounds for the target probability.

The standard solution to the measure selection problem is entropy maximization: select the measure with maximal entropy within the set of measures satisfying the constraints. In [2] we have presented a radically different approach: based on a view of constraints as statistical data, we have developed a framework for measure selection by way of statistical maximal likelihood estimation. This approach requires to specify for each probability measure $\boldsymbol{p}$ on the finite state space a distribution $f_{\boldsymbol{p}}$ on the set of all constraints, so that for given constraints $c_1, \ldots, c_n$ one can maximize the likelihood $\prod f_{\boldsymbol{p}}(c_i)$ of the observed constraints as a function of $\boldsymbol{p}$. Key problem here is to identify a suitable family $\{f_{\boldsymbol{p}} \mid \boldsymbol{p}\}$ of distributions, so that the resulting maximum likelihood selection rule is well-behaved. In [2] a particular such family was proposed based on invariance and robustness considerations.

References

- [1] M. Jaeger. Automatic derivation of probabilistic inference rules. *International Journal of Approximate Reasoning*, 28(1):1–22, 2001.
- [2] M. Jaeger. Constraints as data: a new perspective on inferring probabilities. In B. Nebel, editor, *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, pages 755–760, Seattle, USA, 2001. International Joint Conferences on Artificial Intelligence, Inc., Morgan Kaufmann.

6.2.3 Fairness and Randomness

Investigator: Manfred Jaeger

Fairness conditions are introduced in the analysis of nondeterministic system models in order to exclude from consideration certain pathological behaviors of the nondeterministic system. When one passes from a nondeterministic system to a corresponding probabilistic system (i.e. a system in which the possible transitions of a nondeterministic system are assigned probabilities) then one finds that fairness conditions are satisfied with probability 1. This is not surprising in view of the fact that fairness conditions can be interpreted as tests of randomness in the sense of Martin-Löf [2]. Motivated by this analogy, we have proposed in [1] the concept of computable fairness which, on the one hand, is a canonical strengthening of previously defined concepts of fairness, and on the other hand, can be seen as a qualitative approximation of Martin-Löf-randomness.

For some properties P of infinite state sequences this approximation is complete in the sense that every computable fair sequence satisfies P iff P holds with probability 1 in a probabilistic system. In [1] it is shown that this equivalence holds for all properties P definable by a deterministic Büchi automaton.

References

- [1] M. Jaeger. Fairness, computable fairness and randomness. In M. Kwiatkowska, editor, *Proceedings of the 2nd International Workshop on Probabilistic Methods in Verification (PROBMIV-99)*, Technical Report CSR-99-8, pages 57–66, Eindhoven, the Netherlands, August 1999. School of Computer Science, University of Birmingham.

[2] P. Martin-Löf. The definition of random sequences. *Information and Control*, 9:602–619, 1966.

7 Software

Along with 'pure' research, the Programming Logics Group is developing various pieces of software in order to demonstrate the practical feasibility of our theoretical concepts.

In this section we describe our progress with regard to these systems, their availability, and (where applicable) their performance as measured against other competitive systems.

7.1 Bliksem 1.12

Investigator: Hans de Nivelle

Bliksem was inherited from the University of Amsterdam, when Hans de Nivelle moved from Amsterdam to Saarbrücken.

The main distinctive feature of Bliksem is its capability to generate so-called explicit proofs. Explicit proofs are proofs in which all details are explicitly reported. Such proofs can be either checked directly, or translated into other calculi, in particular type theory or higher order natural deduction. This is needed for the integration of first-order automated theorem proving into interactive theorem provers.

Most of the existing theorem provers are able to output some form of proofs, at least the clauses that were kept by the prover, and used in the derivation of the empty clause. However this is not enough, since literals resolved upon, and term positions used for paramodulation/rewriting can be still ambiguous.

Generating explicit proofs turned out harder than expected. Explicit proof generation was not anticipated in the original design or Bliksem. Initially Bliksem had the same style of proof output as other theorem provers. In the given situation, it seemed the best decision not to change the inference engine, but to add a proof reconstruction module. The motivations for this decision were the following: Full logging of details costs some time. Since only a fraction of the generated clauses eventually will occur in the proof, it was thought that it was better to do costly operations on a small set of clauses, than to do cheaper operations on a much larger set of clauses. It was realized that filling in the details of a proof step could take exponential time in the size of the clauses involved, but it was assumed that this in practice would not be a real problem. Another argument in favor of proof reconstruction was that proof reconstruction allows modularity. The proof reconstructor can be seen as a small theorem prover, which uses its own small set of rules. These rules are in principle independent of those used by the main inference engine. In this way the proof output format has some independence of the calculus used by the main inference engine. It may be possible to add rules to the inference engine, without changing the proof reconstructor. Unfortunately, it turned out that proof reconstruction is surprisingly hard. After doing some experiments, we have concluded that proof reconstruction is the inferior approach. We have been able to find heuristics that work in most of the cases, but we do not believe that proof reconstruction will be sufficiently robust in the end. We concluded that it is better to give up this approach, and to implement full logging.

Bliksem 1.12 comes with three additional Prolog programs, that handle the proof processing. The program **translate.pro** translates explicit proofs into type theory. This experimental program turned out successful. Big proofs can be efficiently and reliably transformed into type theory proofs.

The program **tt.pro** is a simple proof checker in Prolog that checks the type theory proofs generated by **translate.pro**. Big proofs can be checked in short time, without difficulties.

Alternatively the output of **translate.pro** can be translated into Coq syntax [1] using **coq.pro**. Coq is an interactive proof checker, based on type theory, developed by Rocquencourt. A translator for Isabelle [2] is planned.

In the near future Bliksem will be completely rewritten, with the following objectives in mind:

- Simple types will be built in. Nearly all proof tasks coming from applications are typed. In theory it is easy to explicitly axiomatize the types. In practice, the prover cannot cope with these explicit axiomatizations.
- The explicit proof reconstruction will be replaced by explicit proof logging.
- All parts of Bliksem, with the exception of the core inference mechanism will be rewritten in Java. Only the core inference mechanism is time critical. The advantage in maintainability is worth the loss in speed.
- The data structures have been redesigned into order to deal with basicness, and unbounded integers. This will make it possible to experiment with more advanced first-order calculi, and with more application-directed calculi that have built in data-types.

References

[1] B. Barras, S. Boutin, C. Cornes, J. Courant, J.-C. Filliâtre, E. Giménez, H. Herbelin, G. Huet, C. M. nez, C. Murthy, C. Parent, C. Paulin-Mohring, A. Saïbi, and B. Werner. The coq proof assistant reference manual, version 6.2.4, 1998. available on `http://pauillac.inria.fr/coq/` .

[2] T. Nipkow and L. Paulson. Isabelle. available on `http://isabelle.in.tum.de/`, 1999.

7.2 SPASS Version 2.0

Investigator: Christoph Weidenbach

Research assistants: Bijan Afshordel, Uwe Brahm, Christian Cohrs, Thorsten Engel, Enno Keen, Christian Theobalt, Dalibor Topić

SPASS is an automated theorem prover for full sorted first-order logic with equality that extends superposition by a sort concept and by a splitting rule for case analysis [15]. It is one of the most powerful automated theorem proving systems currently available. For example, it won several prices at CADE theorem prover system competitions [10] and shows to be a useful tool in various application areas [13, 12, 5].

Our current focus is to strengthen SPASS even further for formal system analysis, e.g., see [8, 6]. To this end it is indispensable to have a strong reduction machinery on the formula as well as on the clausal level, to make the prover more robust against "small" changes to the problem and to increase the flexibility of its implementation. That's exactly what we did during the past two years, resulting in SPASS version 2.0:

- Integration of stronger reduction mechanisms: contextual rewriting [3], the terminator [2], conditioned atom definitions [1].
- Enhancements to the SPASS search engine: more robust input clause ordering, stronger reduction oriented selection mechanisms, robust precedence computation, lazy reduction loop.
- Improvements to the code: elimination of almost all global variables, highly improved documentation, automated installation procedure based on GNU configure.

In the sequel we will explain our implementation of contextual rewriting in more detail. Most of the above mentioned items are discussed in two chapters of the new *Handbook of Automated Reasoning* [9, 14].

Although contextual rewriting is a rather old concept, refining and instantiating and actually implementing this reduction rule such that it performs well in relevant application areas is a difficult task.

If $\Gamma_1 \rightarrow \Delta_1, s \approx t$ is a given clause, then contextual rewriting[8] replaces $C = \Gamma_2, E[s']_p \rightarrow \Delta_2$ with $\Gamma_2, E[p/t\sigma] \rightarrow \Delta_2$ (reducing $s' = s\sigma$ to $t\sigma$), provided that the literals in $\Gamma_1 \rightarrow \Delta_1$ can be shown to be logically implied by $\Gamma_2 \rightarrow \Delta_2$ and all other given clauses smaller than C.

Although this semantic test can be easily approximated recursively, e.g., by considering only syntactically smaller clauses with respect to variable renaming and by replacing the models relation by recursive application of contextual rewriting and congruence closure, the result is still "too strong". For example, the Ackermann function can be computed this way, meaning that such an implementation in SPASS could already cause the system to spend too much time in one contextual rewriting sequence. Therefore, we decided to introduce even weaker restrictions that still give many benefits in practice, but can be much more efficiently computed. Our restrictions to the rule are: (i) $s\sigma = s'$, (ii) $s \succ t$, (iii) $s \approx t$ is strictly maximal in $\Gamma_1 \rightarrow \Delta_1, s \approx t$, (iv) for any term $t' \neq s$ in $\Gamma_1 \rightarrow \Delta_1, s \approx t$, $s \succ t'$, (v) $\models \mathit{fred}(\Gamma_2 \rightarrow A)$ for all atoms $A \in \Gamma_1\sigma$, (vi) $\models \mathit{fred}(A \rightarrow \Delta_2)$ for all atoms $A \in \Delta_1\sigma$.

The conditions (ii)–(iv) guarantee the necessary ordering restrictions, while the conditions (v) and (vi) force the context of $s \approx t$ to be valid. The function *fred* (recursively) applies the reduction rules of SPASS to the constructed subclauses before they are checked as tautologies. This is a weaker but also much more efficient variant of the above mentioned syntactic test. The tautology check itself is highly efficiently integrated into SPASS by the implementation of a variant of the well-known congruence closure algorithm [7, 5]. The performance of *fred*, i.e., which reduction rules are actually tested/applied, determines the strength as well as the cost for testing and applying contextual rewriting.

One nice result of this enhancement is that SPASS is now able to finitely saturate Balbiani's [4] formulation of a certain fragment of the Euclidean geometry [5], proving its consistency. Currently we are investigating examples from formal system analysis, e.g. using SPASS as the reasoning backend in static analysis of list-processing C programs [11], to get an appropriate mix of reduction rules for the *fred* function in this application domain. We want to significantly improve the results presented by Fischer et. al. [8].

References

[1] B. Afshordel, T. Hillenbrand, and C. Weidenbach. First-order atom definitions extended. 2001. Submitted.

[2] G. Antoniou and H. J. Ohlbach. Terminator. In A. Bundy, editor, *Proceedings of 8th International Joint Conference on Artificial Intelligence, IJCAI-83*, pages 916–919, 1983.

[3] L. Bachmair and H. Ganzinger. Rewrite-based equational theorem proving with selection and simplification. *Journal of Logic and Computation*, 4(3):217–247, 1994. Revised version of Max-Planck-Institut für Informatik technical report, MPI-I-91-208, 1991.

[4] P. Balbiani. Equation solving in geometrical theories. In N. Dershowitz and N. Lindenstrauss, editors, *Proceedings of the 4th international workshop on conditional and typed rewriting systems*, volume 968 of *LNCS*. Springer, 1995.

[8] We show a simplified non-sorted version and only the left variant of the inference rule.

- [5] C. Brinker. Geometrisches Schließen mit SPASS. Masters thesis, Universität des Saarlandes, January 2000.

[6] W. Charatonik, D. McAllester, D. Niwinski, A. Podelski, and I. Walukiewicz. The horn mu-calculus. In *Proceedings 13th IEEE Symposium on Logic in Computer Science, LICS'98*, pages 58–69. IEEE Computer Society Press, 1998.

[7] P. J. Downey, R. Sethi, and R. E. Tarjan. Variations on the common subexpression problem. *Journal of the ACM*, 27(4):758–771, 1980.

[8] B. Fischer, J. Schumann, and G. Snelting. Deduction-based software component retrieval. In W. Bibel and P. H. Schmitt, editors, *Automated Deduction - A Basis for Applications*, volume 3 of *Applied Logic*, chapter 11, pages 265–292. Kluwer, 1998.

- [9] A. Nonnengart and C. Weidenbach. Computing small clause normal forms. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 6, pages 335–367. Elsevier, Amsterdam, the Netherlands, January 2001.

[10] G. Sutcliffe. The cade-16 atp system competition. *Journal of Automated Reasoning*, 24(3):371–396, April 2000.

[11] D. Topić. Static analysis of list-processing c programs. Diplomarbeit, Universität des Saarlandes and Max-Planck-Institut für Informatik, Saarbrücken, Germany, 2001. Supervisors: A. Podelski, C. Weidenbach; in preparation.

- [12] J. Veit. Formal fairness proofs for optimistic contract signing protocols. Masters thesis, Universität des Saarlandes, October 1999.
- [13] C. Weidenbach. Towards an automatic analysis of security protocols in first-order logic. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction (CADE-16)*, volume 1632 of *Lecture Notes in Artificial Intelligence*, pages 378–382, Trento, Italy, 1999. Springer.
- [14] C. Weidenbach. Combining superposition, sorts and splitting. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 27, pages 1965–2013. Elsevier, Amsterdam, the Netherlands, January 2001.
- [15] C. Weidenbach, B. Afshordel, U. Brahm, C. Cohrs, T. Engel, E. Keen, C. Theobalt, and D. Topic. System description: SPASS version 1.0.0. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction (CADE-16)*, volume 1632 of *Lecture Notes in Artificial Intelligence*, pages 314–318, Trento, Italy, 1999. Springer.

7.3 WALDMEISTER

Investigator: Thomas Hillenbrand

WALDMEISTER is a prover for unit equational deduction [4]. Its theoretical basis is unfailing completion in the sense of [2] with refinements towards ordered completion [1]. The development was undertaken at the University of Kaiserslautern by A. Buch and Th. Hillenbrand, who had started the project, and A. Jaeger and B. Löchner. With the arrival of Th. Hillenbrand at the MPI, it will jointly continue here and in Kaiserslautern.

The system is one of the strongest provers in its class, as demonstrated repeatedly in last years competitions at the Conference on Automated Deduction. It is built on solidly engineered algorithms and data structures, especially indexing and space saving techniques, and it features automated selection of reduction ordering and heuristic guidance of the proof search.

Current research is aimed at the integration of new theoretical developments into the existing system. This often poses substantial problems, as run-time and proof search behavior tend to be affected by the new components in an unforeseen way.

As part of these efforts, and as an internship project for a visiting student from India, the basicness restriction [3] has been implemented in the WALDMEISTER system, and its effects on the proof search have been studied. First results are promising already in that a more or less uniform reduction of derivation steps necessary for proof completion is obtained in several problem domains.

References

[1] J. Avenhaus, T. Hillenbrand, and B. Löchner. On using ground joinable equations in equational theorem proving. In P. Baumgartner and H. Zhang, editors, *Proceedings of the Third International Workshop on First-Order Theorem Proving*, Fachberichte Informatik 5/2000, pages 33–43. Universität Koblenz-Landau, 2000. To appear also in the *Journal of Symbolic Computation*.

[2] L. Bachmair, N. Dershowitz, and D. A. Plaisted. Completion without failure. In H. Aït-Kaci and M. Nivat, editors, *Resolution of Equations in Algebraic Structures*, volume 2, chapter 1, pages 1–30. Academic Press, 1989.

[3] L. Bachmair, H. Ganzinger, C. Lynch, and W. Snyder. Basic paramodulation. *Information and Computation*, 121(2):172–192, 1995.

[4] T. Hillenbrand, A. Jaeger, and B. Löchner. System description: WALDMEISTER — improvements in performance and ease of use. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction*, volume 1632 of *LNAI*, pages 232–236. Springer-Verlag, 1999.

7.4 DMC and GCA: Deductive Model Checking

Investigator: Georg Jung

We are continuing our implementation work on our prototype tool DMC (Deductive Model Checker) and our experiments with fixpoint acceleration techniques (widening and narrowing).

DMC is built on top of the SICSTUS-Prolog system, which allows us to use its built-in constraint solvers (such as for linear real arithmetic) and to easily extend it e.g. with new widening rules. These widening rules avoid the loss of precision through (automated) reasoning on the program text.

We have built two frontends for DMC. One translates graph format representations of embedded systems with real-time constraints into DMC-syntax. This has enabled us to compare the performance of our tool with Uppaal, a specialized tool for verifying timed systems built in Uppsala and Aalborg.

GCA is a second frontend, translating ANSI-standard *C* programs into DMC constraint logic programs. Additionally, it infers the specification of a special temporal property, namely the array bound integrity at all program points with array accesses. We have analyzed all stand-alone procedures of a widely used library of arithmetical algorithms; we have been able to compute the exact preconditions for array bound violations, for example `dim c < m` in the function `float chebev(float c[], int m)` (Chebichev approximation) and `dim data <= n` for the function `void avevar(float data[], int n)` (average).

7.5 Emotion: An Evolutionary Algorithm for Optimizing Photo Mask Layout for Grey-Tone Lithography

Investigator: Jörn Hopf

In the last two years we continued working on applications of evolutionary algorithms (a term which collects together evolutionary programming, evolutionary strategies and genetic algorithms) to optimization of photo masks for the lithographic process.

With this work, techniques for the optimization of photo mask layout for silicon micro machining are being developed. Silicon micro machining could use equipment and processes of IC-technology (photo lithography and dry-etching) to produce micro-mechanical systems and highly miniaturized three-dimensional micro structures on silicon, as for example sensors, actuators and lenses. In particular the problem of laying out an appropriate photo mask under consideration of several production constraints is solved.

The problem can be described more formally in the following way: Given is a surface function $f(x, y)$ to approximate and a set $\mathbf{M} = \{m_1, \ldots, m_n\}$, $n \in \mathbb{N}$ of all mask elements. Wanted is a set $\mathbf{M}^* \subset \mathbf{M}$ of elements which consider several production constraints and achieve the best possible approximation of $g^*(x, y) = \sum_{m_i \in \mathbf{M}^*} G_i(x, y)$ respecting $f(x, y)$, where $G_i(x, y)$ represents the grey value generated by the element m_i placed at the coordinates (x, y). Since a calculation of an optimal solution $g^*(x, y)$ is not possible in an analytic way, an algorithm was studied which allows an approximation of $g(x, y) = \sum_{m_i \in \mathbf{M}'} G_i(x, y)$ to $f(x, y)$ with $\mathbf{M}' \subset \mathbf{M}$.

A drawback of the current state of this technology is that photo masks are structured as fixed arrays of 'rasters', since the general optimization problem is computationally intractable. Today, as projects grow more ambitious, this is becoming a greater problem. The evolutionary algorithm in our program *Emotion* (*e*volutionary *m*ask *o*ptimiz*ation*) is a first step towards solving this problem.

8 Journal and Conference Activities

8.1 Editorial positions

The editorial activities of staff of the programming logics group are as follows:

Harald Ganzinger is an editor of the following journals:

- *Information Processing Letters*
- *Journal of Automated Reasoning*
- *Theory of Computing Systems*
- *Discrete Mathematics and Theoretical Computer Science*

Andreas Podelski is an editor of the *International Journal on Software Tools for Technology Transfer*

8.2 Conference Activities

The participation of the members of the programming logics group in the organization of various conferences is as follows.

Program committees

Hans de Nivelle was a member of the program committee of the *17th International Conference on Automated Deduction* (CADE 2000).

Harald Ganzinger is or was a member of the program committee of

- *Third International Conference on Foundations of Software Science and Computation Structures* (FOSSACS 2000)
- *8th International Conference on Algebraic methodology and software technology* (AMAST 2000)
- *17th International Conference on Automated Deduction* (CADE 2000)
- *15th Annual IEEE Symposium on Logic in Computer Science* (LICS 2000)
- *13th International Conference on Theorem proving in higher order logics* (TPHOLs 2000)
- *7th International Conference on Logic for Programming and Automated Reasoning* (LPAR 2000)
- *First International Joint Conference on Automated Reasoning* (IJCAR 2001)
- *14th International Conference on Theorem proving in higher order logics* (TPHOLs 2001)
- *8th International Conference on Logic for Programming and Automated Reasoning* (LPAR 2001)
- *19th Annual Symposium on Theoretical Aspects of Computer Science* (STACS 2002)
- *17th Annual IEEE Symposium on Logic in Computer Science* (LICS 2002)
- *18th International Conference on Automated Deduction* (CADE 2002)

Andreas Podelski is or was a member of the program committees of

- *17th Annual IEEE Symposium on Logic in Computer Science* (LICS 2002)

- *Fourth International Symposium on Practical Aspects of Declarative Languages* (PADL 2002)
- *ACM SIGPLAN Conference on Programming Language Design and Implementation* (PLDI 2001)
- *17th International Conference on Logic Programming. Paphos* (ICLP 2001)
- *Third International ACM Conference on Principles and Practice of Declarative Programming* (PPDP 2001)
- *Tenth French Conference on Logic and Constraint Programming* (JFPLC'01)
- *7th International Conference on Tools and Algorithms for the Construction and Analysis of Systems* (TACAS 2001)
- *First International Joint Conference on Automated Reasoning* (IJCAR 2001)
- *28th International Colloquium on Automata, Languages and Programming* (ICALP 2001)
- *Computational Logic Conference 2000* (CL 2000)
- *9th European Symposium on Programming* (ESOP 2000)
- *6th International Symposium on Static Analysis* (SAS 99)

Viorica Sofronie-Stokkermans was a member of the program committee of the *Fifth International Conference on Artificial Intelligence and Symbolic Computation* (AISC 2000)

Christoph Weidenbach was a member of the program committee of

- *International Workshop on First Order Theorem Proving* (FTP'2000)
- *Workshop Logisches Programmieren* (WLP'2000)

Organization of workshops and conferences

Andreas Podelski is or was on the steering/organizing committee of

- the *European Joint Conferences on Theory and Practice of Software* (ETAPS)
- *Specification, Analysis and Validation for Emerging Technologies* (SAVE'01)
- the *Workshop on Verification and Computational Logic* (VCL'01)
- the *Schloss Ringberg Seminar: Model Checking and Program Analysis*
- the *Electronic Tool Integration Platform Day* (Satellite Event of ETAPS'2001)
- the *International Workshop on Verification and Computational Logic*(VCL'00)

Christoph Weidenbach is on the steering committee of the *International Workshop on First Order Theorem Proving* (FTP).

9 Teaching Activities

Members of the programming logic group regularly teach undergraduate and graduate courses at the Universität des Saarlandes and on invitation at other universities. They also supervise diploma students in the preparation of their theses.
Key: L – Lectures, LE – Lectures and exercises, S – Seminar, FoPra – Project class.

Winter Semester 1999/2000

AK der Theoretischen Informatik 3 V. Sofronie-Stokkermans – L (at the Technical University Vienna)
Informatik für Hörer aller Fakultäten II A.Podelski – LE
Constraint Programming for Verification A.Podelski – L (at Ecole Jeunes Chercheurs en Programmation, ENS-Lyon)
Implementierung von Software C. Weidenbach – LE
Stochastik für Informatiker M. Jaeger – LE
Logik für Informatiker H. Ganzinger – LE

Summer Semester 2000

Informatik für Hörer aller Fakultäten I A.Podelski – LE

Winter Semester 2000/2001

Automata Theory and Applications W. Charatonik, J-M. Talbot – LE
Informatik für Hörer aller Fakultäten II U. Waldmann – LE
Computer Algebra F. Eisenbrand – LE
Logik für Informatiker H. Ganzinger – LE

Summer Semester 2001

Programmiersprachen A.Podelski, H.Ganzinger – LE
Universelle Algebra und Verbandstheorie V. Sofronie-Stokkermans – LE
IT Projektmanagement C. Weidenbach – LE
Bayessche Netzwerke M. Jaeger – LE

Diploma students

Jörg Veit, *Formal Fairness Proofs for Optimistic Contract Signing Protocols*
Christof Brinker, *Geometrisches Schließen mit SPASS*
Dalibor Topić, *Statische Softwareanalyse mit SPASS*
Ingo Grenner, *Cutting plane algorithms*

10 Dissertations and Habilitations

10.1 Doctorates

Completed

H. Baumeister, *Relations Between Abstract Datatypes Modeled as Abstract Datatypes* (May 1999)
F. Eisenbrand, *Gomory-Chvtal Cutting Planes and the Elementary Closure of Polyhedra* (July 2000)
U. Hustadt, *Resolution-Based Decision Procedures for Subclasses of First-Order Logic* (November 1999)
C. Meyer, *Soft Typing for Clausal Inference Systems* (October 1999)
S. Mukhopadhyay, *A Uniform Constraint-based Framework for the Verification of Infinite State Systems* (June 2001)
J. Stuber, *Superposition Theorem Proving for Commutative Algebraic Theories* (December 1999)
M. Tzakova, *Hybrid Languages* (January 2000)

In Progress

W. Backes, *Program Analysis Based on the Extended Register Transfer Language (XRTL)*
T. Hillenbrand, *Combination of Decision Procedures with General Inference Systems for First-Order Logic*
J. Hopf, *Photomaskenlayout für eine 3D-Grauton-Lithographie als kombinatorisches Optimierungsproblem* (submitted)
P. Maier, *Constraint-based Compositional Verification*

10.2 Habilitations

Completed

C. Weidenbach

In Progress

W. Charatonik

11 Grants and Cooperations

Procope Project "Program Analysis and Model Checking"

Description

The goal of the project to exploit the potential of synergie that lies in the complementary expertise of the two communities of program analysis and model checking. The target of the project is the verification of temporal properties of infinite-state systems.

Technical Data

Starting date:	January 1, 2000
Duration:	2 years + 1 year renewal
Funding:	DAAD, CNRS
Staff at MPI f. Informatik:	Witold Charatonik Giorgio Delzanno Georg Jung Andreas Podelski

Partners

Ecole Normale Superieure (Patrick Cousot and his group)

12 Publications

The following list contains all papers by members of the programming logics group that were published or accepted for publication during the period covered by this report. This includes papers the contents of which has already been described in previous reports. These are listed here without being discussed in sections 3–7.

Books and monographs

[1] H. Ganzinger, editor. *Proceedings of the 16th International Conference on Automated Deduction (CADE-16)*, volume 1632 of *Lecture Notes in Artificial Intelligence*, Berlin, Germany, 1999. Springer.

[2] H. Ganzinger, D. McAllester, and A. Voronkov, editors. *Proceedings of the 6th International Conference on Logic for Programming and Automated Reasoning (LPAR-99)*, volume 1705 of *Lecture Notes in Artificial Intelligence*, Berlin, Germany, 1999. Springer. Earlier version: Technical Report MPI-I-1999-2-003, Max-Planck-Institut für Informatik, Saarbrücken.

Journals and chapters

[1] P. Anderson and D. A. Basin. Program development schemata as derived rules. *Journal of Symbolic Computation*, 30(1):5–36, 2000.

[2] A. Ayari and D. A. Basin. A higher-order interpretation of deductive tableau. *Journal of Symbolic Computation*, 31(5):487–520, May 2001.

[3] L. Bachmair and H. Ganzinger. A theory of resolution. In J. A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 2, pages 19–99. Elsevier, Amsterdam, the Netherlands, 2001.

[4] D. A. Basin and S. Friedrich. Modeling a hardware synthesis methodology in Isabelle. *Formal Methods in Systems Design*, 15(2):99–122, September 1999.

[5] D. A. Basin and H. Ganzinger. Automated complexity analysis based on ordered resolution. *Journal of the ACM*, 48(1):70–109, 2001.

[6] D. A. Basin and B. Krieg-Brückner. Formalization of the development process. In E. Astesiano, H.-J. Kreowski, and B. Krieg-Brückner, editors, *Algebraic foundations of systems specification*, pages 521–562. Springer, Berlin, Germany, 1999.

[7] D. A. Basin and S. Matthews. Logical frameworks. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, second edition*. Reidel, Dordrecht, The Netherlands. To appear.

[8] P. Blackburn and M. Tzakova. Hybrid languages and temporal logic. *Logic Journal of the IGPL*, 7(1):27–54, January 1999. Revised Version of MPI-I-98-2-006.

[9] A. Bockmayr and F. Eisenbrand. Cutting planes and the elementary closure in fixed dimension. *Mathematics of Operations Research*, 26(2):304–312, May 2001.

[10] A. Bockmayr, F. Eisenbrand, M. Hartmann, and A. S. Schulz. On the Chvátal rank of polytopes in the 0/1 cube. *Discrete Applied Mathematics*, 98(1/2):21–27, 1999.

[11] A. Bockmayr and V. Weispfenning. Solving numerical constraints. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 12, pages 751–842. Elsevier, Amsterdam, the Netherlands, January 2001.

[12] R. Booth. The lexicographic closure as a revision process. *Journal of Applied Non-Classical Logics*. To appear.

[13] M. Bozzano, G. Delzanno, M. Martelli, V. Mascardi, and F. Zini. Logic programming and multi-agent systems: a synergic combination for applications and semantics. In K. Apt, V. Marek, M. Truszczynski, and D. Warren, editors, *The Logic Programming Paradigm: a 25-Year Perspective*, Artificial Intelligence, chapter 1, pages 5–32. Springer, Berlin, Germany, January 1999.

[14] J. Buchmann and F. Eisenbrand. On factor refinement in number fields. *Mathematics of Computation*, 68(225):345–350, January 1999.

[15] M. Bugliesi, G. Delzanno, L. Liquori, and M. Martelli. Object calculi in linear logic. *Journal of Logic and Computation*, 10(1):75–104, April 2000.

[16] H. de Nivelle. Deciding the E-plus class by an a posteriori, liftable order. *Annals of Pure and Applied Logic*, 88(1):219–232, 2000.

[17] H. de Nivelle, U. Hustadt, and R. A. Schmidt. Resolution-based methods for modal logics. *Logic Journal of the IGPL*, 8(3):265–292, May 2000.

[18] A. Degtyarev, Y. Gurevich, P. Narendran, M. Veanes, and A. Voronkov. Decidability and complexity of simultaneous rigid E-unification with one variable and related results. *Theoretical Computer Science*, 243(1/2):167–184, July 2000.

[19] G. Delzanno, D. Galmiche, and M. Martelli. A specification logic for concurrent object-oriented programming. *Mathematical Structures in Computer Science*, 9(3):253–286, 1999.

[20] F. Eisenbrand. On the membership problem for the elementary closure of a polyhedron. *Combinatorica*, 19(2):297–300, 1999.

[21] F. Eisenbrand. Short vectors of planar integral lattices via continued fractions. *Information Processing Letters*, 79:121–126, 2001.

[22] C. G. Fermüller, A. Leitsch, U. Hustadt, and T. Tammet. Resolution decision procedures. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 25, pages 1793–1849. Elsevier, Amsterdam, the Netherlands, January 2001.

[23] H. Ganzinger, U. Hustadt, C. Meyer, and R. A. Schmidt. A resolution-based decision procedure for extensions of K4. In M. Zakharyaschev, K. Segerberg, M. de Rijke, and H. Wansing, editors, *Advances in Modal Logic, Volume 2*, volume 119 of *CSLI Lecture Notes*, chapter 9, pages 225–246. CSLI, Stanford, USA, 2001.

[24] H. Ganzinger, F. Jacquemard, and M. Veanes. Rigid reachability: The non-symmetric form of rigid E-unification. *International Journal of Foundations of Computer Science*, 11(1):3–27, 2000.

[25] Y. Gurevich and M. Veanes. Logic with equality: Partisan corroboration and shifted pairing. *Information and Computation*, 152(2):205–235, 1999.

[26] U. Hustadt and R. A. Schmidt. An empirical analysis of modal theorem provers. *Journal of Applied Non-Classical Logics*, 9(4):479–522, 1999.

[27] L. Iturrioz and V. Sofronie-Stokkermans. SHn-algebras (Symmetric Heyting algebras of order n). In L. Iturrioz, E. Orlowska, and E. Turunen, editors, *COST Action 15 (Many-Valued Logics for Computer Science Applications) ATLAS of Many-Valued Structures*, number 75 in Department of Information Technology. Mathematics Report, pages 1–11. Tampere University of Technology (Mathematics), Tampere, January 2000.

[28] M. Jaeger. On the complexity of inference about probabilistic relational models. *Artificial Intelligence*, 117:297–308, 2000.

[29] M. Jaeger. Automatic derivation of probabilistic inference rules. *International Journal of Approximate Reasoning*, 28(1):1–22, 2001.

[30] M. Jaeger. Complex probabilistic modeling with recursive relational bayesian networks. *Annals of Mathematics and Artificial Intelligence*, 2001. To appear.

[31] M. R. K. Krishna Rao. Some characteristics of strong innermost normalization. *Theoretical Computer Science*, 239(1):141–164, 2000.

[32] M. R. K. Krishna Rao. Some classes of prolog programs inferable from positive data. *Theoretical Computer Science*, 241(1/2):211–223, 2000. Special issue for ALT '96.

[33] J. Lang, L. W. N. van der Torre, and E. Weydert. Utilitarian desires. *Autonomous Agents and Multi-Agent systems*, 2001. To appear.

[34] J. Levy and M. Veanes. On the undecidability of second-order unification. *Information and Computation*, 159:125–150, 2000.

[35] M. Lu and J. Wu. On theorem proving in annotated logics. *Journal of Applied Non-Classical Logics*, 10(2):121–143, June 2000.

[36] P. Madden, A. Bundy, and A. Smaill. Recursive program optimization through inductive synthesis proof transformation. *Journal of Automated Reasoning*, 22(1):65–115, 1999. Also available as Research Report MPI-I-94-239, Max-Planck-Institut für Informatik, Saarbrücken.

[37] D. Makinson and L. W. N. van der Torre. Input-output logics. *Journal of Philosophical Logic*, 29(4):383–408, August 2000.

[38] S. Matthews and D. A. Basin. Structuring metatheory on inductive definitions. *Information and Computation*, 162(1/2):80–95, 2000.

[39] M. Müller, J. Niehren, and A. Podelski. Ordering constraints over feature trees. *Constraints*, 5(1/2):7–41, January 2000.

[40] A. Nonnengart, H. J. Ohlbach, and D. M. Gabbay. Encoding two-valued non-classical logics in classic logic. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 2, chapter 21, pages 1403–1486. Elsevier, Amsterdam, the Netherlands, January 2001.

[41] A. Nonnengart and A. Szalas. A fixpoint approach to second-order quantifier elimination with applications to correspondence theory. In E. Orlowska, editor, *Logic at Work: Essays Dedicated to the Memory of Helena Rasiowa*, volume 24 of *Studies in Fuzziness and Soft Computing*, page 18. Springer, Berlin, Germany, January 1999.

[42] A. Nonnengart and C. Weidenbach. Computing small clause normal forms. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 6, pages 335–367. Elsevier, Amsterdam, the Netherlands, January 2001.

[43] R. A. Schmidt. Decidability by resolution for propositional modal logics. *Journal of Automated Reasoning*, 22(4):379–396, May 1999.

[44] R. A. Schmidt. Relational grammars for knowledge representation. In M. Böttner and W. Thümmel, editors, *Variable-Free Semantics*, Artikulation und Sprache, pages 162–180. secolo Verlag, Osnabrück, Germany, 2000.

[45] V. Sofronie-Stokkermans. Duality and canonical extensions of bounded distributive lattices with operators and applications to the semantics of non-classical logics. Part I. *Studia Logica*, 64(1):93–132, February 2000.

[46] V. Sofronie-Stokkermans. Duality and canonical extensions of bounded distributive lattices with operators and applications to the semantics of non-classical logics. Part II. *Studia Logica*, 64(2):151–172, March 2000.

[47] V. Sofronie-Stokkermans. Priestley duality for SHn-algebras and applications to the study of Kripke-style models for SHn-logics. *Multiple-Valued Logic - An International Journal*, 5(4):281–305, 2000.

[48] V. Sofronie-Stokkermans. Some properties of Kleene algebras. In L. Iturrioz, E. Orlowska, and E. Turunen, editors, *COST Action 15 (Many-Valued Logics for Computer Science Applications) ATLAS of Many-Valued Structures*, number 75 in Department of Information Technology. Mathematics Report, pages 1–7. Tampere University of Technology (Mathematics), Tampere, January 2000.

[49] V. Sofronie-Stokkermans. Automated theorem proving by resolution for finitely-valued logics based on distributive lattices with operators. *Multiple-Valued Logic - An International Journal*, 6(3/4):289–344, March 2001.

[50] J.-M. Talbot. The $\exists\forall^2$ fragment of the first-order theory of set constraints is π_1^0-hard. *Information Processing Letters*, 74(1/2):27–33, 2000.

[51] J.-M. Talbot, P. Devienne, and S. Tison. Generalized definite set constraints. *Constraints*, 5(1/2):161–202, January 2000.

[52] L. W. N. van der Torre and Y.-H. Tan. Contextual deontic logic: violation contexts and factual defeasibility. In M. Cavalcanti, editor, *Formal Aspects in Context*, Applied Logic Series, pages 173–192. Kluwer, Dordrecht, The Netherlands, January 1999.

[53] L. W. N. van der Torre and Y.-H. Tan. Contrary-to-duty reasoning with preference-based dyadic obligations. *Annals of Mathematics and Artificial Intelligence*, 27:49–78, January 1999.

[54] L. W. N. van der Torre and Y.-H. Tan. Diagnosis and decision making in normative reasoning. *Journal of Artificial Intelligence and Law*, 7:51–67, 1999.

[55] L. W. N. van der Torre and Y.-H. Tan. An update semantics for deontic reasoning. In P. McNamara and H. Prakken, editors, *Norms, Logics and Information Systems: New Studies on Deontic Logic and Computer Science*, volume 49 of *Frontiers in Artificial Intelligence and Applications*, pages 73–90. IOS Press, Amsterdam, the Netherlands, January 1999.

[56] L. W. N. van der Torre and E. Weydert. Parameters for utilitarian desires in a qualitative decision theory. *Applied Artificial Intelligence*, 14:285–301, January 2001.

[57] M. Veanes. Farmer's theorem revisited. *Information Processing Letters*, 74(1/2):47–53, April 2000.

[58] S. Vorobyov. $\forall\exists^5$-equational theory of context unification is undecidable. *Theoretical Computer Science*. To appear.

[59] S. Vorobyov. The undecidability of the first-order theories of one step rewriting in linear canonical systems. *Information and Computation*. To appear.

[60] A. Voronkov. The ground-negative fragment of first-order logics is π_2^p-complete. *The Journal of Symbolic Logic*, 64(3):984–990, September 1999.

[61] Y. Wang. Randomness, stochasticity, and approximation. *Theory of Computing Systems*, 32:517–529, 1999.

[62] C. Weidenbach. Spass v0.95tptp. *Journal of Automated Reasoning*, 23(1):21–21, 1999.

[63] C. Weidenbach. Combining superposition, sorts and splitting. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 27, pages 1965–2013. Elsevier, Amsterdam, the Netherlands, January 2001.

[64] E. Weydert. How to revise ranked probabilities. In W. Horn, editor, *Proceedings of the 14th European Conference on Artificial Intelligence (ECAI-00)*, volume 54 of *Frontiers in Artificial Intelligence and Applications*, pages 38–42, Berlin, Germany, August 2000. IOS Press.

[65] E. Weydert. Thoughts on evolving science and belief. In *Proceedings of the Workshop "Scientific Reasoning in AI and Philosophy of Science" at the 14th European Conference on Artificial Intelligence (ECAI-00)*, pages 62–69, Berlin, Germany, 2000. ECAI.

[66] J. Wu. First-order polynomial based theorem proving. In X.-S. Gao and D. Wang, editors, *Mathematics Mechanizations and Applications*, pages 273–294. Academic Press, San Diego, USA, January 2000.

Conference proceedings

[1] P. Argon, G. Delzanno, S. Mukhopadhyay, and A. Podelski. Model checking for communication protocols. In *Proceedings of the 28th Annual Conference on Current Trends in Theory and Practice of Informatics*, Lecture Notes in Computer Science, Piestany, Slovak Republic, 2001. Springer. To appear.

[2] W. Backes and S. Wetzel. New results on lattice basis reduction in practice. In W. Bosma, editor, *Proceedings of the 4th International Algorithmic Number Theory Symposium (ANTS-IV)*, volume 1838 of *Lecture Notes in Computer Science*, pages 135–152, Leiden, Netherlands, 2000. Springer.

[3] W. Backes and S. Wetzel. Lattice basis reduction with dynamic approximation. In S. Näher and D. Wagner, editors, *Proceedings of the 4th Workshop On Algorithm Engineering (WAE-2000)*, volume 1982 of *Lecture Notes in Computer Science*, pages 63–73, Saarbrücken, Deutschland, 2001. Max-Planck-Institut für Informatik.

[4] T. Ball, A. Podelski, and S. K. Rajamani. Boolean and cartesian abstraction for model checking C programs. In T. Margaria and W. Yi, editors, *Proceedings of the 7th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2001)*, volume 2031 of *Lecture Notes in Computer Science*, pages 268–283, Genova, Italy, April 2001. Springer.

[5] M. Bozzano, G. Delzanno, M. Martelli, V. Mascardi, and F. Zini. Multi-agent systems development as a software engineering enterprise. In G. Gupta, editor, *Proceedings of the 1st International Workshop on Practical Aspects of Declarative Languages (PADL-99)*, volume 1551 of *Lecture Notes in Computer Science*, pages 46–60, San Antonio, Texas, USA, October 1999. Compulog Americas and Association for Logic Programming (ALP), Springer.

[6] W. Charatonik. Directional type checking for logic programs: Beyond discriminative types. In G. Smolka, editor, *Proceedings of the 8th European Symposium on Programming (ESOP-00)*, volume 1782 of *Lecture Notes in Computer Science*, pages 72–87, Berlin, Germany, March 2000. Springer.

[7] W. Charatonik, S. Dal-Zilio, A. D. Gordon, S. Mukhopadhyay, and J.-M. Talbot. The complexity of model checking mobile ambients. In F. Honsell and M. Miculan, editors, *Foundations of Software Science and Computation Structures. Proceedings of the 4th International Conference (FOSSACS-01). Held as Part of the Joint European Conferences on Theory and Practice of Software (ETAPS-01)*, volume 2030 of *Lecture Notes in Computer Science*, pages 152–167, Genova, Italy, April 2001. Springer.

[8] W. Charatonik, A. Podelski, and J.-M. Talbot. Paths vs. trees in set-based program analysis. In *Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL-00)*, pages 330–337, Boston, Massachusetts, USA, January 2000. ACM.

[9] W. Charatonik and J.-M. Talbot. The decidability of model checking mobile ambients. In *Computer Science Logic, 15th Annual Conference of the EACSL*, Lecture Notes in Computer Science, Paris, France, September 2001. European Association for Computer Science Logic (EACSL), Springer. To appear.

[10] V. Cortier, H. Ganzinger, F. Jacquemard, and M. Veanes. Decidable fragments of simultaneous rigid reachability. In J. Wiedermann, P. van Emde Boas, and M. Nielsen, editors, *Proceedings of the 26th International Colloquium on Automata, Languages and Programming (ICALP-99)*, volume 1644 of *Lecture Notes in Computer Science*, pages 250–260, Prague, Czech Republik, 1999. Springer. Full version of this paper is available as MPI Research Report MPI-I-1999-2-004.

[11] H. de Nivelle. An overview of resolution decision procedures. In M. Faller, S. Kaufmann, and M. Pauly, editors, *Formalizing the Dynamics of Information*, volume 91 of *CSLI Publications*, pages 115–130, Stanford University, Palo Alto, USA, 2000. Center for the Study of Language and Information.

[12] H. de Nivelle, C. Areces, and M. de Rijke. Prefixed resolution: A resolution method for modal and description logics. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction (CADE-16)*, volume 1632 of *Lecture Notes in Artificial Intelligence*, pages 187–201, Trento, Italy, July 1999. Springer.

[13] H. de Nivelle, M. Bezem, and D. Hendriks. Automated proof construction in type theory using resolution. In D. McAllester, editor, *Proceedings of the 17th International Conference on Automated Deduction (CADE-17)*, volume 1831 of *Lecture Notes in Artificial Intelligence*, pages 148–163, Carnegie Mellon University, Pittsburgh, PA, USA, June 2000. Springer.

[14] H. de Nivelle and I. Pratt-Hartmann. A resolution-based decision procedure for the two-variable fragment with equality. In R. Gore, A. Leitsch, and T. Nipkow, editors, *Proceedings of the 1st International Joint Conference on Automated Reasoning (IJCAR-2001)*, volume 2083 of *Lecture Notes in Artificial Intelligence*, pages 211–225, Siena, Italy, July 2001. Springer.

[15] G. Delzanno and A. Podelski. Model checking in CLP. In R. Cleaveland, editor, *Proceedings of the 5th International Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS-99)*, volume 1579 of *Lecture Notes in Computer Science*, pages 223–239, Amsterdam, The Netherlands, January 1999. European Association for Programming Languages and Systems (EAPLS); European Association for Theoretical Computer Science (EATCS), Springer.

[16] G. Delzanno and J.-F. Raskin. Symbolic representation of upward-closed sets. In S. Graf and M. I. Schwartzbach, editors, *Proceedings of the 6th Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS-00); Held as Part of the European Joint Conferences on the Theory and Practice of Software (ETAPS-00)*, volume 1785 of *Lecture Notes in Computer Science*, pages 426–440, Berlin, Germany, 2000. Springer.

[17] F. Eisenbrand and G. Rote. Fast 2-variable integer programming. In B. Gerards and K. Aardal, editors, *Proceedings of the 8th Conference on Integer and Combinatorial Optimization (IPCO-01)*, volume 2081 of *Lecture Notes in Computer Science*, pages 78–89, Utrecht, Netherlands, June 2001. Springer.

[18] F. Eisenbrand and G. Rote. Fast reduction of ternary quadratic forms. In J. Silverman, editor, *Proceedings of the 1st Conference on Lattices and Cryptography (CALC-2001)*, volume 2146 of *Lecture Notes in Computer Science*, pages 99–111, Brown University, Providence, USA, 2001. Springer.

[19] F. Eisenbrand and A. S. Schulz. Bounds on the Chvátal rank of polytopes in the 0/1-cube. In G. Cornuéjols, R. E. Burkard, and G. J. Woeginger, editors, *Proceedings of the 7th Conference on Integer Programming and Combinatorial Optimization (IPCO-99)*, volume 1610 of *Lecture Notes in Computer Science*, pages 137–150, Graz, Austria, 1999. Springer.

[20] H. Ganzinger. Relating semantic and proof-theoretic concepts for polynomial time decidability of uniform word problems. In D. A. Williams, editor, *Proceedings of the 16th IEEE Symposium on Logic in Computer Science (LICS-01)*, pages 81–90, Boston, USA, 2001. IEEE Computer Society.

[21] H. Ganzinger and H. de Nivelle. A superposition decision procedure for the guarded fragment with equality. In G. Longo, editor, *Proceedings of the 14th Annual IEEE Symposium on Logic in Computer Science (LICS-99)*, pages 295–303, Trento, Italy, 1999. IEEE Computer Society, IEEE Computer Society.

[22] H. Ganzinger and D. McAllester. A new meta-complexity theorem for bottom-up logic programs. In R. Goré, A. Leitsch, and T. Nipkow, editors, *Proceedings of the 1st International Joint Conference on Automated Reasoning (IJCAR-2001)*, volume 2083 of *Lecture Notes in Artificial Intelligence*, pages 514–528, Siena, Italy, 2001. Springer.

[23] H. Ganzinger, C. Meyer, and M. Veanes. The two-variable guarded fragment with transitive relations. In G. Longo, editor, *Proceedings of the 14th Annual IEEE Symposium on Logic in Computer Science (LICS-99)*, pages 24–34, Trento, Italy, 1999. IEEE Computer Society, IEEE Computer Society.

[24] H. Ganzinger and R. Nieuwenhuis. Constraints and theorem proving. In H. Comon, C. Marché, and R. Treinen, editors, *Contraints in Computational Logics, International Summer School (CCL-99)*, volume 2002 of *Lecture Notes in Computer Science*, pages 159–201, Gif-sur-Yvette, France, 2001. Springer.

[25] H. Ganzinger, R. Nieuwenhuis, and P. Nivela. Context trees. In R. Nieuwenhuis, P. Nivela, R. Goré, A. Leitsch, and T. Nipkow, editors, *Proceedings of the 1st International Joint Conference on Automated Reasoning (IJCAR-2001)*, volume 2083 of *Lecture Notes in Artificial Intelligence*, pages 242–256, Siena, Italy, 2001. Springer.

[26] H. Ganzinger and V. Sofronie-Stokkermans. Chaining techniques for automated theorem proving in many-valued logics. In *Proceedings of the 30th IEEE International Symposium on Multiple-Valued Logic (ISMVL-00)*, pages 337–344, Portland, Oregon, 2000. IEEE Computer Society Technical Committee on Multiple-Valued Logic; Oregon Center for Advanced Technology Education, IEEE Computer Society.

[27] U. Hustadt and R. A. Schmidt. Maslov's class K revisited. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction (CADE-16)*, volume 1632 of *Lecture Notes in Artificial Intelligence*, pages 172–186, Trento, Italy, 1999. Springer.

[28] U. Hustadt and R. A. Schmidt. On the relation of resolution and tableaux proof systems for description logics. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, volume 1, pages 110–115, Stockholm, Sweden, 1999. International Joint Conference on Artificial Intelligence, Inc., The Scandanavian AI Societies, Morgan Kaufmann.

[29] M. Jaeger. Constraints as data: a new perspective on inferring probabilities. In B. Nebel, editor, *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI-01)*, pages 755–760, Seattle, USA, 2001. International Joint Conferences on Artificial Intelligence, Inc., Morgan Kaufmann.

[30] P. Maier. A set-theoretic framework for assume-guarantee reasoning. In F. Orejas, P. G. Spirakis, and J. van Leeuwen, editors, *Proceedings of the 28th International Colloquium on Automata, Languages and Programming (ICALP 2001)*, volume 2076 of *Lecture Notes in Computer Science*, pages 821–834, Crete, Greece, July 2001. Springer.

[31] S. Mukhopadhyay and A. Podelski. Beyond region graphs: Symbolic forward analysis of timed automata. In R. Ramanujam and V. Raman, editors, *Proceedings of the 19th Conference on Foundations of Software Technology and Theoretical Computer Science (FST&TCS-99)*, volume 1738 of *Lecture Notes in Computer Science*, pages 232–244, Chennai, India, December 1999. IARCS, Springer.

[32] S. Mukhopadhyay and A. Podelski. Model checking for timed logic processes. In J. W. Lloyd, V. Dahl, U. Furbach, M. Kerber, K.-K. Lau, C. Palamidessi, L. M. Pereira, Y. Sagiv, and P. J. Stuckey, editors, *Proceedings of the 1st International Conference on Computational Logic*, volume 1861 of *Lecture Notes in Computer Science*, pages 598–612, Imperial College, London, UK, 2000. Springer.

[33] S. Mukhopadhyay and A. Podelski. Accurate widenings and boundedness properties of timed systems. In A. Zamulin, M. Broy, and D. Bjorner, editors, *Perspectives of System Informatics: 4th International Andrei Ershov Memorial Conference*, Lecture Notes in Computer Science, Novosibirsk, Russia, 2001. Springer. To appear.

[34] S. Mukhopadhyay and A. Podelski. Constraint database models characterizing timed bisimilarity. In I. Ramakrishnan, editor, *Proceedings of the 3rd International Symposium on Practical Aspects of Declarative Languages*, volume 1990 of *Lecture Notes in Computer Science*, pages 245–258, Las Vegas, USA, March 2001. Compulog Americas, Springer.

[35] F. Nielson and H. Seidl. Control-flow analysis in cubic time. In *Proceedings of the 10th European Symposium on Programming (ESOP-01) held as part of the Joint European Conferences on Theory and Practice of Software (ETAPS-01)*, volume 2028 of *Lecture Notes in Computer Science*, pages 252–268, Genova, Italy, 2001. Springer.

[36] A. Podelski. Model checking as constraint solving. In J. Palsberg, editor, *Proceedings of the 7th International Symposium on Static Analysis (SAS-00)*, volume 1824 of *Lecture Notes in Computer Science*, pages 221–237, Santa Barbara, California, USA, 2000. Springer.

[37] A. Podelski, W. Charatonik, and M. Müller. Set-based failure analysis for logic programs and concurrent constraint programs. In S. D. Swierstra, editor, *Programming Languages and Systems: Proceedings of the 8th European Symposium on Programming (ESOP-99)*, volume 1576 of *Lecture Notes in Computer Science*, pages 177–192, Amsterdam, The Netherlands, March 1999. Springer.

[38] V. Sofronie-Stokkermans. On the universal theory of varieties of distributive lattices with operators: Some decidability and complexity results. In H. Ganzinger, editor, *Proceedings of the 16th International Conference on Automated Deduction (CADE-16)*, volume 1632 of *Lecture Notes in Artificial Intelligence*, pages 157–171, Trento, Italy, 1999. Springer.

[39] V. Sofronie-Stokkermans. Representation theorems and automated theorem proving in non-classical logics. In *Proceedings of the 29th IEEE International Symposium on Multiple-Valued Logic (ISMVL-99)*, pages 242–247, Freiburg im Breisgau, Germany, May 1999. IEEE Computer Society, IEEE Computer Society.

[40] V. Sofronie-Stokkermans. Resolution-based theorem proving for non-classical logics based on distributive lattices with operators. In J. Cachro and K. Kijania-Placek, editors, *Proceedings of the 11th International Congress of Logic, Methodology and Philosophy of Science. Volume of abstracts*, pages 481–481, Cracow, Poland, August 1999. The Faculty of Philosophy, Jagellonian University.

[41] V. Sofronie-Stokkermans. On unification for bounded distributive lattices. In D. McAllester, editor, *Proceedings of the 17th International Conference on Automated Deduction (CADE-17)*, volume 1831 of *Lecture Notes in Artificial Intelligence*, pages 465–481, Pittsburgh, Pennsylvania, USA, 2000. Springer.

[42] V. Sofronie-Stokkermans. Representation theorems and the semantics of (semi)lattice-based logics. In *Proceedings of the 31st IEEE International Symposium on Multiple-Valued Logics*, pages 125–134, Warsaw, Poland, 2001. IEEE Computer Society, Institute of Computer Science Polish Academy of Sciences, IEEE Computer Society.

[43] V. Sofronie-Stokkermans and K. Stokkermans. Modeling interaction by sheaves and geometric logic. In G. Ciobanu and G. Paun, editors, *Proceedings of the 12th International Symposium Fundamentals of Computation Theory (FCT-99)*, volume 1684 of *Lecture Notes in Computer Science*, pages 512–523, Iasi, Romania, 1999. Springer.

[44] J. Stuber. Theory path orderings. In P. Narendran and M. Rusinowitch, editors, *Proceedings of the 10th International Conference on Rewriting Techniques and Applications (RTA-99)*, volume 1631 of *Lecture Notes in Computer Science*, pages 148–162, Trento, Italy, 1999. Springer.

[45] J. Stuber. Deriving theory superposition calculi from convergent term rewriting systems. In L. Bachmair, editor, *Proceedings of the 11th International Conference Rewriting Techniques and Applications (RTA-00)*, volume 1833 of *Lecture Notes in Computer Science*, pages 229–245, Norwich, UK, 2000. Springer.

[46] J.-M. Talbot. On the alternation-free Horn mu-calculus. In M. Parigot and A. Voronkov, editors, *Proceedings of the 7th International Conference on Logic for Programming and Automated Reasoning (LPAR-2000)*, volume 1955 of *Lecture Notes in Artificial Intelligence*, pages 418–435, La Reunion, France, 2000. Springer.

[47] J.-M. Talbot, J. Niehren, and M. Müller. Entailment of atomic set constraints is PSPACE-complete. In G. Longo, editor, *Proceedings of the 14th Annual IEEE Symposium on Logic in Computer Science (LICS-99)*, pages 285–294, Trento, Italy, 1999. IEEE Computer Society, IEEE Computer Society.

[48] M. Tzakova. Tableau calculi for hybrid logics. In N. V. Murray, editor, *Proceedings of the International Conference TABLEAUX'99 - Automated Reasoning with Analytic Tableaux and Related Methods*, volume 1617 of *Lecture Notes in Artificial Intelligence*, pages 278–292, Saratoga Springs, NY, USA, 1999. Springer.

[49] L. W. N. van der Torre. Defeasible goals. In A. Hunter and S. Parsons, editors, *Proceedings of the 5th European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty (ECSQARU-99)*, volume 1638 of *Lecture Notes in Artificial Intelligence*, pages 374–385, London, 1999. Springer.

[50] L. W. N. van der Torre and Y.-H. Tan. Rights, duties and commitments between agents. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99)*, volume 2, pages 1239–1244, Stockholm, Sweden, 1999. IJCAII and the Scandinavian AI Societies, Morgan Kaufmann.

[51] L. W. N. van der Torre and Y.-H. Tan. An update semantics for defeasible obligations. In K. B. Laskey and H. Prade, editors, *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence (UAI-99)*, pages 631–638, Stockholm, Sweden, 1999. Association for Uncertainty in AI (AUAI), Morgan Kaufmann.

[52] U. Waldmann. Cancellative superposition decides the theory of divisible torsion-free abelian groups. In H. Ganzinger, D. McAllester, and A. Voronkov, editors, *Proceedings of the 6th International Conference on Logic for Programming and Automated Reasoning (LPAR-99)*, volume 1705 of *Lecture Notes in Artificial Intelligence*, pages 131–147, Tbilisi, Georgia, 1999. Springer. Max-Planck-Institut für Informatik, Saarbrücken.

[53] U. Waldmann. Superposition and chaining for totally ordered divisible abelian groups (extended abstract). In R. Goré, A. Leitsch, and T. Nipkow, editors, *Proceedings of the 1st International Joint Conference on Automated Reasoning (IJCAR-2001)*, volume 2083 of *Lecture Notes in Artificial Intelligence*, pages 226–241, Siena, Italy, 2001. Springer.

[54] E. Weydert. How to revise ranked probabilities. In W. Horn, editor, *Proceedings of the 14th European Conference on Artificial Intelligence (ECAI-00)*, volume 54 of *Frontiers in Artificial Intelligence and Applications*, pages 38–42, Berlin, Germany, August 2000. IOS Press.

[55] E. Weydert. Rankings in flux. In J. Cunningham and D. Gabbay, editors, *Proceedings of the 3rd International Conference on Formal and Applied Practical Reasoning (FAPR-00)*, pages 37–48, London, United Kingdom, September 2000. Imperial College Department of Computing.

[56] E. Weydert. Rankings we prefer: a minimal construction semantics for default reasoning. In *Proceedings of the 6th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU-01)*, Toulouse, France, September 2001. Springer. To appear.

Workshop proceedings

[1] R. Booth. The lexicographic closure as a revision process. In C. Baral and M. Truszczynski, editors, *Proceedings of the 8th International Workshop on Non-Monotonic Reasoning (NMR 2000)*, 2000. Extended version to appear in "Journal of Applied Non-Classical Logics".

[2] G. Delzanno, J. Esparza, and A. Podelski. Constraint-based analysis of broadcast protocols. In J. Flum and M. Rodríguez-Artalejo, editors, *Proceedings of the 13th International Workshop on Computer Science Logic (CSL-99), 8th Annual Conference on the EACSL*, volume 1683 of *Lecture Notes in Computer Science*, pages 50–66, Madrid, Spain, 1999. Springer.

[3] F. Eisenbrand and A. Bockmayr. Combining logic and optimization in cutting plane theory. In H. Kirchner and C. Ringeissen, editors, *Proceedings of the Workshop on Frontiers of Combining Systems (FROCOS-2000)*, volume 1794 of *Lecture Notes in Artificial Intelligence*, pages 1–17, Nancy, France, March 2000. Springer.

[4] M. Göbel and P. Maier. Three remarks on comprehensive Gröbner and SAGBI bases. In V. G. Ganzha, E. W. Mayr, and E. V. Vorozhtsov, editors, *Proceedings of the 3rd Workshop on Computer Algebra in Scientific Computing (CASC-2000)*, pages 191–202, Samarkand, Uzbekistan, 2000. Springer.

[5] M. Jaeger. Fairness, computable fairness and randomness. In M. Kwiatkowska, editor, *Proceedings of the 2nd International Workshop on Probabilistic Methods in Verification (PROBMIV-99)*, Technical Report CSR-99-8, pages 57–66, Eindhoven, the Netherlands, August 1999. School of Computer Science, University of Birmingham.

[6] J. Lang, L. W. N. van der Torre, and E. Weydert. Two kinds of conflicts between desires (and how to resolve them). In *Proceedings of the AAAI Symposium on Game and Decision Theoretic Agents*, pages 1–5, Stanford, USA, March 2001. AAAI.

[7] S. Mukhopadhyay and A. Podelski. Model checking for timed logic processes. In J. W. Lloyd, V. Dahl, U. Furbach, M. Kerber, K.-K. Lau, C. Palamidessi, L. M. Pereira, Y. Sagiv, and P. J. Stuckey, editors, *Proceedings of the 1st International Conference on Computational Logic*, volume 1861 of *Lecture Notes in Computer Science*, pages 598–612, Imperial College, London, UK, 2000. Springer.

[8] V. Sofronie-Stokkermans. Priestley representation for distributive lattices with operators and applications to automated theorem proving. In J. Vaz de Carvalho and I. Ferreirim, editors, *Dualities, Interpretability and Ordered Structures*, pages 43–54, Lisbon, Portugal, 1999. Centro de Álgebra da Universidade de Lisboa.

[9] V. Sofronie-Stokkermans. Resolution-based theorem proving for SHn-logics. In R. Caferra and G. Salzer, editors, *Automated Deduction in Classical and Non-Classical Logic (Selected Papers of FTP'98)*, volume 1761 of *Lecture Notes in Artificial Intelligence*, pages 268–282, Vienna, Austria, 2000. Springer.

[10] S. Vorobyov. Subtyping functional+nonempty record types. In G. Gottlob, E. Grandjean, and K. Seyr, editors, *Proceedings of the 12th International Workshop on Computer Science Logic (CSL-98), Annual Conference on the EACSL*, volume 1584 of *Lecture Notes in Computer Science*, pages 285–297, Brno, Czech Republic, 1999. European Association for Computer Science Logic, Springer.

[11] E. Weydert. JZBR - iterated belief change for conditional ranking constraints. In H. Rott, C. Albert, G. Brewka, and C. Witteveen, editors, *Proceedings of the 4th Dutch-German Workshop on Nonmonotonic Reasoning Techniques and their Applications (DGNMR-99)*, pages 57–66, Amsterdam, The Netherlands, 1999. ILLC Scientific Publications.

[12] E. Weydert. Thoughts on evolving science and belief. In *Proceedings of the Workshop "Scientific Reasoning in AI and Philosophy of Science" at the 14th European Conference on Artificial Intelligence (ECAI-00)*, pages 62–69, Berlin, Germany, 2000. ECAI.

[13] E. Weydert. Rankings we prefer: a rational minimal construction semantics for default reasoning. In *Proceedings of the 5th Dutch-German Workshop on Nonmonotonic Reasoning Techniques and their Applications (DGNMR-01)*, pages 1–12, Potsdam, Germany, April 2001. Institut für Informatik, Universität Potsdam.

[14] E. Weydert and L. van der Torre. Risk parameters for utilitarian desires. In J. Bell, editor, *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-99), Workshop on Practial Reasoning and Rationality*, pages 48–54, Stockholm, Sweden, 1999. IJCAI.

Theses

[1] H. Baumeister. *Relations between Abstract Datatypes modeled as Abstract Datatypes.* Phd thesis, Universität des Saarlandes, November 1999.

[2] C. Brinker. Geometrisches Schließen mit SPASS. Masters thesis, Universität des Saarlandes, January 2000.

[3] F. Eisenbrand. *Gomory-Chvátal Cutting planes and the Elementary Closure of Polyhedra.* Phd thesis, Universität des Saarlandes, July 2000.

[4] M. Fischmann. On Applications of Decidable Object-Oriented Type Theory. Masters thesis, Universität des Saarlandes, 1999.

[5] I. Grenner. Die Erzeugung von Schnittebenen mit maximalem Verletzungsgrad und deren Einsatz im Branch-and-Cut Verfahren. Masters thesis, Universität des Saarlandes, 2001.

[6] U. Hustadt. *Resolution-Based Decision Procedures for Subclasses of First-Order Logic.* Phd thesis, Universität des Saarlandes, March 1999.

[7] C. Meyer. *Soft Typing for Clausal Inference Systems.* Phd thesis, Universität des Saarlandes, 1999.

[8] S. Mukhopadhyay. *A Uniform Constraint-based Framework for the Verification of Infinite State Systems.* Phd thesis, Universität des Saarlandes, May 2001.

[9] C. Schmolzi. Kooperative koevolutionäre Entwicklung von Fuzzy-Reglern. Masters thesis, Universität des Saarlandes, March 1999.

[10] J. Stuber. *Superposition Theorem Proving for Commutative Algebraic Theories.* Phd thesis, Universität des Saarlandes, December 1999.

[11] M. Tzakova. *Hybrid Languages.* Phd thesis, Universität des Saarlandes, 1999.

[12] J. Veit. Formal fairness proofs for optimistic contract signing protocols. Masters thesis, Universität des Saarlandes, October 1999.

Technical reports

[1] A. Bockmayr and F. Eisenbrand. Cutting planes and the elementary closure in fixed dimension. Research Report MPI-I-1999-2-008, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, December 1999.

[2] W. Charatonik, S. Dal-Zilio, A. D. Gordon, S. Mukhopadhyay, and J.-M. Talbot. The complexity of model checking mobile ambients. Technical Report MSR-TR-2001-03, Microsoft Research, Microsoft Corporation, Redmond, USA, May 2001.

[3] H. de Nivelle. Translation of resolution proofs into higher order natural deduction using type theory. Research Report MPI-I-2001-2-004, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, August 2001.

[4] G. Delzanno and J.-F. Raskin. Symbolic representation of upward-closed sets. Research Report MPI-I-1999-2-007, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, November 1999.

[5] F. Eisenbrand. Short vectors of planar lattices via continued fractions. Research Report MPI-I-2000-2-001, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, May 2000.

[6] U. Furbach, H. Ganzinger, R. Hasegawa, and D. Kapur, editors. *Deduction*, number 232 in Dagstuhl-Seminar-Report, Wadern, Germany, 2000. IBFI.

[7] P. Maier. A set-theoretic framework for assume-guarantee reasoning. Research Report MPI-I-2001-2-002, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, April 2001.

[8] A. Nonnengart. A deductive model checking approach for hybrid systems. Research Report MPI-I-1999-2-006, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, November 1999.

[9] V. Sofronie-Stokkermans. On the universal theory of varieties of distributive lattices with operators: Some decidability and complexity results. Research Report MPI-I-2001-2-005, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, 2001.

[10] S. Vorobyov. New lower bounds for the expressiveness and the higher-order matching problem in the simply typed lambda calculus. Research Report MPI-I-1999-3-001, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, July 1999.

[11] S. Vorobyov. Experiments with iterative improvement algorithms on completely unimodel hypercubes. Research Report MPI-I-2001-2-003, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, June 2001.

[12] U. Waldmann. Superposition and chaining for totally ordered divisible abelian groups. Research Report MPI-I-2001-2-001, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, April 2001.

Part V

The Computer Graphics Group

1 Personnel

Director:
Prof. Dr. Hans-Peter Seidel

Research Associates:
Michael Goesele (since Sept. 1999)
Dr. Jörg Haber (since July 1999)
Dr. Wolfgang Heidrich (April 1999 - Sept. 2000)
Jan Kautz (since Sept. 1999)
Priv.-Doz. Dr. Leif Kobbelt (April 1999 - Dec. 2000)
Hendrik Lensch (since Sept. 1999)
Dr. Karol Myszkowski (since June 2000)
Christian Rössl (since Sept. 1999)
Annette Scheel (since April 1999)
Hartmut Schirmacher (since April 1999)
Dr. Robert Schneider (July 1999 - June 2001)
Dr. Marc Stamminger (April 1999 - Nov. 2000)

Postdocs:
Dr. Philippe Bekaert (since Jan. 2001)
Dr. Sung Woo Choi (since March 2000)
Dr. Sherif Ghali (Sept. 1999 - Aug. 2001)
Dr. Vlastimil Havran (since Sept. 2001)
Dr. Ioannis Ivrissimtzis (since Oct. 2001)
Dr. Ulrich Schwanecke (March 2000 - Feb. 2001)
Dr. Cyril Soler (Nov. 1999 - Nov. 2000)
Dr. Hitoshi Yamauchi (since April 2000)
Dr. Frank Zeilfelder (Okt. 1999 - Sept. 2000)

Ph.D. students:
Stephan Bischoff (Oct. 1999 - March 2001)
Mario Botsch (May 1999 - Dec. 2000)
Stefan Brabec (since April 1999)
Katja Daubert (since April 1999)
Won-Ki Jeong (since April 2001)
Kolja Kähler (since Sept. 1999)
Ming Li (since April 2000)
Takehiro Tawara (since July 2000)
Christian Theobalt (since April 2001)
Jens Vorsatz (since April 1999)

Secretary:
Sabine Budde (since June 1999)

2 Visitors

Since June 1999, a total of 53 researchers have visited our group.

Kolja Kähler Katrix Inc. Princeton, USA	20/07/99	
Sven Havemann TU Braunschweig Braunschweig, Germany	22/07/99	
Karol Myszkowski University of Aizu Aizu-Wakamatsu, Japan	02/09/99	
Andreas Pomi vrcom GmbH Darmstadt Germany	14/09/99	
Stephen Mann University of Waterloo Waterloo, Canada	20/09/99	- 23/09/99
Jean Michele Dischler Université de Limoges Limoges, France	23/09/99	
Alexander Keller Universität Kaiserslautern Kaiserslautern, Germany	27/09/99	
Nicolas Holzschuch Université Nancy Nancy, France	01/10/99	
Pere-Pau Vazques Universitat de Girona Girona, Spain	10/10/99	- 04/12/99
George Drettakis IMAGIS- GRAVIR Monbonnot, France	03/02/00	- 04/02/00
Xavier Granier IMAGIS- GRAVIR Montbonnot, France	03/02/00	- 04/02/00
François Sillion IMAGIS- GRAVIR Montbonnot, France	03/02/00	- 04/02/00

Ignacio Martin Campos Universitat de Girona Girona, Spain	02/02/00	- 04/02/00
Frederic Pérez Cazorla Universitat de Girona Girona, Spain	02/02/00	- 04/02/00
Juan Roberto Jiménez Pérez Universitat de Girona Girona, Spain	02/02/00	- 04/02/00
Xavier Pueyo Universitat de Girona Girona, Spain	02/02/00	- 04/02/00
Michael Wand Universität Paderborn Paderborn, Germany	29/02/00	
Robert Shakespeare Indiana University Bloomington, USA	22/03/00	- 26/03/00
Markus Gross ETH Zürich Zürich, Switzerland	07/04/00	
Ulf Labsik Friedrich-Alexander-Universität Erlangen-Nürnberg Erlangen, Switzerland	26/04/00	
James Stewart University of Toronto Toronto, Canada	17/05/00	-19/05/00
Andreas Hubeli ETH Zürich Zürich, Switzerland	24/05/00	
Andrzej Lukaszewski University of Wroclaw Wroclaw, Poland	25/05/00	- 07/06/00
Willian Mark Stanford University Stanford, USA	10/06/00	- 18/06/00

Michael McCool University of Waterloo Waterloo, Canada	16/06/00	- 18/06/00
Oleg Davydov Justus-Liebig-Universtität Gießen Gießen, Germany	08/08/00	- 09/08/00
Roberto Scopigno Instituto CNUCE Pisa, Italy	11/09/00	- 15/09/00
Hiroyuki Akamine University of Aizu Aizu-Wakamatsu, Japan	19/09/00	- 31/01/01
Andreas Weber Fraunhofer Institut IGD Darmstadt, Gemany	27/09/01	
Ron Goldman Universität Kaiserslautern Kaiserslautern, Gemany	23/10/01	- 24/10/01
Geraldine Morin Universität Kaiserslautern Kaiserslautern, Gemany	23/10/01	- 24/10/01
Andrzej Wojdala Orad, Inc. Szczecin, Poland	25/10/01	
Hartmut Prautzsch Universität Karlsruhe Karlsruhe, Gemany	10/11/00	
Lars Linsen Universität Karlsruhe Karlsruhe, Gemany	13/12/00	
Reynald Dumont Cornell University Ithaca, USA	11/01/01	
Norbert Mayer Max-Planck-Institut für Strömungsforschung Göttingen, Germany	28/02/01	

Ron Goldman Universität Kaiserslautern Kaiserslautern, Gemany	02/02/01	
Andrzej Lukaszewski University of Wroclaw Wroclaw, Poland	01/03/01	- 31/08/01
Marco Tarini University of Pisa Pisa, Italy	01/03/01	- 31/08/01
Chang-Hun Kim Korea University Seoul, Korea	01/09/00	
Koung-Hee Kim Ewha Womans University Seoul, Korea	08/02/01	
Pere-Pau Vazques Universitat de Girona Girona, Spain	11/03/01	- 24/03/01
Akira Fujimoto Integra Inc. Tokyo, Japan	22/03/01	
Mark Pauly ETH Zürich Zürich, Switzerland	11/04/01	
Alexander Belyaev University of Aizu Aizu-Wakamatsu, Japan	03/05/01	
Yutaka Ohtake University of Aizu Aizu-Wakamatsu, Japan	03/05/01	
Vincent Scheib University of North Carolina Chapel Hill, USA	23/05/01	- 27/05/01
Oliver Staadt ETH Zürich Zürich, Switzerland	06/06/01	

Vadim Shapiro University of Wisconsin-Madison Madison, USA	25/06/01	- 27/06/01
Reinhard Klein Rheinische Friedrich-Wilhelms-Universität Bonn, Germany	02/07/01	
Volker Blanz Albert-Ludwigs-Universität Freiburg, Germany	04/07/01	
Vlastimil Havran Czech Technical University Prague, The Czech Republic	31/07/01	- 01/08/01
Kirill Dmitriev Keldysh Institute of Applied Mathematics Moscow, Russia	20/05/01	- 20/11/01

3 Group Organization

Our research is currently organized into the following eight research areas, each having its own small group of coordinators:

- Mesh Processing (C. Rössl, H.-P. Seidel)
- Free-Form Surfaces, Subdivision Surfaces, and Shape Analysis (J. Haber, H.-P. Seidel)
- Facial Modeling and Animation (J. Haber)
- Model Acquisition with Realistic Reflection Properties (M. Goesele, H. Lensch)
- Image-Based Rendering and Motion from Video (H. Schirmacher)
- Realistic Hardware-Supported Shading and Lighting (J. Kautz)
- Global Illumination (K. Myszkowski, A. Scheel)
- Perception Issues in Rendering and Animation (K. Myszkowski)

The coordinators coordinate the work in their areas and together with Hans-Peter Seidel form the AG4 steering committee. The steering committee meets on a weekly basis (Tuesday, 3 pm) and discusses all group related issues. In particular, it addresses topics such as recruiting, guests and seminars, teaching, project acquisition, mid-term and long-term strategic planning.

The whole group meets twice a week for the

- AG4 lab meeting (Thursday, 1 pm), where organizational issues are discussed and important news is distributed by the members of the steering committee, and the
- AG4 graphics colloquium (Wednesday, 1pm), where people from within AG4 and the computer graphics group at Saarland University as well as visitors present their ongoing work to the group and to other interested people.

In addition we also have weekly *rendering group meetings* (technical issues in the area of material acquisition, image-based rendering, and hardware rendering) and *geometric modeling meetings* (technical issues in the areas of mesh processing and free form surfaces).

Apart from these weekly meetings, there are several meetings and discussion groups that take place frequently, but not on a totally regular basis, such as paper discussion groups that discuss papers of special interest, especially immediately preceding major conference events; global illumination meetings where issues related with this research area are discussed, usually in cooperation with people from the graphics group at Saarland University; internship and practical course meetings where all people involved in internships or FoPras meet and discuss; and last but not least meetings dedicated to single projects, such as 3D scanning and material acquisition, photo studio, 3D visualization room, conference organization, and several others.

4 Mesh Processing

The use of triangle meshes for the representation of highly complex geometric objects has become the de facto standard in most computer graphics applications. Flexible and effective algorithm have been developed which combine results from approximation theory, numerical analysis, and differential geometry and apply them to the discrete setting of polygonal meshes.

Representing a given (real or virtual) surface geometry by a triangle mesh is usually an approximation process. Hence there is no unique polygonal 3D–model but the density and distribution of sample points and the specific way how these samples are connected by triangles provide many degrees of freedom. For efficient storage and modeling with polygonal meshes, we have to choose a specific instance among those many possible models. Various requirements such as controlling the complexity to adapt the models to the available hardware resources (cf. Section4.1) as well as optimizing the shape or connectivity of the triangles (cf. Section4.2) emerge and are an integral part when processing triangle meshes.

Our goal in mesh processing is to generalize well-established spline techniques to unstructured triangle meshes such that splines can be substituted by polygonal meshes in applications such as multiresolution editing, noise removal, hole-filling, or surface blending (cf. Section 4.3 and Section4.4).

Finally, we discuss surface reconstruction with an emphasis on interactive point cloud triangulation of very large data sets (Section 4.5), feature detection for reverse engineering (Section 4.6), NC milling simulation (Section 4.7), and enhancements of the marching cubes algorithm (Section 4.8).

4.1 Mesh Simplification

Investigators: Leif P. Kobbelt, Kolja Kähler, Jens Vorsatz

Incremental mesh simplification is one of the standard techniques used in computer graphics to control the complexity/quality of triangle meshes and to generate differently coarse approximations to one input mesh. This is particularly interesting for oversampled data sets typically emerging from laser-range scans or for over-tessellated data such as iso-surfaces extracted from medical volume scans. We have developed an incremental mesh reduction scheme [1] that provides fine grained control over the approximation error while being linear in time in the number of vertices. We use the well–known half–edge collapse as our atomic decimation operation but we differ from previously reported methods in that we use a sequence of oriented bounding-boxes to track the decimation error instead of retaining the complete vertex information of the original model. This reduces storage costs and computation time while still providing a reliable upper bound for the deviation from the original data. We have also proposed different error accumulation strategies which makes the algorithm adaptable to different application scenarios. We used the half–edge collapse due to its simplicity, local control, and full reversibility and oriented bounding boxes to achieve tight error bounds. We also investigate different strategies for propagating and merging these error volumes.

The main features of our approach are:

- The error metric measures the real world geometric deviation from the source model.
- The runtime is linear in the number of vertices so that the algorithm does not slow down towards the end of the decimation process.
- The memory requirements are lower than for previously reported schemes. This makes our scheme applicable for larger models.

- The tradeoff between meeting exact error bounds and optimized runtime/storage–costs can be controlled by the user.

References

- [1] J. Vorsatz, K. Kähler, L. Kobbelt, and H.-P. Seidel. Linear time mesh simplification with reliable error-bounds. In B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling, and Visualization (VMV-00)*, pages 145–152, Saarbrücken, Germany, November 2000.

4.2 Remeshing

Investigators: Leif P. Kobbelt, Jens Vorsatz, Christian Rössl

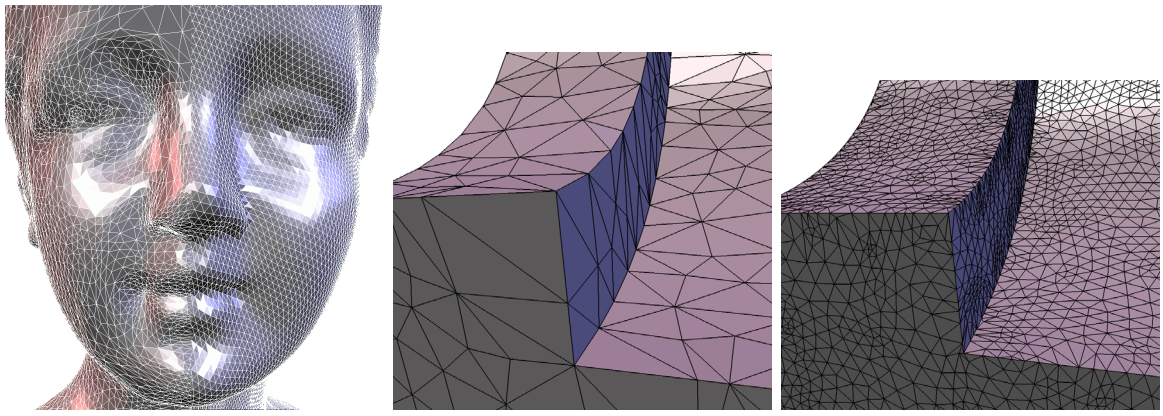

Figure 4.1: *The left-most figure illustrates classical remeshing. The left side of the bust's face shows the original connectivity, the semi-regular mesh on right side is the result of the shrink–wrapping process. Center and right show the fan model at two different resolutions (2.8K△, 40K△ resp.) obtained by our feature sensitive remeshing algorithm.*

Triangle meshes of *arbitrary* connectivity are the most versatile surface representation in computer graphics and CAGD since they enable the description of surfaces with arbitrary shape and topology. However, especially in the context of multiresolution representation and modeling, many algorithms require a specific connectivity, namely *subdivision-connectivity*. This special type of connectivity is generated by successively applying a subdivision operator (cf. Section 5.3) to a coarse base mesh. The resulting mesh is said to be *semi-regular* since each submesh topologically emerging from one coarse triangle has the connectivity of a regular grid. The implicitly defined connectivity established on a coarse base mesh and the direct availability of multiresolution semantics gives rise to many techniques exploiting this representation as the following enumeration shows.

Compression/progressive transmission Based on a wavelet decomposition, a compressed approximation within a given error tolerance can be achieved by suppressing small wavelet coefficients. Moreover such a wavelet representation can easily be transmitted progressively.

Multiresolution editing A combination of subdivision and 'smoothing operations' enables multiresolution editing where efficient modifications of the global shape can be performed while surface detail is preserved.

Parameterization Each submesh *(subdivision–patch)* can be parameterized naturally by assigning barycentric coordinates to the vertices. Combining the local parameterizations of the subdivision–patches yields a global parameterization. Texturing is just one application of such a parameterization.

Level–of–detail control Standard rendering libraries are able to display objects at various levels of detail, that is they display a coarse approximation, if the object is far away and switch to a finer one, if the viewer zooms in. The different subdivision levels naturally support this feature.

The so called *classical remeshing* is the process which converts a triangle mesh with arbitrary connectivity into one having subdivision-connectivity. Just as the wide family of subdivision techniques, remeshing starts with a coarse base mesh and introduces an exponential number of vertices with a split operation. But instead of applying a smoothing rule to the vertices, the remeshing process places the vertices to capture finer detail information and thus approximates the original shape. This is done until a prescribed tolerance is reached.

A well established conversion scheme is the MAPS algorithm [2], which uses mesh decimation and successive harmonic maps to generate a parameterization and thus provides the basis to build a semi-regular remesh. Since MAPS generates the base mesh automatically, it is hard to control it's quality in terms of triangle shape and triangle count. In order to custom tailor the base mesh, we presented a completely different approach to the remeshing problem for genus–zero objects [1]. The physical model behind the algorithm is the process of *shrink wrapping* where a plastic membrane is wrapped around an object and shrunk either by heating the material or by evacuating the air from the space in between the membrane and the object's surface. At the end of the process, the plastic skin provides an exact imprint of the given geometry. To simulate the shrink wrapping process, we approximate the plastic membrane by a semi–regular mesh $\mathcal{S}_m$. Two forces are applied to its vertices. An attracting force pulls them towards the surface. A relaxing force is applied in order to optimize the local distortion energy and to avoid folding. This ensures an even distribution of the vertices. The attracting part is realized by a projecting operator $\mathbf{P}$ that projects $\mathcal{S}_m$'s vertices on $\mathcal{M}$. The relaxing is done by applying a relaxing operator $\mathbf{U}$ to all vertices in $\mathcal{S}_m$, that iteratively balances the vertex distribution.

Beyond 'classical remeshing', we have also investigated remeshing in a more general context. General means that we do not restrict ourselves to the generation of semi–regular meshes, but we generate meshes of arbitrary connectivity in order to prepare a given mesh for a specific application. We take a triangle mesh and resample it in a way that the new tessellation still approximates the same geometric shape but additionally satisfies some quality requirements. Remeshing artifacts are a fundamental problem when converting a given geometry into a triangle mesh. We have proposed a new remeshing technique that is sensitive to features [3]. In our algorithm, the resolution of the mesh is iteratively adapted by a global restructuring process which additionally optimizes the connectivity in a first step. Then a particle system approach evenly distributes the vertices across the original geometry. To exactly find the features we extend this relaxation procedure by an effective mechanism to attract the vertices to feature edges. The attracting force is imposed by means of a hierarchical curvature field and does not require any thresholding parameters to classify the features.

References

- [1] L. Kobbelt, J. Vorsatz, U. Labsik, and H.-P. Seidel. A shrink wrapping approach to remeshing polygonal surfaces. In P. Brunet and R. Scopigno, editors, *Proceedings of the 20th Annual Conference ot*

the European Association of Computer Graphics (Eurographics-99), volume 18 of *Computer Graphics Forum*, pages 119–130, Milan,Italy, September 1999. Blackwell.

[2] A. W. F. Lee, W. Sweldens, P. Schröder, L. Cowsar, and D. Dobkin. MAPS: Multiresolution adaptive parameterization of surfaces. In *SIGGRAPH 98 Conference Proceedings*, pages 95–104, 1998.

- [3] J. Vorsatz, C. Rössl, L. Kobbelt, and H.-P. Seidel. Feature sensitive remeshing. In A. Chalmers and T.-M. Rhyne, editors, *Computer Graphics Forum, Proceedings of Eurographics '01*, volume 20 of *Computer Graphics Forum*, pages 393–401, Manchester, Great Britain, September 2001. Blackwell.

4.3 Discrete Fairing

Investigators: Robert Schneider, Leif P. Kobbelt

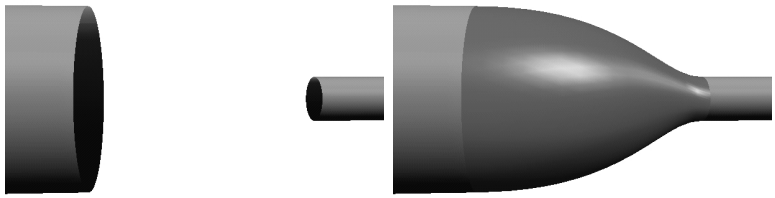

Figure 4.2: *Typical fields of application where fairing problems satisfying certain interpolation constraints have to be solved. The figures show how two objects can be connected by constructing a blending surface with the HMCS method.*

While surface representations based on smoothly joining polynomial segments have dominated the geometric modeling world for a long time, polygonal descriptions are likely to become the future standard. This requires us to be able to manipulate meshes with techniques analogous to the ones that are available for their smooth piecewise polynomial counterparts. An important technique is the creation or manipulation of surfaces in a way that the result is functionally pleasing or looks aesthetic to the human eye. This is known as the surface fairing problem. Our main focus is the development of fairing methods that are independent of an underlying parameterization. Typical fields of application where such problems arise are

- N-sided hole filling
- Surface blending
- Multiresolution modeling

We have present a new intrinsic fairing technique that can be applied to planar polygonal curves as well as to triangular meshes (cf. [2]). This approach can handle general intrinsic problems, but is especially simple for some important classes. In the case of surfaces this fact leads us to a new surface class that is an interesting alternative to minimal energy surfaces which have most often been used in creating high quality surfaces up to now. In addition to this intrinsic fairing algorithm, we have also presented a mesh fairing technique that is closely related to the linear fairing approach presented by Kobbelt et al. [1], but produces a mesh quality that is clearly superior to the results of the original method. All approaches exploit multigrid techniques to speed up the construction process. In detail, our contribution can be summarized as follows:

- We presented a new discrete fairing technique that can be used to create curves and surfaces of high quality.
- We introduced a new surface class, the harmonic mean curvature surfaces (HMCS). Interestingly, this surface class doesn't seem to have appeared in the geometric modeling literature for fairing purposes up to now. We showed that this surface class is a very interesting alternative to minimal energy surfaces (MES).
- Our fairing approach becomes especially simple and efficient for some special classes. In the case of curves this includes the very popular clothoid splines and Ohlin's A5-splines; for surfaces the construction algorithm becomes especially simple for HMCS. For these classes our construction algorithm has many advantages compared to the curvature flow approach, e.g., it is much faster and leads to vertex update steps that arise automatically from the construction algorithm.
- Our algorithm can be used to create curves and surfaces that are defined as solutions of general intrinsic differential equations. In the curve case, we showed that our method can be used to create discrete minimal variation curves (MVCs). We have not seen a construction method for MVCs in a polygonal representation so far.
- In addition to completely intrinsic mesh fairing techniques, we also presented a simple mesh fairing algorithm that improves the fairing method developed by Kobbelt et al. [1]. While sharing many advantages of the original approach, it leads to a shape quality that is much closer to the quality of intrinsic fairing methods.

References

[1] L. Kobbelt, S. Campagna, J. Vorsatz, and H.-P. Seidel. Interactive Multi-Resolution Modeling on Arbitrary Meshes. In *SIGGRAPH '98 Conference Proceedings*, pages 105–114, 1998.

- [2] R. Schneider and L. Kobbelt. Geometric fairing of irregular meshes for free-form surface design. *Computer Aided Geometric Design*, 18(4):359–379, May 2001.

4.4 Mesh Modeling

Investigators: Leif P. Kobbelt, Jens Vorsatz

In [2] we have generalized powerful multiresolution techniques based on semi-regular meshes to meshes with arbitrary connectivity. Our major observation is that the hierarchy of nested spaces which is the structural core element of most multiresolution algorithms can be replaced by a sequence of intermediate meshes emerging from the application of incremental mesh decimation. Performing local frame coding already enables effective algorithms to extract multiresolution information from unstructured meshes. In combination with discrete fairing techniques (cf. Section4.3) we are able to reduce noise from a geometrically specified frequency band in a multiresolution decomposition. Putting mesh hierarchies, local frame coding and multilevel smoothing together gives us a flexible and intuitive paradigm for interactive detail preserving mesh modifications (cf. Fig. 4.3). In [3] we have improved this concept in two aspects. We have extended the two-band decomposition to a robust and intuitive multiband decomposition. On the other hand we have generalized the local frame concept to a more general detail encoding concept using a Phong-type normal field. What is common to all these approaches is the fact that the multiresolution representation always imposes a fixed global hierarchical structure. This is a severe restriction since extreme edits of

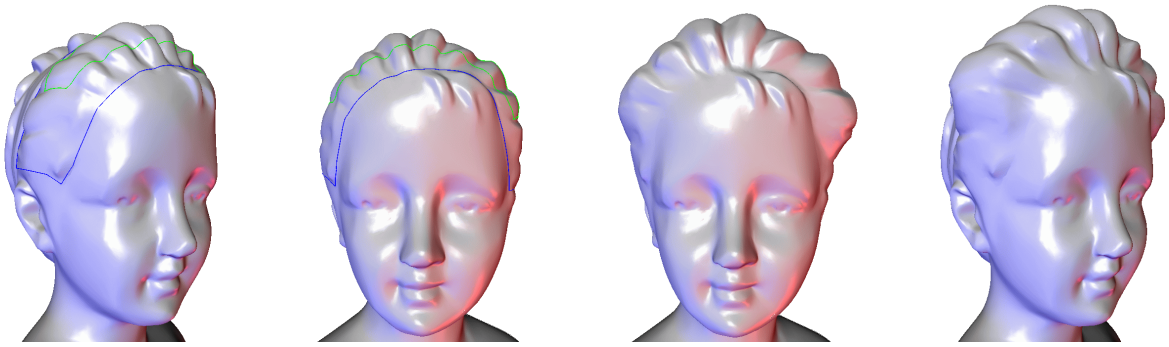

Figure 4.3: *Multiresolution editing of a bust model. The area within the blue line is decomposed into two frequency-bands. The user changes the handle polygon (green) and thus changes the low-frequency surface on top of which the recorded detail based on a Phong-type normal field is reconstructed (two right images).*

the surface geometry causes strong changes in the surface metric. In [1] we propose a multiresolution representation that no longer requires such a global structure. Instead we represent the detail information implicitly by the geometric difference between independent surfaces. The detail information is evaluated by shooting rays in normal direction from one surface to the other without assuming a consistent parameterization or tessellation. In the context of shape deformation we have proposed a dynamic mesh structure which adapts the connectivity during the modification in order to maintain a prescribed mesh quality. Combining these two techniques leads to an efficient mechanism which enables extreme deformations while preserving the mesh from degenerating and reconstructing the detail in a stable and intuitive way.

References

- [1] L. Kobbelt, T. Bareuther, and H.-P. Seidel. Multiresolution shape deformations for meshes with dynamic vertex connectivity. In M. Gross and F. R. A. Hopgood, editors, *Proceedings of the European Association for Computer Graphics 21st Annual Conference (Eurographics-00)*, volume 19 of *Computer Graphics Forum*, pages 249–260, Interlaken, Switzerland, 2000. Blackwell.

  [2] L. Kobbelt, S. Campagna, J. Vorsatz, and H.-P. Seidel. Interactive Multi-Resolution Modeling on Arbitrary Meshes. In *SIGGRAPH '98 Conference Proceedings*, pages 105–114, 1998.

- [3] L. Kobbelt and J. Vorsatz. Multiresolution hierarchies on unstructured triangle meshes. *Computational Geometry*, 14(1-3):5–24, 1999.

4.5 Interactive Point Cloud Triangulation by Virtual Range Scanning

Investigators: Mario Botsch, Christian Rössl, Leif Kobbelt

The amount of geometry data obtained from acquisition hardware is typically huge. Point clouds with millions or even tens of millions of samples are no exception if the surface of a non-trivial geometric object is to be represented. This makes some preprocessing mandatory in order to reduce the input complexity of subsequent steps while observing a prescribed approximation tolerance.

Many algorithms have been proposed for generating meshes from point clouds. The motivation for investigating a new approach was the need for a robust system that is able to process highly complex and completely unstructured input data and produces high-quality triangles meshes. In

contrast to other algorithms we chose an interactive approach that gives the user fine control over the meshing process using the intuitive and powerful *virtual range scanner* metaphor:

The concept behind the user interface is to simulate a virtual 3D scanning device which is more flexible than a real 3D scanner [2]. The user chooses the optimal viewing direction for the processed object on the screen. Then a snapshot (a virtual 3D scan) is taken. Arbitrary subregions of this scan can be masked out to cut away unwanted parts. Taking a 3D snapshot means reading out the z-buffer contents and un-projecting it back to 3-space. Hence, this operation can be performed by the graphics hardware and is accomplished very fast. In a postprocessing step the captured range image can be smoothed, holes can be fixed etc. Then the actual patch is merged with the already existing ones using a zippering [4] like approach. After a few iterations of this procedure we end up with a globally consistent model of the given object. Figure 4.4 shows surfaces that have been reconstructed from unorganized-point clouds by applying this process.

While interactiveness may not be appropriate for every application this new approach gives the user very powerful control over the mesh generation while still being intuitive. Another advantage is the great flexibility of this approach. Every object that can be rendered to the screen can be captured, i.e. we can also handle *hybrid* data combining point clouds and e.g. CAD data easily into a single triangle mesh.

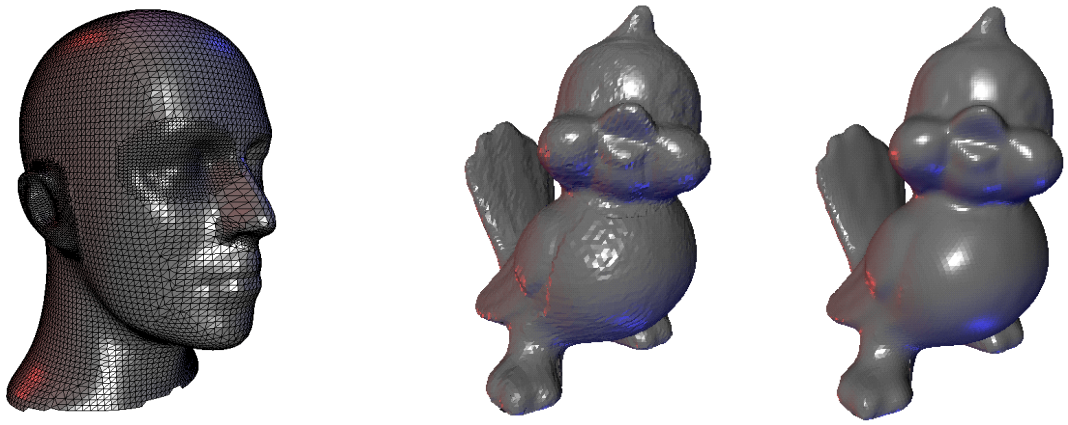

Figure 4.4: *Examples of virtual 3D scanning. Left: the resolution of the model can be controlled locally per patch by simply zooming the point cloud. Right: The reconstructed tweety model before and after smoothing.*

As every sampling technique *virtual range scanning* suffers from aliasing. This can be avoided by choosing the viewing directions appropriately: by aligning sharp features of the object to the sample grid aliasing can be suppressed effectively.

This method of aligning to the sampling has been generalized and automated. Where necessary, we allow curved grids that trace to the feature regions of the surface. Following [3] the user can either manually layout this grid by specifying a surface polygon as backbone or let the system trace this backbone automatically guided by local principal curvature directions. In either case, only view interaction is necessary as ribs are constructed automatically from the backbone resulting in a *fishbone* like structure that is aligned to the local feature and will be used as sampling grid [1].

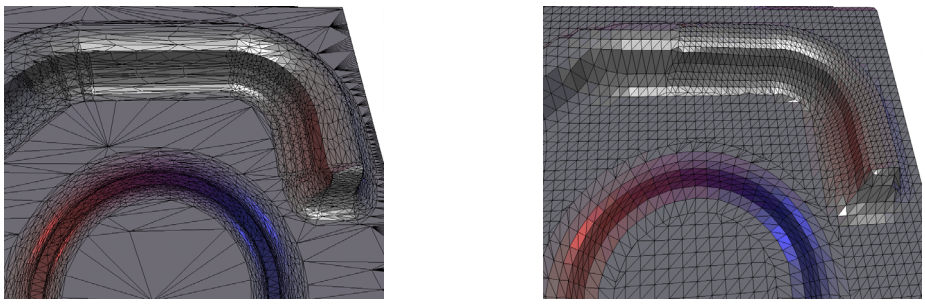

Figure 4.5: *Resampling of an object with curved feature lines (left). The right mesh was obtained by using sample grids that follow the feature. In the upper left part the standard uniform grid (aligned to the x-axis) was used; here, one can clearly notice aliasing artifacts.*

References

- [1] M. Botsch, C. Rössl, and L. Kobbelt. Feature sensitive sampling for interactive remeshing. In B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling, and Visualization (VMV-00)*, pages 129–136, Saarbrücken, Germany, November 2000.
- [2] L. Kobbelt and M. Botsch. An interactive approach to point cloud triangulation. In M. Gross and F. R. A. Hopgood, editors, *Proceedings of the European Association for Computer Graphics 21st Annual Conference (Eurographics-00)*, volume 19 of *Computer Graphics Forum*, pages 479–487, Interlaken, Switzerland, 2000. Eurographics, Blackwell.
- [3] C. Rössl and L. Kobbelt. Line-art rendering of 3D-models. In B. A. Barsky, Y. Shinagawa, and W. Wang, editors, *Proceedings of the 8th Pacific Conference on Computer Graphics and Applications (PG-00)*, pages 87–96, Hong Kong, China, 2000. IEEE Computer Society.

[4] G. Turk and M. Levoy. Zippered polygon meshes from range images. In *SIGGRAPH 94 Conference Proceedings*, pages 311–318, 1994.

4.6 Feature Detection on Surfaces for Reverse Engineering

Investigators: Christian Rössl, Leif Kobbelt

Reverse engineering is the process of converting geometric data into a representation suitable for CAD systems. In spite of the growing popularity of piecewise linear triangle meshes still piecewise smooth, continuous surface representations are the de facto standard in CAD. For the conversion of an arbitrary triangle mesh that may be acquired from a real world object, it is crucial to separate surface regions that can be replaced by one smooth surface patch, such as splines or parts of geometric primitives like cylinders, spheres, etc.

It seems appropriate to use discrete curvature information [1] for guiding this segmentation step. But as second order derivatives are involved, special care is required for handling stochastical noise in the input data.

We use a robust thresholding for separating feature regions. Morphological operators that are well known in digital image processing have been adapted to work on triangle meshes. By applying the dilation and erosion noise is effectively reduced from the initial feature regions but they are still too wide as to be useful as patch boundaries. A new skeletonization scheme is used to shrink the feature regions to feature lines [2, 3] that are suitable for partitioning the surface. Figure 4.6 shows steps in this segmentation process. All operators are kept simple and can be implemented very efficiently on triangle meshes. The obtained structural information on the surface can also be used in other applications where feature detection techniques are needed.

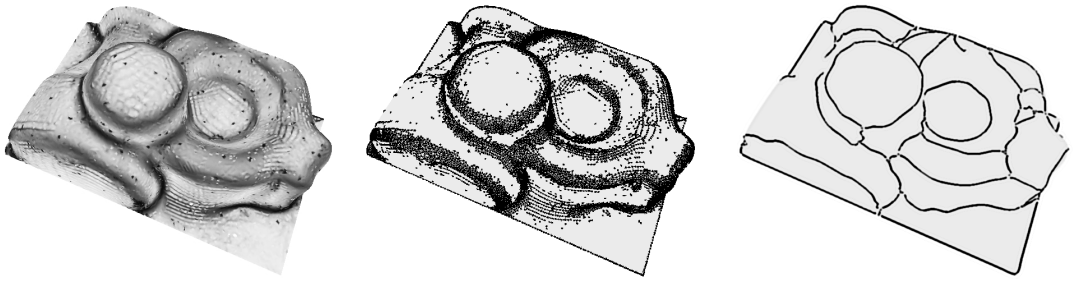

Figure 4.6: *Feature detection on a benchmark model. The maximum curvature (left) is thresholded to obtain the initial feature regions (middle). The (smoothed) feature lines are then extracted from these regions by skeletonization (right).*

References

- [1] C. Rössl and L. Kobbelt. Approximation and visualization of discrete curvature on triangulated surfaces. In B. Girod, H. Niemann, and H.-P. Seidel, editors, *Proceedings of the 4th Conference on Vision, Modeling, and Visualization (VMV-99)*, pages 339–346, Erlangen, Germany, 1999. SFB 603 and Graduate Research Center, in coorperation with IEEE Signal Processing Society and Gesellschaft für Informatik GI, infix.
- [2] C. Rössl, L. Kobbelt, and H.-P. Seidel. Extraction of feature lines on triangulated surfaces using morphological operators. In A. Butz, A. Krüger, and P. Olivier, editors, *Smart Graphics (AAAI Spring Symposium-00)*, volume 00-04 of *Technical Report / SS / American Association for Artificial Intelligence*, pages 71–75, Stanford, USA, 2000. AAAI Press.
- [3] C. Rössl, L. Kobbelt, and H.-P. Seidel. Recovering structural information from triangulated surfaces. In *Mathematical Methods in CAGD: Oslo 2000 Proceedings*, Oslo, Norway, 2001. Vanderbilt University Press.

4.7 NC Milling Simulation

Investigators: Ulrich Schwanecke, Leif Kobbelt

As part of an industrial cooperation a special NC milling process was simulated in order to test NC programs without the immense cost for raw material. The goal was to construct a very detailed triangle mesh representation of the milling result. This investigation lead to an algorithm for approximate envelope reconstruction for moving solids [1].

The standard analytic methods for calculating envelops become extremely complex when the topology of the swept volume gets more complicated. Moreover, these methods are not able to handle self intersections of the swept volume boundary that occur for even fairly simple sweeps.

We developed a new algorithm for reconstructing a polygonal approximation of the envelope of a swept solid that automatically generates the topological correct solution within a prescribed error tolerance. The resulting envelope surface is guaranteed to be a manifold. The algorithm is robust in a way that we do not have to take special attention to – possibly bad conditioned – surface-surface intersections. It is scalable as it allows to calculate coarse approximations very fast while better approximations can be obtained by investing more computation time and memory.

This is accomplished by determining a subdivision of the space the swept volume lives in and using linear sweeps between discrete time steps to approximate the target surfaces rather than calculating algebraic expressions of the swept solid. With the help of an octree based data structure a

polygonization of this surface can be determined in a very efficient manner. By dynamically updating the octree in each time step we successively construct the whole piecewise linear approximation of the swept volume. A surface resulting from a simulated milling process is shown in Figure 4.7.

References

- [1] U. Schwanecke and L. Kobbelt. Approximate envelope reconstruction for moving solids. In *Mathematical Methods in CAGD: Oslo 2000 Proceedings*, Oslo, Norway, 2001. Vanderbilt University Press.

4.8 Extended Marching Cubes Algorithm

Investigators: Ulrich Schwanecke, Mario Botsch, Leif Kobbelt

The representation of geometric objects based on volumetric data structures has advantages in many geometric applications that require, e.g., fast surface interrogation or boolean operations such as intersection and union. However, surface based algorithms like shape optimization (fairing) or free-form modeling often need a topological manifold representation where neighborhood information within the surface is explicitly available.

Consequently it is necessary to find effective conversion algorithms to generate explicit surface descriptions for the geometry which is implicitly defined by a volumetric data set. The de facto standard algorithm for this purpose is Marching Cubes [2] since over a decade. It takes samples of the volume data on a regular grid as input and constructs a triangle mesh as output.

Due to this sampling we often observe severe aliasing artifacts at sharp features on the extracted surfaces. We developed a new algorithm [1] that performs feature sensitive sampling while keeping the simple algorithmic structure of the standard Marching Cubes algorithm.

In order to achieve this we propose an enhanced representation of the discrete distance field: instead of using a *scalar* distance value for each grid point of a uniform spatial grid, we store *directed* distances in x, y, and z direction. This allows us to find more accurate surface samples while the generation and handling of this enhanced distance field turns out to be no more complicated than the processing of scalar fields. Furthermore, we presented the *Extended Marching Cubes* algorithm that detects those grid cells through which a sharp feature (edge or corner) of the considered surface passes. Based on the local distance field information and its gradient, additional sample points lying on the feature are computed and inserted into the mesh. This significantly reduces alias while introducing only a minor computational overhead. Figure 4.7 shows results from both the standard and the extended Marchings Cubes algorithm.

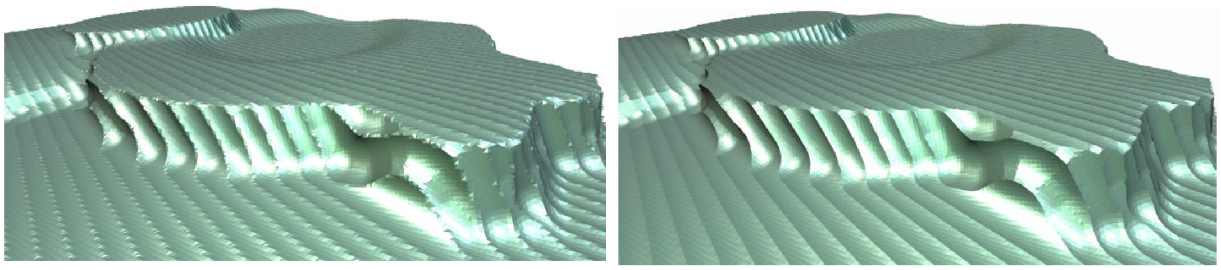

Figure 4.7: *Result of a NC milling simulation. The left image shows the surface extracted by the standard Marching Cubes algorithm, the right image shows the extended MC surface. The sharp ridges are better visible due to the clearly reduced alias.*

References

- [1] L. Kobbelt, M. Botsch, U. Schwanecke, and H.-P. Seidel. Feature sensitive surface extraction from volume data. In *Computer Graphics (SIGGRAPH-01): Conference proceedings*, pages 57–66, Los Angeles, USA, 2001. ACM Press.

[2] W. E. Lorensen and H. E. Cline. Marching cubes: a high resolution 3D surface construction algorithm. In *Computer Graphics (SIGGRAPH 87 Conference Proceedings)*, volume 21, pages 163–170, 1987.

5 Free-Form Surfaces, Subdivision Surfaces, and Shape Analysis

We investigate interpolation and approximation methods for bivariate splines (smooth piecewise polynomial functions defined w.r.t. triangulations in the plane) and subdivision surfaces. These spaces are a natural generalization of univariate splines and therefore of considerable interest in approximation theory and numerical analysis.

In addition, we examine methods for shape analysis such as, e.g., the medial axis transform and an approximation of the discrete cosine transformation, which are used for feature extraction of planar shapes.

5.1 Scattered Data Approximation

Investigators: Jörg Haber, Frank Zeilfelder

We have developed an efficient method to automatically compute a smooth approximation of large functional scattered data sets given over arbitrarily shaped planar domains [2, 1]. Our approach is based on the construction of a C^1-continuous bivariate cubic spline and our method offers optimal approximation order. Both local variation and non-uniform distribution of the data are taken into account by using local polynomial least squares approximations of varying degree (see Figure 5.8). Since we only need to solve small linear systems and no triangulation of the scattered data points is required, the overall complexity of the algorithm is linear in the total number of points. Numerical examples dealing with several real world scattered data sets with up to millions of points demonstrate the efficiency of our method. The resulting spline surface is of high visual quality and can be efficiently evaluated for rendering and modeling. In our implementation we achieve real-time frame rates for typical fly-through sequences and interactive frame rates for recomputing and rendering a locally modified spline surface.

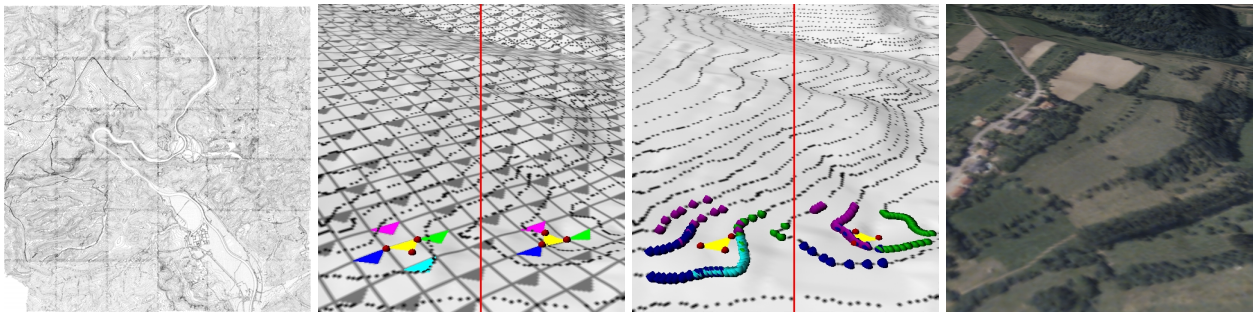

a) data distribution b) two triangle configs c) assoc. data points d) surface with texture

Figure 5.8: *Stages of the approximation process. a) Distribution of 736,577 scattered data points from the area of the* Saarschleife *between Merzig and Mettlach. b) Perspective view onto the surface with projected spline grid and data distribution. Two different sets of triangles (blue, magenta, green, cyan) are shown on the left and on the right. The Bézier points over these triangles are computed from the given scattered data points by local least squares approximation. Figure c) depicts the corresponding scattered data points (color coded) that have been used in this local least squares fitting for each triangle. Once the Bernstein-Bézier representation over the blue, magenta, green, and cyan triangles is available, the Bernstein-Bézier representation over the yellow triangles can be computed from the surrounding triangles using the C^1 smoothness conditions. d) Final result with applied texture.*

References

- [1] J. Haber. Konstruktion und Darstellung glatter Approximationsflächen für große Mengen irregulär verteilter Daten. *Informatik Spektrum*, 24(4):258–259, Aug. 2001.
- [2] J. Haber, F. Zeilfelder, O. Davydov, and H.-P. Seidel. Smooth Approximation and Rendering of Large Scattered Data Sets. In *Proceedings of IEEE Visualization 2001*, San Diego, USA, October 2001. IEEE Computer Society. To appear.

5.2 Bivariate Splines

Investigator: Frank Zeilfelder

Main questions in the field of bivariate splines are as follows (cf. [3, 10]): What is the dimension? What is the approximation order? How can explicit interpolation schemes for these spaces be constructed? These questions yield to complex mathematical problems, in particular, when the degree approaches the smoothness of the space. This is in complete contrast to the univariate theory where characterizing results concerning interpolation and approximation exist (for instance, in [7], we gave the first complete characterization of so-called strongly unique uniform best approximation for periodic splines of even dimension). On the other hand, these splines are powerful tools for many applications.

One case of particular interest is the space of C^1 cubics. This space has a complex structure: it is not known if its dimension always coincides with Schumaker's lower bound. In [1], we develop an algorithm for constructing Lagrange and Hermite interpolation sets for C^1 cubics defined w.r.t. general classes of natural triangulations built up of nested polygons whose vertices are connected by line segments.

Inductive methods to construct Lagrange and Hermite interpolation for splines of degree $q \geq 2r+1$, where $r \in \{1, 2\}$ is the smoothness, simultaneously with the underlying triangulations were given in [4]. In these methods the interpolating splines can be computed locally by passing from triangle to triangle while solving only small linear systems which stands in contrast to earlier global approaches where large system have to be solved.

Since recently, we investigate local Lagrange interpolation by bivariate splines. Although there is a vast literature on local Hermite interpolation and on non-local Lagrange interpolation no results of this type were known, except for the case of continuous (linear) splines. One main difficulty of constructing such interpolation schemes is that the known local Hermite interpolation schemes cannot be transformed into local Lagrange interpolation schemes on the whole triangulation. On the other hand, local Lagrange interpolation schemes are important for instance in scattered data fitting problems, since in contrast to Hermite interpolation methods, only function values (and no (orthogonal) derivatives) have to be computed approximately by using local methods.

In [9], we develop the first local Lagrange interpolation scheme for C^1 cubics on triangulations whose interior vertices have degree six. The method yield optimal approximation order while the number of interpolation conditions is considerably smaller than in the classical schemes. In the meantime, we developed methods (cf. [6, 5]) for constructing local Lagrange interpolation sets for C^1-splines of degree $q \geq 2$ on arbitrary triangulations Δ. The investigations show that for C^1 quadratics, the triangulation Δ has to be refined by applying the so-called Powell-Sabin split to every triangle, while for higher degree splines the triangles of Δ have to be colored in a suitable way. The construction of local Lagrange interpolation sets for C^1 splines on (classes of) triangulated quadrangulation is more complex, since no splits are allowed. The first local Lagrange interpolation scheme for C^1 cubics on so-called checkerboard quadrangulation was given in [2]. Local Lagrange

interpolation by C^1 splines on general (classes of) given triangulated quadrangulations yield to coloring problems of special type and is currently under study.

The algorithmic complexity of our interpolation methods is linear in the number of triangles, and Bernstein-Bézier techniques can be applied to compute the polynomial pieces. Numerical tests with up to 10^6 interpolation points treating test functions as well as real world data show that the interpolation methods work efficiently.

In [8], we consider the problem of constructing a Hermite interpolation operator for bivariate polynomial splines of degree $q \geq 3r+2$ and arbitrary smoothness r, whose approximation error is bounded above by Kh^{q+1}, where h is the maximal diameter of the triangles in Δ and the constant K only depends on the smallest angle of the triangulation. This problem was brought up about ten years ago. The purpose of this work is to develop the first Hermite type interpolation scheme for $S_q^r(\Delta)$, $q \geq 3r+2$, that possesses optimal approximation order in the sense that the corresponding constant K is independent of near-degenerate edges and near-singular vertices. In our approach we argue directly with nodal functionals and develop a new method different from quasi-interpolation that has its roots in the idea of weak interpolation. Since the fundamental functions corresponding to our interpolation scheme are minimally supported, we show in particular that the optimal approximation order can be achieved by using minimally supported splines. Our result is derived since we manage to exploit linear dependencies between directional derivatives of smooth functions to suppress the instability of the minimally supported fundamental functions in the neighborhood of a near-singular vertex.

References

- [1] O. Davydov, G. Nürnberger, and F. Zeilfelder. Cubic spline interpolation on nested polygon triangulations. In A. Cohen, C. Rabut, and L. L. Schumaker, editors, *Curve and Surface Fitting, Saint-Malo 1999*, Innovations in Applied Mathematics, pages 161–170, Saint Malo, June 2000. Vanderbilt University Press.

- [2] G. Nürnberger, L. L. Schumaker, and F. Zeilfelder. Local lagrange interpolation by bivariate cubic C^1 splines. In T. Lyche and et. al., editors, *Proceedings of the Conference on Curves and Surfaces*, Oslo, Norway, 2001. Vanderbilt University Press.

- [3] G. Nürnberger and F. Zeilfelder. Developments in bivariate spline interpolation. *Journal of Computational and Applied Mathematics*, 120(1/2):125–152, 2000.

- [4] G. Nürnberger and F. Zeilfelder. Interpolation by spline spaces on classes of triangulations. *Journal of Computational and Applied Mathematics*, 119(1/2):347–376, May 2000.

[5] G. Nürnberger and F. Zeilfelder. Lagrange interpolation by bivariate C^1-splines with optimal approximation order, 2001. submitted for publication.

- [6] G. Nürnberger and F. Zeilfelder. Local lagrange interpolation on powell-sabin triangulations and terrain modelling. In W. Haussmann and et. al., editors, *Proceedings of the Conference on Multivariate Approximation*, Witten, Germany, 2001. Wiley-VCH.

- [7] F. Zeilfelder. Strong unicity of best uniform approximation from periodic spline spaces. *Journal of Approximation Theory*, 99(2):1–29, 1999.

- [8] F. Zeilfelder, O. Davydov, and G. Nürnberger. Bivariate spline interpolation with optimal approximation order. *Constructive Approximation*, 17(2):181–208, May 2001.

- [9] F. Zeilfelder and G. Nürnberger. Local lagrange interpolation by cubic splines on a class of triangulations. In K. Kopotun, T. Lyche, and M. Neamtu, editors, *Proceedings of the Conference on Trends in Approximation Theory*, pages 341–350, Nashville, USA, June 2001. Vanderbilt University.

- [10] F. Zeilfelder and H.-P. Seidel. Splines over triangulations. In G. Farin, J. Hoschek, M.-S. Kim, and D. Abma, editors, *The Handbook of Computer Aided Design*, chapter Splines over Triangulations. Elsevier Science, Amsterdam, Netherlands, 2001.

5.3 Subdivision Surfaces

Investigators: Leif Kobbelt, Stephan Bischoff, Robert Schneider, Ulrich Schwanecke

Subdivision schemes have become increasingly popular in recent years because they provide a uniform and efficient way to describe smooth curves and surfaces. Their beauty lies in their elegant mathematical formulation and their simple implementation. The real strengths of subdivision schemes are revealed in the bivariate case as they are easily able to handle surfaces of arbitrary topology while automatically maintaining continuity properties. At present almost everything one can do with traditional NURBS-based systems can also be achieved by subdivision techniques. A popular example for their application are animations in computer generated movies where subdivision techniques have become a standard tool [1].

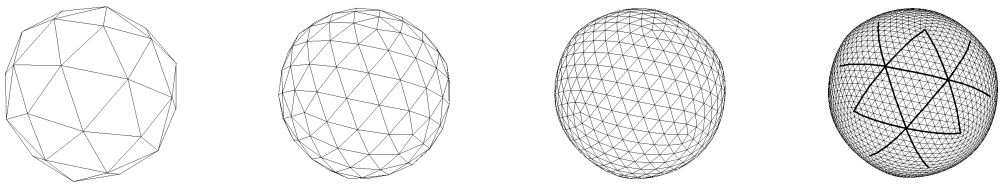

Figure 5.9: *Example of a subdivision process.*

We continued a long tradition of research on subdivision techniques ranging from stationary subdivision schemes [2] and variational subdivision [3] to applications e.g. for ray tracing [4]:

A new stationary subdivision scheme was presented with the $\sqrt{3}$ scheme that generates (almost everywhere) C^2 surfaces and allows adaptive refinement. We developed an algorithm for efficient evaluation of Loop subdivision surfaces aiming at a future hardware implementation. Subdivision techniques can also be exploited for efficient storage and transmission of geometry, a robust transfer protocol for multimedia applications was presented in this context. In the univariate case, research on subdivision schemes focussed on Hermite subdivision.

References

[1] T. DeRose, M. Kass, and T. Truong. Subdivision surfaces in character animation. In *SIGGRAPH 98 Conference Proceedings*, pages 85–94, 1998.

[2] L. Kobbelt. Interpolatory subdivision on open quadrilateral nets with arbitrary topology. *Computer Graphics Forum*, 15(3):409–420, 1996.

[3] L. Kobbelt. A variational approach to subdivision. *Computer Aided Geometric Design*, 13(8):743–761, 1996.

[4] L. Kobbelt, K. Daubert, and H.-P. Seidel. Ray tracing of subdivision surfaces. *Eurographics Rendering Workshop*, pages 69–80, 1998.

5.3.1 $\sqrt{3}$ Subdivision

$\sqrt{3}$-subdivision [1, 2] is a new stationary subdivision scheme that produces C^2 limit surfaces everywhere except for the extraordinary vertices where it is C^1. The stencils for the subdivision rules

have minimum size and maximum symmetry. A 1–to–3 split is used for refinement (Figure 5.10) in contrast to the common 1–to–4 split operator that most other subdivision schemes on triangle meshes are based on. This results in a considerably slower topological refinement compared to the usual dyadic split operation, i.e. less triangles are needed to represent the same geometry with a given approximation tolerance.

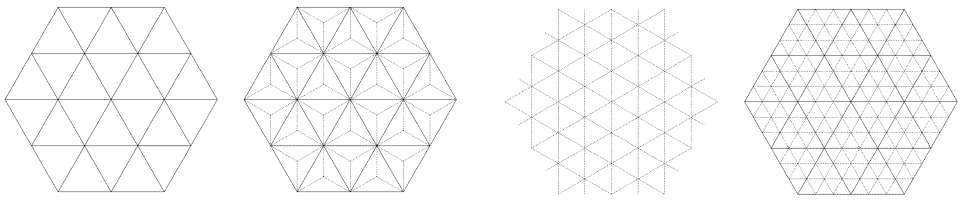

Figure 5.10: *$\sqrt{3}$-scheme: original mesh (left), after inserting barycenters (center left) and after edge flipping (center right). Note that two $\sqrt{3}$-subdivision steps result in a tri-section of the original triangle edges (right), hence the name of the scheme.*

The tri-adic split is a simple operation: In a first step every original triangle is split into three triangles by inserting a new vertex at its barycenter. In the second step the original edges are flipped, yielding a 30 degree rotated triangle mesh. Applying the $\sqrt{3}$-subdivision scheme twice leads to a uniform refinement with tri-section of every original edge while two dyadic splits would quad-sect every original edge (hence the name $\sqrt{3}$-subdivision).

We developed a new general technique for the convergence analysis that also applies to the analysis of other subdivision schemes. For regular meshes, a new matrix formulation is used that is much easier to handle than the standard generating functions notation. This approach requires only a few matrix computations that can easily be performed with the help of some standard software tools.

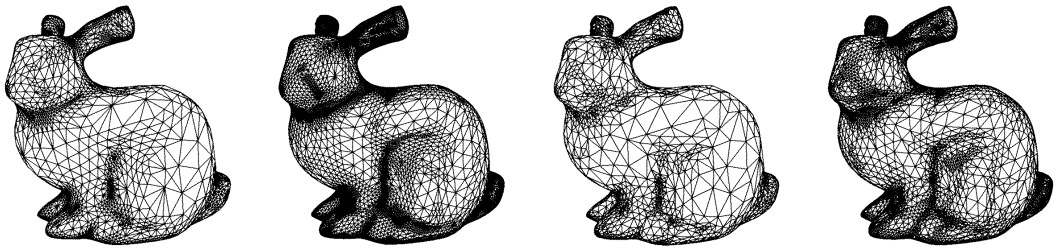

Figure 5.11: *Adaptive refinement based on red-green triangulation with Loop subdivision (left, center left) and based on the $\sqrt{3}$-refinement (center right, right). While the same stopping criterion is used, the Loop meshes have 10072 and 28654 triangles while the $\sqrt{3}$-meshes only have 7174 and 20772 triangles.*

The complexity of the refined meshes increases exponentially in the number of triangles for any subdivision scheme. This is why adaptive refinement strategies are applied (Figure 5.11). They insert new vertices only in regions where geometric detail is expected, flat regions of the surface are sufficiently well approximated by large triangles. $\sqrt{3}$-subdivision seems better suited for adaptive refinement than than other schemes for several reasons: First, the slower refinement reduces the average over-tessellation in case the stopping criterion is slightly failed for a coarse triangle and the result of the result of the refinement falls significantly below the threshold. Second, the localization is better than for dyadic refinement and no temporary triangle fans are necessary to keep the mesh

consistent. In fact, the consistency preserving adaptive refinement can be implemented in a simple recursive procedure.

References

- [1] L. Kobbelt. $\sqrt{3}$-subdivision. In K. Akeley, editor, *Computer Graphics (SIGGRAPH-00): Conference proceedings*, volume 2000 of *Computer Graphics Proceedings, Annual Conference Series*, pages 103–112, New Orleans, USA, 2000. ACM SIGGRAPH, ACM Press.
- [2] L. Kobbelt and H.-P. Seidel. sqrt(3)-subdivision and forward adaptive refinement. In *Proceedings of the Israelian/Korean workshop on Geometric Modeling and Computer Graphics*, pages 245–252, 1999.

5.3.2 Efficient Evaluation of Subdivision Schemes

Current implementations treat subdivision surfaces usually as high level representations of surface geometry. As a consequence, before rendering the surface is converted to an ordinary triangle mesh by refining the given control mesh sufficiently. A natural question that arises is whether subdivision surfaces can be used as a new rendering primitive, i.e., as a basic piece of geometry that is passed to the graphics sub-system directly. This way, only the control mesh would be transfered to the graphics hardware, and the evaluation procedure is moved beyond the graphics API and ideally should be implemented in hardware.

The naive approach of first generating a refined mesh and passing its triangles to the graphics pipeline is not appropriate due to the memory requirements that increase exponentially with the refinement depth. There has been considerable investigation in this topic, most recently [3, 2]. In all proposed solutions the storage requirements still depend on the subdivision depth to which the original control mesh has to be refined.

We developed a novel algorithm that is independent of the subdivision depth and that requires a small and constant amount of memory [1]. We focused on the Loop subdivision scheme as it is currently the most commonly used scheme in computer graphics since it works on arbitrary triangle meshes and generates curvature continuous surfaces in the limit.

Our algorithm exploits the knowledge about the piecewise polynomial structure of Loop subdivision surfaces. It can be modified to work for the Doo/Sabin and the Catmull/Clark schemes as well, but it cannot be generalized to non-polynomial schemes. A forward difference technique is applied for efficiently computing uniform samples on the limit surface.

The simple structure of our algorithm enables a scalable degree of hardware implementation: The central procedure uses only (and exactly) 23 register variables. An extension for computing the exact surface normals would increase the constant memory requirements to 29 variables. The algorithm is fast, we only need four operations for each triangle. Additional speed-up can be obtained by executing these operations in parallel or by caching intermediate results. Both, the algorithmic structure and the programming interface are kept simple.

References

- [1] S. Bischoff, L. Kobbelt, and H.-P. Seidel. Towards hardware implementation of Loop subdivision. In S. N. Spencer, editor, *Proceedings of 2000 SIGGRAPH/Eurographics Workshop on Graphics Hardware*, pages 41–50, Interlaken, Switzerland, August 2000. ACM SIGGRAPH, ACM Press.
- [2] K. Müller and S. Havemann. Subdivision surface tesselation on the fly using a versatile mesh data structure. In *Computer Graphics Forum, Proc. Eurographics 2000*, 2000.
- [3] K. Pulli and M. Segal. Fast rendering of subdivision surfaces. In *Eurographics Rendering Workshop 1996*, pages 61–70, 1996.

5.3.3 Progressive Transmission of Subdivision Surfaces

While current graphics hardware can process very large numbers of triangles, this processing is restricted to rendering the surface geometry on a local computer. But there is a number of applications that require highly detailed geometric models to be broadcast to a remote client over a low-bandwidth network connection like the Internet. Examples are virtual reality or applications for e-commerce.

With the naive sequential transmission of geometric data the user has to wait until the transmission is completely finished before a reasonable surface can be displayed. Thus, progressive transmission is mandatory. The standard solution for this is using progressive meshes [2], where an initial, coarse shape is transmitted first and can be displayed immediately. With more detail information being received, the mesh can be progressively refined until the complete model is recovered.

An alternative method for progressive transmission is to restrict to special triangle meshes with so-called subdivision connectivity. Such surface representations can be obtained from remeshing (cf. Sec. 4.2) arbitrary meshes. By exploiting this structure wavelet-like methods can be applied which allows efficient encoding of the geometric detail. We use this type of surface representation for efficient and robust transmission [4] of geometric models. Subdivision connectivity or semi-regular triangle meshes have been subject of more research, most recently by [3, 1, 5].

We developed a multi-media communication protocol that enables the client to display a smooth approximation of high visual quality of the surface even if not all detail information has been received yet. This is accomplished by stationary subdivision on a coarser level of detail that is already present at the client side. Different time steps of a transmission are illustrated by Figure 5.12. The protocol must be robust in a way that it tolerates permutation of the input data. This is due to the fact that the underlying network protocol for real-time transmission (e.g. UDP) cannot guarantee the correct ordering of data packets. This problem is solved by introducing a special indexing scheme on the semi-regular mesh structure with one index sent per detail coefficient.

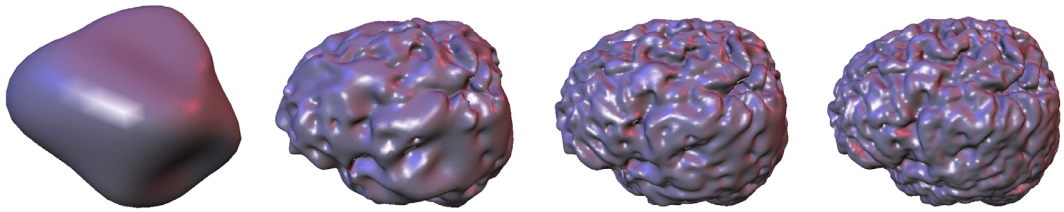

Figure 5.12: *Progressive transmission of a complex data set (90K△). From left:* 0%, 1%, 3% *and* 10% *of the detail coefficients are included.*

References

[1] I. Guskov, K. Vidimce, W. Sweldens, and P. Schröder. Normal meshes. In *SIGGRAPH 2000 Conference Proceedings*, pages 95–102, 2000.

[2] H. Hoppe. Progressive meshes. In *SIGGRAPH 96 Conference Proceedings*, pages 99–108, 1996.

[3] A. Khodakovsky, P. Schröder, and W. Sweldens. Progressive geometry compression. In *SIGGRAPH 00 Conference Proceedings*, pages 271–278, 2000.

[4] U. Labsik, L. Kobbelt, R. Schneider, and H.-P. Seidel. Progressive transmission of subdivision surfaces. *Computational Geometry: Theory and Applications*, 15(1-3):25–39, 2000.

[5] A. Lee, H. Moreton, and H. Hoppe. Displaced subdivision surfaces. In *SIGGRAPH 00 Conference Proceedings*, pages 85–94, 2000.

5.3.4 Hermite Subdivision Schemes

Approximation and interpolation of curves and surfaces from point sets is a challenging topic of research. In practical applications the given points are often assigned derivatives leading to Hermite elements. Approximation of Hermite-data by standard B-Spline curves and surfaces leads to highly non-linear problems. Their linearization produces new problems such as dependence on parameterization or the need for iterative methods. We focused our research on univariate Hermite-subdivision schemes. A stationary C^2 Hermite-subdivision scheme for second order Hermite-data was developed as well as the theory for the analysis of such schemes [1].

References

- [1] U. Schwanecke and B. Jüttler. A B-spline approach to Hermite subdivision. In A. Cohen, C. Rabut, and L. L. Schumaker, editors, *Curve and surface fitting, Saint-Malo 1999*, Innovations in Applied Mathematics, pages 385–392, Saint-Malo, June 2000. Vanderbilt University Press.

5.4 Medial Axis Transform and Shape Analysis

Investigator: Sung Woo Choi

The medial axis (transform), sometimes called the skeleton, is a thinned version of a given planar shape, which extracts the overall feature of the shape. Specifically, the *medial axis*, $\mathbf{MA}(\Omega)$, of a shape Ω is defined by the set of all the centers of the maximal circles contained in Ω, and the *medial axis transform*, $\mathbf{MAT}(\Omega)$, is defined by the set of all the pairs of the medial axis points and the radii of the corresponding maximal circles. Under some general condition on the shape, called the *normality* condition, the medial axis transform has a finite graph-like structure, and is homotopically equivalent to the given shape [1]. Due to these advantages, the medial axis transform has been one of the most widely used tools for shape analysis.

But the major difficulty in using the medial axis transform is its notorious instability: even when the boundary of the shape changes slightly the resulting change in the medial axis transform can be very large. To overcome this, the *hyperbolic Hausdorff distance* was introduced in [3], which has its motivation from the Minkowski metric on the hyperbolic space. It was shown that the transform $\mathbf{MAT}$: {planar shapes} $\rightarrow$ {medial axis transforms} is almost an isometry, where the target space is endowed with the hyperbolic Hausdorff distance.

A normal domain Ω is called *injective*, if the boundary $\partial\Omega$ has no sharp corners, and the medial axis $\mathbf{MA}(\Omega)$ has no end points which is a center of the maximal circle in Ω osculating the boundary. When the sharp corners are allowed, Ω is called *weakly injective*. Suppose we perturb a normal domain Ω to another normal domain Ω'. When Ω is injective, it is shown in [2] that the one-sided Hausdorff distance of $\mathbf{MA}(\Omega)$ (*resp.*, $\mathbf{MAT}(\Omega)$) with respect to the perturbed $\mathbf{MA}(\Omega')$ (*resp.*, $\mathbf{MAT}(\Omega')$) is bounded *linearly* with the Hausdorff distance between Ω and Ω'. This result was extended to the case when Ω is weakly injective in [4] and [6]. Moreover it was shown that the weakly injective domains are the largest class for which the linear bound is possible. The case when Ω is a general normal domain was investigated in [5]. In this case, it was shown that in the worst case the above one-sided Hausdorff distance for the medial axis (transform)s is bounded by the 1/3-th power of the Hausdorff distance between Ω and Ω'.

Additionally, in [7] and [8] a new efficient method for the detection of the artificial scene changes in MPEG video was introduced, which is fast since it extracts the scene change information directly from the MPEG encoded data by using an approximation of the discrete cosine transform.

References

[1] H. I. Choi, S. W. Choi, and H. P. Moon. Mathematical Theory of Medial Axis Trasform. *Pacific J. Math.*, 181(1):57–88, Nov. 1997.

- [2] S. W. Choi and S.-W. Lee. Stability analysis of medial axis transform under relative Hausdorff distance. In A. Sanfeliu, J. J. Villanueva, M. Vanrell, R. Alquzar, T. Huang, and J. Serra, editors, *Proceedings of the 15th International Conference on Pattern Recognition (ICPR-00)*, pages 139–142, Barcelona, Spain, September 2000. IEEE Computer Society.
- [3] S. W. Choi and H.-P. Seidel. Hyperbolic hausdorff distance for medial axis transform. *Graphical Models*, 2001. To appear.
- [4] S. W. Choi and H.-P. Seidel. Linear one-sided stability of mat for weakly injective domain. Research Report MPI-I-2001-4-004, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, June 2001.
- [5] S. W. Choi and H.-P. Seidel. One-sided stability of MAT and its applications. In T. Ertl, B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Vision, Modeling, and Visualization 2001*, Stuttgart, Germany, November 2001. To appear.
- [6] S. W. Choi and H.-P. Seidel. One-sided stability of medial axis transform. In B. Radig and S. Florczyk, editors, *Proceedings of Pattern Recognition, 23rd DAGM Symposium*, volume 2191 of *Lecture Notes in Computer Science*, pages 132–139, München, Germany, 2001. Springer.
- [7] Y.-M. Kim, S. W. Choi, and S.-W. Lee. Fast scene change detection using direct feature extraction from MPEG compressed videos. In A. Sanfeliu, J. J. Villanueva, M. Vanrell, R. Alquzar, T. Huang, and J. Serra, editors, *Proceedings of the 15th International Conference on Pattern Recognition (ICPR-00)*, number 3, pages 178–181, Barcelona, Spain, September 2000. IEEE Computer Society.
- [8] S.-W. Lee, Y.-M. Kim, and S. W. Choi. Fast scene change detection using direct feature extraction from MPEG compressed videos. *IEEE Transactions on Multimedia*, 2(4):240–254, December 2000.

6 Facial Modeling and Animation

We have developed a system for photo-realistic facial modeling and animation, which includes several tools that facilitate necessary tasks such as mesh processing, texture registration, and assembling of facial components (see also Figure 6.13). The resulting head model reflects the anatomical structure of the human head including skull, skin, and muscles. Semi-automatic generation of high-quality models from scan data for physics-based animation becomes possible with little effort. A state-of-the-art speech synchronization technique is integrated into our system, resulting in realistic that can be rendered at real-time frame rates on current PC hardware.

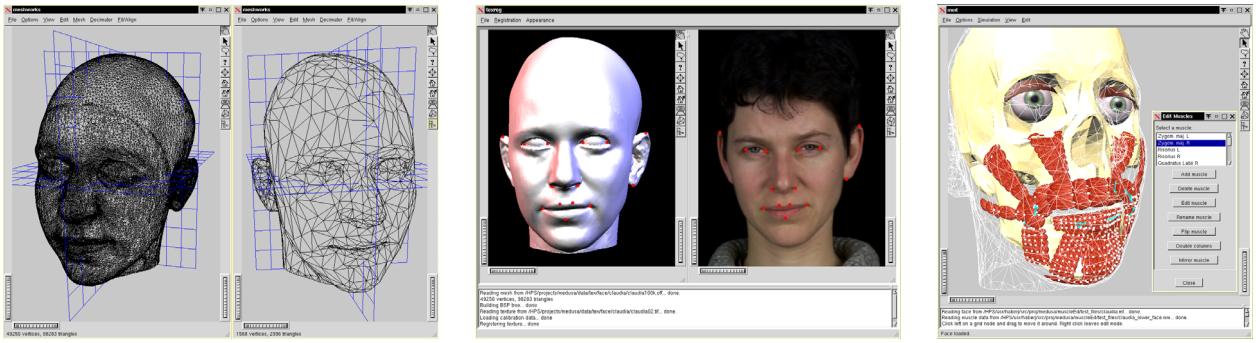

a) b) c)

Figure 6.13: *Tools of our facial modeling and animation system. a) Two views of our mesh decimation and alignment tool. b) GUI of our texture registration tool. c) GUI of our muscle editor.*

We aim at the construction of realistic animatable head models from real human individuals. Geometry and images are acquired in high resolution and converted to textured polygonal geometry, simplified according to the requirements of the targeted hardware. On the lowest level, we use a physics-based approach with muscle contraction values as animation parameters. The construction of the virtual head model is based on the human anatomy, and we have developed tools to automate the task of adapting and linking this model to the actual face geometry.

6.1 Geometry-Based Muscle Modeling

Investigators: Kolja Kähler, Jörg Haber

Recent developments in the accurate simulation of human faces make use of anatomical knowledge, which is incorporated into the model [3]. In the general context of virtual characters, the so-called *anatomically-based modeling* aims for construction of animatable models from bones, muscles, and skin [4]. For human faces, however, this strategy is unsuitable if the target geometry is already given. Since the muscles of the face lie closely underneath the skin and have a great influence on the shape and appearance of the surface, it is difficult to model the skin from skull and muscles in such a way that the result bears close resemblance with the target face. On the other hand, surface geometry can easily be acquired using for instance a range scanner. Thus, our approach is to adapt the muscle geometry to the prescribed facial geometry.

To facilitate this, we have devised and implemented a muscle model and methods for muscle construction that allow to easily create animatable facial models from given face geometry [2]. Using our editing tool (see Figures 6.13 c and 6.14), one can interactively specify coarse outlines of the muscles, which are then automatically created to fit the face geometry. Muscles for a new geometry can easily be created in a fully automatical fashion from stored outlines.

Our muscle model incorporates different types of muscles such as linear muscles, sheet muscles, and sphincters (see Figure 6.15) as well as the effects of bulging and intertwining muscle fibers. The influence of muscle contraction onto the skin is simulated using a mass-spring system that connects the skull, muscle, and skin layers of our model. We are currently working on the development of an even more realistic muscle model for the *orbicularis oris* that is capable of showing all the subtle details needed for proper speech synchronization [1].

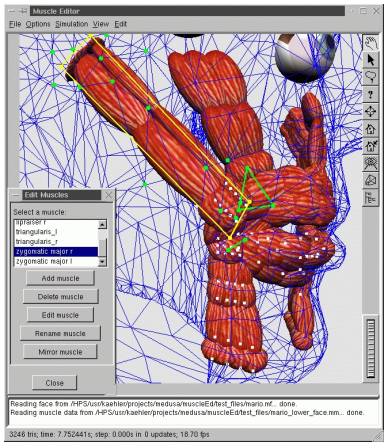

Figure 6.14: *Visual information while editing a muscle: the muscle grid of the currently edited muscle (yellow); the skin vertices influenced by this muscle (green dots); muscle control points attached to the jaw (white dots); merged muscle segments (connected by green lines).*

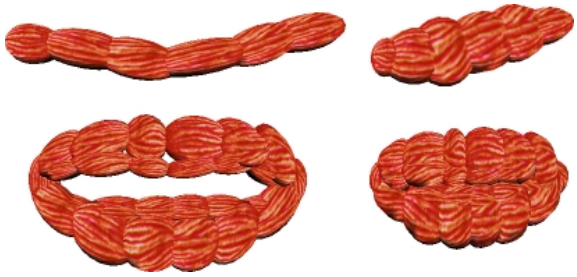

Figure 6.15: *Relaxed (left) and contracted (right) muscles: a single fiber (top) and a sphincter (bottom) modeling the* orbicularis oris.

References

- [1] J. Haber, K. Kähler, I. Albrecht, H. Yamauchi, and H.-P. Seidel. Face to Face: From Real Humans to Realistic Facial Animation, 2001. submitted for publication.
- [2] K. Kähler, J. Haber, and H.-P. Seidel. Geometry-based Muscle Modeling for Facial Animation. In B. Watson and J. W. Buchanan, editors, *Proceedings Graphics Interface 2001*, pages 37–46, Ottawa, Ontario, June 2001. Canadian Human-Computer Communications Society, Morgan Kaufmann.

[3] Y. Lee, D. Terzopoulos, and K. Waters. Realistic Modeling for Facial Animations. In *Computer Graphics (SIGGRAPH 95 Proceedings)*, pages 55–62, Aug. 1995.

[4] J. Wilhelms and A. Van Gelder. Anatomically Based Modeling. In *Computer Graphics (SIGGRAPH 97 Proceedings)*, pages 173–180, Aug. 1997.

6.2 Texturing

Investigators: Hitoshi Yamauchi, Jörg Haber

To generate a texture for a head model, we take several photographs of the person's head using different, uncalibrated camera positions [1]. The number of photographs we use varies from four to eight, depending on the desired quality of the combined texture and the time available for texture registration. All photographs are taken with a high resolution digital camera under diffuse

illumination. During the photo session, the facial expression of the person should resemble the neutral expression during the scanning process.

The photographs are registered and combined into a single texture suitable for OpenGL rendering similarly to the approach in [2], which is based on the camera calibration technique developed by Tsai [3]. In contrast to the method proposed in [2], we do not rely on the property that every vertex of a face mesh is bound to an input texture. Especially in the regions inside and behind the ears, some vertices typically remain unbound. For each unbound vertex v in the mesh, we perform a region growing over the topological neighborhood of v and interpolate the texture coordinates of those vertices within that region R_v, which are bound to the same and most frequently used texture. If the texture binding within R_v is evenly distributed among several textures, we examine remaining *valid* texture bindings of the vertices in R_v (cf. [2, Sect. 3] for a definition of *valid*) and interpolate the texture coordinates within the prevailing texture. If there is still no preferred texture among all texture bindings, we randomly choose one of the most frequently used textures. The interpolated texture coordinates of vertex v need to be verified to lie within $[0, 1]^2$. If this is not the case, another frequently used texture is chosen for interpolation. Figure 6.13 b) shows a snapshot of the GUI of our texture registration tool, where corresponding points have been interactively selected on both the 3D geometry and the 2D input photograph.

References

- [1] J. Haber, K. Kähler, I. Albrecht, H. Yamauchi, and H.-P. Seidel. Face to Face: From Real Humans to Realistic Facial Animation, 2001. submitted for publication.

[2] C. Rocchini, P. Cignoni, C. Montani, and R. Scopigno. Multiple Textures Stitching and Blending on 3D Objects. In D. Lischinski and G. Ward Larson, editors, *Rendering Techniques '99 (Proc. 10th Eurographics Workshop on Rendering)*, pages 119–130, 1999.

[3] R. Y. Tsai. An Efficient and Accurate Camera Calibration Technique for 3D Machine Vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 364–374, June 1986.

6.3 Head Modeling Using Subdivision Surfaces

Investigators: Won-Ki Jeong, Jörg Haber, Kolja Kähler

To generate individual human head models, we acquire geometric information of a human face using a 3D range scanner. Resulting from the range scans we obtain a point cloud, which is used to fit a generic head model by employing an energy optimization technique. The generic head model is represented by a subdivision surface that can easily be parameterized for texturing. Moreover, we can overcome imperfection of the range scan data like, for instance, noise or holes in the data by fitting our generic head model. Figure 6.16 shows different stages of the process.

We have developed a displaced subdivision surface reconstruction algorithm for known topology [1] and we are currently integrating it into our facial modeling system. The displaced subdivision surface [3] is a new mesh representation that defines a detailed mesh with a displacement map over a smooth domain surface, and it is appropriate for our face model structure since it is a compact mesh representation and has regular structure for compression and multiresolution analysis. The overall process of our modeling system is as follows. First, we deform a generic head model to minimize a given total energy functional consisting of four different energy terms: distance functional, constraint functional, smoothness functional, and sampling functional. In this way, we minimize not only the distance from points to the mesh but also enhance the surface quality due

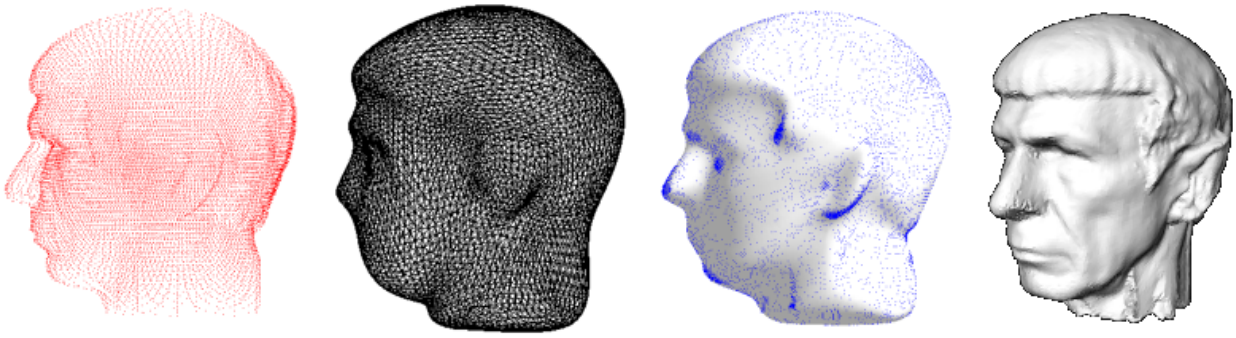

Figure 6.16: *Different stages of our head modeling process. Left to right: point cloud resulting from 3D range scanning; subdivision surface fit to the point cloud; displacement vectors of the displaced subdivision surface shown in blue; final result.*

to the sampling energy term that is defined by the distance between each sampling ray to the nearest input point. Next, we subdivide our generic model until we get a proper resolution using the Loop subdivision scheme. Finally, we sample surface detail in the direction of each vertex normal. That can be done by heuristic determination [1], which is much faster than local triangulation of points. The output of our algorithm is a 3D approximation of the input point cloud represented as a displaced subdivision surface, i.e. its structure is piecewise-regular and can it be represented by a simple coarse control mesh with scalar displacement values.

To efficiently render the head model, we have developed a method to adaptively refine an irregular triangle mesh as it deforms in real-time [2]. The method increases surface smoothness in regions of high deformation by splitting triangles in a fashion similar to one or two steps of Loop subdivision. The refinement is computed for an arbitrary triangle mesh and the subdivided triangles are simply passed to the rendering engine, leaving the mesh itself unchanged. The algorithm can thus be easily plugged into existing systems to enhance visual appearance of animated meshes. The refinement step has very low computational overhead and is easy to implement.

References

- [1] W.-K. Jeong and C.-H. Kim. Direct reconstruction of displaced subdivision surface from unorganized points. In *Proceedings of the 9th Conference on Pacific Conference on Computer Graphics and Applications(PG'01)*, Tokyo, Japan, 2001. IEEE Computer Society.
- [2] K. Kähler, J. Haber, and H.-P. Seidel. Dynamic Refinement of Deformable Triangle Meshes for Rendering. In H. H. S. Ip, N. Magnenat-Thalmann, R. W. Lau, and T.-S. Chua, editors, *Proceedings Computer Graphics International 2001*, pages 285–290, Hong Kong, 2001. IEEE Computer Society.

[3] A. Lee, H. Moreton, and H. Hoppe. Displaced Subdivision Surfaces. In *Computer Graphics (SIGGRAPH 2000 Proceedings)*, pages 85–94, Aug 2000.

7 Model Acquisition with Realistic Reflection Properties

The use of realistic models for all components of images synthesis is a fundamental prerequisite for photorealistic rendering. The generation of these models in a manual process often becomes infeasible as the demand for visual complexity increases steadily. We have developed a system for the acquisition of high-quality 3D models of real world objects representing the object's geometry as well as its reflection properties. The generated model allows for rendering under arbitrary viewing and lighting conditions and realistically reproduces the appearance of the original object (see Figure 7.17).

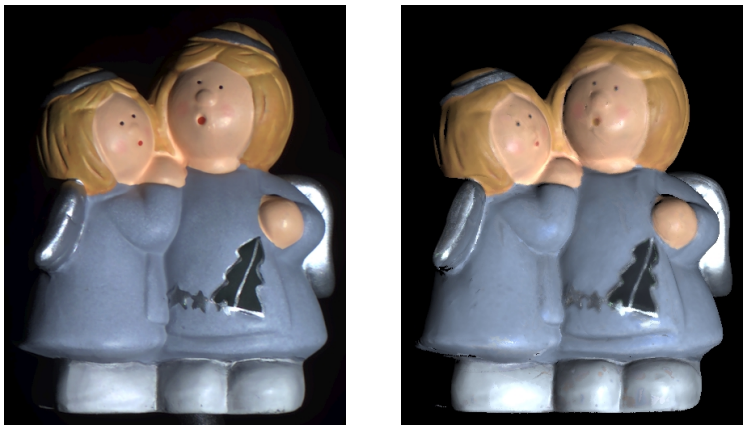

Figure 7.17: *Left: Photograph of real object. Right: Model with acquired BRDF rendered from the same view with similar lighting direction.*

To obtain the complete 3D model a number of different steps are necessary. We acquire the geometry of the object with a 3D scanner. In order to capture the reflection properties (BRDF) we take several high dynamic range images of the object, lit by a point-light source in a controlled environment. For all images we recover the camera position and orientation relative to the geometric model as well as the light source position. Combining all the acquired data it is possible to determine the spatially varying BRDFs of the object, i.e., a specific BRDF for each surface point.

7.1 Digital Cameras as Measuring Devices

Investigators: Michael Goesele, Hendrik P.A. Lensch, Wolfgang Heidrich

In contrast to most traditional devices for measuring reflection properties such as gonioreflectometers digital cameras are highly parallel. The sensor captures the irradiance on each sensor element on the image plane simultaneously. Given a suitable environment such a camera can be used to measure the reflection properties of real objects. We built a digital photo studio for measurement purposes tailored to our specific needs in acquiring realistic input data [4]. It consists of a digital camera, almost point shaped light sources, a dark room with low reflection, and resourceful computing equipment.

As light sources we use HMI lamps from K5600. Their photometric properties (near daylight spectrum) are well suited for measurement purposes. The light is emitted from a small spark which serves as a very good approximation to a point light source. Additional reflectors allow for more diffuse illumination conditions.

The environment has a large influence on all measurements and unaccounted indirect light can

lead to large errors. We try to minimize this influence by using dark, diffuse reflecting material (black felt, matt paint) to cover the inside of the photo studio and the equipment.

The most important requirements for the camera were hereby high flexibility, faithful measurement results, and the ability to control as much settings as possible remotely from a computer. We use a Kodak DCS 560 Professional Digital Camera for our experiments which combines the flexibility of an analog SLR camera with high resolution images (6 million pixels per image). It can be controlled remotely via the IEEE 1394 FireWire interface and custom control software can be written using a SDK provided by the manufacturer.

Due to the limited dynamic range of the camera it is not possible to capture the full dynamic range of most scenes with a single image. However, by combining several images with different exposure times into a single image it is possible to capture scenes with almost arbitrary dynamic range [1, 5]. We extended the approach of Robertson et al. [5] to include color calibration into the high dynamic range generation [2] (see Section 7.5). In addition, images with long exposure times (several seconds) contain a high amount of noise. As these exposure times are common in our applications, we developed a new noise removal technique [3] (see Section 7.4).

References

[1] P. Debevec and J. Malik. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proceedings of SIGGRAPH 97*, pages 369–378, August 1997.

- [2] M. Goesele, W. Heidrich, and H.-P. Seidel. Color calibrated high dynamic range imaging with ICC profiles. In *Proceedings of the 9th Color Imaging Conference*, Scottsdale, USA, 2001.

- [3] M. Goesele, W. Heidrich, and H.-P. Seidel. Entropy-based dark frame subtraction. In *Proceedings of PICS 2001: Image Processing, Image Quality, Image Capture, Systems Conference*, pages 293–298, Montreal, Canada, April 2001. The Society for Imaging Science and Technology (IS&T).

- [4] M. Goesele, H. Lensch, W. Heidrich, and H.-P. Seidel. Building a photo studio for measurement purposes. In B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling, and Visualization (VMV-00)*, pages 231–238, Saarbrücken, Germany, November 2000.

[5] M. A. Robertson, S. Borman, and R. L. Stevenson. Dynamic Range Improvement Through Multiple Exposures. In *Proceedings of the 1999 International Conference on Image Processing (ICIP-99)*, pages 159–163. IEEE, oct. 1999.

7.2 Silhouette-Based Registration

Investigators: Hendrik P.A. Lensch, Wolfgang Heidrich

Given high dynamic range images of an object with known lighting conditions, the next step towards a complete 3D model is to acquire its 3D geometry. A 3D scanner, e.g. a structured light or computer tomography scanner, is used to sample the object's surface which results in a triangular mesh approximating the original object.

To combine both types of data, the 3D geometry and the 2D images, it is necessary to register each image with the 3D mesh. Knowing the exact camera parameters for one image, the position and orientation of the camera with respect to the object as well as the intrinsic parameters like the focal length, allows to map the image as a texture onto the 3D geometry as can be seen in Figure 7.18.

Several techniques can be applied to obtain the camera parameters. A set of corresponding points in the image and on the 3D model can be selected manually from which the parameters can be directly calculated using camera calibration methods like [4]. Other approaches like [3, 2]

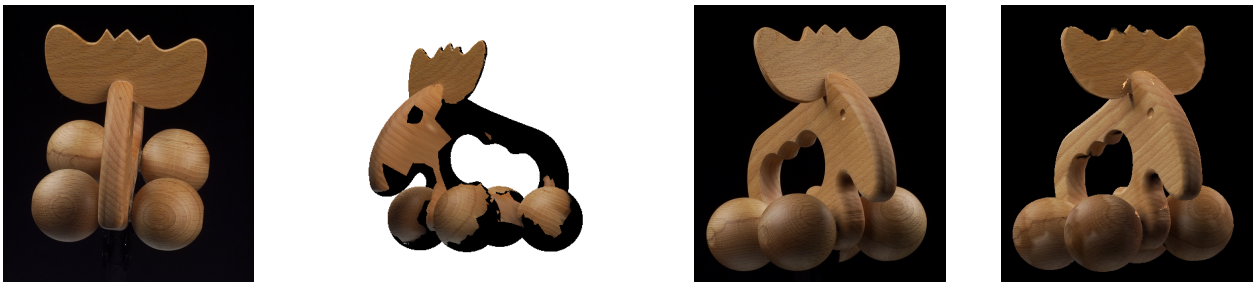

a) b) c) d)

Figure 7.18: *a) Photograph b) Photograph mapped onto 3D geometry model. c) and d) Comparison of another photo with the 3D model textured with 15 images, rendered from the same view as photograph.*

compare the outline of a 2D image to the outline of the projected 3D geometry using the current set of camera parameters. By minimizing the distance between points on one outline and the nearest points on the second outline the optimal set of camera parameters can be obtained. Unfortunately, computing the outline of one view of the object is quite time consuming, since the object has to be rendered and the outline has to be extracted.

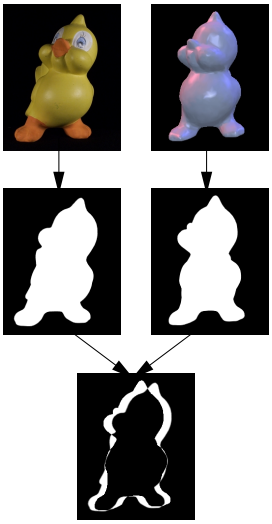

Figure 7.19: *Measuring the difference between the photo (left) and one view of the model (right) by the area occupied by the XOR-ed foreground pixels.*

Our approach [1] uses silhouettes instead of outlines to make the comparison easier and faster. To further speed up the process we make use of current graphics hardware features.

In a preprocessing step the silhouettes of the input images are generated by segmentation into foreground and background pixels. To obtain the silhouette of one view of the object the object is simply rendered in monochrome color. Both silhouettes are then combined using the logical XOR-operation (see Figure 7.19). The effect of this operation is that only those pixels remain that are covered by exactly one of the silhouettes, that is where the silhouettes differ. The number of remaining pixels directly yields a measure for the distance between both silhouettes.

Note that all operations, rendering the 3D geometry, performing the XOR-operation on the silhouettes and counting the number of remaining pixels by calculating a histogram, can easily be carried out in hardware allowing for efficient evaluation of the similarity measure. Based on this similarity measure we can optimize for the camera parameters for each image yielding the individual camera transformations.

References

- [1] H. Lensch, W. Heidrich, and H.-P. Seidel. Automated texture registration and stitching for real world models. In B. A. Barsky, Y. Shinagawa, and W. Wang, editors, *Proceedings of the 8th Pacific Conference*

on Computer Graphics and Applications (PG-00), pages 317–326, Hong Kong, China, October 2000. IEEE Computer Society.

[2] K. Matsushita and T. Kaneko. Efficient and Handy Texture Mapping on 3D Surfaces. *Computer Graphics Forum*, 18(3):349–358, September 1999.

[3] P. Neugebauer and K. Klein. Texturing 3D Models of Real World Objects from Multiple Unregistered Photographic Views. *Computer Graphics Forum*, 18(3):245–256, September 1999.

[4] R. Tsai. A versatile camera calibration technique for high accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses. *IEEE Journal of Robotics nd Automation*, 3(4), August 1987.

7.3 Measuring Spatially Varying Reflection Properties

Investigators: Hendrik P.A. Lensch, Jan Kautz, Michael Goesele, Wolfgang Heidrich, Hartmut Schirmacher

Although quite realistic results can be obtained by just mapping images of the diffusely lit object onto the 3D geometry artifacts will always be visible since highlights do not move when viewed from different positions, nor can the model be adapted to varied lighting conditions. These effects can only be reproduced realistically if the reflection properties of the object's surface have been reconstructed.

The measuring of real-world materials has recently received a lot of attention in the computer graphics community. For example, Marschner et al. [5] measured the BRDF (bi-directional reflectance distribution function) of a homogeneous material by taking a small number of pictures of a curved sample with varying light source positions. Since all pixels show the same material with different normals, i.e. different angles with respect to the viewing and lighting direction, a dense sampling of the BRDF is obtained when combining the data of all pixels. In our first approach [6] we also measured homogeneous samples and fit a BRDF model to the acquired samples. Debevec et al. [1] measure a different BRDF for each point on the surface visible from one view by densely sampling the incident light direction which unfortunately requires around 200 images and a special setup. Our method [2, 3] is similar to the method of Marschner et al. in combining the data of different points on the surface showing the same material, but we are also dealing with spatially varying BRDFs as Debevec does.

In order to capture the reflection properties we take a relatively small number (around 20) of high dynamic range images of the object lit by a point-light source. Both, the position of the camera and of the light source are changed for each image. We recover the camera position and orientation as well as the light source position relative to the geometric model for all images.

For every point on the object's surface we collect all available data from the different views in a data structure called *lumitexel*. It contains the position of the surface point, its normal, and a list of radiance samples together with their viewing and lighting directions.

Since a single lumitexel does not carry enough information to reliably fit a BRDF model to the radiance samples, we first determine clusters of lumitexels belonging to similar materials. Starting with a single cluster containing all lumitexels, the parameters of an average BRDF are fitted using the Levenberg-Marquardt algorithm. From this, two new sets of parameters are generated by varying the fitted parameters along the direction of maximum variance, yielding two slightly separated BRDFs. The lumitexels of the original cluster are then assigned to the nearest of these BRDFs, forming two new clusters. A stable separation of the materials in the clusters is obtained by repeatedly fitting BRDFs to the two clusters and redistributing the original lumitexels. Further splitting isolates the different materials until the number of clusters matches the number of materials of the object as illustrated in Figure 7.20.

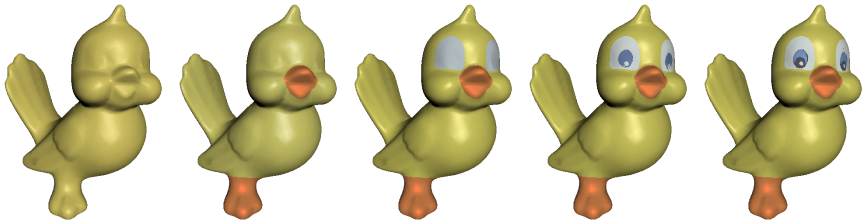

Figure 7.20: *The clustering process at work. In every image a new cluster was created.*

After the clustering we still have the same reflection behavior assigned to all lumitexels in one cluster. However, small features on the surface and smooth transition between materials can only be represented if every lumitexel is assigned its own BRDF.

In our algorithm, this BRDF is a linear combination of the BRDFs recovered by the clustering procedure. This can be represented by a set of basis BRDFs for the entire model plus a set of weighting coefficients for each lumitexel. An optimal set of weighting coefficients minimizes the error between the measured radiance and the weighted radiance values obtained by evaluating the basis BRDFs for the specific viewing and lighting directions. To recover the coefficients we compute the least square solution of the corresponding system of equations using singular value decomposition.

This method allows for accurately shaded, photorealistic rendering of complex solid objects from new viewpoints under arbitrary lighting conditions with relatively small acquisition effort.

For a complete overview over the entire framework see [4], the handout for a tutorial we held at the Web3D 2001 conference.

Figure 7.21: *Three models rendered with spatially varying BRDFs, which were acquired with our reconstruction method.*

References

[1] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the Reflectance Field of a Human Face. In *Computer Graphics (SIGGRAPH 2000 Proceedings)*, pages 145–156, July

2000. ISBN 1-58113-208-5.

- [2] H. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based reconstruction of spatially varying materials. In S. Gortler and K. Myszkowski, editors, *Rendering Techniques: Proceedings of the 12th Eurographics Workshop on Rendering*, pages 104–115, London, Great Britain, 2001. Springer.

- [3] H. P. A. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based reconstruction of spatially varying materials. Research Report MPI-I-2001-4-001, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, March 2001.

- [4] H. P. A. Lensch, J. Kautz, M. Goesele, and H.-P. Seidel. A framework for the acquisition, processing, transmission, and interactive display of high quality 3d models on the web. Research Report MPI-I-2001-4-002, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, March 2001. Also published as Tutorial Notes for the Web3D Conference 2001, Paderborn, Germany.

[5] S. Marschner, S. Westin, E. Lafortune, K. Torrance, and D. Greenberg. Image-based BRDF Measurement Including Human Skin. In *10th Eurographics Workshop on Rendering*, pages 131–144, June 1999.

- [6] H. Schirmacher, W. Heidrich, M. Rubick, D. Schiron, and H.-P. Seidel. Image-based BRDF reconstruction. In B. Girod, H. Niemann, and H.-P. Seidel, editors, *Proceedings of the 4th Conference on Vision, Modeling, and Visualization (VMV-99)*, pages 285–292, Erlangen, Germany, November 1999. SFB 603 and Graduate Research Center, in coorperation with IEEE Signal Processing Society and Gesellschaft für Informatik GI, Infix.

7.4 Noise Removal

Investigators: Michael Goesele and Wolfgang Heidrich

Noise due to dark current is a serious limitation for taking images with long exposure times with modern digital cameras [1]. On the physical level charge is accumulated in a sensor element of a CCD chip although no light is hitting its surface leading to noisy images. As the dark current is mainly caused by defects on the CCD chip the spatial distribution of the resulting noise has a fixed pattern and is therefore predictable. However, its amplitude depends strongly on the temperature during the exposure.

We have developed a technique that effectively removes the dark current noise from even highly corrupted images [2]. A scaled dark frame is subtracted from a noisy image. The correct scaling factor is found by an optimization technique that minimizes the entropy of the resulting image. Figure 7.22 shows an example of the ceaning process.

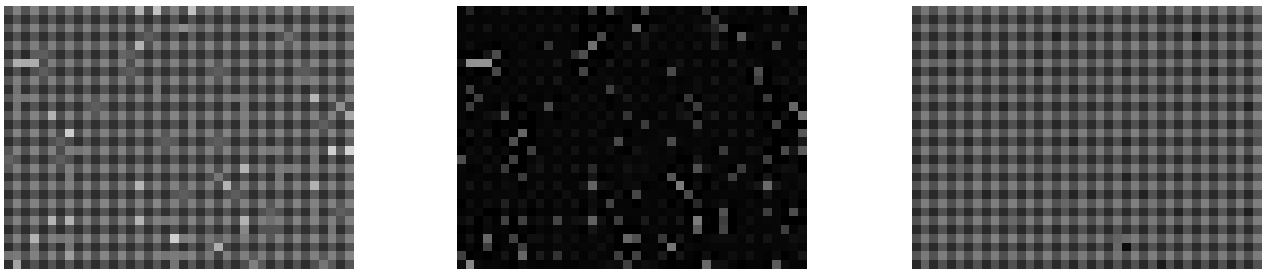

Figure 7.22: *Excerpts from a noisy test image taken with 25 s exposure time. Left: original version of the image. Middle: captured dark frame. Bottom: cleaned version for an optimal scaling factor. The checkerboard pattern is due to a color filter array on the CCD chip.*

In contrast to general noise removal techniques we are not only able to detect noisy pixels and interpolate their value from neighboring pixels but we can also reconstruct the correct pixel values as long as the pixel is not overexposed.

References

[1] R. L. Baer. CCD Requirements for Digital Photography. In *Proceedings of IS&T's PICS Conference*, pages 26–30, March 2000.

- [2] M. Goesele, W. Heidrich, and H.-P. Seidel. Entropy-based dark frame subtraction. In *Proceedings of PICS 2001: Image Processing, Image Quality, Image Capture, Systems Conference*, pages 293–298, Montreal, Canada, April 2001. The Society for Imaging Science and Technology (IS&T).

7.5 High Dynamic Range Imaging

Investigators: Michael Goesele and Wolfgang Heidrich

The dynamic range of standard analog or digital cameras is too low to capture most real scenes. Many details are lost in dark regions, bright regions or even both. Several groups including Debevec et al. [2] and Robertson et al. [4] have therefore developed methods to combine multiple images with different exposure times into a single image in order to capture the full dynamic range of a scene. They first recover the response curve of the imaging system and linearize the image data. In a second step a weighted average of the pixel values from the set of images is computed taking into account their respective exposure times.

Our approach [3] integrates high dynamic range imaging in the workflow of traditional, industry standard color management systems. An ICC profile is computed for the specific camera and lighting conditions [1, 5]. It can be used to convert the input images into the linear CIE XYZ color space where all high dynamic range calculations are performed. After additional processing steps the output images can be transformed into an arbitrary target color space using another ICC profile. We analyzed the linearity of the converted input images and showed that it is sufficient for this purpose. If all additional processing steps preserve the linearity of the system we are able to generate color calibrated output while working with high dynamic range images.

References

[1] Specification ICC.1:1998-09, File Format for Color Profiles. Available from http://www.color.org, 1998.

[2] P. Debevec and J. Malik. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proceedings of SIGGRAPH 97*, pages 369–378, August 1997.

- [3] M. Goesele, W. Heidrich, and H.-P. Seidel. Color calibrated high dynamic range imaging with ICC profiles. In *Proceedings of the 9th Color Imaging Conference*, Scottsdale, USA, 2001.

[4] M. A. Robertson, S. Borman, and R. L. Stevenson. Dynamic Range Improvement Through Multiple Exposures. In *Proceedings of the 1999 International Conference on Image Processing (ICIP-99)*, pages 159–163. IEEE, oct. 1999.

[5] D. Wallner. Building ICC Profiles - the Mechanics and Engineering, 2000.

8 Image-Based Rendering and Motion from Video

The general idea of image-based rendering (IBR) is to use images as input data (e.g. real-world photographs), and to generate synthetic images by recombination and modification of these input images, avoiding the very difficult step of first acquiring consistent models of the geometry, material, and lighting in the scene.

The IBR research activities in AG4 have focused strongly on image-based scene viewing, especially *light field rendering*. The following sections describe the two main branches of our IBR research during the past two years: light field acquisition, representation, and rendering (Section 8.1), and applications and extensions of light field techniques (Section 8.2). Finally, motion from video is discussed in Section 8.3.

8.1 Efficient Acquisition, Representation, and Rendering of Light Fields

Investigators: Hartmut Schirmacher, Wolfgang Heidrich, and Li Ming

Light fields have been introduced to the computer graphics community in 1996, and have since then inspired a large body of innovative research. A light field represents the 4D distribution of light leaving or entering a bounded region in 3D space. This approach is also called "ray database", since it relies on a sampled representation of all possible rays through this 3D region. A *Lumigraph* is a light field with additional geometric information. This information is used to reduce the number of rays needed for a high-quality reconstruction of arbitrary views of the light field. A discrete light field can also be imagined as a 2D collection of 2D images, with two dimensions representing virtual "eye points" corresponding to these input images, and the other two dimensions representing the pixel position within each image [5].

The research in our group covers nearly all aspects of light field and Lumigraph processing, starting from the acquisition of light fields from synthetic scenes (Section 8.1.1), representation of Lumigraphs as images with per-pixel depth information, and efficient rendering from such representations (Section 8.1.2), alternative and efficient parameterizations for real-word applications (Section 8.1.3) and finally a complete system that can capture, transmit, and render light fields from dynamic scenes on the fly (Section 8.1.4).

8.1.1 Adaptive Acquisition from Synthetic Scenes

In our previous work on Lumigraph refinement [3] we demonstrated that it is possible to refine coarse Lumigraphs and obtain a high rendering quality if the set of images contains sufficient information about all visible features of the scene. Since the acquisition of Lumigraph from a synthetic scene usually involves large amounts of computational effort (sometimes up to several hours for acquiring a single image with a ray tracer or similar high-end renderers), it is imperative to acquire only those images that are really needed for the final reconstruction.

This can be achieved using a modified 3D image warping as an error estimator for the reconstruction quality [7]. First, we acquire an initial, very coarse, Lumigraph (e.g. 2×2 eye points / images). We create a triangulation of the eye points and estimate the warping error for a number of reasonable candidate eye points, e.g. for the centers of the triangle edges. We choose the eye point with the largest error, acquire the corresponding image by ray tracing, and adapt the triangulation and error estimate accordingly. The algorithm stops after the desired number of images is acquired, or the error of the next candidate drops below a user-defined threshold.

The error estimate mainly takes into account two different factors: the difference of the colors that will be interpolated to obtain the final pixel color, and the number of source pixels that map

onto the same target pixel. This way it is relatively easy to detect holes, blending errors, and pixels for which no conservative color estimate can be derived (e.g. if there is only information from one source pixel).

Our technique provides a fully automatic adaptive acquisition of synthetic Lumigraphs that will produce the (in our sense) optimal Lumigraph for a given number of images. The method could also be extended to select the most appropriate images from a sequence, e.g. when creating Lumigraphs from video streams.

8.1.2 Interactive Lumigraph Rendering Through Warping

As we have shown in previous work [3], warping can efficiently produce intermediate Lumigraph images of high quality. But the Lumigraph refinement method results in a very large light field that consumes enormous amounts of memory and graphics hardware resources for interactive display.

Rather than warping in a preprocessing step, we introduced a technique that allows to employ warping for viewing Lumigraphs directly at interactive frame rates [8]. For every novel view, we first partition the output image into regions using a triangulation of the eye point plane. For each triangle, we determine from which reference images we want to extract and interpolate the color information.

Next, we have to determine the source region in each affected reference image that has to be reprojected in order to fill the corresponding target region (a triangle fan) in the output image. We compute conservative bounds on the source regions by taking into account the pixel flow into that target region. Then, we reproject only the necessary pixels from each source image to obtain the interpolated final image. This approach can be seen as an efficient solution to the per-triangle-fan *inverse warping* problem. This warped image is generated on the original Lumigraph image plane, and finally reprojected into the actual view by the use of hardware-supported texture mapping.

The technique allows interactive viewing of relatively sparse Lumigraphs in very high quality, achieving frame rates of around 5–7 frames/second on an sgi O2 machine. The performance of the algorithm is nearly independent of the number of reference images in the Lumigraph. Furthermore, the warping-based rendering can directly use per-pixel depth information, as provided by most existing depth sensors for real-world applications. Figure 8.23 demonstrates the reconstruction quality and performance for synthetic scenes with 6–8bit of artificially quantized depth information per pixel. The graph on the right of Figure 8.23 shows the weak, sublinear dependency between rendering complexity (number of pixels to be warped) and the number of input images, which is a very big advantage compared to other warping techniques.

8.1.3 Free-Form Light Fields

Since the introduction of light fields to computer graphics in 1996 the most commonly used parameterization has been the so-called *two-plane* parameterization, mainly because its implementation is easy and very efficient. Some authors have proposed alternatives, such as spherical, or sphere-plane and sphere-sphere parameterizations [2]. However, for acquiring light fields from the real world (e.g. using camera arrays) it is very hard to position and calibrate the cameras in such way that they match the desired parameterization directly. This means that the data from the input views usually has to be resampled (or *rebinned*) before it can be used for efficient light field rendering.

We have shown a simple and efficient way for rendering arbitrary views from so-called *free-form light fields* [10], employing a convex free form camera surface (formed by the optical centers of the input cameras) and a set of arbitrarily oriented camera planes (the different image planes of the

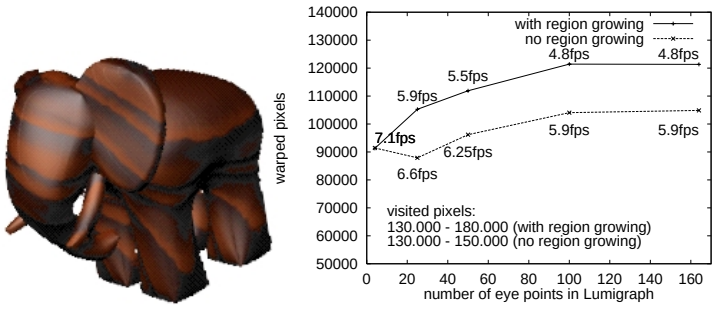

Figure 8.23: Lumigraph view reconstructed using our warping technique at approx. 5–7 frames/second on an sgi O2 machine. The plot shows that the rendering complexity depends only sublinearly on the number of input images.

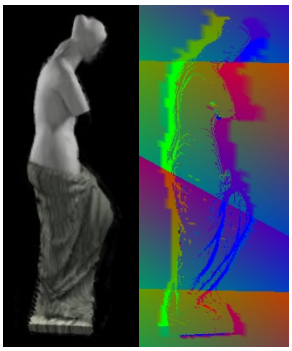

Figure 8.24: Reconstructed view and occlusion correction for polygon-based free form light field rendering.

input cameras). This way directionally varying real-world imagery can be displayed without intermediate resampling steps, and yet rendering of free form light fields can be performed as efficiently as for two-plane-parameterized light fields using texture mapping graphics hardware. Comparable to sphere-based parameterizations, a single free form light field can represent all possible views of the scene without the need for multiple slabs, and it allows for relatively uniform sampling. Furthermore, we extend the rendering algorithm to account for occlusion in certain input views. We also show that the two-plane parameterization (with one or multiple light slabs) is a special case of the free form light field. In contrast to previous as well as very recent work on *unstructured* light field rendering [4, 1], using a convex camera mesh allows to use a static eye point triangulation that is consistent over time and space.

We applied our method to synthetic and real-world datasets with and without additional geometric information and compared the resulting rendering performance and quality to two-plane-parameterized light field rendering. It turned out that in the worst case (with massive adaptive depth correction) free form parameterizations require twice the rendering time than comparable two-plane parameterization.

8.1.4 On-the-Fly Lumigraphs from Dynamic Scenes

Bringing together the advantages from our free form parameterization (Section 8.1.3) and our flexible per-pixel depth representation (Section 8.1.2), we were able to put together a complete image-based capturing and rendering system that can capture input views, reconstruct per-pixel depth information, transmit only the necessary data to the rendering entity, and render desired views of the scene on the fly [9]. The system exhibits two prominent features:

1. **On-the-fly processing.** For real-world scenes, by the use of some computer vision techniques, video camera pairs can deliver color data together with a dense depth map at multiple frames per second. This data will be consumed by the Lumigraph renderer immediately as well. No lengthy preprocessing is involved throughout the whole system.

2. **Distribution and Scalability.** In order to obtain more information about the 3D scene, multiple cameras at different positions must be used. We distribute the data acquisition tasks of these cameras to several computers (see Figure 8.26). Such a client-server paradigm scales well for more cameras and spares more computation power for the renderer. In contrast to

other real-time image-based capturing and rendering systems [6], our rendering complexity mainly depends on the size of the output image and increases sublinearly with the number of input images [8]. In addition, our distribution scheme also provides a unified interface for different types of input data/sensors.

Figure 8.25 shows our so-called *Lumi-Shelf*, a simple bookshelf equipped with some FireWire cameras that we used for our real-world experiments. The whole system is able to acquire and render scenes at 1 – 10 frames per second, depending mainly on the image resolution (320×240 in our experiments) and on the performance of the depth reconstruction technique. Rendering quality is inherently determined by the correctness and fineness of the reconstructed depth information. As a result, synthetic datasets produce high-quality rendering, whereas rendering quality for real world data coming from the video cameras suffers from incorrect and missing depth information. However, there is still great potential to improve the quality of depth maps from real data. Hence we hope to suppress the rendering artifacts significantly in the near future.

Figure 8.25: The Lumi-Shelf, consisting of six FireWire cameras mounted in a bookshelf.

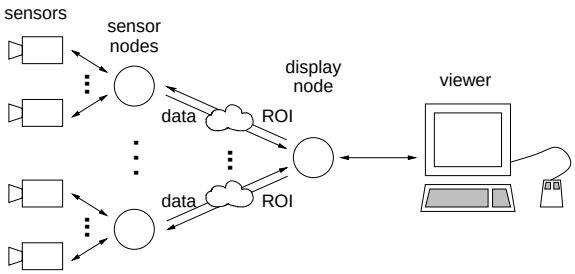

Figure 8.26: On-the-fly Lumigraph acquisition, processing, and rendering system overview. The processing is distributed among *sensor nodes* and a single *rendering node*.

References

[1] C. Buehler, M. Bosse, L. McMillan, S. Gortler, and M. Cohen. Unstructured lumigraph rendering. In *Proc. SIGGRAPH 2001*, pages 425–432. Addison Wesley/ACM Press.

[2] E. Camahort, A. Lerios, and D. Fussell. Uniformly sampled light fields. In *Rendering Techniques '98 (Proc. Eurographics Workshop on Rendering)*, pages 117–130. Springer, 1998.

- [3] W. Heidrich, H. Schirmacher, H. Kück, and H.-P. Seidel. A warping-based refinement of lumigraphs. In V. Skala, editor, *Proceedings of the 7th International Conference in Central Europe on Computer Graphics, Visualization and Interactive Digital Media (WSCG-99)*, pages 102–109, Plzen, Czech Republic, February 1999.

[4] B. Heigl, M. Pollefeys, J. Denzler, and L. van Gool. Plenoptic modeling and rendering from image sequences taken by a hand-held camera. In *Proc. 21. Symposium für Mustererkennung (DAGM '99)*, pages 94–101, 1999.

- [5] M. Magnor and W. Heidrich. Image-based rendering. In B. Girod, G. Greiner, and H. Niemann, editors, *Principles of 3D Image Analysis and Synthesis*, chapter 5.4, pages 232–241. Kluwer Academic Publishers, Boston, Massachusetts, January 2000.

[6] W. Matusik, C. Buehler, R. Raskar, S. J. Gortler, and L. McMillan. Image-based visual hulls. *Proc. SIGGRAPH 2000*, pages 369–374, 2000. ISBN 1-58113-208-5.

- [7] H. Schirmacher, W. Heidrich, and H.-P. Seidel. Adaptive acquisition of lumigraphs from synthetic scenes. In P. Brunet and R. Scopigno, editors, *Proceedings of the 20th Annual Conference ot the European Association of Computer Graphics (Eurographics-99)*, volume 18 of *Computer Graphics Forum*, pages 151–160, Milano, Italy, 1999. Eurographics Association, Blackwell.
- [8] H. Schirmacher, W. Heidrich, and H.-P. Seidel. High-quality interactive lumigraph rendering through warping. In S. Fels and P. Poulin, editors, *Proceedings of Graphics Interface 2000 (GI-00)*, pages 87–94, Montreal, Canada, May 2000. CHCCS, Canadian Information Processing Society.
- [9] H. Schirmacher, L. Ming, and H.-P. Seidel. On-the-fly processing of generalized lumigraphs. In A. Chalmers and T.-M. Rhyne, editors, *Proceedings of Eurographics 2001*, volume 20 of *Computer Graphics Forum*, pages 165–173, Manchester, UK, 2001. Eurographics Association, Blackwell.
- [10] H. Schirmacher, C. Vogelgsang, H.-P. Seidel, and G. Greiner. Efficient free form light field rendering. In T. Ertl, B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proc. 6th International Fall Workshop on Vision, Modeling, and Visualization 2001 (VMV01)*, Stuttgart, Germany, 2001. To appear.

8.2 Extensions and Applications of Light Fields

Investigators: Stefan Brabec, Hendrik P.A. Lensch, Wolfgang Heidrich

In addition to the "core" light field research presented in the previous sections, we also worked on extensions and applications of the light field approach. Section 8.2.1 describes a way to efficiently represent complex light sources as light fields, whereas in Section 8.2.2 a light field-like structure is used to store geometric and material properties of a scene, rather than the actual light distribution.

8.2.1 Canned Light Sources

To avoid costly global illumination algorithms, we pre-compute the outgoing light field of a given lamp and store away these radiance values in a four-dimensional data structure [1, 2]. This discretized lighting information can later be used to reconstruct the radiance emitted by the light source at every point in every direction. Instead of generating a global illumination solution, a *canned light source* could also be generated by resampling measured data or simply by computing radiance values to achieve special effects.

A canned light source can easily be integrated in standard global illumination algorithms, e.g. ray-tracing or radiosity. Due to the huge amount of data, speed-up techniques such as shadow maps or adaptive downsampling become mandatory for real world applications. Figure 8.27 (left) shows the illumination caused by a point source located in the focal point of a parabolic reflector.

In addition to global illumination methods, it is also possible to use computer graphics hardware (projective texture mapping, multi-pass rendering) for reconstructing the illumination from a Canned Light Source. Figure 8.27 (right) shows a scene consisting of 500 polygons and 40 textures, lit by a canned light source resembling a slide projector. The focal plane is approximately located at the projection screen. There, the image is very sharp, while it is somewhat blurred on the wall behind. This scene can be rendered at interactive frame rates ($\approx 20\mathsf{fps}$) on modern graphics hardware.

8.2.2 Light Field Techniques for Reflections and Refractions

Until recently, only a few algorithms exist which allow rendering of reflections and refractions of arbitrary shaped objects at interactive frame rates. One common way to deal with reflections off curved surfaces is using environment maps. This technique is supported by most contemporary graphics hardware. But environment maps restrict the spatial settings of the scene and sometimes

even produce visible artifacts. Additionally, this approach is not suited to handle refractions in thick, solid objects which refract light rays more than once.

Light fields have been introduced to store the illumination of complex objects from different viewpoints. Based on this idea we store view dependent geometrical information (instead of illumination) that can be used to look up the correct illumination [4, 3]. The geometrical data may consist of precomputed texture coordinates into an environment map or into another light field. Geometry and illumination are strictly decoupled by this step. Hardware support is possible for all stages of this technique through existing extensions of the OpenGL rendering pipeline. Figure 8.28 shows a glass teapot that reflects and refracts its environment, rendered with our technique.

References

- [1] S. Brabec. Lumigraphen von Lichtquellen. Masters thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg, March 1999.

  [2] W. Heidrich, J. Kautz, P. Slusallek, and H.-P. Seidel. Canned lightsources. *Eurographics Rendering Workshop 1998*, pages 293–300, June 1998. ISBN 3-211-83213-0. Held in Vienna, Austria.

- [3] W. Heidrich, H. Lensch, M. Cohen, and H.-P. Seidel. Light field techniques for reflections and refractions. In D. Lischinski and G. W. Larson, editors, *Rendering Techniques '99: Proceedings of the 10th Eurographics Workshop on Rendering (EGRW-99)*, pages 187–196, Granada, Spain, June 1999. Eurographics, Springer.

- [4] H. Lensch. Techniques for hardware-accelerated light field rendering. Masters thesis, Friedrich-Alexander-Universität Erlangen-Nürnberg, July 1999.

8.3 Vision-Based Capture of Highly Articulated Body Motion

Investigator: Christian Theobalt

Research in this field focuses on the visual interpretation and analysis of video streams showing a human actor performing a complex motion sequence [12]. The goal is to extract the 3D pose and body configuration (e.g. joint angles) and to visualize the tracking data with an animation model. In the last two decades many different approaches for capturing human motion data were proposed which in most cases impose strong constraints on the scene setup to keep the problem feasible. The industry offers real-time motion capture systems which depend on the tracked person wearing special clothes or markers. Hence, there is a strong interest in non-intrusive inexpensive motion

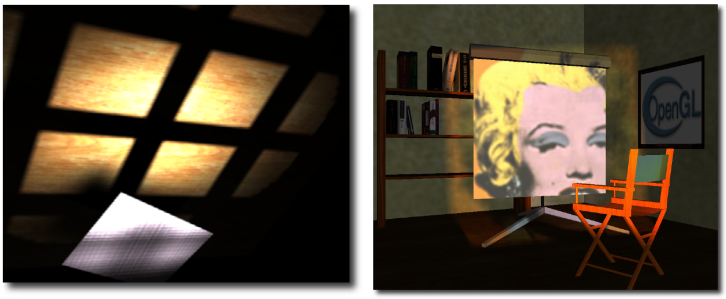

Figure 8.27: Two "canned light sources" representing slide projectors (left: ray tracing; right: OpenGL rendering). Note the correct simulation of sharp and blurred image regions, depending on the focal distance.

Figure 8.28: Reflection and refraction in a glass teapot, represented by a modified light field structure.

capture systems. This early stage of the project only allows to present preliminary results with the first components of a capture system prototype and an outline of intended future research.

Figure 8.29: *Tracking model fitted to a video frame showing a person walking*

The task of a human body motion capture system can be divided into three main subtasks : initialization, tracking and pose estimation. The initialization consists of fitting the model to the first frame and adapting it to the observed person's shape, this is either done automatically [1] but mostly by human interaction [2]. Tracking is the identification of corresponding body parts in subsequent frames of a video stream as well as the separation of the person from the background. To achieve this some systems use image features such as edges or contours [6, 11] whereas more recent approaches make use of spatiotemporal information such as optical flow [2]. In estimation the tracked body parts are matched to a model of the observed person from which a new body configuration and 3D pose are derived. Simple tracking models used by other researchers reach from purely iconic image features such as edges or statistical representations of image regions to identify body parts without modeling the body as a kinematic chain [5] to silhouette-based approaches.

The simplest body models also representing the kinematic properties are stick-figure models where the human body is reduced to a simplified skeleton of joints connected by line segments [3]. Volumetric approaches extend this simple skeleton by modeling limbs using ellipsoids [2] or cylinders [6] and more recently also tapered superquadrics [4], whereas the most complex multi-layered approaches also model muscle tissue using isosurfaces and B-spline patch surfaces on top [9].

Our current tracking system prototype is designed to process video streams recorded off-line from one or several camera views. No camera calibration is needed and the initialization procedure requires the user to manually mark joint locations in the initial frame. A hierarchical volumetric kinematic chain body model using ellipsoids or elliptical cylinders for the limbs represents the tracked person which allows to efficiently handle self-occlusion and supports the idea of using the same model for tracking and visualization. The model is fit to the first frame by minimizing an energy function describing the position error between marked and projected joint locations using non-linear optimization for the inverse kinematics [10, 13], the body dimensions are also approximated to the real anthropometric data of the observed person (see Figure 8.29).

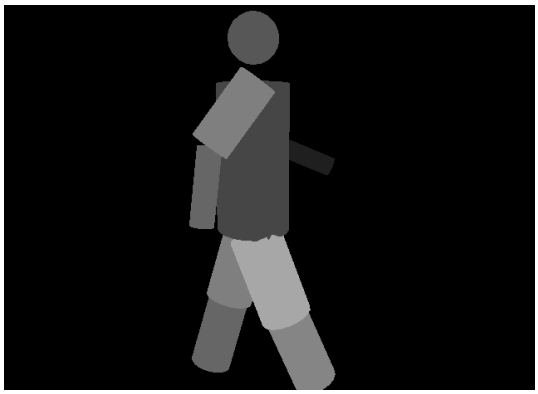

Figure 8.30: *Model with support regions for different body parts*

The projected model segments represent support regions (see Figure 8.30) allowing the grouping of image pixels into coherent blobs for different body parts. Using optical flow [7] and a twist representation of the joint rotations [8, 2] a differential approach is applied to estimate body pose and joint angle changes between two consecutive video frames. The use of region-based features like

blobs for tracking was shown to be more robust than boundary features and particularly suitable for real-time applications [2, 1, 5].
Further work will deal with the development of a tracking software and hardware environment allowing to process multiple video streams in a distributed system in real-time. The main focus will lie on the integration of efficient and robust feature matching and representations of the kinematic properties with optimization algorithms which allow the system to run online. Ways to automatically initialize the system and label body parts in the first frame as well as improved methods of adaption to the real anthropometric dimensions will need more attention. Furthermore, the current basic kinematic human model shall be extended to a multi-layered one also modeling tissue (and muscles) which might be adapted to the real person in combination with tracking. Different ways to cluster image pixels to regions have been proposed, and more recent probabilistic approaches using mixture models, different image statistics and the integration of machine learning algorithms (e.g. EM) into model adjustment shall be developed further. The integration of boundary feature and region feature tracking into a hybrid system appears to be promising and will also be tested in this context.

References

[1] A. Azarbayejani and A. Pentland. Real-time self-calibrating stereo person tracking using 3-d shape estimation from blob features. In *ICPR96*, page C8A.6, 1996.

[2] C. Bregler and J. Malik. Tracking people with twists and exponential maps. In *CVPR98*, pages 8–15, 1998.

[3] Z. Chen and H. Lee. Knowledge-guided visual perception of 3d human gait from a single camera sequence. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(2):336–342, 1992.

[4] D. Gavrila and L. Davis. 3d model-based tracking of humans in action: A multi-view approach. In *CVPR96*, pages 73–80, 1996.

[5] I. Haritaoglu, D. Harwood, and L. Davis. W4: Who? When? Where? What? A Real Time System for Detecting and Tracking People. In *AFGR98 (Tracking and Segmentation of Moving Figures)*, 1998.

[6] D. Hogg. Model-based vision : a program to see a walking person. *Image and Vision Computing*, 1(1):5–20, 1983.

[7] B. Horn and B. Schunck. Determining optical flow. *AI*, 17:185–203, 1981.

[8] R. M. Murray, Z. Li, and S. S. Sastry. *A mathematical introduction to robotic manipulation.* CRC Press, 1994.

[9] R. Plankers and P. Fua. Tracking and modeling people in video sequences. *CVIU*, 81(3):285–302, March 2001.

[10] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical recipes in C : the art of scientific computing.* Cambridge Univ. Press, 2nd ed. edition, 1994.

[11] K. Rohr. Incremental recognition of pedestrians from image sequences. In *CVPR93*, pages 8–13, 1993.

[12] D. Terzopoulos. Visual modeling for computer animation: Graphics with a vision. *Computer Graphics*, 33(4):42–45, November 1999.

[13] J. Zhao and N. Badler. Inverse kinematics positioning using nonlinear programming for highly articulated figures. *ACM TOG*, 13:313–336, October 1994.

9 Realistic Hardware-Supported Shading and Lighting

With fast 3D graphics becoming more and more available even on low end platforms, the focus is beginning to shift towards higher quality rendering.

In our research we develop new algorithms for high quality shading and lighting using computer graphics hardware. The research is mostly concerned with algorithms for generating various local shading and lighting effects and for visualizing global illumination solutions. In particular, we discuss algorithms for shadows, bump mapping, alternative material models, mirror reflections and glossy reflections off curved surfaces.

9.1 Accelerated Local Illumination

Investigator: Jan Kautz

9.1.1 Bump Mapping with Anisotropic Shift-Variant BRDFs

Blinn [1] has shown how wrinkled surfaces can be simulated by only perturbing the normal vector, without changing the underlying surface itself. The perturbed normal is then used for the lighting calculations instead of the original normal. This technique is generally called bump mapping. A simple lighting model such as the Phong model or the Blinn-Phong model is usually used for the lighting calculations. Usually the same simple models are used for hardware accelerated bump mapping [2], which also only allow to vary the diffuse and specular coefficient on a per-pixel basis.

We have developed a technique [3] that evaluates (any analytical) BRDF on a per-pixel basis using these two capabilities at interactive rates, even if the local surface coordinate frame and the material properties (i.e the parameters of the BRDF) are allowed to vary per pixel.

Figure 9.31: *Bump mapped sphere with anisotropic BRDF, running at 30Hz.*

The basic idea we use is simple. A given analytical lighting model is broken up into complex functions that cannot be computed by the graphics hardware. These functions are sampled and stored in texture maps. Each of the functions depends on some parameters which need to be computed from surface and material properties, such as the local coordinate frame, roughness of the surface, and so on. For every pixel of the surface these properties are stored in texture maps. The computation of the parameters of the complex functions depends on the used reflectance model, but requires only simple math, such as dot-products and multiplications, which can be done in hardware, since dependent texturing is used to perform all complex operations. The computed parameters are then used to index into the sampled complex functions using a dependent texture lookup, which effectively evaluates them. Finally the evaluated functions are recombined, depending on how they were broken up, e.g. using multiplication or addition. See Figure 9.31 for an example.

9.1.2 Hardware Accelerated Displacement Mapping

Displacement mapping is an effective technique to add detail to a polygon-based surface model while keeping the polygon count low. For every pixel on a polygon a value is given that defines the displacement of that particular pixel along the normal direction effectively encoding a height field.

Displacement mapping recently made its way into hardware accelerated rendering using standard features. The basic technique was introduced by Schaufler [5] in the context of warping for layered impostors. This algorithm encodes the displacement in the α-channel of a texture. It then draws surface-aligned slices through the volume defined by the maximum displacement. The α-test is used to render only the appropriate parts of every slice. Occlusions are handled properly by this method. This algorithm works well only for surface-aligned slices. At grazing angles it is possible to look through the slices.

We have developed an enhanced method [4] that supports arbitrary slicing planes, allowing orthogonal slicing directions or screen-space aligned slicing commonly used in volume rendering, avoiding artifacts when looking from grazing angles. The main idea is fairly simple. It is based on the observation that every pixel in the volume defined by the maximum displacement can be assigned a distance h_α from the base surface. A pixel belongs to the displacement if its distance h_α is smaller or equal to the displacement d_α encoded in the alpha channel of the displacement texture. We do this test by using the following alpha test: $d_\alpha - h_\alpha > 0$. If true, then the pixel is part of the displacement.

Figure 9.32: *Displacement mapped torus running at 30Hz.*

On the one hand, we use this new method to render traditional displacement mapped objects; see Figure 9.32. This works at interactive rates even for large textures and displacements employing current graphics hardware.

On the other hand this technique can also be used to accelerate image-based rendering techniques, such as the image-based depth objects [4] or the visual hull.

References

[1] J. Blinn. Simluation of Wrinkled Surfaces. In *Proceeding SIGGRAPH 1978*, pages 286–292, Aug. 1978.

- [2] W. Heidrich and H.-P. Seidel. Realistic, hardware-accelerated shading and lighting. In A. Rockwood, editor, *Computer Graphics (SIGGRAPH-99): Conference proceedings*, pages 171–178, Los Angeles, USA, August 1999. ACM Siggraph, ACM.

- [3] J. Kautz and H.-P. Seidel. Towards interactive bump mapping with anisotropic shift-variant BRDFs. In S. N. Spencer, editor, *Proceedings of 2000 SIGGRAPH/Eurographics Workshop on Graphics Hardware*, pages 51–58, Interlaken, Switzerland, August 2000. ACM SIGGRAPH, ACM Press.

- [4] J. Kautz and H.-P. Seidel. Hardware accelerated displacement mapping for image based rendering. In B. Watson and J. W. Buchanan, editors, *Graphics Interface 2001*, pages 61–70, Ottawa, Canada, 2001. Canadian Human-Computer Communications Society, Morgan Kaufmann.

[5] G. Schaufler. Per-Object Image Warping with Layered Impostors. In *9th Eurographics Rendering Workshop*, pages 145–156, June 1998.

9.2 Accelerated Global Illumination

Investigators: Katja Daubert, Wolfgang Heidrich, Jan Kautz, Hendrik Lensch

9.2.1 Illuminating Micro Geometry Based on Precomputed Visibility

Many researchers have been arguing that geometry, bump maps, and BRDFs present a hierarchy of detail that should be exploited for efficient rendering purposes. In practice however, this is often not possible, as no techniques for computing indirect illumination for bump maps existed, while BRDFs and geometry based rendering both took this phenomenon into account. The most costly part when accounting for indirect lighting is computing the visibility. In [4] we show how to precompute and store this information for height fields in a set of textures. The information can then be used during rendering, e.g. in a software renderer, to speed up many different Monte Carlo algorithms. To compute the image in Figure 9.33 (above) we incorporated our method as a bump map shader into a software ray tracer. Note that nearly all surfaces in the image are bump mapped.

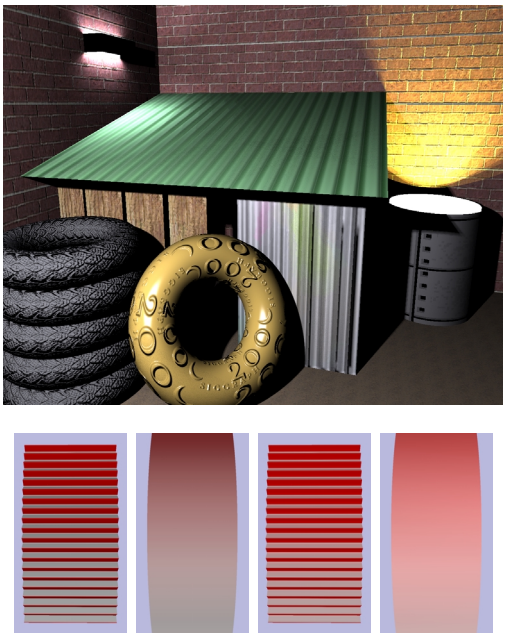

(a) (b) (c) (d)

Figure 9.33: *Top: Software: bump map shader with indirect illumination (and shadows, see Section 9.3.4). Bottom: illuminated height field and BRDF, (a)+(b): only direct lighting (c)+(d) also indirect lighting.*

By using a variant of Monte Carlo algorithms called the Method of Dependent tests we can write the indirect lighting computation as a series of SIMD operations, which compute the indirect lighting for all points on the height field simultaneously. This algorithm can easily be mapped onto hardware, the only requirements being a way of rendering bump maps and dependent texturing. As a benefit we can compute indirect illumination extremely efficiently (i.e. 100 reflections per seconds for a 64x64 height field), making it possible to rapidly compute samples for BRDFs or other higher dimensional data structures. In Figure 9.33 (below) we compare the results of computing the illumination in a height field with the corresponding BRDF. Figure 9.33(a)+(b) show the results for only direct illumination while Figure 9.33(c)+(d) also include indirect illumination. Compared to (a) the height field in (c) is generally brighter and clearly exhibits color bleeding of the red color onto the white faces. Both phenomena are also apparent in the corresponding BRDF.

9.2.2 Efficient Cloth Modeling and Rendering

Rendering cloth is a challenging task, because many different phenomena need to be taken into account. For instance the appearance of cloth is highly dependent on the light and viewing direction, which is due to occlusion and self-shadowing. Often the individual stitches and weaves can be resolved, which means we need a reflection model which can take the spatial dimensions as well as the four dimensions for viewing direction and light direction into account. The model we propose [2] point samples the spatial dimensions into a 2D array, making each array entry a 4D function. Each entry then consists of Lafortune lobes (dependent on light and view) combined with and a view-dependent lookup table. The parameters of this proposed model are acquired in the following way:

Given the micro geometry of a stitch or weave and a set of pairs of viewing and light directions we compute the lighting, including self-shadowing and indirect lighting, of the micro geometry for each pair of directions. We then sample radiance data, normals and transparency values which we use to fit the parameters of each entry in the BRDF array.

Once we have all parameters we can apply the computed spatially variant reflection model to a given cloth. By exploiting OpenGL texture mapping we can easily find out which BRDF array entry to evaluate for each pixel on the garment. Results of applying the reflection model of knitted fabric to a sweater model can be seen in Figure 9.34.

As texture based techniques are prone to aliasing, it is essential, that they lend themselves to mip-mapping. We can build several mip-mapping levels by grouping the data for four neighboring array entries in the acquisition phase and then fitting this grouped data to generate BRDF array entries for the next level. Again we can use OpenGL hardware texturing to determine which entry in which level to evaluate.

Figure 9.34: *Reflection model of knitted fabric applied to sweater model.*

9.2.3 Environment Maps

Environment maps [1] are a widely used technique to approximate reflections in interactive rendering. Recently environment maps have been introduced as a means to render glossy reflections, e.g. [3]. These methods prefilter an environment map with a given fixed reflection model, which can take up to a few hours for large environment maps.

We have developed two methods [6] to speed up this computation. One is a fast hierarchical prefiltering method, that can be used to compute all known types of prefiltered environment maps. And secondly, we have developed a hardware-accelerated filtering method that prefilters environment maps for certain reflectance models [5] at interactive rates even on low-end workstations.

Filtering an environment map corresponds to applying a space-variant filter, which corresponds to a 2D slice through a BRDF. The hierarchical prefiltering method first builds a mip-map of the environment map as well as for all the filter kernels, which vary per pixel. Then the filter is applied to the coarsest possible level of the environment mip-map. This level is found by checking the differences to the next finer mip-map levels of the filter kernel. If it is smaller than some threshold, we use this mip-map level of the filter kernel (and the corresponding environment mip-map level). This hierarchical method has sublinear complexity in contrast to the brute force method, which has linear complexity in the number of touched pixels.

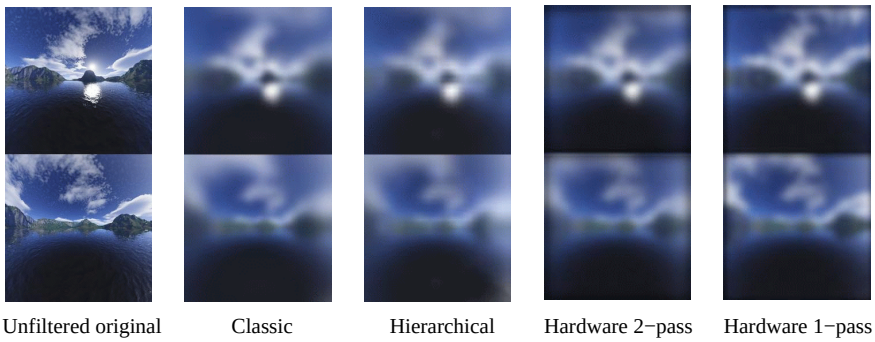

Figure 9.35: *Comparison of the different filtering methods. From left to right: Unfiltered, classic method, new hierarchical method, hardware accelerated method with (2 and 1 pass(es).*

The hardware accelerated filtering method works only with radially symmetric filter kernels, which is for example the case for the Phong model. Then we represent the environment map as a parabolic map, which has the property of mapping a circular filter over the sphere to a circular filter in the parabolic map, only the radius of the mapped circular filter changes per-pixel. We will overcome this problem by generating two prefiltered environment maps, one with the smallest and one with the largest necessary filter size. Then we blend between both prefiltered environment maps, the blending weights can be precomputed since they only depend on the position in the parabolic map. The filtering with the space-invariant filters can be done with OpenGL, and is supported in hardware. See Figure 9.35 for a comparison between the different filtering methods.

References

[1] J. Blinn and M. Newell. Texture and Reflection in Computer Generated Images. *Communications of the ACM*, 19:542–546, 1976.

- [2] K. Daubert, H. Lensch, W. Heidrich, and H.-P. Seidel. Efficient cloth modeling and rendering. In S. Gortler and K. Myszkowski, editors, *Rendering Techniques 2001, Proceedings of the 12th Eurographics Workshop on Rendering*, pages 63–70, London, Great Britain, 2001. Eurographics, Springer.

[3] N. Greene. Environment Mapping and Other Applications of World Projections. *IEEE Computer Graphics & Applications*, 6(11):21–29, November 1986.

- [4] W. Heidrich, K. Daubert, J. Kautz, and H.-P. Seidel. Illuminating micro geometry based on precomputed visibility. In K. Akeley, editor, *Computer Graphics (SIGGRAPH-00): Conference proceedings*, volume 2000 of *Computer Graphics Proceedings, Annual Conference Series*, pages 455–464, New Orleans, USA, July 2000. ACM Siggraph, ACM Press.

- [5] J. Kautz and M. McCool. Approximation of glossy reflection with prefiltered environment maps. In S. Fels and P. Poulin, editors, *Proceedings of Graphics Interface 2000 (GI-00)*, pages 119–126, Montreal, Canada, May 2000. CHCCS, Canadian Information Processing Society.

- [6] J. Kautz, P.-P. Vzquez, W. Heidrich, and H.-P. Seidel. A unified approach to prefiltered environment maps. In B. Proche and H. Rushmeier, editors, *Rendering Techniques 2000, Proceedings of the 11th Eurographics Workshop on Rendering*, Springer Computer Science, pages 185–196, Brno, Czech Republic, June 2000. Eurographics, Springer.

9.3 Interactive Shadow Computation

Investigators: Stefan Brabec, Katja Daubert, Wolfgang Heidrich, Jan Kautz

One of the most problematic tasks in computer graphics is the accurate and efficient computation of shadows caused by point, extended, or area light sources. Although there have been enormous efforts in this specific area, only a small subset of algorithms are really appropriate for interactive rendering applications. Recent developments in graphics hardware led to systems that now allow very sophisticated task to be performed by the underlying graphics hardware. Utilizing these programmable features we have developed new algorithms that are able to perform shadow computations at interactive rates.

9.3.1 Extended Light Maps

Based on the traditional *shadow mapping* technique [8], which computes a shadow mask to determine lit and shadowed pixels in the scene, we have developed a combined light- and shadow-mapping algorithm to amortize the cost of normal shadow mapping [1]. The benefit of this approach, is that nearly all stages of the graphics pipeline can be used during the shadow map generation phase, which results in better image quality and a reduced number of rendering passes. Using these *extended* light maps, we are able to pre-compute shading calculations during the shadow map generation phase which saves valuable resources. Figure 9.36 shows a simple example were we computed all lighting and texturing during the generation phase. For each of the two light sources we obtain a RGB image and depth values encoded in the alpha channel. These pre-computed extended light maps are then applied for the final image just like projective textures.

Figure 9.36: *Scene illuminated by two extended light maps.*

9.3.2 Hardware-accelerated Filtering of Shadow Maps

Since the shadow map algorithm operates in image space this means that the method relies on a sampling scheme to process the scene geometry, therefore making it somehow independent of the underlying scene complexity. Since sampling is the fundamental principal of all rasterization-based graphics architectures shadow mapping is among the first choices for real-time rendering.

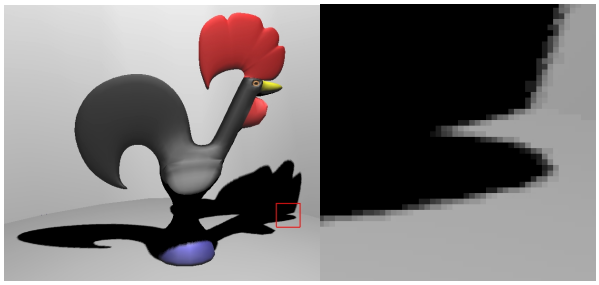

Figure 9.37: *Filtered shadows using fast PCF.*

However, sampling methods all suffer from undersampling artifacts when it comes to higher frequency parts and the same is true for the shadow map algorithm. As discussed by Reeves et al. [6], sampling problems occur during the generation phase as well as when performing the actual shadow test. The first can be solved by increasing the image resolution and by using stochastic sampling. In order to resolve artifacts during the shadow test, Reeves proposes a filtering method called *percentage closer filtering* (PCF) for which we developed a hardware-accelerated variant [2].

Given a 2×2 footprint the four components can simply be stored using the red, green, blue, and alpha channel of the texture image. This means that we have to render the scene four times where in each pass only one color channel is enabled for writing and the image plane is jittered. Given such a packed shadow map it is relatively easy to adapt the shadow test since we only have to extend the one component scheme to four components and compute the percentage of shadowing in a final step.

This algorithm works very well for footprints of size 2×2. If it comes to larger filter sizes, e.g. 3×3 or 4×4, the algorithm needs to be split up into parts of a maximum of four components per texel. The resulting contributions can then be summed up using the accumulation buffer.

A faster way of performing percentage closer filtering can be achieved if we try to retain the effective resolution of the shadow map and use a filtering scheme that softens shadow boundaries by just looking at adjacent texels to compute the shadow mask. This can be done by rendering the

shadow map simultaneously into all four color channel and perform the pixel packing by setting up a convolution filter that assembles a center pixel using depth values taken from the 2×2 region. Figure 9.37 show the results of our fast filtering approach.

9.3.3 Fast Computation of Soft Shadows caused by Extended Light Sources

We present a modification to the shadow map algorithm that allows us to render soft shadows for linear light sources of a high visual fidelity with a very small number of light source samples [3]. This algorithm is well suited for both software and hardware rendering.

We assume that the visibility term can be separated from the local illumination part, and that the latter is smooth enough to be represented by very few light source samples. The task is then to reconstruct the visibility term with a high quality, while only using a small number of light source samples.

Figure 9.38: *Example scene.*

We integrate this visibility information by using a two-channel shadow map, called *soft shadow map*, generated for each sample position on the light source. These shadow maps not only contain information about *which* object points are visible, but also a percentage value that describes *how much* of the whole linear light source can be seen by that point.

In order to achieve interactive frame rates we use a linear approximation for the transition in the penumbra regions and use image processing techniques to determine the location of penumbra regions. This is done by detecting discontinuities in the depth maps and transforming this information to the visibility channel of adjacent sample points. By triangulating depth discontinuities we can utilize graphics hardware to do the linear penumbra approximation and resolve visibility.

Figure 9.38 shows the result of a simple scene illuminated by a linear light source. In this example, only two sample points at the ends of the light source were used to calculate soft shadows.

9.3.4 Shadow Ellipses – Efficient Shadows for Bump Maps

Shadows are not only important for adding realism to macroscopic scenes, they are also essential when rendering micro geometry. Up to now, however, methods for computing shadows in bump maps are not widely used. Our algorithm [4] was developed simultaneously to Sloan et al.'s [7] who mapped the horizon map algorithm by Nelson Max [5] to graphics hardware. Our method relies on a data structure called shadow ellipse, which is computed from visibility textures described in 9.2.1. For each point on the height field and several directions this data structure holds the information whether another height field point is visible in this direction or not.

As shown in the left images in Figure 9.39, we project these directions to the tangent plane and fit the unoccluded directions (yellow) with an ellipse (red). The shadow test then consists of first projecting the light direction to the tangent plane and mapping it to the ellipse coordinate system, then checking whether the transformed light direction is outside the ellipse (shadow) or not. This shadow test can either be integrated into a software renderer's bump map shader (see Figure 9.33 (top)), but also lends itself to hardware implementation, yielding bump map shadows in real time. Figure 9.39 (right) shows a frame from a hardware implementation in which we used shadow ellipses to produce shadows on a terrain in real time.

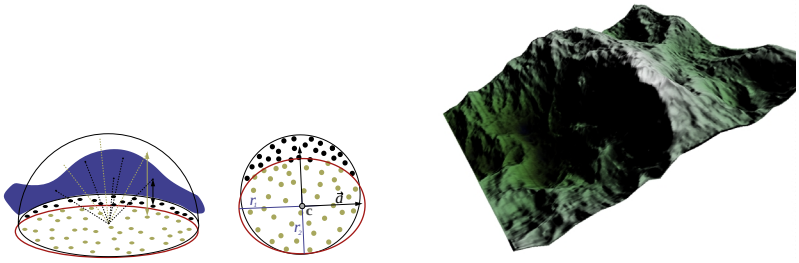

Figure 9.39: *Left: Computation of the shadow ellipse. Right: Example.*

References

- [1] S. Brabec and H.-P. Seidel. Extended light maps. In Q. Mehdi and N. Gough, editors, *GAME-ON 2000 / 1st International Conference on Intelligent Games and Simulation*, pages 10–13, London, United Kingdom, November 2000. Society for Computer Simulation International.
- [2] S. Brabec and H.-P. Seidel. Hardware-accelerated rendering of antialiased shadows with shadow maps. In H. H. S. Ip, N. Magnenat-Thalmann, R. W. H. Lau, and T.-S. Chua, editors, *Proceedings Computer Graphics International 2001*, pages 209–214, Hong Kong, China, July 2001. IEEE Computer Society.
- [3] W. Heidrich, S. Brabec, and H.-P. Seidel. Soft shadow maps for linear lights. In B. Proche and H. Rushmeier, editors, *Rendering Techniques 2000, Proceedings of the 11th Eurographics Workshop on Rendering*, Springer Computer Science, pages 269–280, Brno, Czech Republic, June 2000. Eurographics, Springer.
- [4] W. Heidrich, K. Daubert, J. Kautz, and H.-P. Seidel. Illuminating micro geometry based on precomputed visibility. In K. Akeley, editor, *Computer Graphics (SIGGRAPH-00): Conference proceedings*, volume 2000 of *Computer Graphics Proceedings, Annual Conference Series*, pages 455–464, New Orleans, USA, July 2000. ACM Siggraph, ACM Press.

[5] N. L. Max. Horizon mapping: shadows for bump-mapped surfaces. *The Visual Computer*, 4(2):109–117, July 1988.

[6] W. T. Reeves, D. H. Salesin, and R. L. Cook. Rendering antialiased shadows with depth maps. In *Computer Graphics (SIGGRAPH '87 Proceedings)*, pages 283–291, July 1987.

[7] P. Sloan and M. Cohen. Hardware Accelerated Horizon Mapping. In *11th Eurographics Workshop on Rendering*, pages 281–286, June 2000.

[8] L. Williams. Casting curved shadows on curved surfaces. In *Computer Graphics (SIGGRAPH '78 Proceedings)*, pages 270–274, Aug. 1978.

10 Global Illumination

Synthesis of realistic images which predict the appearance of the real world has many applications including architecture and interior design, illumination engineering, environmental assessment, special effects and film production, along with many others. Due to costly global illumination computation, which is required for the prediction of appearance, physically-based rendering still remains the domain of research laboratories, and is rarely used in industrial practice. The main goal of our research is to analyze problems and provide solutions towards making global illumination affordable in practical applications.

In Section 10.1 efficient solutions which handle complex geometry and multiple light sources are proposed. Real-time lighting computation is not affordable for complex environments, thus, techniques of lighting storage and real-time reconstruction using pre-calculated results are developed (Section 10.2). Since many algorithms make simplifying assumptions about the underlying physical model in order to achieve gains in rendering performance, a validation procedure for testing lighting simulation accuracy and image quality is proposed (Section 10.3). Also, to check the requirement of appearance predictability imposed on the developed algorithms, the rendered images are compared against the corresponding real-world views. In Section 10.4 we outline major guidelines for currently developed solution which enables handling complex scenes with surfaces featuring arbitrary light scattering characteristics. In Section 10.5 the problem of visibility computation in the global illumination framework is discussed.

10.1 Thrifty Final Gather for Radiosity

Investigators: Annette Scheel and Marc Stamminger

Finite Element methods like Hierarchical Radiosity are well suited to the computation of the light distribution in mostly diffuse scenes, but the resulting mesh is often far from optimal to accurately represent illumination. Shadow boundaries are hard to capture in the mesh, and the illumination may contain artifacts due to light transports at different mesh hierarchy levels.

To render a high quality image a final gather reconstruction step is usually done after the radiosity process, which re-evaluates the illumination integral for each pixel. This final gather step is very expensive; as a stochastic rule about 1000 samples per pixel are needed to achieve a good image quality.

We drastically reduce the time spent for the final gather step by exploiting spatial and directional coherence information taken from the radiosity solution [1]. Additionally to the radiosity estimate for each sender and receiver, other information about the light transport between each sender and receiver can be stored already during the radiosity process in so-called links. This includes for example the gradient of the illumination, or the presence of occluders. Based on these data we determine a set of senders which need re-evaluation, where an importance driven sampling is used. The contribution of the other senders is only evaluated exactly at the vertices of the radiosity mesh and then interpolated inside the triangles.

Experiments show that usually a large number of senders can be interpolated and the total number of samples is reduced to only 30% which leads to faster computation times. Another advantage of the method is its scalability to the precision of the underlying radiosity solution. The better the radiosity solution is, the less effort will be spent for the final gather. Figure 10.40 shows a result of the new final gather method for a scene containing significant indirect light which is hard to capture by other final gather techniques.

Figure 10.40: Left: Result of the final gather step. Right upper small image: number of samples (red: visib., green: formfactor). Lower small image: radiosity solution.

References

- [1] A. Scheel, M. Stamminger, and H.-P. Seidel. Thrifty final gather for radiosity. In S. Gortler and K. Myszkowski, editors, *Rendering Techniques: Proceedings of the 12th Eurographics Workshop on Rendering*, pages 1–12, London, 2001. Springer.

10.2 Corrective Texturing

Investigators: Marc Stamminger, Jörg Haber, and Hartmut Schirmacher

Realism in computer graphics can be obtained in two ways. In photorealistic rendering, synthetic images are produced that look 'real'. Computation of such images usually requires minutes or hours. But realism can also be created through interactivity, even if the world the user explores looks merely synthetic.

The gap between photorealistic and interactive rendering is still large. The rapid development of graphics hardware makes interactivity possible for complex scenes, even on cheap PC graphics boards. Nevertheless, just the inclusion of shadows as the most basic global illumination effect usually destroys interactivity. On the other hand, even high-end multiprocessor machines are not fast enough to achieve interactivity using ray tracing for complex scenes.

We try to bring together interactivity and ray tracing by a hybrid approach [3]. In our system a user moves through a world interactively. Views are generated using graphics hardware, with all known limitations in realism. Whenever computation power is available on local or remote processors, we acquire ray tracing samples in parallel to the rendering of the current view. The resulting samples are then used to generate *corrective textures* that represent the error between the hardware generated images and the ray tracing results. These corrective textures are mapped onto the objects during hardware rendering, correcting the hardware image towards the ray tracing solution (see Figure 10.41).

Several issues have to be addressed by this approach:

- Assignment of textures: We do not assign a texture to each scene triangle, but to entire groups that form an object, e.g. a cup, a chair, or a book. For this purpose a scene hierarchy is constructed [2, 1]. The hierarchy level to which the textures are assigned depends on the screen size of the object and on its depth range within the current view.

- Texture mapping: Corrective textures are projected onto objects by point projection from the current view point. In this way, ambiguities between ray samples and texture pixels are avoided.

- Insertion of samples: Samples are splatted into the textures using a Voronoi diagram creation method. Blending is used to blur the Voronoi cell boundaries.

- Sample generation: Ray samples are concentrated to regions that exhibit the largest error. The samples can be generated on parallel processors or remote machines.

- Aging of textures: Textures are reused from frame to frame, but the splatted samples age according to the camera movement. Old and thus uncertain samples are likely to be overwritten by new ones.

In our implementation, we render a radiosity solution with graphics hardware and use a standard distribution ray tracer to compute the corrective textures. Even for rather complex test scenes, we achieve display rates of several frames per second.

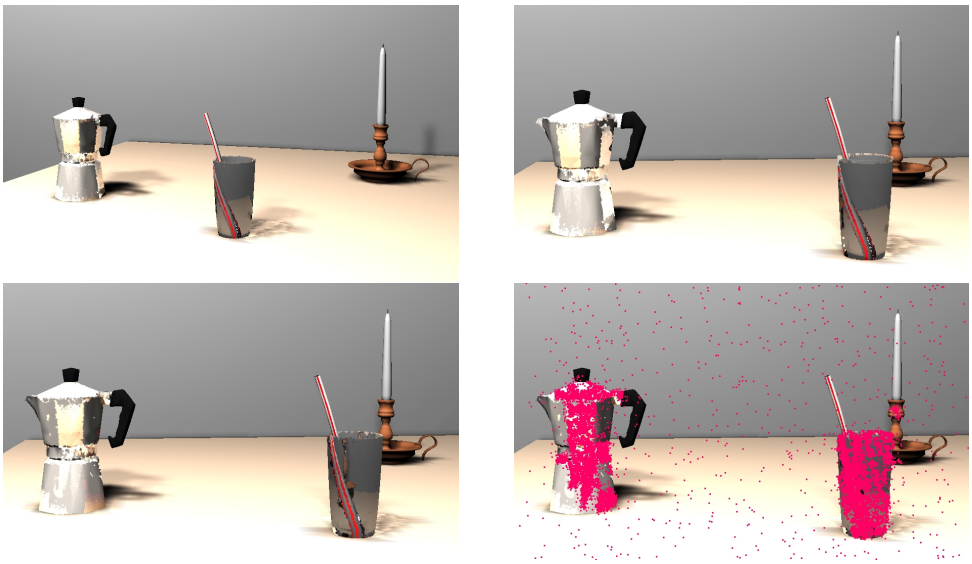

Figure 10.41: *Snapshots from an interactive session. Upper left: almost converged rendering of a scene after a few seconds. Upper right: new viewpoint, with the same corrective textures. Lower left: after 10,000 new samples have been shot. Lower right: positions of the new samples.*

References

- [1] J. Haber, M. Stamminger, and H.-P. Seidel. Customized Object Hierarchies. In B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling, and Visualization (VMV-00)*, pages 255–262, Saarbrücken, Germany, November 2000.

- [2] J. Haber, M. Stamminger, and H.-P. Seidel. Enhanced Automatic Creation of Multi-Purpose Object Hierarchies. In B. A. Barsky, Y. Shinagawa, and W. Wang, editors, *Proceedings of the 8th Pacific Conference on Computer Graphics and Applications (PG-00)*, pages 52–61;437, Hong Kong, China, October 2000. IEEE Computer Society.

- [3] M. Stamminger, J. Haber, H. Schirmacher, and H.-P. Seidel. Walkthroughs with Corrective Texturing. In B. Peroche and H. Rushmeier, editors, *Rendering Techniques 2000, Proceedings of the 11th Eurographics Workshop on Rendering*, pages 377–388;422, Brno, Czech Republic, 2000. Springer.

10.3 Validation of Global Illumination and Rendering Solutions

Investigator: Karol Myszkowski

It is relatively easy to use commodity rendering techniques to create great looking images; however, it is much more difficult to create images that match the appearance of a real environment [3]. The basic precondition to achieving this goal is physically-based lighting simulation, which is a computationally demanding problem. To make computation tractable in practical applications, many simplifying assumptions are usually introduced to underlying physical models. Because analytic evaluation of such simplifications and interactions between them is generally impractical, the correctness of a given technique must be checked experimentally by a comparison of simulation results to some reference data. For example, the distribution of illumination at some predefined points derived analytically or measured experimentally can be used to validate the lighting simulation part of a rendering algorithm. An effective way to test complete rendering algorithms, including the tone-reproduction techniques used for displaying images on the CRT device (refer to Section 11.3), is a direct comparison of the appearance of virtual and real-world images as seen by the human observer [2, 3]. Unfortunately, such experimental validation was almost never performed for existing global illumination solutions, which makes it difficult to compare their efficiency, or even test their implementation correctness.

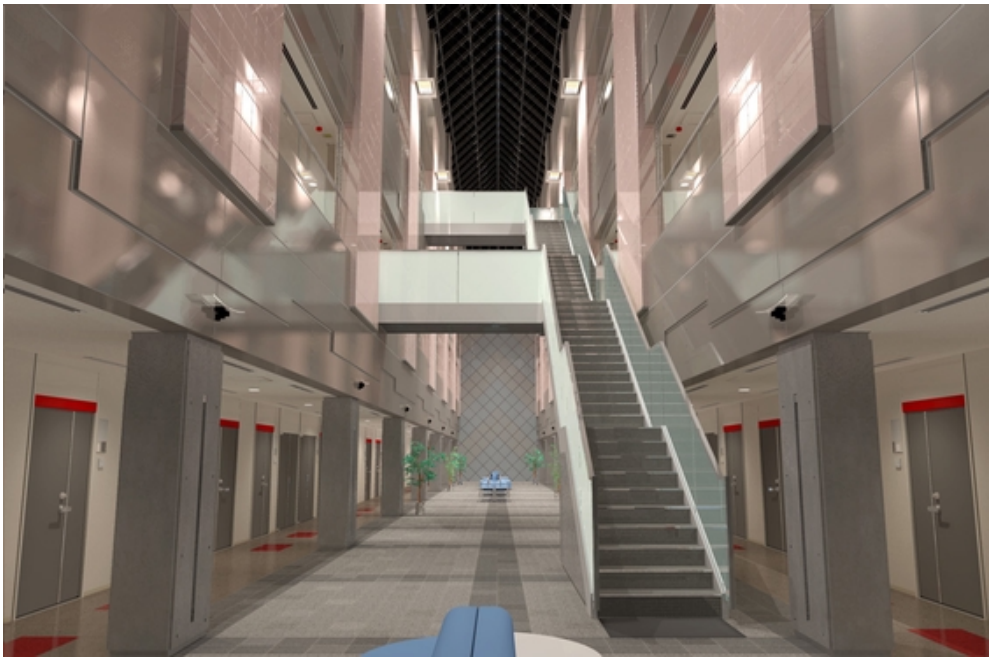

Figure 10.42: Rendered image of the atrium.

One of the reasons that such validation experiments are rarely performed in practice is lack of standardized, robust, non-trivial, and easily accessible test data. We have developed a complete set of data describing an atrium at the University of Aizu (shown in Figure 10.42) which fulfill these requirements. We disseminate this data through the Internet for public use `http://www.mpi-sb.mpg.de/resources/atrium`. Special care was taken to reconstruct a vast majority of the atrium details. The model is built of almost 700,000 polygons. The reflectance characteristics of the most important construction materials were experimentally measured, and BRDFs of high sampling density were obtained. Also, the lighting distribution on the floor was measured at selected sample points. The goal of the experiments with these data was twofold [1]:

- Estimation of the lighting simulation accuracy through comparison with measurement data;
- Evaluation of computer images fidelity in respect to the real world environment.

Our global illumination solution [4] successively passed the standardized set of tests [3] with the atrium data.

References

- [1] F. Drago and K. Myszkowski. Validation proposal for global illumination and rendering techniques. *Computers & Graphics*, 25(3):511–518, 2001.

  [2] G. W. Meyer, H. E. Rushmeier, M. F. Cohen, D. P. Greenberg, and K. E. Torrance. An experimental evaluation of computer graphics imagery. *ACM Transactions on Graphics*, 5(1):30–50, January 1986.

- [3] K. Myszkowski and T. L. Kunii. A case study towards validation of global illumination algorithms: Progressive hierarchical radiosity with clustering. *The Visual Computer*, 16(5):271–288, 2000.

- [4] V. Volevich, K. Myszkowski, A. Khodulev, and E. Kopylov. Using the visual differences predictor to improve performance of progressive global illumination computations. *ACM Transactions on Graphics*, 19(2):122–161, April 2000.

10.4 Towards Interactive Predictive Global Illumination in Large 3-Dimensional Environments

Investigator: Philippe Bekaert

Although much progress has been made in recent years, current illumination simulation algorithms still suffer from several shortcomings, that prevent their common use in many application areas. Current algorithms either lead to high computation times, they have enormous storage requirements for large environments, they may unexpectedly fail to produce meaningful results when applied to real-world data, and/or their usage requires expert knowledge.

Based on a recent analysis of the state-of-the-art in the field, this project focuses on the development of more reliable, rapid and automatic illumination simulation algorithms along the following lines:

Hybrid Finite Element/Monte Carlo Strategies Pure finite elements approaches appear to suffer from reliability and accuracy problems. Current Monte Carlo methods on the other hand, converge only very slowly. We develop new algorithms in which a finite element illumination representation is not what will be presented as the result, but in which it serves to reduce the variance of a robust Monte Carlo method;

Hierarchical illumination representation with memory/speed trade-off New hierarchical spatial data structures and refinement criteria are developed that allow efficient illumination storage and queries in such a hybrid approach. The data structures allow a memory-speed trade-off, in order to be able to handle very large environments without excessive storage requirements;

Bi-directional approach In order to reliably handle all possible lighting situations, we investigate algorithms that operate both from the light sources and from the observer in tightly intertwined phases. These phases are combined in a way that avoids certain problems with nowadays bi-directional algorithms;

Sequential Monte Carlo sampling We focus on algorithms that perform the calculations in stages, in which previous results are used in order to optimize random sampling later on. Doing so in novel ways, we expect to achieve a significantly higher error reduction rate than the usual $1/\sqrt{N}$ rate of Monte Carlo methods;

Memory coherence and parallelism Computing hardware increasingly suffers from a gap between processor and memory speed. In order to execute rapidly, our algorithms will be organised in such a way that small, but high speed, memory caches are used optimally. The exploitation of memory coherence will also allow parallel implementation with less communication overhead between processing elements.

10.5 Analytically Correct Shadows

Investigator: Sherif Ghali

A long-standing problem in Computer Graphics is the efficient and robust computation of analytically–correct shadow boundaries. Such boundaries are useful both as a preprocessing tool for radiosity meshing as well as for alias–free shadows by other rendering techniques. Strides have been taken in the pursuit of this problem by tackling a number of related subproblems discussed below.

Efficient Computation of the Occlusion Graph: The *occlusion graph*, developed and used independently in graphics [7] and in geometry [6], is a superset of the order relationships provided by a BSP tree. The occlusion graph gives a complete set of occlusion relationships between elements of a scene in 3D. We have investigated methods for computing the edges of the occlusion graph. An efficient algorithm has been designed and implemented.

Image Vectorization: Image magnification and minification are standard image manipulation techniques. Both techniques can be thought of as operating in raster or image space. We consider a related problem, image *vectorization*, where a single transformation results in images that can be magnified at will by using solely scale and translate operations. The vectorized image fills a rectangular region in the plane and is in turn capable of being rasterized at an arbitrary resolution.

Embedding Automatic Shadows in Scene Graphs: Current scene graphs (Inventor, Java3D, VRML) leave the task of computing shadows to the user. We designed and implemented a complete scene graph that includes a novel node visitor whose task is to extract and render correct shadows without assistance from the user or client [1].

Computation of Object Space Shadows: The computation of shadows by reduction to visibility has long been known. This reduction is widely utilized in computing raster shadows [9] and in computing analytical shadows [2]. By sacrificing the ability to compute shadows on curved objects and by insisting on computing true or analytical shadows, we identify the visibility map as a valuable data structure in shadow computation. The edges in the visibility map identify analytically the shadow edges in a scene [4]. Our approach does not produce fragmented scene geometry such as the one resulting from Shadow Volume BSP trees. We have also described techniques for computing analytical shadows in the presence of linear and area light sources [5]. We showed how to reduce analytical shadows under extended light sources to analytical visibility. We also showed that such reduction requires a solution to two problems traditionally studied in computer vision, the computation of points of change of visibility, or critical points, and the computation of edges of change of visibility, or the aspect graph. A heuristic for computing the latter has been described.

Object Space Visibility: In a classical taxonomy of visibility algorithms [8], Sutherland et al. classify the object space visibility algorithms known in 1974. This taxonomy has been updated to include new algorithmic categories that correspond to the solutions reported since then [3]. In keeping with the original taxonomy, only practical algorithms were considered.

References

- [1] D. Blythe, S. Ghali, L. Kettner, and H. Sowizral. Migrating to an object-oriented graphics api (course 15). In S. Spencer, editor, *Computer Graphics (SIGGRAPH-00): Course Notes CD-ROM*. ACM SIGGRAPH, New York, USA, July 2000.

[2] N. Chin and S. Feiner. Near real–time shadow generation using BSP trees. In *SIGGRAPH '89*, pages 99–106, Aug. 1989.

- [3] S. Ghali. Object space visibility. In L. Harrison, editor, *Computer Graphics (SIGGRAPH-01): Course Notes CD-ROM*. ACM SIGGRAPH, New York, USA, August 2001.

- [4] S. Ghali, E. Fiume, and H.-P. Seidel. Object-space, connectivity-preserving, shadow computation. In B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling, and Visualization (VMV-00)*, pages 281–289, Saarbrücken, Germany, November 2000.

- [5] S. Ghali, E. Fiume, and H.-P. Seidel. Shadow computation: A unified perspective. In *European Association for Computer Graphics 21st Annual Conference (Eurographics-00), State of The Art Reports*, pages 1–17, Aire-la-Ville, Switzerland, August 2000. Eurographics.

[6] M. T. Goodrich. A polygonal approach to hidden–line and hidden–surface elimination. *CVGIP: Graph. Models Image Process.*, 54(1):1–12, 1992.

[7] J. Snyder and J. Lengyel. Visibility sorting and compositing for image–based rendering. Technical Report MSR–97–11, Microsoft Research, April 1997.

[8] I. E. Sutherland, R. F. Sproull, and R. A. Schumacker. A characterization of ten hidden-surface algorithms. *ACM Computing Surveys*, 6(1):1–55, Mar. 1974.

[9] L. Williams. Casting curved shadows on curved surfaces. *Computer Graphics (SIGGRAPH '78 Proceedings)*, 12(3):270–274, Aug. 1978.

11 Perception Issues in Rendering and Animation

Realistic rendering of high-quality images and animation sequences involves the solution of the global illumination problem which is computationally hard (refer to Section 10). The overall problem complexity results in difficulties in formulating proper solution error metrics that could be used to guide the computation efficiently and prevent spending too much effort on measurably negligible scene details. The computation in current rendering algorithms is usually guided by the energy-based metrics of solution errors, which do not necessarily correspond to the visible improvements of the image quality. Since the human observer is the final judge of the quality of rendered images, the computation should be focused only on those image features that are readily perceivable under given viewing conditions. In our research, we consider perceptual issues to improve the efficiency of global illumination computation, rendering, and image display.

In Section 11.1 we present two successful applications of an early vision model (the visual path beginning with the retina and ending with the visual cortex V1 is considered) to guide the off-line computation of high-quality animation frames. In Section 11.2 we show some applications of a visual attention model (involving higher level perceptual and cognitive elements) in the context of high-quality interactive rendering scenario. In Section 11.3 we describe our perception-motivated tone-reproduction solution for image display, which overcomes numerous physical limitations of the display device. In Section 11.4 we study the effects of lossy image compression on the image quality and improve the well-known LBG algorithm to reduce the codebook error significantly.

11.1 Perceptually Guided Solutions for High-Quality Rendering of Animation Sequences

Investigators: Karol Myszkowski, Takehiro Tawara, and Hiroyuki Akamine

The traditional approach to computer animation relies on rendering of every single frame separately, which means that the eye sensitivity variations resulting from the temporal considerations cannot be properly accounted. Effectively, computational efforts can be easily wasted on processing image details that cannot be perceived in the animated sequence. In this context, a global approach to rendering involving both spatial and temporal dimensions appears promising and relatively unexplored research direction.

In our research, we consider applications of perception-based video quality metrics in the context of computer generated animation sequences. Our goal is to improve animation rendering performance by introducing simpler (less costly) rendering methods of in-between frames without affecting the animation quality as perceived by the human observer [6]. We developed a framework for fully automatic, perception-based guidance of the in-between frames computation, which reduces the number of pixels computed using costly ray tracing, and seamlessly (in terms of perception of animated sequences) combine them with pixels derived using inexpensive Image-Based Rendering (IBR) techniques. We have found also two very useful applications of the Pixel Flow obtained as a by-product of IBR processing. The first application is to estimate the spatio-velocity Contrast Sensitivity Function which made possible incorporation of temporal factors into our perceptually-informed image quality metric. The second is to perform the spatio-temporal antialiasing with motion-compensated filtering based on image processing principles (in contrast to traditional antialiasing techniques used in computer graphics). We integrated all these techniques into a balanced animation rendering system.

The central part of the rendering system is our Animation Quality Metric (AQM). The following procedure is called recursively for all animation segments. For a sequence of frames located between

two keyframes, the AQM is used to find perceivable differences between the keyframes warped (reprojected) to the middle inbetween frame. Based on the AQM prediction the following decisions are made:

- whether to split a sequence of inbetween frames by inserting one extra keyframe in the sequence middle,
- which glossy and transparent objects lead to perceivable animation artifacts when derived using IBR techniques (pixels depicting such objects must be ray traced).

As the framework of our AQM we decided to expand the perception-based visible differences predictor for static images proposed by Eriksson el al. [2]. We found that responses of this predictor are very robust, and its architecture was suitable for incorporation of the spatio-velocity CSF derived in Kelly's psychophysical experiments. To account for target tracking by the eye movements we incorporated the extensions of the Kelly's model that have been proposed by Daly [1]. The velocity data required by the CSF model was derived from the Pixel Flow computed using IBR methods [5] as a by-product of pixel reprojection - thus, our AQM metric takes advantage of data that are readily available for synthetic images. For the natural image sequences the Optical Flow can also be derived, but is more costly and usually of far less accuracy. For dynamic environments we assume that the motion of animated objects is fully compensated by the smooth pursuit eye motion, which leads to the high sensitivity of the AQM for such objects. This assumption is justified by the fact that moving objects are very important attractors of the visual attention [4] which means that efficiency of the eye tracking for such objects is very high. All those simplifications result in the AQM processing time of about 4 seconds for a pair of compared frames. Fig. 11.43 illustrates the processing flow of the AQM.

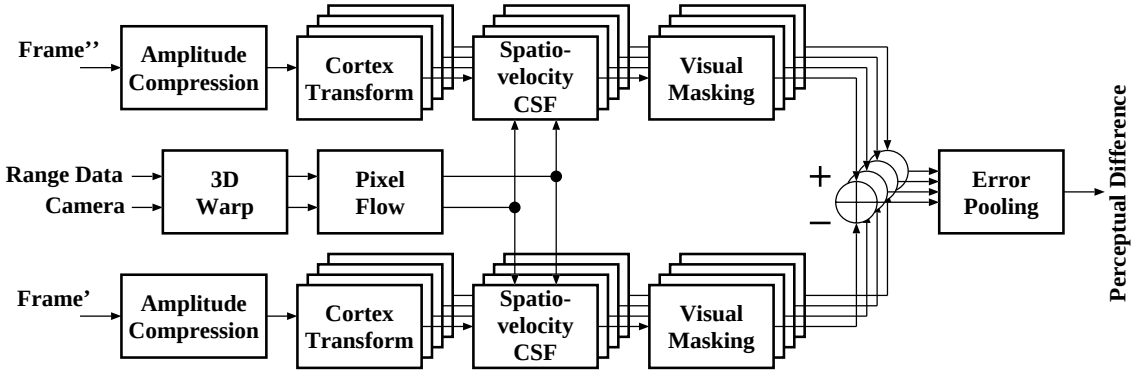

Figure 11.43: Animation Quality Metric (AQM)

More details about this work can be seen on our Web page under the URL:
`http://www.mpi-sb.mpg.de/resources/aqm/`

We found another successful application of AQM to improve the performance of global illumination computation for dynamic environments [7]. Global illumination is an important visual cue, which greatly improves the appearance of rendered images [3]. Our global illumination solution is based on stochastic photon tracing [8] and takes advantage of temporal coherence of lighting distribution, by processing photons both in the spatial and temporal domains. The AQM is used to keep noise inherent in stochastic methods below the sensitivity level of the human observer. As a result a perceptually-consistent quality across all animation frames is obtained. Furthermore, the

computation cost is reduced compared to the traditional approaches operating solely in the spatial domain.

Figures 11.44 show an example animation frame that was computed with/without temporal processing under the same computation time. Notice the noise in Figure 11.44b which becomes even more pronounced in the animation context due to the lack of temporal coherence. The noise was successfully eliminated in Figure 11.44a.

a)  b) 

Figure 11.44: Example frame with temporal processing a) and without temporal processing b) under the same computation time.

More details about this work can be seen on our Web page under the URL:
`http://www.mpi-sb.mpg.de/resources/aqm/dynenv/`

References

[1] S. Daly. Engineering observations from spatiovelocity and spatiotemporal visual models. In *Human Vision and Electronic Imaging III*, pages 180–191. SPIE Vol. 3299, 1998.

[2] R. Eriksson, B. Andren, and K. Brunnstrom. Modelling of perception of digital images: a performance study. In *Human Vision and Electronic Imaging III*, pages 88–97. Proc. of SPIE Vol. 3299, 1998.

[3] D. Greenberg. A framework for realistic image synthesis. *Communications of the ACM*, 42(8):43–53, August 1999.

[4] L. Itti. *Models of Bottom-Up und Top-Down Visual Attention.* Ph.D. Thesis, California Institute of Technology, 2000.

[5] L. McMillan. *An Image-Based Approach to 3D Computer Graphics.* Ph.D. thesis, North Carolina University, Chapel Hill, 1997.

- [6] K. Myszkowski, P. Rokita, and T. Tawara. Perception-based fast rendering and antialiasing of walk-through sequences. *IEEE Transactions on Visualization and Computer Graphics*, 6(4):360–379, October 2000.

- [7] K. Myszkowski, T. Tawara, H. Akamine, and H.-P. Seidel. Perception-guided global illumination solution for animation rendering. In E. Fiume, editor, *Computer Graphics (SIGGRAPH-01): Conference proceedings*, pages 221–230, Los Angeles, CA, USA, 2001. Association of Computing Machinery (ACM), ACM.

• [8] V. Volevich, K. Myszkowski, A. Khodulev, and E. Kopylov. Using the visual differences predictor to improve performance of progressive global illumination computations. *ACM Transactions on Graphics*, 19(2):122–161, April 2000.

11.2 Perceptually Guided Corrective Splatting

Investigators: Jörg Haber, Karol Myszkowski, and Hitoshi Yamauchi

Despite of the progress in (graphics) hardware development, interactive walkthroughs in photometrically complex scenes are still a challenging application, if high quality rendering is required. Since full ray tracing at interactive rates is usually impossible, we render a precomputed global illumination solution using graphics hardware and use remaining computational power to correct the appearance of non-diffuse objects on-the-fly [1]. To obtain the best image quality as perceived by a human observer within a limited amount of time for each frame, we control corrective computation of non-diffuse objects according to a computational model of visual attention. See also Figure 11.45.

Our multi-threaded implementation consists of one *rendering thread* T_R, one *analyzer thread* T_A, and one or more *ray tracing threads* T_{RT_i}. T_R renders the scene using OpenGL hardware. Non-diffuse objects are rendered simultaneously into the stencil buffer with a unique ID code. The frame buffer is analyzed by T_A, and the resulting order of corrections is passed to the T_{RT_i} via a priority queue. Every sample that is computed by one of the T_{RT_i} is sent back to T_R, where it is splatted into the frame buffer using a stencil test to ensure that only the corresponding (non-diffuse) object will be affected by this correction.

To select and order the non-diffuse objects that need to be corrected, we extend the state-of-the-art model of attention developed by Itti *et al.* This bottom-up model performs very well for static images and has been soundly validated by its developers in many demanding applications. Every input image is decomposed into a set of channels (intensity, color components), which are further decomposed into eight frequency scales each using Gaussian pyramids. By applying center-surround differences and across-scale normalization, a *saliency map* is generated, which encodes the saliency of objects into grey levels. To take into account the user's volitional focus of attention during interactive walkthroughs, we additionally consider task-driven factors such as distance from image center and pixel coverage for non-diffuse objects. Weighted blending between these top-down components and the saliency obtained from the Itti model results in a *visual importance priority* assigned to each non-diffuse object.

We use a hierarchical image-space sampling scheme to control ray tracing and splat the generated point samples with a square footprint. Due to the layout of the sampling scheme, the resulting image converges progressively to a ray traced solution if the viewing parameters remain unchanged. Moreover, a sample cache is used to enhance visual appearance if the time budget for correction has been too low for some frame. We measure the validity of the cached samples based on the deviation of the dot product between the surface normal in the hit point and the current viewing direction with respect to the value of the dot product during sample generation. Valid samples are then reprojected into the current view.

In our current implementation, we obtain frame rates of about 10 fps on an `sgi` Onyx3 for interactive walkthroughs in scenes with up to 100,000 primitives of various photometrical complexity such as mirrors, glasses, and glossy objects. We found the predictions of our visual attention model to be usually in good agreement with the user's fixations. We would like to extend our approach using some level-of-detail and occlusion culling techniques to speed up rendering for scenes of higher geometric complexity.

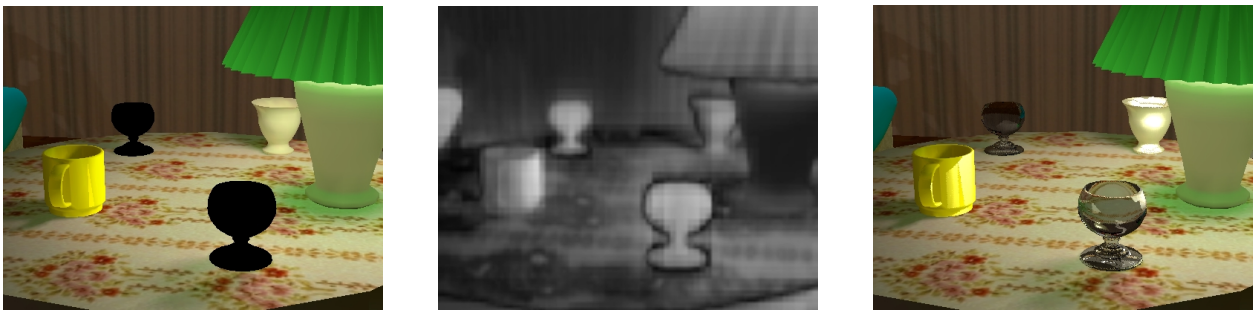

Figure 11.45: *Different stages of our correction process. Left: An input image showing a view-independent global illumination solution computed in a preprocessing step. Middle: The saliency map resulting from visual attention processing of the input image. Grey levels encode the saliency of objects. Right: A fully converged solution is obtained after corrections have been splatted according to the results of our attention model for about 5 seconds.*

References

- [1] J. Haber, K. Myszkowski, H. Yamauchi, and H.-P. Seidel. Perceptually Guided Corrective Splatting. In A. Chalmers and T.-M. Rhyne, editors, *Computer Graphics Forum (Proceedings of Eurographics 2001)*, volume 20 of *Computer Graphics Forum*, pages 142–152, Manchester, UK, September 2001. Eurographics, Blackwell.

11.3 Tone-Reproduction for Interactive Walkthroughs

Investigators: Annette Scheel and Marc Stamminger

The goal of global illumination algorithms is to compute radiance values which closely match the radiance values of the corresponding real world's scene (if this scene would exist). The result of such a simulation is usually displayed on output devices such as a CRT, a powerwall (with stereo projection for a 3D impression), or a CAVE (which is a cube with 5 projection walls). This step requires a conversion from radiance values to rgb-values which can be displayed by the output device. At this point, we encounter a problem: the range of luminances that the output devices can display is much smaller than the range which exists in the real world. For example, outdoor scenes can range from 10^{-6}cd/m^2 in moonlight to 10^{8}cd/m^2 in sunlight, whereas a CRTs dynamic range is only $1 - 100\text{cd/m}^2$.

Tone-Mapping or -Reproduction algorithms try to map the 'world' luminance range to the display luminance range such that the image still is perceived like its pendant in the real world. To achieve this ambitious task, most tone-mapping operators try to mimic the functioning of the visual system, or at least some aspects of it, as for example perception of contrast, or brightness. The behavior of the visual system depends heavily on the adaptation. This is the mechanism that allows the eye to function over a wide range of illumination, although its photoreceptors can cope only with a relatively small range of luminances. Most tone-mapping operators compute an approximation of the state of adaptation and try to reproduce the image like the adapted eye would see it, thereby exploiting the compression already done by the eye.

So far, tone-mapping operators have only been applied to single static images. However, for interactive walks through globally illuminated scenes, tone-mapping is equally necessary. The difference is that in a walkthrough the eye of the observer constantly adapts to the new light conditions which also requires an adjustment of the tone-mapping. This behavior is already known

from video cameras (compare Figure 11.46).

We provided a solution for tone-mapping of interactive walkthroughs [1] which repeats the following three steps for each frame:

- Determine the state of adaptation of the new view.
- Based on this adaptation, compute a new tone-mapping curve.
- Create a texture containing the mapping from world to display luminances and apply it.

The state of adaptation of the human eye is mainly influenced by a field of 1 degree around the focal point. Because it is reasonable to assume that the person operating the interactive walkthrough looks at the center of the window and pans objects he wants to observe to the center, we use a center-weighted average of all visible scene points to determine the adaptation. For tone-mapping then any operator developed for still images can be used which is based on adaptation and outputs only a single tone-mapping curve. Finally this tone-mapping curve is stored in a texture. Alternatively, it would be possible to change all colors of the scene directly according to the tone-mapping, but this would be too time-consuming. The use of the tone-mapping textures allows us to integrate tone-mapping without loosing interactivity. The bottom row of Figure 11.46 shows three frames from an interactive session. Similarly to the video camera, the brightness level is raised when looking into the dark room such that at least some features become visible. In contrast, if the camera pans towards the bright room, everything becomes darker.

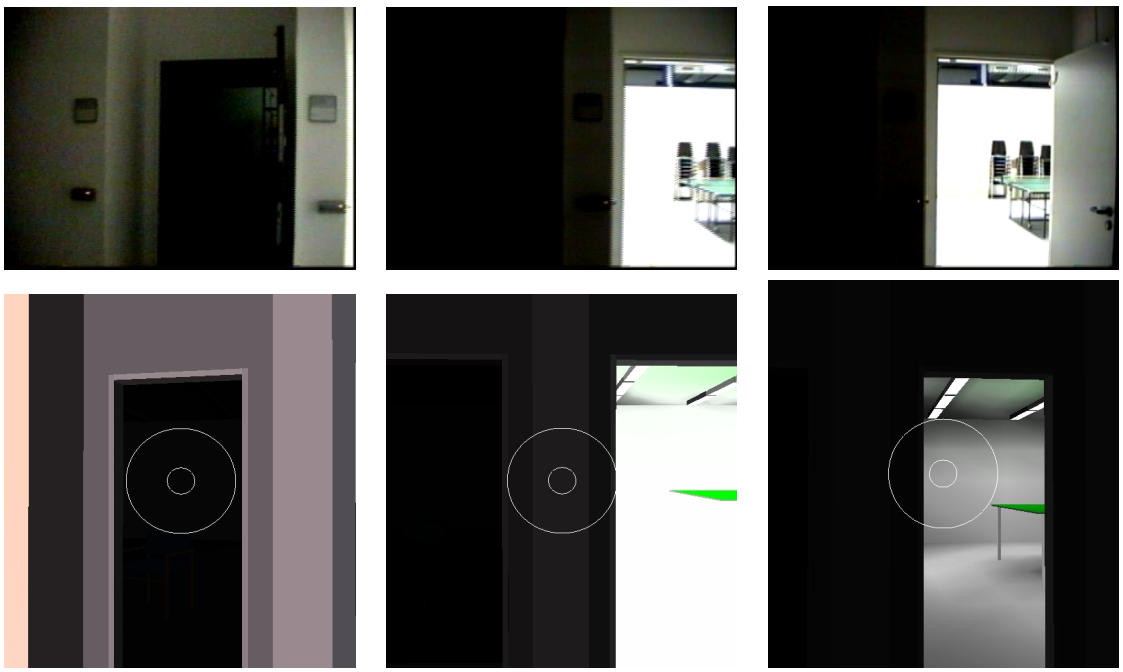

Figure 11.46: Top row: real world images from a video camera, bottom row: interactive viewing of a similar artificial scene with tone-mapping.

References

- [1] A. Scheel, M. Stamminger, and H.-P. Seidel. Tone reproduction for interactive walkthroughs. In M. Gross and F. R. A. Hopgood, editors, *Proceedings of the European Association for Computer Graphics 21st Annual Conference (Eurographics-00)*, volume 19 of *Computer Graphics Forum*, pages 301–311, Interlaken, Switzerland, August 2000. EUROGRAPHICS, Blackwell.

11.4 Lossy Image Compression

Investigator: Jörg Haber

We have developed a framework for evaluating and comparing the quality of various lossy image compression techniques based on a multiresolution decomposition of the image data [1, 2]. Much attention is paid to the interdependencies of the individual steps of such compression techniques, i.e. the *transformation*, *quantization*, and *coding* of the image data. A modification of the well-known LBG algorithm for the generation of codebooks in vector quantization has been published in [3]. Our algorithm, denoted as the ILBG algorithm, reduces the codebook error of the LBG algorithm drastically in typical applications. In our experiments we were able to achieve up to 75 % reduction of the codebook error in only a few additional iteration steps. In the context of lossy image compression this error reduction leads to an increase of 2–3 dB of the peak-signal-to-noise-ratio (PSNR) in turn.

References

- [1] J. Haber. *Konstruktion und Implementierung eines neuen Verfahrens zur Kompression von Bilddaten.* Phd thesis, Technische Universität München, München, May 1999. ISBN: 3-89675-576-5.
- [2] J. Haber and H.-P. Seidel. A framework for evaluating the quality of lossy image compression. Research Report MPI-I-1999-4-001, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, December 1999.
- [3] J. Haber and H.-P. Seidel. Using an Enhanced LBG Algorithm to Reduce the Codebook Error inVector Quantization. In *Proceedings of Computer Graphics International 2000 (CGI-00)*, pages 99–104, Geneva, Switzerland, June 2000. IEEE Computer Society.

12 Software

As part of the research process, several libraries, development tools, and application frameworks have been developed by members of the group. In this section we describe some of them that evolved to a level where it was appropriate to either distribute them as *open source* projects or let members of other research institutes benefit from software that had been developed in our group.

12.1 TMK

Investigators: Hartmut Schirmacher and Stefan Brabec

12.1.1 Overview

`tmk` [3] is a tool that embeds the functionality of `make` in the scripting language `Tcl` [1] in a very simple and convenient way. Furthermore, `tmk` allows higher levels of abstraction via *modules* and a flexible *configuration* framework. In addition to using `tmk` simply as a replacement for `make` [4], the users can create *projects* with global methods, objects, and options, and extend or modify the globally defined tasks using per-directory control files similar to the traditional `Makefile` concept.

The design of `tmk` has been driven by the demand for two things: a simple system for managing larger software projects without having platform- or site-specific code in each `Makefile`, and a scripting environment that is combined with the core functionality of `make`. As a common basis for achieving both goals, we have chosen to embed `make`-like functions into `Tcl`. Additionally, the `tmk` core natively supports architecture-dependent output, multi-directory processing, and things such as exception/exclusion handling.

On top of the `tmk` core, we have added additional abstraction layers by a module mechanism and a centralized configuration system. Through this, it is possible to remove any platform- or configuration-specific code from the control files. Figure 12.47 illustrates the hierarchy of components that build up `tmk`.

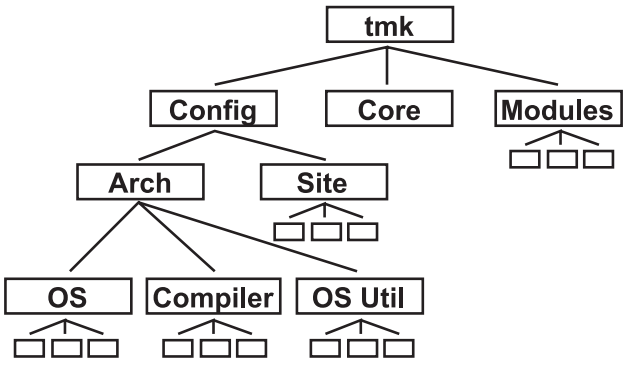

Figure 12.47: The components that build up `tmk`. The configuration is split into two major branches to separate architecture-specific configuration from site-related issues like package existence and installation paths. The CONFIG/ARCH branch is split further to distinguish between operating-system support and things like compiler environments and OS-specific helper tools.

12.1.2 The tmk core

The core functionality of `tmk` was designed to embed the functionality of `make` into `Tcl`. The control files for `tmk` are simply `Tcl` scripts, called `TMakefiles`, and define, either implicitly or explicitly, how to create a target from a set of source files (or primary dependencies). Primary dependencies alone are used to select the appropriate rule (if there are multiple candidates) for each target, whereas secondary dependencies simply define additional preconditions before the dependent target can be built.

In addition to these most basic features, `tmk` also has a way of handling exceptions and exclusions. An *exclusion* means that some target will not be built and will not appear as dependency in any rule. An *exception* temporarily overrides the values of some variables for just some targets. Exceptions also allow to replace the rule completely by a different one.

12.1.3 Modules and Configuration

On top of the core, `tmk` has a *module* mechanism that allows to globally store rules, options, and procedures for certain classes of tasks. Modules are explicitly requested in the control files in order to allow the user to choose the right set of methods for the specific task, and they are parameterized through global or namespace-relative variables.

Site-dependent variables (e.g. installation paths) are not defined inside the module, but rather in the appropriate *site-config* files that are processed by `tmk`'s central *configuration system*. Similar to this, `tmk` reads *arch-config* files that define architecture-dependent options, like for example a procedure for how to call the compiler and linker for a certain task.

12.1.4 TMK Versions and Updates

In August 2000, the `tmk` web site [2] became online and the first beta release V0.9 was available for the public. At its current state, `tmk` runs under IRIX, Solaris, Linux, FreeBSD as well as under Windows 98, 95, NT, and 2000, including support for many of the commonly used compilers under each operating system.

Although the public version of `tmk` has not been updated so far, the MPII-internal version of `tmk` is being maintained and extended continuously, since it is used excessively for nearly all projects within our working group (AG4). The first public V1.0 release is planned for the end of 2001.

References

[1] J. K. Ousterhout. *Tcl and the Tk Toolkit.* Addison-Wesley, 1994.

- [2] H. Schirmacher and S. Brabec. Tmk home page. `http://www.tmk-site.org`.

- [3] H. Schirmacher and S. Brabec. tmk - a multi-site, multi-platform system for software development. In C. Zerbst, editor, *Proc. First European Tcl/Tk User Meeting*, Hamburg, 2000. Technische Universität Hamburg-Harburg.

[4] R. Stallman and R. McGrath. GNU Make. `http://www.gnu.org/software/make`.

12.2 AG4 Shared Projects

Since the computer graphics group (AG4) was founded in 1999, great efforts have been undertaken to develop well-designed libraries that can easily be combined and extended. Some of these libraries and projects are used and developed as shared projects with other sites, as for example the University of Erlangen, the University of Aachen, and the University of British Columbia (Canada).

12.2.1 Shared Projects System

Investigators: Hartmut Schirmacher and Stefan Brabec

From a system point of view, our shared projects focus on the following points:

- hierarchical and transparent project structure
- versioning and central repository
- shared source code basis
- shared object code and executable code
- platform-independence and support for multiple platforms in parallel
- semi-automatic documentation
- easy start and minimal resource requirements for students

The source code is maintained in a central *CVS* [1] repository that provides a number of features for tracking changes, defining versions, and merging code fragments that have been changed concurrently by multiple users. The essential parts of this code tree are checked out from the repository and compiled on the most important platforms (currently IRIX and Linux) every night by an automatic system based on `tmk` (see 12.1). This results in a complete and always-up-to-date *shared version* of the complete source code, libraries, and executables. Should the compilation fail for a certain project, the system restores the previous day's state and sends error report mails to the project administrators as well as to everybody who committed changes to these projects during the last days.

If someone wishes to use a centrally maintained library, she or he simply specifies the project-relative library directory, e.g. `IBR/image/img_core`, in the local *TMakefile*. `tmk` will then search for this library in the person's project directory, and take the shared instance of the library if it is not provided by the user. This way, library locations are completely transparent to the user and need not to be specified in the project *Makefiles*. Specifically, a new user does not need to maintain any source or object code of the shared projects in her/his directory, except if she/he wants to change that code.

If a user compiles a project or application, the generated object code will always be stored in a platform- and codelevel-dependent subdirectory, such that compilation on multiple platforms will simply generate multiple subdirectories below the source code directory. `tmk` will automatically separate all platform-specific things and also use the right version of the shared code basis.

The web-based documentation of the code is generated by using *doxygen* [2], a freely available system for semi-automatic extraction of documentation information from C/C++ source code. The documentation is updated every night in the same way (and with the same methods) as it is done for the code.

References

[1] Concurrent versioning system (cvs) home page. `http:://www.cvshome.org`

[2] Doxygen home page. `http://www.doxygen.org`

12.2.2 *base* Project

Investigators: Hartmut Schirmacher, Stefan Brabec, Hendrik Lensch, Christian Rössl, Marc Stamminger, et al.

The AG4 base library is a collection of useful general-purpose classes, functions, macros, etc:

- Vectors, matrices, and helpers.
- Reference Counting and Smart Pointers.
- Files and Filenames, Path Lists.
- Value Representation as Strings.
- Key/Value Parameters and Parameter Lists.
- Applications with decoupled OpenGL output.
- Graphical user interfaces (*Qt* widgets with extended functionality).

Single or multiple features of the base project are widely used within the group and by the projects described below. Like all our projects, *base* is organized as a number of mostly independent directories/libraries.

12.2.3 Image Based Rendering Code Base

Investigators: Hartmut Schirmacher, Hendrik Lensch, Wolfgang Heidrich, Michael Goesele, Martin Hochstraßer, Jan Uschok, et al.

This IBR code base contains some basic types and functions for tasks in the context of image-based modeling and rendering. Of course the most central concept in this context is that of an *image*. An image is a collection of multiple *buffers*, each of which is a regular array of scalars or vectors of a certain type, one vector or scalar representing one layer of a *pixel* in the image. For example, an image can consist of a RGB-byte-valued *color* buffer plus a floating point-valued *depth* buffer. Additionally, every buffer may have a list of parameters (key/value pairs) that are represented and stored (I/O) transparently. Images and buffers can be stored in a proprietary and portable file format (called *IBRraw*), or they can (with some limitations) be mapped to well-known image formats such as *PPM*, *PNG*, *TIFF*, etc. The IBR project provides a general plug-in mechanism for I/O subroutines, the so-called *loaders* and *writers*.

Operating on images, the IBR project defines generic *filter* classes. A filter takes one or multiple images as input, and produces one or many images as output. It is parameterized via a generic string-based parameter interface (using the *base* project's key/value and parameter interface), and a graphical user interface can be generated automatically for every filter. Meanwhile a lot of different filters have been implemented for different kinds of tasks, such as standard image processing (color manipulation, arithmetic operations on pixels, cropping, resizing, blurring, etc.), reconstruction of colors from sensors (e.g. we have implemented an alternative color reconstruction filter for our Kodak DCS professional digital camera), color space conversion, type conversion (short/float/byte

values, packing, coding, compression), buffer-specific routines such as conversion between different types of depth value representations, and many more.

Furthermore, the IBR project also contains routines for handling large collections of images memory that is shared by multiple processes, and for caching images in shared memory. There is also a number of routines for mapping images to textures, generating image hierarchies, and similar things. Last, but not least, there are some command line tools as well as a new graphical framework for using images and filters via their generic interfaces.

On top of these data structures and algorithmic framework, many different projects have been carried out, and the image interface is also used in further shared projects such as *Antelao* (see Sec. 12.2.6).

12.2.4 Geometric Modeling Utilities

Investigators: Kolja Kähler, Christian Rössl, and Jens Vorsatz

The *Geometric Modeling Utilities (GMU)* library provides support for geometry processing with focus on handling polygonal meshes. The core of the library is a set of generic data structures for the representation and manipulation of triangle meshes. The library is written in `C++`, and uses `template` constructs for parameterizing generic classes like the mesh data structures in order to adapt them to specific applications in an intuitive and efficient way. Current versions support the IRIX MIPSPro and the GNU compilers. *GMU* is organized in several modules:

- low-level math code (e.g. matrix computations, optimization)
- low-level system dependent code (e.g. general purpose IO, parallel processing)
- input/output of triangle meshes (multiple file formats, transparent to the user, easily expandable)
- core triangle mesh data structures
- several supporting modules for high level operations such as mesh decimation, subdivision, parameterization, filtering, discrete differential geometry
- scene graph for visualization of arbitrary data (many existing nodes e.g. for visualizing triangle meshes, easily expandable)
- user interface/widgets for rendering and exploring a scene graph (similar to *Open Inventor* widgets, based on the *Qt* library)

GMU uses the *base* library, *OpenGL* for rendering and *Qt* for user interfaces. It is entirely built by *tmk* and provides documentation and examples.

12.2.5 OGLutil - OpenGL utility library

Investigator: Stefan Brabec

The *OpenGL* utility library (*OGLutil*) started as a collection of helpful functions often needed when developing *OpenGL* applications. Most of the functions provided by this library can be used to debug applications by examining the current state of various *OpenGL* stages (texturing, imaging, buffer configuration, etc.).

One of the most helpful features of *OGLutil* is the extensions parser. Extensions are vendor-specific extensions to the standard *OpenGL* API that provide extended functionality, such as NVIDIA's programmable texture and vertex stages.

Since these extensions are provided by an entry function which returns the function pointer by name, developers have to deal with a great overhead checking and requesting function pointers for specific *OpenGL* extensions. Another problem is that most vendors do not provide *C/C++* include files rather than only the core specifications. *OGLutil* solves theses problems by extracting all necessary information directly from the specification files and automatically generates code for all entry points and definitions. This way a programmer can use extensions without the additional overhead of writing extension-specific setup routines.

Apart from these features, *OGLutil* also includes several utility functions (ported from NVIDIA's OpenGL SDK) to support high-level programming for vertex and texture shaders as well as register combiners. We are currently in contact with NVIDIA to discuss copyright-issues due to a planned public release of *OGLutil*.

OGLutil in its current state supports *Linux*, *IRIX* and *Windows* platforms, which makes cross-platform development of *OpenGL* applications an easy task.

12.2.6 Antelao

Investigators: Marc Stamminger, Annette Scheel, Katja Daubert, Jörg Haber, Hartmut Schirmacher, Hitoshi Yamauchi, Jörg Pütz, Oliver Schwinn, and Markus Weber

Antelao is a rendering system designed for the development and testing of new lighting algorithms. It provides a basis which most algorithms need:

- Scene description. Antelao can read different scene description formats (radiance and mgf, an own antelao file format, and inventor (via a converter)) and builds an internal scene graph. The scene graph can contain basic objects (sphere, polygon, cylinder, triangle, etc.) and more complex ones like Indexed Face Sets.

- Ray Casting Methods. Several applications need to cast rays, for example for Ray Tracing, visibility tests and Monte Carlo Methods in general. Antelao contains intersection tests for all kinds of objects in the scene graph. Acceleration structures like Grid, BSP tree or Octree have also been implemented.

- Camera. The camera can be positioned and moved, and its viewing angle can be changed interactively. Two camera modes exist, examine and walk. In the examine mode, the camera rotates around a fixed point. The walk modus corresponds to an observer walking through the scene and turning his head. Additionally, for animations camera paths can be pre-defined.

- OpenGL Viewer. If a scene is loaded it can be viewed with an OpenGL Viewer. If no lighting algorithm has been applied yet, OpenGL lighting is used to give a first impression of the scene. There are several viewing modes like wireframe, backface culling, and textured / untextured.

- Rendering. To render a chosen camera view either sequential sampling, adaptive sampling, or Directional Coherence Maps can be used. Oversampling is available as well.

Several lighting simulation algorithms and improvements for interactive viewing and single high quality images have already been developed inside the antelao framework:

- Ray Tracing. Caustics can be obtained via Photon Maps.

- Hierarchical Radiosity with Clustering.
- Three Point Transport.
- Hierarchical Radiosity Particle Tracing.
- Corrective Texturing for glossy effects in interactive viewing.
- Final Gather for Finite Element methods.

12.2.7 Tone-Mapping

Investigators: Annette Scheel, Marc Stamminger, and Martin Hochstraßer

This software implements a variety of tone-mapping algorithms. Additionally, it provides some means for image analysis (false color images, luminance histogram) and conversion between different image spaces. It contains the following tone-mapping operators:

- Adjustable Linear Tone-Mapping (min. and max. value and gamma correction)
- Ward 1991
- Tumblin and Rushmeier 1993
- Ward 1997

13 Journal and Conference Activities

13.1 Journal Positions

Hans-Peter Seidel is on the editorial board of the following journals:
International Journal of Shape Modeling (since 2001),
Computer Aided Geometric Design (since 1999),
The Visual Computer (since 1999),
Graphical Models (since 1995), and
Computer Graphics Forum (since 1993).

Leif Kobbelt was co-editor of a special issue on Subdivision Surfaces in *Computer Aided Geometric Design* (2001).

Karol Myszkowski is on the editorial board of the journal *Machine Graphics & Vision* (since 1998).

13.2 Conference and Workshop Positions

13.2.1 Membership in program committees

Philippe Bekaert:

- 7th Spring Conference on Computer Graphics 2001 (SCCG'01) (Budmerice Castle, Slovakia)
- Winter School on Computer Graphics 2001 (WSCG'01) (Pilsen, Czech Republic)

Wolfgang Heidrich:

- 11th Eurographics Workshop on Rendering 2000 (Brno, Czech Republic)
- GameOn 2000 (London, UK)
- 12th Eurographics Workshop on Rendering 2001 (London, UK)
- ACM SIGGRAPH 2001 (Los Angeles, USA)
- Eurographics/Siggraph Workshop on Graphics Hardware 2001 (Los Angeles, USA)
- The International Game Technology Conference 2001 (GTEC'01), (Hongkong, China)

Leif Kobbelt:

- Mini symposium on Mathematical Methods for Curves and Surfaces - Conference on Mathematical Methods for Curves and Surfaces 2000, (Oslo, Norway) - chair
- ACM SIGGRAPH 2000 (New Orleans, USA)
- EUROGRAPHICS 2000 (Interlaken, Switzerland)
- Pacific Graphics 2000 (Hongkong, China)
- Subdivision in Geometric Modeling and Computer Graphics 2000 (Dagstuhl, Germany) - co-chair
- Vision, Modeling, and Visualization 2000 (VMV'00) (Saarbrücken, Germany)
- ACM SIGGRAPH 2001 (Los Angeles, USA)
- Pacific Graphics 2001 (Tokyo, Japan) - co-chair
- EUROGRAPHICS 2001 (Manchester, UK)

Karol Myszkowski:

- 10th Eurographics Workshop on Rendering 1999 (Granada, Spain)
- 11th Eurographics Workshop on Rendering 2000 (Brno, Czech Republic)
- 12th Eurographics Workshop on Rendering 2001 (London, UK) - program co-chair
- 7th Spring Conference on Computer Graphics 2001 (SCCG'01) (Budmerice Castle, Slovakia)
- 7th International Conference Advanced Computer Systems 2001 (ACS '01) (Szczecin, Poland)
- The International Game Technology Conference 2001 (GTEC'01), (Hongkong, China)
- The International Game Technology Conference 2002 (GTEC'02), (Hongkong, China)

Hans-Peter Seidel:

- Pacific Graphics 1999 (Seoul, Korea) - program chair
- Vision, Modeling, and Visualization 1999 (VMV'99) (Erlangen, Germany) - program chair
- Intern. Conf. on Comp. Aided Design and Comp. Graphics (CAD/CG'99) (Shanghai, China)
- EUROGRAPHICS 1999 (Milano, Italy)
- 4th International Conference on Curves and Surfaces 1999 (Saint-Malo, France)
- 10th Eurographics Workshop on Rendering 1999 (Granada, Spain)
- Computer Graphics International 1999 (Canmore, Canada)
- International Workshop on Volume Graphics 1999 (Swansea, UK)
- Shape Modeling International 1999 (SMI'99) (Aizu Wakamatsu, Japan)
- Winter School on Computer Graphics 1999 (WSCG'99) (Pilsen, Czech Republic)
- Image Synthesis and Interactive 3D Graphics 2000 (Dagstuhl, Germany) - co-chair
- Vision, Modeling, and Visualization 2000 (VMV'00) (Saarbrücken, Germany) - conference chair
- Pacific Graphics 2000 (Hongkong, China)
- Eurographics Workshop on Parallel Graphics and Visualization 2000 (Girona, Spain)
- Winter School on Computer Graphics 2000 (WSCG'00) (Pilsen, Czech Republic)
- EUROGRAPHICS 2000 (Interlaken, Switzerland)
- Geometric Modeling and Processing 2000 (GMP'00) (Hongkong, China)
- Vision, Modeling, and Visualization 2001 (VMV'01) (Stuttgart, Germany) - program chair
- 5th ACM Symposium on Solid Modeling 2001 (Ann Arbor, Michigan, USA) - conference chair
- Afrigraph 2001 (Cape Town, South Africa)
- Sibgraphi 2001 (Florianopolis, Brazil)
- Pacific Graphics 2001 (Tokyo, Japan)
- DAGM '01 (Munich, Germany)
- EUROGRAPHICS 2001 (Manchester, UK)
- Computer Graphics International 2001 (Hongkong, China)

- International Workshop on Volume Graphics 2001 (VG'01), (SUNY Stony Brook, USA)
- 7th Spring Conference on Computer Graphics 2001 (SCCG'01) (Budmerice Castle, Slovakia)
- Winter School on Computer Graphics 2001 (WSCG'01) (Pilsen, Czech Republic)
- The International Game Technology Conference 2001 (GTEC'01), (Hongkong, China)
- EUROGRAPHICS 2002 (Saarbrücken, Germany) - program co-chair
- 6th ACM Symposium on Solid Modeling 2002 (Saarbrücken, Germany) - conference co-chair
- International Conference on Vision, Graphics and Image Processing 2002 (ICVGIP'02) (Ahmedabad, India)
- Geometric Modeling and Processing 2002 (GMP'02) (Tokyo, Japan)
- 5th International Conference on Curves and Surfaces 2002 (Saint-Malo, France)
- Shape Modelling International 2002 (SMI'02), (Banff, Canada)
- 7th International Conference on 3D Web Technology 2002 (Web—3D'02), (Tempe, USA)
- Winter School on Computer Graphics 2002 (WSCG'02) (Pilsen, Czech Republic)
- The International Game Technology Conference 2002 (GTEC'02), (Hongkong, China)

13.2.2 Membership in organizing committees

Stefan Brabec:

- Vision, Modeling, and Visualization 2000 (VMV'00) (Saarbrücken, Germany)

Katja Daubert:

- Vision, Modeling, and Visualization 2000 (VMV'00) (Saarbrücken, Germany)

Christian Rössl:

- Vision, Modeling, and Visualization 2000 (VMV'00) (Saarbrücken, Germany)
- 6th ACM Symposium on Solid Modeling 2002 (Saarbrücken, Germany)

Hartmut Schirmacher:

- Vision, Modeling, and Visualization 2000 (VMV'00) (Saarbrücken, Germany)
- EUROGRAPHICS 2002 (Saarbrücken, Germany)

Jens Vorsatz:

- Web3D 2001 Conference (WEB3D'01) (Paderborn, Germany)
- 6th ACM Symposium on Solid Modeling 2002 (Saarbrücken, Germany)

13.3 Invited Talks and Tutorials

13.3.1 Invited Talks

Philippe Bekaert:

- Weighted importance sampling: theory and applications in radiosity, Universität Kaiserslautern, Germany, April 2001.
- Monte Carlo methods for radiosity, INRIA/Loria Nancy, France, June 2001.
- An extensible scene description language, Universitat de Girona, Spain, July 2001.

Sung Woo Choi:

- Medial Axis Transform: Shape, Algorithm, and Stability, Invited Lecture, The Tenth International Colloquium on Numerical Analysis and Computer Science with Applications, Plovdiv, Bulgaria, August, 2001.

Sherif Ghali:

- Raster-Free Computer Graphics, Universität Tübingen, Germany, March 2000.
- Shadow Computation: A Unified Perspective, Eurographics 2000 STAR, Interlaken, Switzerland, August 2000.
- Shadow Computation: Algorithmic and Systems Challenges, Högskolan i Gävle, Sweden, March 2001.
- Analytical Shadows in Computer Graphics, University of Alberta, Canada, May 2001.
- Analytical Visibility and Shadows in Computer Graphics, University of Ottawa, Canada, June 2001.
- Analytical Visibility and Shadows in Computer Graphics, Université de Montréal, Canada, June 2001.

Wolfgang Heidrich:

- Raster-based Rendering, Universität Tübingen, Germany, January 2000.
- Image-based Rendering at the MPI Informatik, Graduiertenkolleg 3D Bildanalyse und -synthese, Universität Erlangen-Nürnberg, Germany, January 2000.
- Interactive Realism through Hardware Acceleration and und Image-based Techniques, Technische Universität Dresden, Germany, January 2000.
- Towards Realistic Materials in Interactive Computer Graphics, The University of British Columbia, Canada, April 2000.
- Towards Realistic Materials in Interactive Computer Graphics, University of Toronto, Canada, April 2000.
- Towards Realistic Materials in Interactive Computer Graphics, University of Utah, Canada, April 2000.
- Hardware Shading: State-of-the-Art and Future Challenges, Keynote Speech, Eurographics/Siggraph Workshop on Graphics Hardware, Interlaken, Switzerland, August 2000.
- Interactive Display of Global Illumination Solutions for Non-Diffuse Environments, Eurographics 2000 STAR, August 2000.

- Hardware Shading: New Developments in the Graphics Pipeline, Invited Lecture, First International Game Technology Conference (GTEC'01), Hongkong, China, January 2001.

Jan Kautz:

- Real-Time BRDF Rendering, GDC Hardcore Technical Seminar, San Francisco, December 1999.
- Image-Based Reconstruction of Spatially Varying Materials, University of Waterloo, Canada, June 2001.
- Advanced Shading for Bump and Displacement Mapping, Universität Bonn, Germany, June 2001.

Leif Kobbelt:

- Efficient representation and modification of polygonal meshes in computer graphics, Gesellschaft für Mathematik und Datenverarbeitung, Bonn, Germany, February 1999.
- Multi-resolution modeling with polygonal meshes, Universität Tübingen, Germany, February 1999.
- Efficient representation and processing of complex 3D models in computer graphics, Universität Darmstadt, Germany, April 1999.
- Efficient representation and processing of complex 3D models in computer graphics, Universität Bonn, Germany, May 1999.
- Efficient representation and processing of complex 3D models in computer graphics, Universität Aachen, Germany, June 1999.
- Free-form modeling with polygonal meshes, Ludwig-Maximilians-Universität, Munich, Germany, June 1999.
- Efficient representation and processing of complex 3D models in computer graphics, Universität Erlangen-Nürnberg, Germany, November 1999.
- New types of subdivision schemes, ETH Zürich, Switzerland, December 1999.
- Efficient processing of complex 3D models in computer graphics, Technische Universität Braunschweig, Germany, May 2000.
- Efficient processing of complex 3D models, Heinz-Maier-Leibnitz award acceptance speach, Universität Jena, Germany, May 2000.
- Geometrical data processing, NATO Science Fellowship Administrator's Meeting (DAAD), Potsdam, Germany, June 2000.
- Efficient processing of complex 3D models in computer graphics, Technische Universität München, Germany, July 2000.
- Multiresolution deformations with dynamic meshes, ETH Zürich, Switzerland, August 2000.

Karol Myszkowski:

- Perception-driven global illumination computation, University Claude Bernard Lyon, France, November 2000.
- Perception-based realistic rendering and animation, Warsaw Institute of Technology, Poland, April 2000.

- Perception-based realistic rendering and animation, Szczecin Technical University, Poland, April 2000.
- Perception-based realistic rendering and animation, University of Applied Sciences, Düsseldorf, Germany, May 2001.
- Perception-based global illumination, rendering, and animation techniques, Fourth International Conference on Human and Computer (HC'2001), Invited Lecture, Aizu Wakamatsu, Japan, September, 2001.

Christian Rössl:

- Using discrete curvature for feature extraction and visualization on polygonal models, University of Magdeburg, June 2000.

Hartmut Schirmacher:

- Lumigraph-Rendering mittels Warping, University of Erlangen-Nürnberg, Germany, March 2000.
- Warping Techniques for Light Fields, University of Waterloo, Canada, May 2000.

Hans-Peter Seidel:

- Mesh Reduction and Interactive Multi-Resolution Modeling on Arbitrary Triangle Meshes, Invited Lecture, CEIG'99, Jaen, Spain, June 1999.
- Hierarchical Methods in Computer Graphics, Universität Tübingen, Germany, June 1999.
- An image is worth more than a thousand words, Lecture Series on Multimedia, RWTH Aachen, Germany, November 1999.
- Mesh Reduction and Interactive Multi-Resolution Modeling on Arbitrary Triangle Meshes, Invited Lecture, 2nd Israel-Korea Bi-National Conference on Geometric Modeling, Seoul, Korea, September 1999.
- Hardware-Shading: State-of-the-Art and Future Challenges, Keynote Speech, Eurographics/Siggraph Workshop on Graphics Hardware, Interlaken, Switzerland, August 2000.
- Efficient Processing of Large 3D Meshes, Invited Section Lecture, DMV-Jahrestagung, Sektion Geometrie, Dresden, Germany, September 2000.
- Efficient Processing of Large 3D Meshes, Invited Lecture, Pacific Graphics 2000, Hongkong, China, October 2000.
- Multiresolution Modeling and Interactive Deformations of Large 3D Meshes, Invited Lecture, IFIP Workshop Deform 2001, Geneve, Switzerland, November 2000.
- Multiresolution Modeling and Interactive Deformations of Large 3D Meshes, Invited Lecture, Game Technology Conference 2001 (GTEC'01), Hongkong, China, January 2001.
- Efficient Processing of Large 3D Meshes, Invited Lecture, 10th International Conference on Approximation Theory, St. Louis, USA, March 2001.
- Efficient Processing of Large 3D Meshes, University of Utah Distinguished Lecture Series, The University of Utah, Salt Lake City, USA, March 2001.
- Efficient Processing of Large 3D Meshes, Invited Lecture, IMA Workshop on Geometric Design, The Institute for Mathematics and its Applications (IMA), Minneapolis, USA, April 2001.

- Efficient Processing of Large 3D Meshes, Invited Lecture, Shape Modeling International 2001 (SMI'01), Genua, Italy, May 2001.
- Efficient Processing of Large 3D Meshes, Invited Lecture, CGC Workshop ETH Zürich, Monte Verita, Italy, May 2001.
- Efficient Processing of Surfaces using Triangle Meshes, Universität Bielefeld, Germany, May 2001.
- A Framework for the Acquisition, Processing, and Interactive Display of High-Quality 3D Models, Invited Lecture, IMA Workshop Computer Graphics, The Institute for Mathematics and its Applications (IMA), Minneapolis, USA, May 2001.
- Multiresolution Modeling and Multiresolution Smoothing of Large 3D Meshes, European Summer School on Principles of Multiresolution in Geometric Modeling, Munich, Germany, August 2001.

Marc Stamminger:

- Finite elements for glossy global illumination, INRIA/Loria Nancy, France, September 1999.
- Glossy global illumination with finite elements, Study Day on Physically Based Rendering, Leuven, Belgium, December 1999.
- Global illumination with finite elements, Universitat de Barcelona, Spain, May 2000.
- Walkthroughs with corrective texturing, Universitat de Girona, Spain, May 2000.

Jens Vorsatz:

- Hierarchical Meshes and Remeshing, Graduate Research Center, University of Erlangen, June 2001.

13.3.2 Conference Tutorials

Philippe Bekaert:

- Stochastic radiosity: doing radiosity without form factors, ACM SIGGRAPH 2001 Tutorial, Los Angeles, USA, August 2001.
- RenderPark: a test-bed system for global illumination, ACM SIGGRAPH 2001 Creative Applications Lab Hands-On Session, Los Angeles, USA, August 2001.

Mario Botsch:

- Geometric Modeling Based on Polygonal Meshes, Eurographics 2000 Tutorial, Interlaken, Switzerland, September 2000.

Stefan Brabec:

- Advanced Lighting for Interactive Applications, Tutorial GTEC'01 Hongkong, China, January 2001.

Sherif Ghali:

- Migrating to an Object-Oriented Graphics API, ACM SIGGRAPH 2000 Tutorial, New Orleans, USA, July 2000.

- Object Space Visibility, ACM SIGGRAPH 2001 Tutorial, Los Angeles, USA, August 2001.

Michael Goesele:

- A Framework for the Acquisition, Processing, Transmission, and Interactive Display of High Quality 3D Models on the Web, Web3d Conference 2001 Tutorial, Paderborn, Germany, February 2001.
- A Framework for the Acquisition, Processing, and Interactive Display of High Quality 3D Models, DAGM 2001 Tutorial, Munich, Germany, September 2001.

Wolfgang Heidrich:

- Advanced Graphics Programming with OpenGL and Extensions, Eurographics 1999 Tutorial, Milano, Italy, September 1999.
- Approaches for Procedural Shading on Graphics Hardware, ACM SIGGRAPH 2000 Tutorial, New Orleans, USA, July 2000.

Kolja Kähler:

- Geometric Modeling Based on Polygonal Meshes, Eurographics 2000 Tutorial, Interlaken, Switzerland, September 2000.

Jan Kautz:

- Advanced Lighting for Interactive Applications, Tutorial GTEC'01 Hongkong, China, January 2001.
- A Framework for the Acquisition, Processing, Transmission, and Interactive Display of High Quality 3D Models on the Web, Web3d Conference 2001 Tutorial, Paderborn, Germany, February 2001.

Leif Kobbelt:

- Subdivision for Modeling and Animation, ACM SIGGRAPH 1999 Tutorial, Los Angeles, USA, August 1999.
- Subdivision for Modeling and Animation, ACM SIGGRAPH 2000 Tutorial, New Orleans, USA, July 2000.
- Geometric Modeling Based on Polygonal Meshes, Eurographics 2000 Tutorial, Interlaken, Switzerland, September 2000.

Hendrik Lensch:

- A Framework for the Acquisition, Processing, Transmission, and Interactive Display of High Quality 3D Models on the Web, Web3d Conference 2001 Tutorial, Paderborn, Germany, February 2001.
- A Framework for the Acquisition, Processing, and Interactive Display of High Quality 3D Models, DAGM 2001 Tutorial, Munich, Germany, September 2001.

Karol Myszkowski:

- Image Quality Metrics, ACM SIGGRAPH 2000 Tutorial, New Orleans, USA, July 2000.

- Seeing is Believing: Reality Perception in Modeling, Rendering, and Animation, ACM SIGGRAPH 2001 Tutorial, Los Angeles, USA, August 2001.

Christian Rössl:

- Geometric Modeling Based on Polygonal Meshes, Eurographics 2000 Tutorial, Interlaken, Switzerland, September 2000.

Robert Schneider:

- Geometric Modeling Based on Polygonal Meshes, Eurographics 2000 Tutorial, Interlaken, Switzerland, September 2000.

Hans-Peter Seidel:

- A Framework for the Acquisition, Processing, Transmission, and Interactive Display of High Quality 3D Models on the Web, Web3d Conference 2001 Tutorial, Paderborn, Germany, February 2001.
- A Framework for the Acquisition, Processing, and Interactive Display of High Quality 3D Models, DAGM 2001 Tutorial, Munich, Germany, September 2001.

Marc Stamminger:

- Advanced Radiosity - Complex Scenes and glossy Reflections, Eurographics 1999 Tutorial, Milano, Italy, September 1999.
- Advanced Lighting for Interactive Applications, Tutorial GTEC'01 Hongkong, China, January 2001.

Jens Vorsatz:

- Geometric Modeling Based on Polygonal Meshes, Eurographics 2000 Tutorial, Interlaken, Switzerland, September 2000.

14 Teaching Activities

Winter Semester 1999/2000

Lectures:
Bildbasierte Modelierung und Bildsynthese (W. Heidrich)
Geometrische Modelierung (H.-P. Seidel)
Komplexe Polygonmodelle (L. Kobbelt)

Project Classes:
Fortgeschrittene Graphikprogrammierung (H.-P. Seidel, M. Stamminger)

Summer Semester 2000

LECTURES:
Flächenmodelierung (H.-P. Seidel, L. Kobbelt, F. Zeilfelder)
Hardware-unterstütztes Rendering (W. Heidrich)

SEMINARS:
Objektmodelierung (H.-P. Seidel, L. Kobbelt)
Advanced Graphics (H.-P. Seidel, P. Slusallek, L. Kobbelt, W. Heidrich, M. Stamminger)

PROJECT CLASSES:
Fortgeschrittene Graphikprogrammierung (H.-P. Seidel, M. Stamminger)

Winter Semester 2000/2001

LECTURES:
Geometrische Modelierung (H.-P. Seidel)
Human Perception Issues in Rendering and Animation (K. Myszkowski)

SEMINARS:
Modelierung und Animation von Gesichtern (J. Haber, K. Kähler, H.-P. Seidel)

PROJECT CLASSES:
3D-Reconstruction aus Bildern, Bildfolgen und 3D-Scans (M. Goesele, H. Lensch, C. Rössl, H.-P. Seidel)

Summer Semester 2001

LECTURES:
Gesichtsmodelierung und -animation (J. Haber)
Interactive Rendering and Animation (K. Myszkowski, J. Kautz)
Fortgeschrittene 3D-Modelierung mit Polygonnetzen (H.-P. Seidel)

SEMINARS:
3D-Rekonstruktion aus Bildern, Bildfolgen und 3D-Scans (M. Goesele, H. Lensch, C. Rössl, H.-P. Seidel)
Advanced Graphics (H.-P. Seidel, P. Slusallek, K. Myszkowski, J. Kautz)

PROJECT CLASSES:
Fortgeschrittene Grafikprogrammierung (J. Kautz, A. Scheel, H. Schirmacher, H.-P. Seidel)

Diploma Theses

During the last two years, the following 10 diploma theses have been completed under guidance of members of our group.

Irene Albrecht: Speech Synchronization for Physically-based Human Face Models, 2001.
Mario Botsch: 3-D Gesichtsmodelierung zur Operationsplanung, 1999.
Stefan Brabec: Lumigraphen von Lichtquellen, 1999.
Jan Kautz: Hardware Rendering with Bidirectional Reflectances, 1999.
Hendrik Lensch: Techniques for Hardware-Accelerated Light Field Rendering, 1999.
Jörg Pütz: Progressives Rendering mit Texturen unter Ausnutzung von Kohärenz, 2000.
Christian Rössl: Semi-Automatische Methoden für die Rekonstruktion von CAD-Modellen aus Punktdaten, 1999.
Pascal Schüler: 3-D Rekonstruktion aus kalibrierten Bildaten mittels Space Carving, 2001.
Takehiro Tawara: Applications of Image-Based Rendering to High-Quality Walkthrough Animation, 2000.
Christian Theobalt: Navigation on a Mobile Robot, 2001.

15 Dissertations, Habilitations, and Offers for Faculty Positions

15.1 Dissertations

Completed:
Robert Schneider: Efficient High-Quality Fairing of Discrete Curves and Surfaces

In preparation:
Stefan Brabec: Shadow Techniques for Real-time and Interactive Applications
Katja Daubert: Illuminating Microgeometry for Rendering Fabric and Textiles
Michael Goesele: Image-based Acquisition Techniques
Won-Ki Jeong: Reconstruction of High-Quality Subdivision Surfaces from Point Clouds
Jan Kautz: Realistic and Interactive Rendering of Arbitrary Materials
Kolja Kähler: A Head Model with Anatomical Structure for Facial Animation
Hendrik Lensch: 3D Object Acquisition including Reflection Properties
Ming Li: Towards Real-Time Image-Based Modeling and Rendering
Christian Rössl: Robust and Efficient Processing of Large Polygonal Meshes
Annette Scheel: Interactive and High-Quality Rendering of Finite Element Global Illumination Solutions
Hartmut Schirmacher: Efficient Acquisition, Representation, and Rendering of Light Fields
Takehiro Tawara: Spatio-Temporal Techniques for Efficient Rendering of Animation Sequences
Christian Theobalt: Efficient Video-based Analysis of Highly Articulated Motion
Jens Vorsatz: Interactive Multiresolution Modeling with Dynamic Triangle Meshes

15.2 Habilitations

Leif Kobbelt

15.3 Offers for Faculty Positions

Wolfgang Heidrich
University of British Columbia, Canada, 2000.

Cyril Soler
INRIA Rhone-Alpes, Grenoble, France, 2000.

Leif Kobbelt
RWTH Aachen, Aachen, Germany, 2001.

Sherif Ghali
University of Alberta, Canada, 2001.

16 Grants and Cooperations

16.1 Projects funded by the European Union (EU)

The following projects are funded by the European Union (EU).

16.1.1 Multiresolution in Geometric Modeling (MINGLE)

The goal of the EU research project MINGLE is the investigation and development of efficient and robust methods for generating, storing, and manipulating multiresolution models based on triangulated data sets. Applications for such techniques include fast rendering (e.g. in real-time flight simulation), editing and compression of both geometric models and images, surface illumination, computer animation, and scientific visualization (such as weather simulations over the earth). The joint activities of the training program include lecture series, workshops, short-term exchange visits, and software management. The work on MINGLE started in 2000, and the project will run until 2003. It involves nine European research teams, representing six different countries, from universities, research centers, and industry.

Partners and group leaders of the project are:

SINTEF Applied Mathematics, Oslo, Norway (Dr. M. Floater, Dr. E. Quak, Prof. T. Lyche)
Tel Aviv University, Tel Aviv, Israel (Prof. N. Dyn, Prof. D. Levin, Prof. D. Cohen-Or)
Munich University of Technology, Munich, Germany (Dr. A. Iske, Dr. J. Prestin)
Israel Institute of Technology, Haifa, Israel (Prof. C. Gotsman)
Laboratoire de Modlisation et Calcul, Grenoble, France (Dr. G. Bonneau, Dr. S. Hahmann)
University of Cambridge Computer Laboratory, Cambridge, United Kingdom (Dr. N. Dodgson, Dr. M. Sabin)
Dept. of Comp. and Inf. Sciences, Univ. of Genova, Genova, Italy (Prof. L. De Floriani, Prof. E. Puppo)
Systems in Motion AS, Oslo, Norway (M. Kintel)
Max-Planck-Institut für Informatik (Prof. H.-P. Seidel, PD Dr. L. Kobbelt, Dr. F. Zeilfelder)

16.1.2 Realistic Simulation of Light in General Environments (SIMULGEN)

The goal of the ESPRIT open long term research project SIMULGEN is to extend global illumination methods so that they can handle all kinds of reflections, including participating media like fog. It has two main objectives: interactive viewing of global illumination solutions and the creation of single high quality images and animations.

The first phase of the project started in 1997 with the goal to investigate the extension of Finite Element methods–which can traditionally only handle diffuse reflections–towards general reflection properties. The first phase was successfully completed and was followed by a second phase which started in 1999 and will have a duration of 36 months. The current research concentrates on a combination of Finite Element methods and Monte Carlo approaches as well as integration of human perception to steer the effort of the computation.

Partners and group leaders of the project are:

University of Girona, Spain (Prof. X. Pueyo, Dr. I. Martin)
University Joseph Fourier de Grenoble, France (Dr. N. Holzschuh, Prof. F.X. Sillion)
LightWork Design Ltd, Sheffield, United Kingdom (Dr. N. Gatenby)

Max-Planck-Institut für Informatik (Prof. H.-P. Seidel, Dr. K. Myszkowski)

16.1.3 Towards Interactive Predictive Photo-Realistic Image Synthesis of Large 3-Dimensional Environments (Marie Curie PostDoc Fellowship)

This project is associated with a Marie-Curie post doctoral fellowship taken by Philippe Bekaert. The project focuses on photo-realistic image synthesis: the generation by computer of images of virtual 3-dimensional scenes that would be indistinguishable from corresponding real scenery, when viewed by a human under corresponding conditions. Current algorithms for photo-realistic image synthesis suffer various shortcomings that prevent their common use in a wide range of application areas nowadays. They make crude assumptions about surface-material properties, they need high computation times, they have excessive storage requirements for real-world model sizes, they are not robust or their usage requires expert knowledge. Based on a recent analysis, new algorithms and data structures will be developed, in which these shortcomings are effectively addressed.

Partners and group leaders of the project are:

Max-Planck-Institut für Informatik (Prof. H.-P. Seidel, Dr. Ph. Bekaert)

16.2 Projects funded by GIF

The following project is funded by the German-Israeli Foundation for Scientific Research and Development (GIF).

16.2.1 Compact Representation and Efficient Processing of Very Large Triangle Meshes

The goal of this research project is to develop compact representations for static, dynamic and progressive meshes and to investigate how good these representations are relative to the best possible. This involves developing effective coding algorithms and establishing lower bounds. Being able to render from compressed form in order to preserve memory-chip bandwidth is an important advantage, which will enable on-chip geometry transformations in next-generation low-end graphics cards.

Partners and group leaders of the project are:

Tel Aviv University, Tel Aviv, Israel (Prof. D. Cohen-Or, Prof. D. Levin)
Technion, Haifa, Israel (Prof. C. Gotsman)
Universität Tübingen, Tübingen, Germany (Prof. W. Strasser)
Universität Bonn, Bonn, Germany (Prof. R. Klein)
RWTH Aachen, Aachen, Germany (Prof. L. Kobbelt)
Max-Planck-Institut für Informatik, Saarbrücken, Germany (Prof. H.-P. Seidel)

16.3 Projects funded by DAAD

The following cooperation is funded by the German Academic Exchange Service (DAAD - Deutscher Akademischer Austauschdienst).

16.3.1 Global Illumination using Simplified Geometry

A large problem with global illumination methods is the high memory cost for storing the large and often highly detailed scene descriptions. One category of global illumination algorithms, the finite element methods, already operates on multiple levels of detail for the scene. The goal of this project was to investigate how well the category of stochastic methods can be combined with simplified geometry. The main focus was on, but not restricted to enhancing photon maps to operate on simplified geometry, thereby comparing efficiency to finite element methods.

Partners and group leaders of the project are:

Max-Planck-Institut für Informatik (Prof. H-P. Seidel)
Universitat de Girona, Spain (Prof. X. Pueyo)
UPC Barcelona, Spain (Prof. P. Brunet)

16.4 Projects funded by DFG

The following projects are funded by the German National Science Foundation (DFG - Deutsche Forschungsgemeinschaft).

16.4.1 Acquisition, Compression, and Transmission of Digital 3D Models and Their Interactive Display

The project is part of the DFG-Schwerpunktprogramm "Verteilte Verarbeitung und Vermittlung digitaler Dokumente". The topic of this project is the development and implementation of techniques and tools for the acquisition, compression, transmission, and display of complex digital 3D models and their integration into digital documents. In particular, we are working on the creation of a hardware/software environment for the acquisition of high quality 3D models, capturing both geometry and material properties.

Partners and group leaders of the project are:

Max-Planck-Institut für Informatik (Prof. Dr. H.-P. Seidel, Dr. W. Heidrich, Dr. J. Haber, Dipl. Inf. H. Lensch)

16.4.2 Robust Transmission and Adaptive Display of Complex 3D Models and 3D Animations for Integration in Digital Documents

This project is also part of the DFG-Schwerpunktprogramm "Verteilte Verarbeitung und Vermittlung digitaler Dokumente". The project focuses on the generalization of successful methods for the progressive transmission of static 3D models to dynamic 3D animations. This will make it possible, e.g., to directly integrate the results of a dynamic simulation into digital documents.

Partners and group leaders of the project are:

Max-Planck-Institut für Informatik (PD Dr. L. Kobbelt)

16.4.3 Graduiertenkolleg 3D Image Analysis and Synthesis

The Graduiertenkolleg (Graduate Research Center) "3D Image Analysis and Synthesis" comprises a program of collaborative research and advanced studies at the University of Erlangen with special emphasis on problems of 3D image acquisition, computer vision, 3D computer graphics, and selected

applications ranging from medicine to manufacturing. The graduate research center supports a PhD thesis on mesh modeling at MPI.

Partners and group leaders of the project are:

University of Erlangen-Nuremberg (Profs. Greiner, Niemann, Fahlbusch, Gerhäuser, Häusler, Hirschfelder, Kalender, Rüde, Seitzer, Steffen)
Max-Planck-Institut für Informatik (Prof. H.-P. Seidel)
Stanford University (Prof. Dr. B. Girod)

16.4.4 Virtual BioLab

The computer graphics group also participates in the "Virtual Biolab". The project – a common proposal of the Max-Planck-Institute for Computer Science (MPI), the Saarland University, and the German Research Center for Artificial Intelligence (DFKI), under the scientific leadership of the Bioinformatics group of AG1 – was one of the five winners of the "Initiative Bioinformatics" of the DFG (see Section 14.2.1). Within the project we will work on the processing and interactive display of very large data sets within the proposed biological information system.

The partners and group leaders are:

Universität des Saarlands (Profs. H. Lenhof, R. Seidel, R. Bernhardt, F. Giffhorn, R. Hartmann, E. Heinzle, J. Hüttermann, J. Jose, C.-M. Lehr, A. Louis, E. Meese, M. Springborg, G. Weikum)
Max-Planck-Institut für Informatik (Profs. T. Lengauer, K. Mehlhorn, H.-P. Seidel)
Deutsches Forschungszentrum für Künstliche Intelligenz (DFKI) (Prof. W. Wahlster)

16.5 Projects funded by BMBF

The following project is funded by the German Ministry of Education, Science, and Technology (BMBF - Bundesminister für Bildung und Forschung).

16.5.1 OpenSG PLUS

OpenSG PLUS is a project that is funded by the German Ministry for Research and Education (BMBF). The duration of the project is from February 2001 to December 2003. In this project 9 german research institutions (universities and independent research groups) from the area of 3D computer graphics are cooperating to develop basic technology for OpenSG (OpenSG (Open Source Scenegraph) a scenegraph-based rendering API), define architectures and data structures together and provide important components through R&D work for new application areas.

Partners and group leaders of the project are:

OpenSG Forum, Darmstadt (Dr. S. Müller, Dipl. Inf. D. Reiners)
Fraunhofer Institut für Graphische Datenverarbeitung, Darmstadt (Dr. C. Busch)
Max-Planck-Institut für Informatik (Dipl. Inf. J. Kautz)
Universität Bonn, Bonn (Prof. Dr. R. Klein)
Universität Stuttgart, Stuttgart (Prof. Dr. T. Ertl)
Universität Tübingen, Tübingen (Prof. Dr. W. Strasser)
RWTH Aachen, Aachen, (Prof. Dr. L. Kobbelt)
TU Braunschweig, Braunschweig (Prof. Dr. D. Fellner)
TU Darmstadt, Darmstadt (Dipl. Inf. M. Alexa)
Zentrum für Graphische Datenverarbeitung, Darmstadt (Dipl. Inf. J. Behr)

16.6 Projects funded by FhG

The following project is funded by Fraunhofer-Gesellschaft (FhG).

16.6.1 High Dynamic Range Images and Interactive Tone-Mapping

This project was done in cooperation with the FhG/IGD and BMW. The goal of the project was the simulation of illumination inside a car at night. The global illumination solution was computed by a radiosity simulation and visualized in a CAVE. The low levels of illumination at night and the high dynamic range of the illumination required a Tone-Mapping step, which was developed at MPI. Furthermore, for validation of the results of the illumination simulation, a method to capture accurate luminance values, so called high dynamic range images, from real car interiors has been developed and tested.

Partners and group leaders of the project are:

FhG, Institut für Graphische Datenverarbeitung, Darmstadt (Dr. S. Müller)
BMW AG, München (T. Harrasim)
Max-Planck-Institut für Informatik (Prof. H.-P. Seidel, Dr. M. Stamminger)

16.7 Cooperations with Industry

The following projects are funded by industry.

16.7.1 Siemens: Detailed Reconstruction of Surfaces from the Simulation of a Milling Process

In this project we investigate the use of voxel representations in the context of NC milling. Since volumetric data is usually sampled on a regular grid with a given step width, we often observe severe alias artifacts at sharp features on the extracted surfaces. As part of this project we have developed a new technique for surface extraction that performs feature sensitive sampling and thus reduces these alias effects while keeping the simple algorithmic structure of the standard Marching Cubes algorithm. Applications of this work range from CSG modeling and NC milling simulation to surface reconstruction and remeshing of polygonal models. This project is funded by Siemens AG, Erlangen.

The partners and group leaders are:

Max-Planck-Institut für Informatik (PD Dr. L. Kobbelt, Prof. Dr. H.-P. Seidel)
Siemens AG, Nürnberg (Dr. W. Friedrich)

16.7.2 BMW: Efficient Generation of and Modeling with Polygonal Meshes for Fluid Flow Simulations

In numerical fluid flow simulations within the car industry, the underlying car geometry is usually represented by a polygonal mesh. Hence this approach usually requires several conversion steps (surface measurement, reverse engineering (reconstruction of a spline based CAD model, tessellation of the resulting spline surface, generation of a consistent mesh representation). The main goal of this project is to reduce the number of conversion steps between different representations by

developing new techniques for mesh reconstruction from measurement data, for mesh editing, and for remeshing. This project is funded by BMW.

The partners and group leaders are:
Max-Planck-Institut für Informatik (PD Dr. L. Kobbelt)
BMW AG, München (Dr. W. Bartelheimer)

16.7.3 ProMePro: Progressive Meshes for rapid and virtual Prototyping

The goal of this project has been the evaluation of state-of-the-art geometry processing algorithms on hierarchical and progressive mesh representations of geometric models that come from rapid and virtual prototyping applications. This includes the identification and elimination of topological artifacts, the reduction of mesh complexity while keeping a certain approximation tolerance, and remeshing for quality enhancement. This project has been funded by the participating partners and Bayern Innovativ.

The partners and group leaders are:

Alphaform, Laser Modell- und Formenbau GmbH, Feldkirchen (Dr. M. Girardet)
Janet, Gesellschaft für interaktive Medien GmbH, München (Dr. W. Posch)
Max-Planck-Institut für Informatik (PD Dr. L. Kobbelt)

16.7.4 Motorola: Network-Integrated Multimedia Middleware

Multimedia is still not well supported in most Unix environments in contrast to other OS-platforms where several multimedia infrastructures are established (e.g. Apple Quicktime, MS DirectShow). However these models adopt a PC-centric approach, where all multimedia processing takes place within a single computer. The network is, at best, used for streaming data transmission.

Since there is a strong trend towards networked multimedia devices (like networked cameras, audio devices, settop-boxes, game consoles, and PCs) these centralized approaches are becoming obsolete. The goal of our work is to develop a multimedia middleware, which considers the network as an integral part and enables the intelligent use of devices distributed across a network.

The goal of this project is to design and implement a network-integrated multimedia infrastructure for Linux as well as other operating systems. Since many low-level parts of such a system already exist (like en-/decoders, de-/compression, multimedia networking APIs, CORBA, and others), our primary goal is the development of an architecture, which integrates these usually isolated modules. This unified architecture will offer a simple and easy to use interface for applications to integrate multimedia functionality. The result of this work will be made available as OpenSource.

The project started in February 2001. Most of the core local functionality has been implemented already, including a DVD player, a DVB/TV player, and many plugins for audio and other multimedia processing. Currently work concentrates on implementing the network layer of the infrastructure, adding more audio and video plugins, and implementing more example applications. We expect a first public release early next year. This project is funded by Motorola, Germany.

Partners and group leaders of the project are:

Universität des Saarlandes, Lehrstuhl Computergraphik (Prof. Ph. Slusallek)
Max-Planck-Institut für Informatik (Prof. H.-P. Seidel)
Motorola Imaging & Entertainment, Munich, Germany

17 Publications

Books

[1] T. Ertl, B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors. *Proceedings of Vision, Modeling, and Visualization (VMV-01)*, Stuttgart, Germany, 2001.

[2] B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors. *Proceedings of Vision, Modeling, and Visualization (VMV-00)*, Saarbrücken, Germany, 2000.

[3] B. Girod, H. Niemann, and H.-P. Seidel, editors. *Proceedings of the 4th Conference on Vision, Modeling, and Visualization (VMV-99)*, Erlangen, Germany, 1999. SFB 603 and Graduate Research Center, in coorperation with IEEE Signal Processing Society and Gesellschaft für Informatik GI, infix.

[4] S. Gortler and K. Myszkowski, editors. *Rendering Techniques 2001*, Springer Computer Science, Wien, New York, August 2001. Eurographics, Springer.

[5] M.-S. Kim and H.-P. Seidel, editors. *Proceedings of Pacific Graphics '99*, Seoul, Korea, 1999. IEEE Computer Society.

[6] K. Myszkowski. *Efficient and Predictive Realistic Image Synthesis.* Oficyna Wydawnicza Politechniki Warszawskiej, Warsaw, Poland, May 2001. ISSN 0137-2343.

In Journals and Refereed Conference Proceedings

[1] P. Bekaert and H.-P. Seidel. A theoretical comparison of Monte Carlo radiosity algorithms. In T. Ertl, B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling and Visualisation (VMV-01)*, Stuttgart, Germany, November 2001. To appear.

[2] S. Bischoff, L. Kobbelt, and H.-P. Seidel. Towards hardware implementation of Loop subdivision. In S. N. Spencer, editor, *Proceedings of 2000 SIGGRAPH/Eurographics Workshop on Graphics Hardware*, pages 41–50, Interlaken, Switzerland, August 2000. ACM SIGGRAPH, ACM Press.

[3] M. Botsch, C. Rössl, and L. Kobbelt. Feature sensitive sampling for interactive remeshing. In B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling, and Visualization (VMV-00)*, pages 129–136, Saarbrücken, Germany, November 2000.

[4] S. Brabec and H.-P. Seidel. Extended light maps. In Q. Mehdi and N. Gough, editors, *GAME-ON 2000 / 1st International Conference on Intelligent Games and Simulation*, pages 10–13, London, United Kingdom, November 2000. Society for Computer Simulation International.

[5] S. Brabec and H.-P. Seidel. Hardware-accelerated rendering of antialiased shadows with shadow maps. In H. H. S. Ip, N. Magnenat-Thalmann, R. W. H. Lau, and T.-S. Chua, editors, *Proceedings Computer Graphics International 2001*, pages 209–214, Hong Kong, China, July 2001. IEEE Computer Society.

[6] S. Campagna, L. Kobbelt, R. Schneider, H.-P. Seidel, and J. Vorsatz. Mesh reduction and interactive multiresolution modeling on arbitrary triangle meshes. In J. Zra, editor, *Proceedings of the 15th Spring Conference on Computer Graphics (SCCG-99)*, pages 34–44, Budmerice, Slovakia, 1999. Comenius University.

[7] S. Campagna, L. Kobbelt, and H.-P. Seidel. Directed edges - a scalable representation for triangle meshes. *Journal of Graphics Tools*, 3(4):1–12, 1999.

[8] S. W. Choi and S.-W. Lee. Stability analysis of medial axis transform under relative Hausdorff distance. In A. Sanfeliu, J. J. Villanueva, M. Vanrell, R. Alquzar, T. Huang, and J. Serra, editors, *Proceedings of the 15th International Conference on Pattern Recognition (ICPR-00)*, pages 139–142, Barcelona, Spain, September 2000. IEEE Computer Society.

[9] S. W. Choi and H.-P. Seidel. Hyperbolic hausdorff distance for medial axis transform. *Graphical Models*, 2001. To appear.

[10] S. W. Choi and H.-P. Seidel. One-sided stability of MAT and its applications. In T. Ertl, B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling, and Visualization (VMV-01)*, Stuttgart, Germany, November 2001. To appear.

[11] S. W. Choi and H.-P. Seidel. One-sided stability of medial axis transform. In B. Radig and S. Florczyk, editors, *Proceedings of Pattern Recognition, 23rd DAGM Symposium*, volume 2191 of *Lecture Notes in Computer Science*, pages 132–139, München, Germany, 2001. Springer.

[12] M. Daehlen, T. Lyche, K. Morken, R. Schneider, and H.-P. Seidel. Multiresolution analysis oder triangles, based on quadratic Hermite interpolation. *Journal of Computational and Applied Mathematics*, 119:97–114, 2000.

[13] K. Daubert, H. Lensch, W. Heidrich, and H.-P. Seidel. Efficient cloth modeling and rendering. In S. Gortler and K. Myszkowski, editors, *Rendering Techniques 2001, Proceedings of the 12th Eurographics Workshop on Rendering*, pages 63–70, London, Great Britain, 2001. Eurographics, Springer.

[14] O. Davydov, G. Nürnberger, and F. Zeilfelder. Cubic spline interpolation on nested polygon triangulations. In A. Cohen, C. Rabut, and L. L. Schumaker, editors, *Curve and Surface Fitting, Saint-Malo 1999*, Innovations in Applied Mathematics, pages 161–170, Saint Malo, June 2000. Vanderbilt University Press.

[15] F. Drago and K. Myszkowski. Validation proposal for global illumination and rendering techniques. *Computers & Graphics*, 25(3):511–518, 2001.

[16] S. Ershov, K. Kolchin, and K. Myszkowski. Rendering pearlescent appearance based on paint-composition modeling. In A. Chalmers and T.-M. Rhyne, editors, *Computer Graphics Forum (Proceedings of Eurographics 2001)*, volume 20 of *Computer Graphics Forum*, pages 227–238, Manchester, UK, September 2001. Eurographics, Blackwell.

[17] D. Fellner, J. Haber, S. Havemann, L. Kobbelt, H. Lensch, G. Müller, I. Peter, R. Schneider, H.-P. Seidel, and W. Straer. Beiträge der Computergraphik zur Realisierung eines verallgemeinerten Dokumentbegriffs. *Informationstechnik und Technische Informatik*, 42(6):8–18, December 2000.

[18] S. Ghali, E. Fiume, and H.-P. Seidel. Object-space, connectivity-preserving, shadow computation. In B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling, and Visualization (VMV-00)*, pages 281–289, Saarbrücken, Germany, November 2000.

[19] M. Goesele, W. Heidrich, and H.-P. Seidel. Color calibrated high dynamic range imaging with ICC profiles. In *Proceedings of the 9th Color Imaging Conference*, Scottsdale, USA, 2001.

[20] M. Goesele, W. Heidrich, and H.-P. Seidel. Entropy-based dark frame subtraction. In *Proceedings of PICS 2001: Image Processing, Image Quality, Image Capture, Systems Conference*, pages 293–298, Montreal, Canada, April 2001. The Society for Imaging Science and Technology (IS&T).

[21] M. Goesele, H. Lensch, W. Heidrich, and H.-P. Seidel. Building a photo studio for measurement purposes. In B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling, and Visualization (VMV-00)*, pages 231–238, Saarbrücken, Germany, November 2000.

[22] M. Goesele and W. Stuerzlinger. Semantic constraints for scene manipulation. In J. Zra, editor, *Proceedings of the 15th Spring Conference on Computer Graphics (SCCG-99)*, pages 140–146, Budmerice, Slovakia, 1999. Comenius University.

[23] J. Haber, K. Myszkowski, H. Yamauchi, and H.-P. Seidel. Perceptually Guided Corrective Splatting. In A. Chalmers and T.-M. Rhyne, editors, *Computer Graphics Forum (Proceedings of Eurographics 2001)*, volume 20 of *Computer Graphics Forum*, pages 142–152, Manchester, UK, September 2001. Eurographics, Blackwell. To appear.

[24] J. Haber and H.-P. Seidel. Using an Enhanced LBG Algorithm to Reduce the Codebook Error inVector Quantization. In *Proceedings of Computer Graphics International 2000 (CGI-00)*, pages 99–104, Geneva, Switzerland, June 2000. IEEE Computer Society.

[25] J. Haber, M. Stamminger, and H.-P. Seidel. Customized Object Hierarchies. In B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling, and Visualization (VMV-00)*, pages 255–262, Saarbrücken, Germany, November 2000.

[26] J. Haber, M. Stamminger, and H.-P. Seidel. Enhanced Automatic Creation of Multi-Purpose Object Hierarchies. In B. A. Barsky, Y. Shinagawa, and W. Wang, editors, *Proceedings of the 8th Pacific Conference on Computer Graphics and Applications (PG-00)*, pages 52–61;437, Hong Kong, China, October 2000. IEEE Computer Society.

[27] J. Haber, F. Zeilfelder, O. Davydov, and H.-P. Seidel. Smooth Approximation and Rendering of Large Scattered Data Sets. In *Proceedings of IEEE Visualization 2001*, San Diego, USA, October 2001. IEEE Computer Society. To appear.

[28] W. Heidrich, S. Brabec, and H.-P. Seidel. Soft shadow maps for linear lights. In B. Proche and H. Rushmeier, editors, *Rendering Techniques 2000, Proceedings of the 11th Eurographics Workshop on Rendering*, Springer Computer Science, pages 269–280, Brno, Czech Republic, June 2000. Eurographics, Springer.

[29] W. Heidrich, K. Daubert, J. Kautz, and H.-P. Seidel. Illuminating micro geometry based on precomputed visibility. In K. Akeley, editor, *Computer Graphics (SIGGRAPH-00): Conference proceedings*, volume 2000 of *Computer Graphics Proceedings, Annual Conference Series*, pages 455–464, New Orleans, USA, July 2000. ACM Siggraph, ACM Press.

[30] W. Heidrich, H. Lensch, M. Cohen, and H.-P. Seidel. Light field techniques for reflections and refractions. In D. Lischinski and G. W. Larson, editors, *Rendering Techniques '99: Proceedings of the 10th Eurographics Workshop on Rendering (EGRW-99)*, pages 187–196, Granada, Spain, June 1999. Eurographics, Springer.

[31] W. Heidrich, H. Schirmacher, H. Kück, and H.-P. Seidel. A warping-based refinement of Lumigraphs. In V. Skala, editor, *Proceedings of the 7th International Conference in Central Europe on Computer Graphics, Visualization and Interactive Digital Media (WSCG-99)*, pages 102–109, Plzen, Czech Republic, February 1999.

[32] W. Heidrich and H.-P. Seidel. Realistic, hardware-accelerated shading and lighting. In A. Rockwood, editor, *Computer Graphics (SIGGRAPH-99): Conference proceedings*, pages 171–178, Los Angeles, USA, August 1999. ACM Siggraph, ACM.

[33] W. Heidrich, R. Westermann, H.-P. Seidel, and T. Ertl. Applications of pixel textures in visualization and realistic image synthesis. In J. Hodgins and J. Foley, editors, *Proceedings of the 1999 ACM Symposium on Interactive 3D Graphics (I3DG-99)*, pages 127–134, Atlanta, USA, April 1999. ACM Press.

[34] W.-K. Jeong and C.-H. Kim. Direct reconstruction of displaced subdivision surface from unorganized points. In *Proceedings of the 9th Conference on Pacific Conference on Computer Graphics and Applications(PG'01)*, Tokyo, Japan, 2001. IEEE Computer Society.

[35] K. Kähler, J. Haber, and H.-P. Seidel. Dynamic Refinement of Deformable Triangle Meshes for Rendering. In H. H. S. Ip, N. Magnenat-Thalmann, R. W. Lau, and T.-S. Chua, editors, *Proceedings Computer Graphics International 2001*, pages 285–290, Hong Kong, 2001. IEEE Computer Society.

[36] K. Kähler, J. Haber, and H.-P. Seidel. Geometry-based Muscle Modeling for Facial Animation. In B. Watson and J. W. Buchanan, editors, *Proceedings Graphics Interface 2001*, pages 37–46, Ottawa, Ontario, June 2001. Canadian Human-Computer Communications Society, Morgan Kaufmann.

[37] K. Kähler, C. Rössl, R. Schneider, J. Vorsatz, and H.-P. Seidel. Efficient processing of large 3d meshes. In A. Pasko and M. Spagnuolo, editors, *Proceedings of the International Conference on Shape Modeling and Applications*, pages 228–239, Genova, Italy, May 2001. IEEE Computer Society.

[38] J. Kautz, W. Heidrich, and H.-P. Seidel. Real-time bump map synthesis. In *Proceedings of the Eurographics/SIGGRAPH Workshop on Graphics Hardware 2001*, pages 109–114, Los Angeles, USA, August 2001. Eurographics/SIGGRAPH, ACM Press.

[39] J. Kautz and M. McCool. Interactive rendering with arbitrary BRDFs using separable approximations. In D. Lischinski and G. W. Larson, editors, *Rendering Techniques '99: Proceedings of the 10th Eurographics Workshop on Rendering (EGRW-99)*, pages 247–260, Granada, Spain, June 1999. Eurographics, Springer.

[40] J. Kautz and M. McCool. Approximation of glossy reflection with prefiltered environment maps. In S. Fels and P. Poulin, editors, *Proceedings of Graphics Interface 2000 (GI-00)*, pages 119–126, Montreal, Canada, May 2000. CHCCS, Canadian Information Processing Society.

[41] J. Kautz and H.-P. Seidel. Towards interactive bump mapping with anisotropic shift-variant BRDFs. In S. N. Spencer, editor, *Proceedings of 2000 SIGGRAPH/Eurographics Workshop on Graphics Hardware*, pages 51–58, Interlaken, Switzerland, August 2000. ACM SIGGRAPH, ACM Press.

[42] J. Kautz and H.-P. Seidel. Hardware accelerated displacement mapping for image based rendering. In B. Watson and J. W. Buchanan, editors, *Graphics Interface 2001*, pages 61–70, Ottawa, Canada, 2001. Canadian Human-Computer Communications Society, Morgan Kaufmann.

[43] J. Kautz, P.-P. Vazquez, W. Heidrich, and H.-P. Seidel. A unified approach to prefiltered environment maps. In B. Proche and H. Rushmeier, editors, *Rendering Techniques 2000, Proceedings of the 11th Eurographics Workshop on Rendering*, Springer Computer Science, pages 185–196, Brno, Czech Republic, June 2000. Eurographics, Springer.

[44] Y.-M. Kim, S. W. Choi, and S.-W. Lee. Fast scene change detection using direct feature extraction from MPEG compressed videos. In A. Sanfeliu, J. J. Villanueva, M. Vanrell, R. Alquzar, T. Huang, and J. Serra, editors, *Proceedings of the 15th International Conference on Pattern Recognition (ICPR-00)*, number 3, pages 178–181, Barcelona, Spain, September 2000. IEEE Computer Society.

[45] L. Kobbelt. Discrete fairing and variational subdivision for freeform surface design. *The Visual Computer*, 16(3/4):142–158, 2000.

[46] L. Kobbelt. $\sqrt{3}$-subdivision. In K. Akeley, editor, *Computer Graphics (SIGGRAPH-00): Conference proceedings*, volume 2000 of *Computer Graphics Proceedings, Annual Conference Series*, pages 103–112, New Orleans, USA, 2000. ACM SIGGRAPH, ACM Press.

[47] L. Kobbelt, T. Bareuther, and H.-P. Seidel. Multiresolution shape deformations for meshes with dynamic vertex connectivity. In M. Gross and F. R. A. Hopgood, editors, *Proceedings of the European Association for Computer Graphics 21st Annual Conference (Eurographics-00)*, volume 19 of *Computer Graphics Forum*, pages 249–260, Interlaken, Switzerland, 2000. Blackwell.

[48] L. Kobbelt and M. Botsch. An interactive approach to point cloud triangulation. In M. Gross and F. R. A. Hopgood, editors, *Proceedings of the European Association for Computer Graphics 21st Annual Conference (Eurographics-00)*, volume 19 of *Computer Graphics Forum*, pages 479–487, Interlaken, Switzerland, 2000. Eurographics, Blackwell.

[49] L. Kobbelt, M. Botsch, U. Schwanecke, and H.-P. Seidel. Feature sensitive surface extraction from volume data. In *Computer Graphics (SIGGRAPH-01): Conference proceedings*, pages 57–66, Los Angeles, USA, 2001. ACM Press.

[50] L. Kobbelt and H.-P. Seidel. sqrt(3)-subdivision and forward adaptive refinement. In *Proceedings of the Israelian/Korean workshop on Geometric Modeling and Computer Graphics*, pages 245–252, 1999.

[51] L. Kobbelt and J. Vorsatz. Multiresolution hierarchies on unstructured triangle meshes. *Computational Geometry*, 14(1-3):5–24, 1999.

[52] L. Kobbelt, J. Vorsatz, U. Labsik, and H.-P. Seidel. A shrink wrapping approach to remeshing polygonal surfaces. In P. Brunet and R. Scopigno, editors, *Proceedings of the 20th Annual Conference ot the European Association of Computer Graphics (Eurographics-99)*, volume 18 of *Computer Graphics Forum*, pages 119–130, Milan,Italy, September 1999. Blackwell.

[53] U. Labsik, L. Kobbelt, R. Schneider, and H.-P. Seidel. Progressive transmission of subdivision surfaces. *Computational Geometry: Theory and Applications*, 15(1-3):25–39, 2000.

[54] S.-W. Lee, Y.-M. Kim, and S. W. Choi. Fast scene change detection using direct feature extraction from MPEG compressed videos. *IEEE Transactions on Multimedia*, 2(4):240–254, December 2000.

[55] H. Lensch, W. Heidrich, and H.-P. Seidel. Automated texture registration and stitching for real world models. In B. A. Barsky, Y. Shinagawa, and W. Wang, editors, *Proceedings of the 8th Pacific Conference on Computer Graphics and Applications (PG-00)*, pages 317–326, Hong Kong, China, October 2000. IEEE Computer Society.

[56] H. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based reconstruction of spatially varying materials. In S. Gortler and K. Myszkowski, editors, *Rendering Techniques: Proceedings of the 12th Eurographics Workshop on Rendering*, pages 104–115, London, Great Britain, 2001. Springer.

[57] H. P. Lensch, J. Kautz, M. Goesele, and H.-P. Seidel. 3D model acquisition including reflection properties. In D. W. Fellner, N. Fuhr, and I. Witten, editors, *Proceedings of the ECDL Workshop Generalized Documents*, Darmstadt, Germany, 2001.

[58] M. Li. Correspondence analysis between the image formation pipelines of graphics and vision. In S. J. Salvador and P. Filiberto, editors, *Proceedings of the IX Spanish Symposium on Pattern Recognition and Image Analysis*, pages 187–192, Benicasim (Castellön), Spain, May 2001. Publications de la Universitat Jaume I.

[59] C. Lürig, L. Kobbelt, and T. Ertl. Hierachical solutions for the deformable surface problem in visualization. *Graphical Models*, 62(1):2–18, 2000.

[60] M. McCool and W. Heidrich. Texture shaders. In A. Kaufman and W. Straer, editors, *Eurographics/SIGGRAPH Workshop on Graphics Hardware*, pages 117–126, Los Angeles, USA, 1999. ACM.

[61] K. Myszkowski and T. L. Kunii. A case study towards validation of global illumination algorithms: Progressive hierarchical radiosity with clustering. *The Visual Computer*, 16(5):271–288, 2000.

[62] K. Myszkowski, P. Rokita, and T. Tawara. Perception-based fast rendering and antialiasing of walkthrough sequences. *IEEE Transactions on Visualization and Computer Graphics*, 6(4):360–379, October 2000.

[63] K. Myszkowski, T. Tawara, H. Akamine, and H.-P. Seidel. Perception-guided global illumination solution for animation rendering. In E. Fiume, editor, *Computer Graphics (SIGGRAPH-01): Conference proceedings*, pages 221–230, Los Angeles, CA, USA, 2001. Association of Computing Machinery (ACM), ACM.

[64] G. Nürnberger, L. L. Schumaker, and F. Zeilfelder. Local lagrange interpolation by bivariate cubic C^1 splines. In T. Lyche and et. al., editors, *Proceedings of the Conference on Curves and Surfaces*, Oslo, Norway, 2001. Vanderbilt University Press.

[65] G. Nürnberger, G. Steidl, and F. Zeilfelder. Explicit estimates for bivariate hierarchical bases. *Communications in Applied Analysis*, 2001.

[66] G. Nürnberger and F. Zeilfelder. Developments in bivariate spline interpolation. *Journal of Computational and Applied Mathematics*, 120(1/2):125–152, 2000.

[67] G. Nürnberger and F. Zeilfelder. Interpolation by spline spaces on classes of triangulations. *Journal of Computational and Applied Mathematics*, 119(1/2):347–376, May 2000.

[68] G. Nürnberger and F. Zeilfelder. Local lagrange interpolation on powell-sabin triangulations and terrain modelling. In W. Haussmann and et. al., editors, *Proceedings of the Conference on Multivariate Approximation*, Witten, Germany, 2001. Wiley-VCH.

[69] J. Rigau, M. Feixas, P. Bekaert, and M. Sbert. View-dependent information theory measures for pixel sampling and scene discretization in flatland. In *Proceedings of the Spring Conference on Computer Graphics*, pages 231–238, Budmerice, Slovak Republic, May 2001. Comenius University.

[70] C. Rössl and L. Kobbelt. Approximation and visualization of discrete curvature on triangulated surfaces. In B. Girod, H. Niemann, and H.-P. Seidel, editors, *Proceedings of the 4th Conference on Vision, Modeling, and Visualization (VMV-99)*, pages 339–346, Erlangen, Germany, 1999. SFB 603 and Graduate Research Center, in coorperation with IEEE Signal Processing Society and Gesellschaft für Informatik GI, infix.

[71] C. Rössl and L. Kobbelt. Line-art rendering of 3D-models. In B. A. Barsky, Y. Shinagawa, and W. Wang, editors, *Proceedings of the 8th Pacific Conference on Computer Graphics and Applications (PG-00)*, pages 87–96, Hong Kong, China, 2000. IEEE Computer Society.

[72] C. Rössl, L. Kobbelt, and H.-P. Seidel. Extraction of feature lines on triangulated surfaces using morphological operators. In A. Butz, A. Krüger, and P. Olivier, editors, *Smart Graphics (AAAI Spring Symposium-00)*, volume 00-04 of *Technical Report / SS / American Association for Artificial Intelligence*, pages 71–75, Stanford, USA, 2000. AAAI Press.

[73] C. Rössl, L. Kobbelt, and H.-P. Seidel. Line art rendering of triangulated surfaces using discrete lines of curvature. In V. Skala, editor, *Proceedings of the 8th International Conference in Central Europe on Computer Graphics, Visualization and Interactive Digital Media (WSCG-00)*, pages 168–175, Plzen, Czech Republic, 2000. EUROGRAPHICS and IFIP WG 5.10, University of West Bohemia.

[74] C. Rössl, L. Kobbelt, and H.-P. Seidel. Recovering structural information from triangulated surfaces. In *Mathematical Methods in CAGD: Oslo 2000 Proceedings*, Oslo, Norway, 2001. Vanderbilt University Press.

[75] A. Scheel, B. Böhm, and H.-P. Seidel. Perception-based radiosity. In B. Girod, H. Niemann, and H.-P. Seidel, editors, *Proceedings of the 4th Conference on Vision, Modeling, and Visualization (VMV-99)*, pages 293–300, Erlangen, Germany, November 1999. SFB 603 and Graduate Research Center, in coorperation with IEEE Signal Processing Society and Gesellschaft für Informatik GI, infix.

[76] A. Scheel, M. Stamminger, J. Pütz, and H.-P. Seidel. Enhancements to directional coherence maps. In N. Magnenat-Thalmann and V. Skala, editors, *Proceedings of the 9th International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision 2001*, pages 403–410, Plzen, Czech Republic, 2001. University of West Bohemia.

[77] A. Scheel, M. Stamminger, and H.-P. Seidel. Tone reproduction for interactive walkthroughs. In M. Gross and F. R. A. Hopgood, editors, *Proceedings of the European Association for Computer Graphics 21st Annual Conference (Eurographics-00)*, volume 19 of *Computer Graphics Forum*, pages 301–311, Interlaken, Switzerland, August 2000. EUROGRAPHICS, Blackwell.

[78] A. Scheel, M. Stamminger, and H.-P. Seidel. Thrifty final gather for radiosity. In S. Gortler and K. Myszkowski, editors, *Rendering Techniques: Proceedings of the 12th Eurographics Workshop on Rendering*, pages 1–12, London, 2001. Springer.

[79] H. Schirmacher and S. Brabec. tmk - a multi-site, multi-platform system for software development. In C. Zerbst, editor, *Proc. First European Tcl/Tk User Meeting*, Hamburg, 2000. Technische Universität Hamburg-Harburg.

[80] H. Schirmacher, W. Heidrich, M. Rubick, D. Schiron, and H.-P. Seidel. Image-based BRDF reconstruction. In B. Girod, H. Niemann, and H.-P. Seidel, editors, *Proceedings of the 4th Conference on Vision, Modeling, and Visualization (VMV-99)*, pages 285–292, Erlangen, Germany, November 1999. SFB 603 and Graduate Research Center, in coorperation with IEEE Signal Processing Society and Gesellschaft für Informatik GI, Infix.

[81] H. Schirmacher, W. Heidrich, and H.-P. Seidel. Adaptive acquisition of Lumigraphs from synthetic scenes. In P. Brunet and R. Scopigno, editors, *Proceedings of the 20th Annual Conference ot the European Association of Computer Graphics (Eurographics-99)*, volume 18 of *Computer Graphics Forum*, pages 151–160, Milano, Italy, 1999. Eurographics Association, Blackwell.

[82] H. Schirmacher, W. Heidrich, and H.-P. Seidel. High-quality interactive Lumigraph rendering through warping. In S. Fels and P. Poulin, editors, *Proceedings of Graphics Interface 2000 (GI-00)*, pages 87–94, Montreal, Canada, May 2000. CHCCS, Canadian Information Processing Society.

[83] H. Schirmacher, M. Li, and H.-P. Seidel. On-the-fly processing of generalized Lumigraphs. In A. Chalmers and T.-M. Rhyne, editors, *Proceedings of Eurographics 2001*, volume 20 of *Computer Graphics Forum*, pages 165–173, Manchester, UK, 2001. Eurographics Association, Blackwell.

[84] H. Schirmacher, C. Vogelgsang, H.-P. Seidel, and G. Greiner. Efficient free form light field rendering. In T. Ertl, B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proc. 6th International Fall Workshop on Vision, Modeling, and Visualization 2001 (VMV01)*, Stuttgart, Germany, 2001. To appear.

[85] R. Schneider and L. Kobbelt. Discrete fairing of curves and surfaces based on linear curvature distribution. In P.-J. Laurent, P. Sablonniere, and L. Schumaker, editors, *Curve and Surface Design, Saint-Malo 1999*, Innovations in Applied Mathematics, pages 371–380, Saint-Malo, France, 2000. Vanderbilt University Press.

[86] R. Schneider and L. Kobbelt. Generating fair meshes with g^1 boundary conditions. In R. Martin and W. Wang, editors, *Proceedings of Geometric Modeling and Processing 2000 (GMP-00)*, pages 251–261, Hongkong, China, 2000. IEEE Computer Society.

[87] R. Schneider and L. Kobbelt. Geometric fairing of irregular meshes for free-form surface design. *Computer Aided Geometric Design*, 18(4):359–379, May 2001.

[88] R. Schneider and L. Kobbelt. Mesh fairing based on an intrinsic PDE approach. *CAD*, 2001. To appear.

[89] R. Schneider, L. Kobbelt, and H.-P. Seidel. Improved bi-laplacian mesh fairing. In *Mathematical Methods in CAGD: Oslo 2000 Proceedings*, Oslo, Norway, 2001. Vanderbilt University.

[90] R. Schneider and H.-P. Seidel. Multiresolution analysis over triangles, based on quadratic Hermite interpolation. *Journal of Computational and Applied Mathematics*, 119(1/2):97–114, 2000.

[91] U. Schwanecke and B. Jüttler. A B-spline approach to Hermite subdivision. In A. Cohen, C. Rabut, and L. L. Schumaker, editors, *Curve and surface fitting, Saint-Malo 1999*, Innovations in Applied Mathematics, pages 385–392, Saint-Malo, June 2000. Vanderbilt University Press.

[92] U. Schwanecke and L. Kobbelt. Approximate envelope reconstruction for moving solids. In *Mathematical Methods in CAGD: Oslo 2000 Proceedings*, Oslo, Norway, 2001. Vanderbilt University Press.

[93] C. Soler and F. Sillion. Hierarchical instantiation for radiosity. In B. Proche and H. Rushmeier, editors, *Rendering Techniques 2000, Proceedings of the 11th Eurographics Workshop on Rendering*, Springer Computer Science, pages 173–184, Brno, Czech Republic, June 2000. Eurographics, Springer.

[94] C. Soler and F. Sillion. Texture-based visibility for efficient lighting simulation. *ACM Transactions on Graphics*, 19(4):302–342, October 2000.

[95] M. Stamminger, K. Daubert, and H.-P. Seidel. Photon tracing with renderman light source shaders. In B. Girod, H. Niemann, and H.-P. Seidel, editors, *Proceedings of the 4th Conference on Vision, Modeling, and Visualization (VMV-99)*, pages 277–284, Erlangen, Germany, November 1999. SFB 603 and Graduate Research Center, in coorperation with IEEE Signal Processing Society and Gesellschaft für Informatik GI, infix.

[96] M. Stamminger, J. Haber, H. Schirmacher, and H.-P. Seidel. Walkthroughs with Corrective Texturing. In B. Peroche and H. Rushmeier, editors, *Rendering Techniques 2000, Proceedings of the 11th Eurographics Workshop on Rendering*, pages 377–388;422, Brno, Czech Republic, 2000. Springer.

[97] M. Stamminger, A. Scheel, X. Granier, F. Perez-Cazorla, G. Drettakis, and F. Sillion. Efficient glossy global illumination with interactive viewing. In I. S. MacKenzie and J. Stewart, editors, *Graphics Interface '99 (GI-99)*, pages 50–57, Kingston, Canada, June 1999. Canadian Human-Computer Communications Society, Morgan Kaufmann.

[98] M. Stamminger, A. Scheel, X. Granier, F. Perez-Cazorla, G. Drettakis, and F. Sillion. Efficient glossy global illumination with interactive viewing. *Computer Graphics Forum*, 19(1):13–26, March 2000.

[99] M. Stamminger, A. Scheel, and H.-P. Seidel. Hierarchical radiosity with global refinement. In B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling, and Visualization (VMV-00)*, pages 263–270, Saarbrücken, Germany, 2000.

[100] V. Volevich, K. Myszkowski, A. Khodulev, and E. Kopylov. Using the visual differences predictor to improve performance of progressive global illumination computations. *ACM Transactions on Graphics*, 19(2):122–161, April 2000.

[101] J. Vorsatz, K. Kähler, L. Kobbelt, and H.-P. Seidel. Linear time mesh simplification with reliable error-bounds. In B. Girod, G. Greiner, H. Niemann, and H.-P. Seidel, editors, *Proceedings of Vision, Modeling, and Visualization (VMV-00)*, pages 145–152, Saarbrücken, Germany, November 2000.

[102] J. Vorsatz and L. Kobbelt. Robust multi-band detail encoding for triangular meshes of arbitrary connectivity. In B. Girod, H. Niemann, and H.-P. Seidel, editors, *Proceedings of the 4th Conference on Vision, Modeling, and Visualization (VMV-99)*, pages 245–252, Erlangen,Germany, November 1999. SFB 603 and Graduate Research Center, in coorperation with IEEE Signal Processing Society and Gesellschaft für Informatik GI, infix.

[103] J. Vorsatz, C. Rössl, L. Kobbelt, and H.-P. Seidel. Feature sensitive remeshing. In A. Chalmers and T.-M. Rhyne, editors, *Computer Graphics Forum, Proceedings of Eurographics '01*, volume 20 of *Computer Graphics Forum*, pages 393–401, Manchester, Great Britain, September 2001. Blackwell.

[104] J. Vorsatz and H.-P. Seidel. Multiresolution modeling and interactive deformation of large 3D meshes. In *Deform Conference*, pages 11–23, Geneva,Switzerland, 2001. Kluwer.

[105] G. Wesche and H.-P. Seidel. FreeDrawer – a free-form sketching system on the responsive workbench. In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology 2001 (ACM VRST'01)*, Banff, Canada, 2001. ? To appear.

[106] R. Westermann, L. Kobbelt, and T. Ertl. Real-time exploration of regular volume data by adaptive reconstruction of iso-surfaces. *The Visual Computer*, 15(2):100–111, 1999.

[107] H. Yamauchi and A. Maeda. Java in the real world - java ha dokomade tukaeru gengoka. In T. Ogawa, editor, *The 41 Programming Symposium*, pages 97–108, Hakone, Japan, January 2000. Information Processing Society of Japan (IPSJ), Information Processing Society of Japan. In Japanese.

[108] H. Yamauchi, A. Maeda, and H. Kobayashi. Developing a practical parallel multi-pass renderer in Java and C++: Toward a Grande application in Java. In *Proceedings of the ACM 2000 Java Grande Conference (JAVA-00)*, pages 126–133, San Francisco, California, USA, June 2000. Association for Computing Machinery (ACM), ACM Press.

[109] F. Zeilfelder. Strong unicity of best uniform approximation from periodic spline spaces. *Journal of Approximation Theory*, 99(2):1–29, 1999.

[110] F. Zeilfelder, O. Davydov, and G. Nürnberger. Interpolation by splines on triangulations. In M. W. Müller, M. D. Buhmann, D. H. Mache, and M. Felten, editors, *New Developments in Approximation Theory. 2nd International Dortmund Meeting (IDoMAT) '98*, volume 132 of *International Series of Numerical Mathematics*, pages 49–70, Witten, Germany, April 1999. University of Dortmund, Germany, Birkhäuser.

[111] F. Zeilfelder, O. Davydov, and G. Nürnberger. Bivariate spline interpolation with optimal approximation order. *Constructive Approximation*, 17(2):181–208, May 2001.

[112] F. Zeilfelder and G. Nürnberger. Local lagrange interpolation by cubic splines on a class of triangulations. In K. Kopotun, T. Lyche, and M. Neamtu, editors, *Proceedings of the Conference on Trends in Approximation Theory*, pages 341–350, Nashville, USA, June 2001. Vanderbilt University.

Book Chapters

[1] D. Blythe, S. Ghali, L. Kettner, and H. Sowizral. Migrating to an object-oriented graphics api (course 15). In S. Spencer, editor, *Computer Graphics (SIGGRAPH-00): Course Notes CD-ROM*. ACM SIGGRAPH, New York, USA, July 2000.

[2] S. Ghali. Object space visibility. In L. Harrison, editor, *Computer Graphics (SIGGRAPH-01): Course Notes CD-ROM*. ACM SIGGRAPH, New York, USA, August 2001.

[3] S. Ghali, E. Fiume, and H.-P. Seidel. Shadow computation: A unified perspective. In *European Association for Computer Graphics 21st Annual Conference (Eurographics-00), State of The Art Reports*, pages 1–17, Aire-la-Ville, Switzerland, August 2000. Eurographics.

[4] W. Heidrich. Interactive display of global illumination solutions for non-diffuse environments. In *European Association for Computer Graphics 21st Annual Conference (Eurographics-00), State of The Art Reports*, pages 1–19, Aire-la-Ville, Switzerland, August 2000. Eurographics, Eurographics.

[5] W. Heidrich. Texture models. In B. Girod, G. Greiner, and H. Niemann, editors, *Principles of 3D Image Analysis and Synthesis*, chapter 5.3, pages 223–232. Kluwer Academic Publishers, Boston, Massachusetts, January 2000.

[6] W. Heidrich, P. Eisert, M. Stamminger, and A. Scheel. Models for illumination and reflection. In B. Girod, G. Greiner, and H. Niemann, editors, *Principles of 3D Image Analysis and Synthesis*, chapter 5.1, pages 204–213. Kluwer Academic Publishers, Boston, Massachusetts, January 2000.

[7] L. Kobbelt, S. Bischoff, M. Botsch, K. Kähler, C. Rössl, R. Schneider, and J. Vorsatz. Geometric modeling based on polygonal meshes. In *Tutorials of the European Association for Computer Graphics 21st Annual Conference (Eurographics-00)*, pages 1–47, Aire-la-Ville, August 2000. European Association for Computer Graphics, Eurographics.

[8] H. P. A. Lensch, M. Goesele, and H.-P. Seidel. A framework for the acquisition, processing, transmission, and interactive display of high quality 3D models. In *Tutorial Notes of the DAGM 2001*, Munich, Germany, September 2001. Also published as Research Report MPI-I-2001-4-005, Max-Planck-Institut fr Informatik, Saarbrcken, Germany.

[9] H. P. A. Lensch, J. Kautz, M. Goesele, and H.-P. Seidel. A framework for the acquisition, processing, transmission, and interactive display of high quality 3D models on the web. In *Tutorial Notes of the Web3D Conference 2001*, Paderborn, Germany, February 2001. Also published as Research Report MPI-I-2001-4-002, Max-Planck-Institut fr Informatik, Saarbrcken, Germany.

[10] M. Magnor and W. Heidrich. Image-based rendering. In B. Girod, G. Greiner, and H. Niemann, editors, *Principles of 3D Image Analysis and Synthesis*, chapter 5.4, pages 232–241. Kluwer Academic Publishers, Boston, Massachusetts, January 2000.

[11] K. Myszkowski. Chapter 4: Perception-driven global illumination and rendering computation, and chapter 6: Perception-driven rendering of high-quality walkthrough animations. In A. McNamara and A. Chalmers, editors, *Image quality metrics (Course 44)*, volume 44 of *ACM Siggraph Course Notes*, chapter 4 and 6, pages 43–59 and 75–81. ACM SIGGRAPH, New York, USA, July 2000.

[12] K. Myszkowski. Perception-based global illumination, rendering, and animation. In A. McNamara and A. Chalmers, editors, *Computer Graphics (Siggraph-01), Course Notes: Seeing is Believing: Reality Perception in Modeling, Rendering and Animation*, volume 21 of *ACM Siggraph 2001, Course Notes*. ACM Siggraph, New York, USA, 2001.

[13] A. Scheel, K. Daubert, and M. Stamminger. Global illumination. In B. Girod, G. Greiner, and H. Niemann, editors, *Principles of 3D Image Analysis and Synthesis*, chapter 5.2, pages 213–223. Kluwer Academic Publishers, Boston, Massachusetts, January 2000.

[14] H.-P. Seidel. Realistic rendering. In B. Girod, G. Greiner, and H. Niemann, editors, *Principles of 3D Image Analysis and Synthesis*, chapter 5, pages 203–241. Kluwer Academic Publishers, Boston, Massachusetts, January 2000.

[15] M. Stamminger, W. Heidrich, J. Kautz, and S. Brabec. Advanced lighting for interactive applications. In *Tutorials of the GTEC'01*, Hong Kong, China.

[16] M. Stamminger, D. Wexler, W. Kresse, N. Holzschuch, and P. H. Christensen. Advanced radiosity: Complex scenes and glossy reflections. In *Tutorials of the European Association for Computer Graphics 20th Annual Conference (Eurographics-99)*, Mailand, Italy. European Association for Computer Graphics, Eurographics.

[17] F. Zeilfelder and H.-P. Seidel. Splines over triangulations. In G. Farin, J. Hoschek, M.-S. Kim, and D. Abma, editors, *The Handbook of Computer Aided Design*, chapter Splines over Triangulations. Elsevier Science, Amsterdam, Netherlands, 2001.

Part VI

Index

Index