

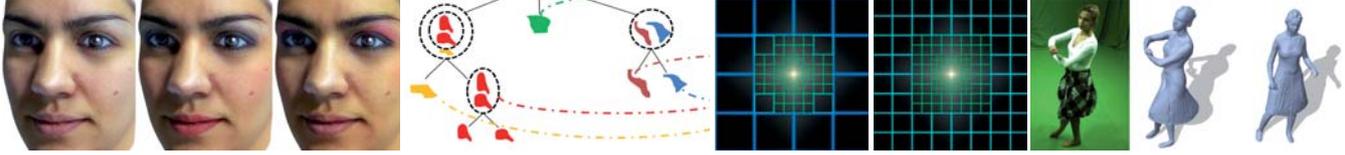
max planck institut
informatik

Report 2009/2010

09 10 11 12 13



max planck institut
informatik



REPORT

Report 2009/2010

Directors

Prof. Dr. Dr. Thomas Lengauer
Prof. Dr. Kurt Mehlhorn
Prof. Dr. Bernd Schiele
Prof. Dr. Hans-Peter Seidel
Prof. Dr. Gerhard Weikum

Scientific Advisory Board

Prof. Dr. Pankaj Kumar Agarwal, Department of
 Computer Science, Duke University, Durham, USA
Prof. Dr. Douglas L. Brutlag, Stanford University,
 School of Medicine, USA
Prof. Dr. Joseph M. Hellerstein, EECS Computer
 Science Division, University of California, Berkeley, USA
Prof. Dr. Yannis E. Ioannidis, Department of Informatics &
 Telecommunications, University of Athens, Greece
Prof. Dr. Friedhelm Meyer auf der Heide, Heinz Nixdorf Institute,
 University of Paderborn, Germany
Prof. Dr. Eugene Myers, Howard Hughes Medical Institute,
 Ashburn, USA
Prof. Dr. Frank Pfenning, Computer Science Department,
 Carnegie Mellon University, Pittsburgh, USA
Prof. Dr. Claude Puech, Laboratoire INRIA, France
Prof. Dr. Éva Tardos, Cornell University, Ithaca, USA
Prof. Dr. Demetri Terzopoulos, Department of Computer
 Science, University of California, Los Angeles, USA

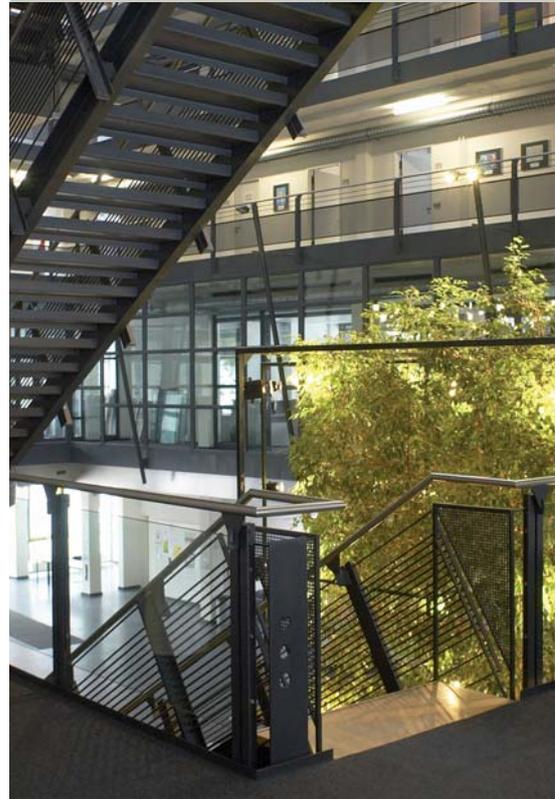
Curatorship Board

Hon. Adv. Prof. (Tsinghua) Dr. Reinhold Achatz,
 Head of Corporate Research
 and Technologies, Siemens AG, Munich
Dr. Siegfried Dais, deputy chairman of the Board of
 Management of Robert Bosch GmbH, Stuttgart
Christiane Götz-Sobel, Head of the editorial department
 of Science and Technology, ZDF, Munich
Dr. Christoph Hartmann, Minister of Economics and
 Science, Saarland
Prof. Dr. Joachim Hertel, CEO of DACOS
 Software GmbH, Saarbrücken
Prof. Dr. Stefan Jähnichen, President of the German
 Computer Science Association (GI), Bonn
Prof. Dr. Henning Kagermann, President of the National
 Academy of Science and Technology (acatech), Munich
Prof. Dr. Volker Linneweber, President of
 Saarland University, Saarbrücken
Prof. Dr. Wolf-Dieter Lukas, German Federal Ministry
 of Education and Research, Bonn
Dr. Nelson Mattos, Vice President Engineering,
 EMEA, Google, Zurich, Switzerland
Prof. Dr. Wolffried Stucky, Institute of Applied
 Informatics and Formal Description Methods (AIFB),
 University of Karlsruhe
Prof. Dr. Margret Wintermantel, President of the
 German Rectors' Conference, Bonn



CONTENTS

7	PREFACE
8	THE MAX PLANCK INSTITUTE FOR INFORMATICS: OVERVIEW
14	DEPARTMENT OVERVIEW
	DEPARTMENTS
14	DEPT . 1 ALGORITHMS AND COMPLEXITY
16	DEPT . 2 COMPUTER VISION AND MULTIMODAL COMPUTING
18	DEPT . 3 COMPUTATIONAL BIOLOGY AND APPLIED ALGORITHMICS
20	DEPT . 4 COMPUTER GRAPHICS
22	DEPT . 5 DATABASES AND INFORMATION SYSTEMS
	RESEARCH GROUPS
24	RG . 1 AUTOMATION OF LOGIC
25	THE MAX PLANCK CENTER
26	THE RESEARCH AREAS
28	UNDERSTANDING IMAGES & VIDEOS
36	BIOINFORMATICS
44	GUARANTEES
52	INFORMATION SEARCH & DIGITAL KNOWLEDGE
62	MULTIMODAL INFORMATION
66	OPTIMIZATION
74	SOFTWARE
82	VISUALIZATION
90	IMPRS-CS
92	NEWS
100	THE INSTITUTE IN FIGURES
102	INFORMATION SERVICES & TECHNOLOGY
108	COOPERATIONS
112	PUBLICATIONS
118	DIRECTIONS



P R E F A C E

P R E F A C E

Every two years, the Max Planck Institute for Informatics publishes a report for the general public. We would like to take this opportunity to present some of the current topics, goals and methods of modern informatics. We introduce the work of our institute and hope to bring you, our readers, closer to the fascinating world of our science.

The Max Planck Institute for Informatics aims to be a lighthouse in research in informatics. We aim to have impact in the following ways: First, through our scientific work, which we disseminate mainly through publications and books, but also in the form of software and internet services. Second, through the training of young scientists, particularly during the doctoral and postdoctoral phases. We are educating future leaders for science and industry. Third, through our role in the field. We initiate and coordinate large research programs, and serve on important committees. Fourth, by attracting talent from within and outside the country. About half of our staff of over 190 scientists comes from abroad. Fifth, through the transfer of our results into industry. These transfers take place through cooperation projects, spin-offs, and people. Sixth, by building a world-class competence center for informatics in cooperation with our partners: Saarland University, the German Research Center for Artificial Intelligence and the Max Planck Institute for Software Systems. We have been very successful in all of these undertakings in recent years.



Our success is becoming visible in the new buildings for the Max Planck Institute for Software Systems, the Center for Bioinformatics, the Intel Visual Computing Institute, the informatics lecture halls, the informatics and mathematics library, and the Cluster of Excellence, all of which are located at the Platz der Informatik.

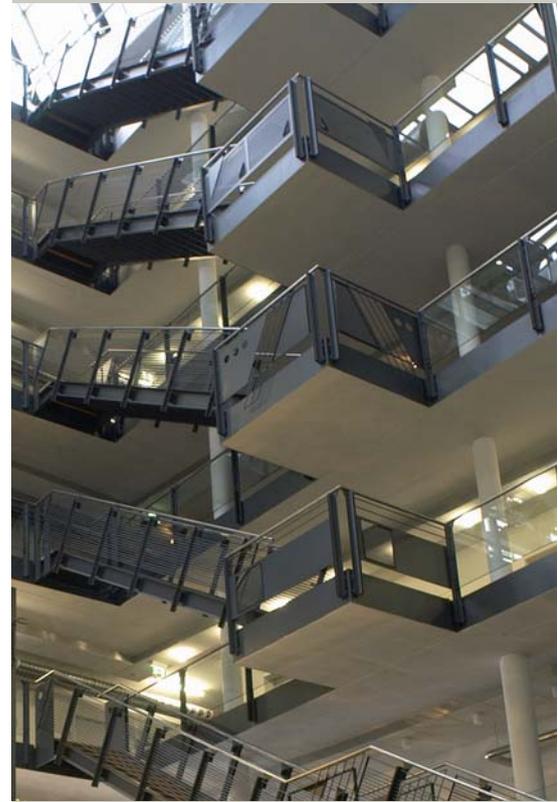
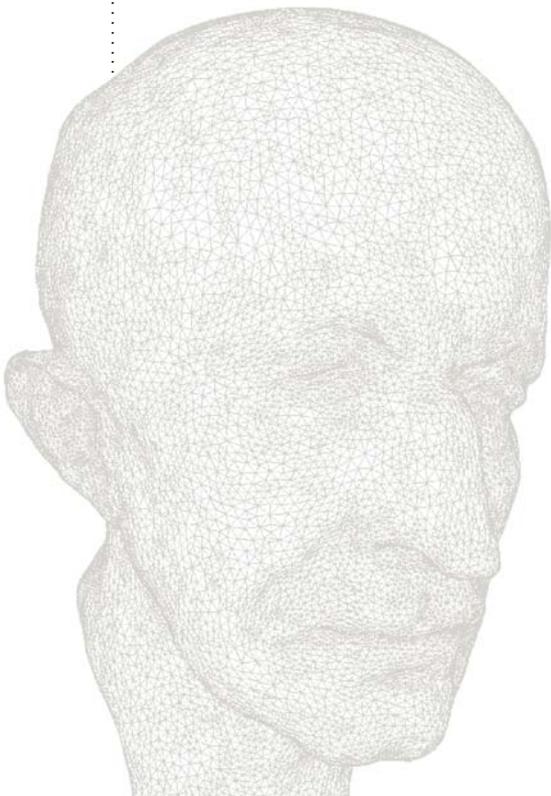
The report is structured as follows: After an overview of the institute and its departments and research groups, we present the main areas of recent work. These topics span several departments and will also be the focus of our work in the next years. The last part of the report contains a selection of recent scientific publications and a compact presentation of the institute through key indicators.

Please enjoy reading this report.

Kurt Mehlhorn *Managing Director*

The Max Planck Institute for Informatics, an Overview

Information technology influences all aspects of our lives. Computer systems, hardware, software, and networks, are among the most complex structures that have been constructed by man. Computational thinking is a new way of studying the universe. Basic research in informatics is needed to cope with this complexity, to lay the foundations for powerful computer systems, and to further develop computational thinking.



OVERVIEW

Basic research in informatics has led to dramatic changes in our everyday lives in recent years. This has become particularly clear in the last two decades: The worldwide web, search engines, compression processes for video and music, and secure electronic banking using cryptographic methods have revolutionized our lives just a few years after their discovery at universities and research institutes.

The Max Planck Society, the leading organization for basic research in Germany, reacted to these challenges by founding the Max Planck Institute for Informatics (MPI-INF) in Saarbrücken in 1990. In 2005, the Max Planck Institute for Software Systems (MPI-SWS) was established with sites in Saarbrücken and Kaiserslautern. There are departments with a strong emphasis on informatics in other institutes of the Max Planck Society as well. The restructuring of the Max Planck Institute for Metal Research into an Institute for Intelligent Systems has further strengthened informatics within the Max Planck Society. Given the importance of the area, the establishment of further institutes for informatics or related areas is desirable.

Goals

The Max Planck Institute for Informatics aims to be a lighthouse in research in informatics. We aim to have impact in the following ways:

First, through our scientific work, which we disseminate mainly through publications and books, but also in the form of software and internet solutions. At the moment, we are concentrating on algorithms for very large, multimodal data. Multimodal means text, speech, images, videos, graphs and high-dimensional data.

Second, through the training of young scientists, particularly in the doctoral and postdoctoral phases, we are educating future leaders for research and business. Over 190 researchers are working at our institute and remain with us on average for three years. In this way, we provide the society with over 60 well-trained young scientists each year.

Third, by our role in the profession. We initiate and coordinate large research programs and serve on important committees, e.g., the “Wissenschaftsrat”. The Institute has played a significant role in forming the Excellence Cluster “*Multimodal Computing and Interaction*” and the “*Graduate School for Computer Science*”.

Fourth, by attracting talent from within and outside the country. Half of the research staff of the institute comes from outside Germany. This strengthens the talent base in Germany and establishes bridges to foreign countries.

Fifth, by transferring our results to industry. These transfers take place through cooperation projects, spin-offs, and people. Intel founded the Intel Visual Computing Institute in 2009 together with the UdS (Saarland University), the DFKI (German Center for Artificial Intelligence), the MPI-SWS, and the MPI-INF. Intel is investing \$12 million in a new research project with its headquarters on the campus of the UdS. The development of future graphics and visual computing technologies are at the core of the center's work. The investment will take place over a period of five years and is the most extensive cooperation to date by Intel with a European university.

Sixth, by building a world-class competence center for informatics in cooperation with our partners, the UdS, the DFKI, and the MPI-SWS.

We have been very successful in all of these endeavors in recent years.

History and Organization

The Max Planck Institute for Informatics was founded in 1990 with Kurt Mehlhorn as the founding director. He has led the “*Algorithms and Complexity*” department since then. Also, Harald Ganzinger was involved from the very beginning and led the “*Logic of Programming*” department until his death in 2004. A third department “*Computer Graphics*” followed in 1999 under the direction of Hans-Peter Seidel. Thomas Lengauer then joined in 2001 to lead the “*Computational Biology and Applied Algorithmics*” department. Gerhard Weikum has led the “*Databases and Information Systems*” department since 2003. In the summer of 2010, the new “*Computer Vision and Multimodal Computing*” department, led by Bernt Schiele, was added.

In addition to the departments, the institute is home to independent research groups. Christoph Weidenbach leads the research group “*Automation of Logic*”, and several research groups, headed by Mario Albrecht, Jan Baumbach, Ivo Ihrke, Meinard Muller, Ralf Schenkel, Robert Strzodka, Martin Theobald, Thorsten Thormählen and Michael Wand have been established within the Cluster of Excellence “*Multimodal Computing and Interaction*” as well as through the cooperation with Stanford University, carry out their work at the institute. The group “*Informatics for Genome Research and Epidemiology*”, led by Alice McHardy, will end in 2011 because Alice has accepted a professorship in Düsseldorf.

Research Topics

Algorithms are our central research subject. An algorithm is a general recipe for solving a class of problems. We enlarge the computational universe through the design of new and better algorithms. We prove the correctness of these algorithms and analyze their performance. We implement them and validate them through experiment. We make them available to the world in the form of software libraries and internet services. We study inherent properties of computation and investigate techniques for robust software design. We apply the algorithms to interesting application problems.

Improved hardware has led to impressive gains in efficiency; improved algorithms can lead, and have led to, even larger gains. Here is an example: The state of hardware and algorithms in 1970 made it possible to calculate an optimal traveling salesman tour through 120 cities (a classic optimization problem and a recognized benchmark for computer performance). Adding one more city multiplies the run time of the classical algorithm by 120 and adding another city multiplies it by 121; the running time grows super-exponentially. Current hardware and the classical algorithm would allow us to solve problems with 135 cities. The combined advancement in hardware and algorithms allows us to find an optimal route through thousands of cities.

The department “*Algorithms and Complexity*” concentrates on the resource requirements of algorithms. The most important resources are running time (*How long must I wait for the result of my computation?*) and space requirement (*Do I have enough storage space for my calculation?*). The group develops new



OVERVIEW

algorithms with improved running time or storage requirement and also studies the basic limits of computation: How much time and space are provably necessary for a computation?

The “*Automation of Logic*” research group investigates logic-based generic procedures for solving “hard” combinatorial and decision problems. Typical applications are the verification of hardware or software as well as optimization problems.

Nowadays, computers are used to model, represent, and simulate aspects of real or virtual worlds. Since the visual sense is a key modality for humans, computer graphics has become a key technology in modern information and communication societies. The department “*Computer Graphics*” researches the entire processing chain from data acquisition via modeling (creation of a suitable scene representation) to image synthesis (generation of views for human consumption). The following scientific challenges emerge from this: For the input side, we want to develop modeling tools for efficient handling and processing of large data flows, and, on the output side, we are seeking new algorithms for fast computation of high-quality views; these algorithms should exploit the capabilities of modern graphics hardware.

The department “*Bioinformatics and Applied Algorithmics*” addresses the potential of computing for the life sciences. The life sciences have an increasing demand for algorithmic support due to the recent large increase in experimental data. Algorithms play a central role in the preparation and configuration of biological experiments and even more in the interpretation of biological data

generated by them. The computer is now an essential tool for biology and medicine. The understanding of biological processes on the molecular level is not possible without sophisticated information processing. Vast amounts of data need to be processed in modern biology, and the biochemical interactions in the living organism are so complex that studying them is hopeless without algorithmic support. Therefore, bioinformatics methods have become essential for modern research on the diagnosis and treatment of illnesses.

The “*Informatics for Genome Research and Epidemiology*” research group has developed new methods for the analysis of genomic sequences for questions of medical and biotechnological relevance.

The “*Databases and Information Systems*” department is dedicated to the topics of search, distribution and organization of data in digital libraries, scientific data collections, and the worldwide web. Our long-term goal is the development of easy-to-use, scalable and precise tools for intelligent searches, which actively support the user in the formulation of queries and in finding relevant information in different data formats. A special characteristic of this research is the automatic extraction of structured information from unstructured sources such as the worldwide web. Our extraction processes combine pattern recognition, linguistic methods, and statistical learning. In this way, the department has created one of the most comprehensive knowledge bases over the past few years and has made it publicly available. In addition, we are developing new methods and software tools for the search and analysis of XML documents, graphics-based RDF data, and very data-intensive internet

archives. We investigate various approaches for the implementation of these methods in distributed computing systems for better scalability.

The newest department “*Computer Vision and Multimodal Computing*” investigates processing and understanding sensor information. Sensors range from relatively simple, e.g., GPS and acceleration sensors, to very powerful sensors, e.g., cameras. They are embedded in more and more devices. Although algorithmic processing of sensor information has advanced considerably, it is still mainly limited to low-level processing. In particular, we are far from being able to fully interpret and understand sensor information. Such sensor understanding is, however, a necessary prerequisite for many areas such as man-machine interaction, the indexing of image and video databases, or for autonomous systems such as robots.

Excellence Cluster Multimodal Computing and Interaction

The Institute plays an important role in the Excellence Cluster “*Multimodal Computing and Interaction*”, which started its work in November 2007. Four of the institute’s directors belong to the 13 principal investigators of the cluster, and Hans-Peter Seidel coordinates the cluster’s work. The cluster has now completed three years of successful work. The following summary of the application is an announcement from 2007 of the cluster’s goals for future work, many of which have already been accomplished:

The last three decades have brought dramatic changes to the way we live and work. We live in a computerized society. The technological basis for these changes are modern computer systems which store, process and transmit data more compactly, more cost-efficiently and faster than ever before. This combination of increasing performance and falling prices for information and communications technology is without precedence.

The challenge is no longer to transmit information, but to organize, understand, and search it in all its forms efficiently, intelligently and robustly, and to create systems which make this possible in a natural and intuitive way.

The cluster is facing these challenges. The term “*multimodal*” in the cluster’s name refers to the human senses especially vision and hearing, the diversity of human expression, and the different types of digital data such as text, voice, images, movies, 3D models and high-dimensional data. The cluster is organized in nine research areas. Four of them, namely text and speech processing, visual computing, algorithmic foundations, and secure autonomous networked systems, are oriented towards basic research. The other five, namely open science web, information processing in the life sciences, large-scale virtual environments, synthetic virtual characters, and multimodal dialogue systems, focus on application-oriented research.

For this research program, the departments for informatics and computer linguistics and phonetics at the UoS and the Max Planck Institute for Informatics, the DFKI, and the Max Planck Institute for Software Systems have all joined forces. Although some groups in these institutions have already been working together successfully for several years, this is the first time that leading researchers from *all* participating institutions have presented a joint research program.

A prominent goal of the cluster is to qualify and promote young scientists. Saarbrücken has acquired a reputation over the years as an “*elite school*” for young scientists. A large part of the budget will be deployed to set up junior research groups. Half of all group leaders are financed by the university and the par-



O V E R V I E W

participating institutions. This financing scheme ensures continuity during and beyond the life cycle of the cluster. A new professorship for “*Image, Video and Multimedia Systems*” will be financed first by the cluster, and then permanently by the university. All participating institutions have open leadership positions in the area of the cluster.

The cluster strengthens the already existing academic excellence in Saarbrücken and establishes it as a leading center for informatics and speech technology.

Publications and Software

The scientific results of the Max Planck Institute for Informatics are distributed through presentations, publications, software, and web services. Our publications appear in the best venues of the computer science field. Most publications are freely available in the institute’s repository. Part of our results is available in the form of downloadable software or as a web service. Examples are CGAL (Computational Geometry Algorithms Library), the GISAID EpiFlu database as well as the clinically-used web service Geno2pheno for HIV therapy support. Publications in the form of software and web services make our results available more directly and to a larger audience than classical publications.

Promotion of Young Scientists

A further goal of the institute is the creation of a stimulating environment for young scientists, in which they can grow, develop their own research programs, and build their own groups. We concentrate on doctoral and postdoctoral training. Our 110 PhD students are trained in cooperation with the Graduate School for Computer Science at Saarland University and the International Max Planck Research School for Computer Science (IMPRS-CS). Our postdoctoral researchers participate in international collaborations such as the “*Max Planck Center for Visual Computing*” (a cooperation with Stanford University in the area of computer graphics), or the “*Indo Max Planck Center for Computer Science*” (a cooperation with leading universities in India), or one of our many EU projects. Our training effort is successful; about 20 PhDs graduate every year and more than 100 alumni from the institute are now professors at institutions all over the world.

Structure of the Report

After a brief introduction to the departments and research groups of our institute, we survey our work by means of representative examples. We group the examples into subject areas, each of which spans at least two departments. This report ends with a presentation of the IMPRS-CS, a presentation of the institute in figures, infrastructure aspects, and a tabular listing of cooperations and publications. Enjoy reading it. :::

Algorithms and Complexity

PROF. DR. KURT MEHLHORN

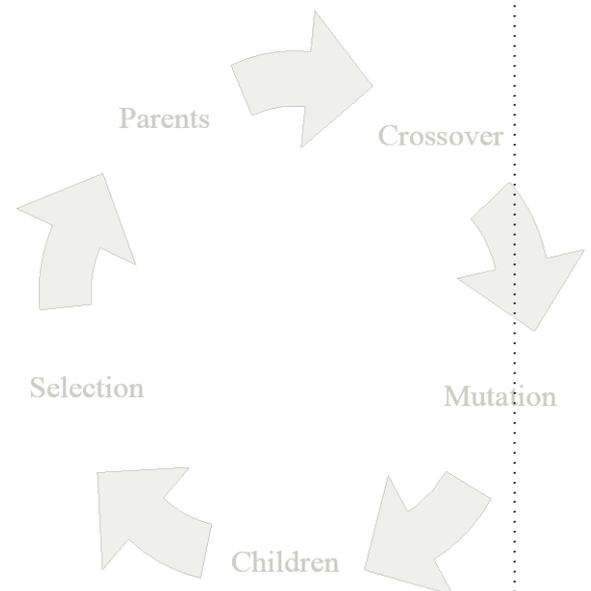
DEPT.1



The group has existed since the founding of the institute. It currently has about 40 staff members and doctoral candidates. Our goals are

- carrying out outstanding basic research in the field of algorithms,
- implementing our fundamental work in demonstrators and generally useful software libraries,
- promoting young scientists in a stimulating work groups environment.

We are successful in all three aspects. We are effective through publications, software and people. We publish in the best journals, present our results at the leading international conferences in the field, our LEDA and CGAL software libraries are used worldwide, and we offer completely new search engine opportunities for efficient and intelligent searches in larger amounts of data. Many former members of the group are in top positions domestically and abroad.



CONTACT

Algorithms and Complexity

Secretary

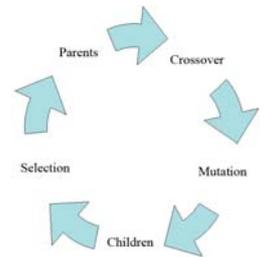
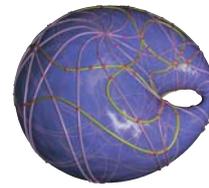
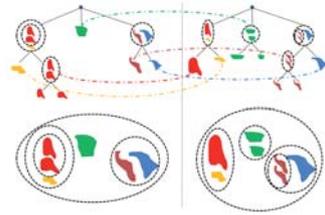
Ingrid Finkler-Paul | Christina Fries

Phone +49 681 9325-1000

Email infi@mpi-inf.mpg.de

chfries@mpi-inf.mpg.de

DEPT. 1



Algorithms are the heart of all software systems. We work on the design and analysis of algorithms in many facets: combinatorial, geometric and algebraic algorithms, data structures and search processes, very different computer models (sequential, parallel, distributed, flat data or hierarchical data), exact and approximate solutions, problem-specific methods and general heuristics, deterministic and randomized solutions, upper and lower bounds, analyses in the worst case and on the average. We develop efficient algorithms for abstract versions of applications problems as well as for concrete applications, e.g., placing wind turbines in a wind farm. A portion of our theoretical insight is used to implement software demonstrators and software libraries; as a part of our practical work, we collaborate with companies.

Outstanding theoretical results in the past two years include new algorithms for provably correct isolation of zeroes, for finding good solutions for the traveling salesman problem in appropriate metric spaces, for load balancing in computer networks and real-time scheduling. Our latest analyses of degree distribution for random planar graphs and our work in complexity theory for randomized search heuristics provide a fundamental understanding for many algorithmic questions.

Outstanding practical results in recent years include our contribution to the CGAL software library, which also makes the handling of non-linear objects possible, the CompleteSearch search engine and its application to one of the most important literature database in computer science, as well as evolutionary algorithms to calculate a yield-optimized placement of wind turbines in wind farms.

Our theoretical work is inspired by our practical work and vice versa. Our theoretical work forms the basis for demonstrators and libraries. For instance, the CompleteSearch Engine is based on new index structures, which are more powerful than known structures, but require no additional space. CGAL algorithms rely on a deep theoretical understanding of algebraic curves and surfaces.

The combination of theoretical and experimental research in algorithms has become a widely accepted research direction. The DFG supports it via its priority program “Algorithm Engineering”.

The group is involved in international projects: the GIF project (geometric computing – with the University of Tel Aviv) and the Indo-German Max Planck Center for Computer Science (IMPECS). In addition, there is a regular international exchange through

various means of support for excellent research. For example, at this time we are hosting four recipients of Humboldt scholarships, a Marie-Curie fellow, and a recipient of a scholarship from the Swiss National Science Foundation. In Germany, we participate in the “Algorithm Engineering” priority program and are part of the trans-regional special research area AVACS (Automatic Verification and Analysis of Complex Systems).

The promotion of young scientists is an integral component of our work. We give lectures at Saarland University, which are addressed not only to students, but also to our doctoral candidates. Part of our training concept is that we continue to employ our students after successful promotion and only after a minimum stay of one year at a research institution abroad. The entirety of these measures ensures that our group members are well-equipped after their stay at the Max Planck Institute to attain attractive positions in industry not limited to research, or to continue their scientific careers at leading universities or research institutes domestically and abroad. ...

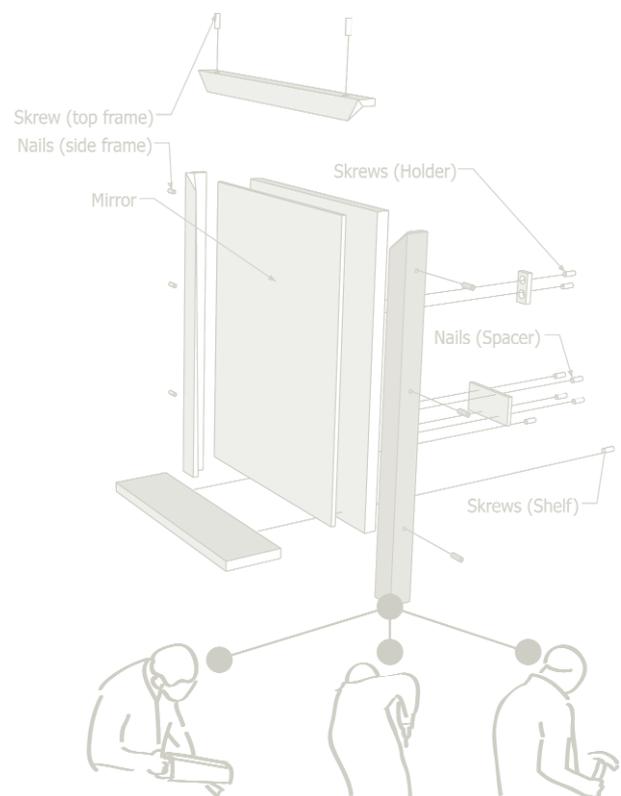
Computer Vision and Multimodal Computing

PROF. DR. BERNT SCHIELE

DEPT.2



The department was founded in 2010 and currently includes 10 scientists. The group's main research areas are computer vision with a focus on object recognition and 3D scene description as well as multi-sensor-based context recognition in the area of ubiquitous and wearable computing.



CONTACT

Computer Vision and Multimodal Computing

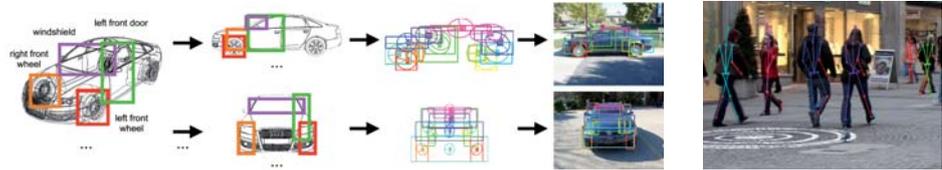
Secretary

Anja Weber

Phone +49 681 9325-2000

Email aweber@mpi-inf.mpg.de

DEPT. 2



Sensors such as cameras, GPS and accelerometers are being increasingly embedded in devices and environments, and they are already helpful in various ways. Computer-controlled processing of sensor information has made enormous progress but is generally limited to simple matters. This means, in particular, that devices and computers which have access to this sensor information do not fully interpret it and thus cannot truly understand their environment. The department is therefore concerned with the understanding of sensor information, using both powerful sensors, such as cameras, and embedded sensors, such as gyroscopes and accelerometers.

In the area of computer vision, the department deals with problems such as object recognition, one of the basic problems of image understanding. In recent years, this area of computer vision has made impressive progress, and the department has played a pioneering role in presenting several innovative approaches. One of these approaches recognizes and segments the object simultaneously, leading to significantly improved results in comparison with standard approaches. Current work has presented approaches where object models are learned directly from computer graphic models, or where speech and image processing is combined in a multimodal learning approach.

Another central theme of the department is people detection and tracking using moving cameras. This problem is

not only scientifically challenging but also has a wide variety of applications, such as image and video understanding, or in robotics and the automobile industry. For example, cars equipped with such a camera may predict the movements of pedestrians and therefore react to their behavior more effectively. The department has developed approaches that robustly detect people and track them over longer periods of time. A recently presented approach not only describes people but also entire 3D scenes, representing a further step towards complete image and scene understanding.

In addition to computer vision, the second central research area is the processing and understanding of multimodal sensor information. The underlying observation here is that an increasing amount of computers and sensors can be found in our environment, in objects and even in our clothing. Context awareness and sensing is often seen as a means of making the computing tasks sensitive to the situation and the user's needs. Ultimately, context awareness may support and enable seamless interaction and communication between humans and computing environments without the need for explicit interaction. In this area, the department has presented approaches to recognize long-term activities and to model personal daily routines. It was also possible to show that a person's interruptibility can be predicted with surprising accuracy using a few sensors embedded in their clothing.

The third research area of the department is machine learning. This plays the important role as a cross-cutting theme, as the other research areas make extensive use of probabilistic modeling and inference techniques. These allow, for example, the modeling of the uncertainties that exist with any sensor processing. In addition, they allow the use of large amounts of data, and can also elegantly integrate previous knowledge. ∴

Computational Biology and Applied Algorithmics

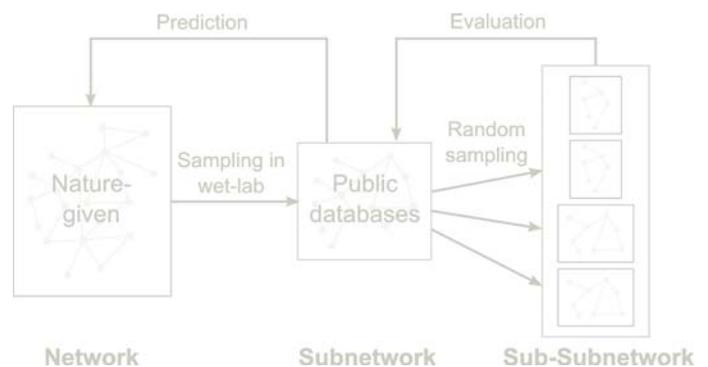
PROF. DR. THOMAS LENGAUER, PH.D.

DEPT. 3



This department has existed since October 2001 and is directed by Prof. Dr. Thomas Lengauer.

The department employs currently about 20 scientists who perform research exclusively in the area of bioinformatics.



CONTACT

Computational Biology and Applied Algorithmics

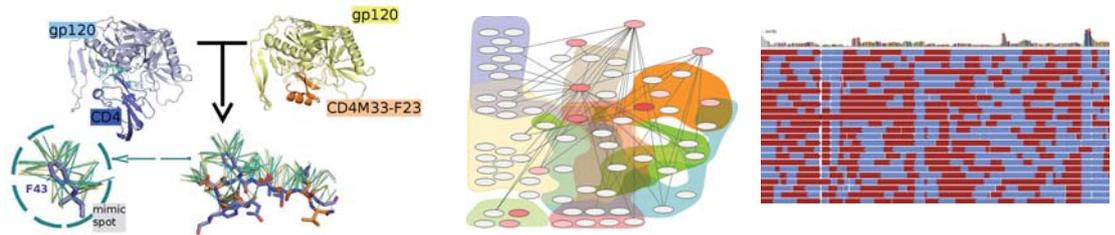
Secretary

Ruth Schnepfen-Christmann

Phone +49 681 9325-3000

Email ruth@mpi-inf.mpg.de

DEPT. 3



The department primarily investigates topics which are directly or indirectly relevant for the diagnosis and therapy of diseases. At the molecular level, disease processes are manifested by aberrations in the biochemical circuitry of an organism. The building blocks of such biochemical networks are usually proteins that interact via binding to each other, to organic molecules or to nucleic acids. In this way, proteins catalyze chemical reactions, regulate the expression of genes and transduce signals within and between cells. The elucidation of the respective molecular functions requires the determination of the three-dimensional structures of the proteins involved and the analysis of relationships between protein structure and protein function (*"Structure Function Relationships of Proteins"*, page 38). In addition, binding processes between biomolecules must be modelled, and complex interaction networks of proteins must be analyzed (*"Regulatory Feedback Networks – How Little Do We Actually Know?"*, page 43).

These methods are also applied to specific case studies of infectious diseases such as AIDS, Hepatitis B (*"Analysis of HBV Resistance"*, page 76) and Hepatitis C (*"Fighting the Hepatitis C Virus"*, page 39) as well as other diseases such as

cancer, neurodegenerative and immunological diseases. While the early recognition of cancer through genetic and so-called epigenetic changes constitute a research focus (*"The Genetic Foundation of Cancer"*, page 42), in other diseases the emphasis lies on the identification and characterization of disease-related proteins (*"Functional Analysis of Medically Relevant Proteins"*, page 40). Regarding the search for optimized treatments for infectious diseases, AIDS plays a special role. With this disease, the Max Planck Institute for Informatics goes even a step further. We analyze the resistances of HI viruses against administered combination drug treatments (*"Analysis of HIV Drug Resistance"*, page 77) as well as other important viral preconditions for an effective patient-specific drug therapy (*"Prediction of HIV Coreceptor Usage"*, page 41).

A substantial part of the method development in the department results in software systems which are used worldwide by many academic, clinical and industrial users. Examples, which are reported on in this 2009/2010 issue, include the field of epigenetics, the analysis of protein function (*"Structure Function Relationships of Proteins"*, page 38) and protein interaction networks as well as the optimization of AIDS

therapies (*"Analysis of HIV Drug Resistance"*, page 77, *"Prediction of HIV Coreceptor Usage"*, page 41). In the area of flu research (*"GISAID"*, page 86) we provide a world-wide-used database.

The department is one of the main pillars of the Bioinformatics Center Saar, an inter-faculty center at Saarland University focusing on teaching and research in the area of bioinformatics. The department is a member of the German Avenir network as well as the European Consortium "EuResist", both of which conduct bioinformatics research for analyzing viral drug resistance, part of the Clinical Research Group 129 of Deutsche Forschungsgemeinschaft on Hepatitis C, as well as of the National Genome Research Network supported by the German Science Ministry. ...

Computer Graphics

PROF. DR. HANS-PETER SEIDEL

DEPT. 4



The Computer Graphics work group was founded in 1999 and includes 40 scientists today.

An important characteristic of the work is the thorough consideration of the entire process chain, from data acquisition through modeling to image synthesis (3D image analysis and synthesis). Typical in this area is the bringing together of very large data sets with the requirement for faster, and more interactive presentation, where possible.



CONTACT

Computer Graphics

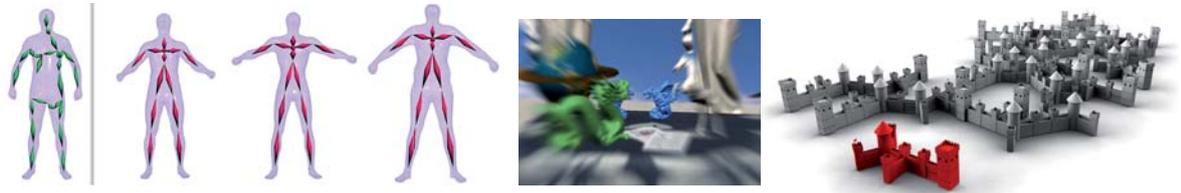
Secretary

Sabine Budde

Phone +49 681 9325-4000

Email budde@mpi-inf.mpg.de

DEPT. 4



Computers are often deployed today to construct segments of the real or a virtual world on the computer, in order to simulate and present the world. Due to the meaning of visual information for people, computer graphics in the previous decade have become a key technology for modern information and communication companies, whose future applications potential through headline words such as multimedia, digital television, telecommunications, virtual reality or 3D internet are merely a partial indication. Typical for this area is bringing together very large data sets with the requirement for faster (if possible, more interactive) visual presentation of results with high image quality. In addition, the user has the intention of interacting as much as possible in an interactive way with his environment.

Due to the above-mentioned challenges, also new approaches are needed from a scientific point of view. An important characteristic in the work group is therefore the thorough consideration of the entire processing chain from data acquisition via modeling (creation of a suitable internal computer scene description) to image synthesis (generation of arbitrary views). These integrated views are necessary in order to adequately use the performance capabilities of modern hardware, the input (image generating processes) as well as the output (graphics hardware). In the meantime, for these integrated views, the term 3D image analysis and synthesis has been developed. As central scientific challenges arising from this development, especially appropriate modeling tools for efficient handling and further processing of data flows from the input side as well as the development of new algorithms for fast and therefore qualitatively high-value pre-

sentation with close linkage with the possibilities and prospects of modern graphics on the output side.

The scientific activities of the Computer Graphics work group are embedded in a series of project activities on a national, European and international level.

Especially important is the jointly founded "Max Plank Center for Visual Computing and Communication", developed with the support of the Max Planck Institute and Stanford University with significant support from the BMBF. It was founded in October 2003 and has been extended in recent years for an additional funding phase. The aim of this bridge-building between the two top locations in Germany and the United States is to strengthen research efforts in this key area of modern information and communications technology. By establishing new exchange mechanisms with attractive return opportunities, there is a significant contribution to training and winning back outstanding young scientists. The direction of the center is in the hands of Professor Bernd Girod (Stanford University) and Professor Hans-Peter Seidel (Max Planck Institute for Informatics)

In addition, the group is significantly involved in the activities of the "Multimodal Computing and Interaction". The excellence cluster was newly formed in 2007 within the framework of an excellence initiative by the German federal and state governments. After the kick-off workshop in November 2008, the second regular inspection of the excellence cluster was held by the advisory board in November 2010. The chairman of the board, Prof. Thomas Gross (ETH Zurich), was particularly impressed by

the interdisciplinary approach of the cluster: "Particularly impressive is that different research groups with different approaches work together very closely within the framework of this cluster. This allows them to achieve much more than isolated individual projects." The scientific coordinator of the excellence cluster is Prof. Hans-Peter Seidel.

An additional important development was the founding of the "Intel Visual Computing Institute" in May 2009. The new research institute is located on the campus and is supported together by Intel, Saarland University, the DFKI, the Max Planck Institute for Informatics and the Max Planck Institute for Software Systems. Prof. H.-P. Seidel is a representative of the Max Planck Institute on the Governance Board of the Institute.

Over the past ten years, more than 30 young scientists from the group have been promoted to professorships domestically and internationally. The group has earned a number of prizes, including young scientist prizes and also the Leibniz Prize of the German Research Society for Professor H.-P. Seidel. With the appointment of Christian Theobalt (Stanford) to a W2 permanent chair at the Institute, the group was further strengthened last year. ...

Databases and Information Systems

PROF. DR.-ING. GERHARD WEIKUM

DEPT. 5



The Department, which is led by Gerhard Weikum, pursues research in five topic areas:

1. Knowledge discovery on the Web with statistics- and logics-based methods for automatic fact extraction from Internet sources like Wikipedia;
2. Text mining for the automatic classification of documents and for the identification of interesting patterns in text corpora, especially in news and web archives over long timescales;
3. Ranking and inference methods for queries for which only the top k answers are important, and for coping with uncertain data (e.g. automatically extracted relations from text);
4. Query processing and optimization of execution plans for efficient search on structured and semistructured data (e.g. in XML or RDF format);
5. Analysis of distributed data, especially in scalable peer-to-peer systems and online communities, for example, in social networks and Web 2.0 media.



CONTACT

Databases and Information Systems

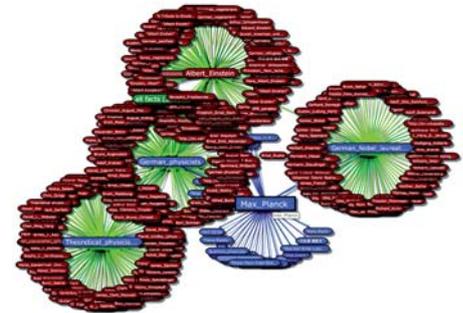
Secretary

Petra Schaaf

Phone +49 681 9325-5000

Email schaaf@mpi-inf.mpg.de

DEPT. 5



A central theme in the scientific work of the department is the automatic gathering of comprehensive knowledge bases from information on the World Wide Web as well as discovering, tracking, and analyzing patterns and individual entities (people, organisations, etc.) and their cross-relationships in dynamic web sources. In the YAGO-NAGA project, a very large knowledge collection called YAGO (*Yet Another Great Ontology*) has been created using fact extraction from Wikipedia and integrating the WordNet taxonomy. For exploration and intelligent search in YAGO, a new type of search engine called NAGA (*Not Another Google Answer*) has been developed. Both use the semantic web data model RDF (*Resource Description Framework*), and the group has developed one of the fastest RDF search engines called RDF-3X (*RDF Triple Express*). A variety of further research topics are addressed within the YAGO-NAGA framework: Inclusion of multilingual information, systematic gathering of multimodal data (e.g. photos of people), analysis of time changes in knowledge, deduction of new relationships on the basis of probable but uncertain knowledge, and more.

The vision that drives this work is the expected convergence of the *Semantic Web* with its formal ontologies and logic-oriented search and inference, the *Social Web* (Web 2.0) with its latent “*Wisdom of the Crowds*”, and the de-facto predominating *Statistical Web* with its statistical approach to data cap-

turing and probabilistic underpinnings of search engines. The web could be the basis of a comprehensive knowledge base, which contains the entire encyclopedic knowledge of mankind in a formal, structured, and machine-readable representation, and enables higher-level programs and intelligent services. The benefits of such a knowledge base would be enormous.

The group is a worldwide leader in its methods of combining logic-based algorithms for structured data sets and statistically-based methods for unstructured textual data. The former fall into the area of database systems (DB) and data analysis, the latter into the area of information retrieval (IR) and search engines. These two directions were historically divided; their connection has become ever more important with increasingly mixed data forms for digital libraries, social online communities, e-science groups and, last but not least, for companies and in the Web itself. The group is a trendsetter in this strategic direction of DB-IR integration.

The group’s research methodology spans the entire spectrum from theoretical foundations to practical systems building for applications and experiments. Many of the prototype systems developed by the group are publicly available as open source software and are used by other research groups around the world. These include, in particular:

1. The XML search engine *TopX*, which has achieved top ranks in the INEX benchmarking series and provides the semantically annotated Wikipedia reference corpus for the INEX competition;
2. The RDF search engine *RDF-3X*, which can search semantic web data and other graph-structured data very efficiently for complex patterns;
3. The software tools that automatically create and maintain the YAGO knowledge base, as well as the YAGO contents itself;
4. The scalable system *PROSPERA*, which extracts new facts from web pages and natural-language texts.

The department is involved in a number of externally funded projects, especially the EU projects “*Living Knowledge*” and “*Longitudinal Analytics of Web Archive Data*” as well as the DFG Excellence Cluster “*Multimodal Computing and Interaction*.” The group has received a Google Research Award. ...

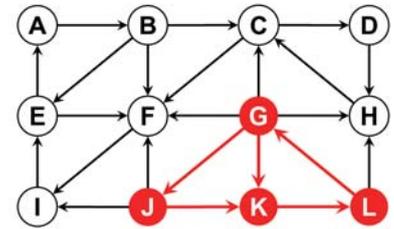
Automation of Logic

PROF. DR. CHRISTOPH WEIDENBACH

RG. 1



RG. 1 Automation of Logic



The independent research group "Automation of Logic", under the direction of Prof. Dr. Christoph Weidenbach, serves the complete pipeline from basic research on new logics to industrially used tools.

Logics are formal languages for calculation with mathematically precise meaning and exact rules for reasoning. A simple example of a logic are linear equation systems that we know from high school, with a solution, that includes the calculation method. Logic was developed at the end of the 19th century in order to exactly describe and calculate general mathematical arguments. Since the invention of computers and information technology, they have been further developed for computer systems so that these could also be described precisely and properties could be formulated, analyzed and ultimately proven. By about the mid-1990s, such analytical and proof efforts required massive manual interaction and therefore a great deal of time. Proving or analyzing computer systems properties is typically "hard", that is, even if there are automatic procedures available, the number of calculation steps grows at least exponentially with the size of the problem. From the mid-90s to date, however, for a number of practically relevant properties, procedures have been developed to work fully automatically in reasonable time. In sum, the application of logic was extended from mathematics to properties of computer systems in the broadest sense. In addition, it has been found that many of the techniques developed are generally successful on hard problems, even if they do not originate from the formal analysis of computer systems and mathematics. The application of analysis and proof techniques has thus once again extended to general hard problems in recent years, e.g., combinatorial optimization problems.

Our research group develops automatic, logic-based procedures. These should be capable of computing solutions for a class of applications within an acceptable amount of time. We are particularly interested in the verification of controls, protocols, safety properties, software and in solutions for decision and optimization problems. In order to keep up with increasingly complex systems and required properties, the formal procedures used today must significantly increase in productivity. This is the goal of our research group, "Automation of Logic".

CONTACT

Automation of Logic

Prof. Dr. Christoph Weidenbach

Secretary Jennifer Müller

Phone +49 681 9325-2900

Email jmueller@mpi-inf.mpg.de

Max Planck Center for Visual Computing and Communication

PROJECT DIRECTOR: PROF. DR. HANS-PETER SEIDEL

The "Max Planck Center for Visual Computing and Communication" was founded in 2003 with substantial support from the BMBF (Federal Ministry for Education and Research). The Max Planck Center brings together the Max Planck Institute for Informatics in Saarbrücken and Stanford University. It has been granted 6.9 million Euros from the BMBF since its founding. Due to the success of the program, in recent years, funds of up to 7.8 million Euros have been approved for a further support phase.

The Max Planck Center for Visual Computing and Communication (MPC-VCC) with corresponding research activities at the Max Planck Institute for Computer Science and at Stanford University was established jointly by the Max Planck Society and Stanford University in October 2003. In 2010 BMBF approved continuation of the center for another five year period. The proposed collaboration has two intertwined goals: Establish a joint research program in the area of visual computing and communication and incorporate a strong career development component to alleviate the shortage of qualified faculty and scientists in information technology in Germany. The Max Planck Center fosters the professional development of a small number of selected outstanding individuals by providing them with the opportunity to work at Stanford University as visiting assistant professors in the area of visual computing and communication for two years and then return to Germany to continue their research as leader of a junior research group at the Max Planck Institute for Computer Science and ultimately as a professor at a German university.

The center is directed jointly by Hans-Peter Seidel (Max Planck Institute for Informatics) and Bernd Girod and Leo Guibas (Stanford).



A successful alumni of the program is Christian Theobalt. He rejoined the Max Planck Institute for Informatics as a tenured faculty member (W2) after his previous appointment as Visiting Assistant Professor at Stanford University. Since December 2010, Christian Theobalt is also an Associate Professor in the Computer Science Department of Saarland University.



CONTACT

Prof. Dr. Hans-Peter Seidel

Secretary Sabine Budde

Phone +49 681 9325-4000

Email budde@mpi-inf.mpg.de

Research Areas

UNDERSTANDING IMAGES & VIDEOS

- 30 People Detection and Pose Estimation
in Challenging Real-World Scenes
- 31 Scalability of Object Class Recognition
- 32 3D Scene Understanding from Monocular Cameras
- 33 Image-based 3D Scene Analysis
- 34 Markerless Reconstruction of Dynamic Scenes

BIOINFORMATICS

- 38 Structure Function Relationships of Proteins
- 39 Fighting the Hepatitis C Virus
- 40 Functional Analysis of Medically Relevant Proteins
- 41 Prediction of HIV Coreceptor Usage
- 42 The Genetic Foundation of Cancer
- 43 Regulatory Feedback Networks – How Little Do We Actually Know?

GUARANTEES

- 46 Dealing with Selfishness in Optimization
- 47 Geometric Algorithms – Exact and Efficient
- 48 Controlled Perturbation – How to Avoid Difficult
Special Situations
- 49 Automated Deduction
- 50 Model Checking for Hybrid Systems
- 51 Modular Proving in Complex Theories

INFORMATION SEARCH & DIGITAL KNOWLEDGE

- 54 Random Telephone Chains – Efficient Communication
in Data Networks
- 55 Search and Mining in Web Archives
- 56 Information Search in Social Networks
- 57 YAGO – a Collection of Digital Knowledge
- 58 URDF – Efficient Reasoning in Uncertain RDF Knowledge Bases
- 59 EnBlogue – What is New and Interesting in Web 2.0?
- 60 Decision Procedures for Ontologies

**MULTIMODAL
INFORMATION**

- 64** : Recognizing Human Activity
65 : Automated Music Processing

OPTIMIZATION

- 68** : The Theory of Bio-inspired Algorithms
69 : Approximation Algorithms for Hard Optimization Problems
70 : Placement of Wind Turbines
71 : Combinatorial Optimization in Information Science
72 : Model-Based Product Configuration

SOFTWARE

- 76** : Analysis of HBV Resistance
77 : Analysis of HIV Drug Resistance
78 : GBACE: Parallel Processing of Adaptive Data
79 : TopX 2.0 – Efficient Search in Digital Libraries
80 : RDF-3X – Fast Searches for Semantic Data

VISUALIZATION

- 84** : HDR – Images and Video with Enhanced Contrast
85 : Advanced Real-Time Rendering
86 : Computational Photography
87 : Statistical Models of People
88 : Correspondences, Symmetries and Automatic Modeling

UNDERSTANDING IMAGES & VIDEOS

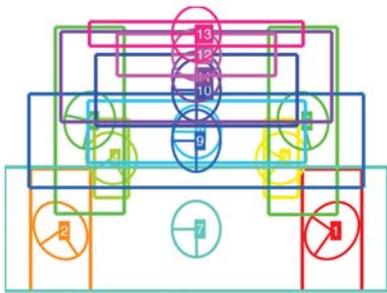
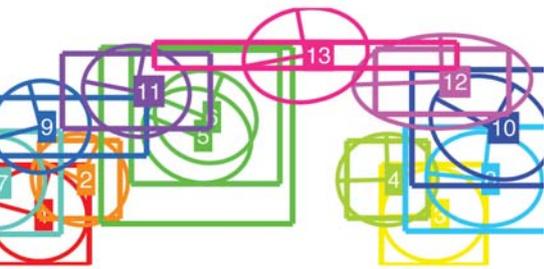
Understanding images and videos is one of the fundamental problems of image processing. The scientific challenges range from the modeling and tracking of people and objects in camera systems to the reconstruction and description of 3D scenes. This area has many applications; these include the animation of people and visualization of 3D scenes, the indexing of image and video material, and also 3D-capturing of the environment for automotive applications. This research area is thus situated at the intersection between computer vision and computer graphics, resulting in numerous opportunities for cooperation within the institute.

Various research groups within the Max Planck Institute for Informatics work on the different aspects of image and video understanding. In the field of modeling people, work is carried out on the reconstruction of animated models from multi-video data, for example. The primary goal here is the ability to model and visualize people in detail and as accurately as possible. Furthermore, methods for people detection and tracking are researched that require monocular cameras only and that are able to detect and track many people simultaneously in complex scenes. Even if these studies pursue fundamentally different targets and use different camera configurations, they benefit from each other, as these areas use similar models and algorithms.

A further area of understanding images and video is the reconstruction and description of 3D scenes. Again, there are different studies within the institute that benefit from each other. For example, 3D scenes are reconstructed from image series to enable a surface description of 3D objects that is highly detailed and accurate. Furthermore, the 3D environment of a moving car is described. Here, a special focus is on the complete identification and description of pedestrians and other road users. All this work shows that the integration and use of 3D scene models can significantly improve results.

One of the fundamental problems of understanding images is the recognition of objects. With the current omnipresence of digital image material, automatic visual object-class recognition techniques are becoming increasingly important. Even if today's approaches can achieve remarkable results, one of the most significant problems remains: How can object models be learned or constructed? Therefore, a central topic of various research work at the institute concerns how such object models can be constructed with as little manual input as possible to enable a broad applicability of today's approaches. ...

CONTRIBUTIONS



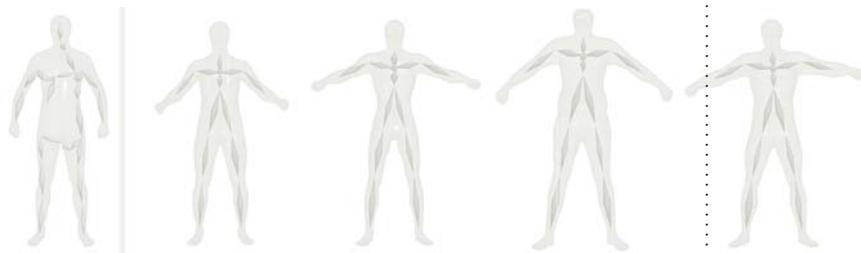
**People Detection and Pose Estimation
in Challenging Real-World Scenes** 30

Scalability of Object Class Recognition 31

3D Scene Understanding from Monocular Cameras 32

Image-based 3D Scene Analysis 33

Markerless Reconstruction of Dynamic Scenes 34



UNDERSTANDING
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH
& DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION

People Detection and Pose Estimation in Challenging Real-World Scenes



Examples of people detection and pose estimation obtained with our approach

Finding and following people is a key technology for many applications such as robotics and automotive safety, human-computer interaction scenarios, or for indexing images and videos from the web or surveillance cameras. At the same time it is one of the most challenging problems in computer vision and remains a scientific challenge for realistic scenes.

We have developed a new approach for detecting people and estimating their poses in complex street scenes with multiple people and dynamic backgrounds. Our approach does not require multiple synchronized video streams but can operate on images from a moving, monocular and uncalibrated camera. Several examples of people detections and estimated body configurations are shown on the pictures. The important challenges addressed in our approach are frequently full and partial occlusions of people, cluttered and dynamically changing backgrounds, and ambiguities in recovering 3D body poses from monocular data.

2D human pose estimation

Several key components contribute to the success of the approach. The first is a novel and generic procedure for people detection and 2D pose estimation that is based on the pictorial structures model and also enables to estimate viewpoints of people from monocular and single images.

Our approach is based on learned local appearance models of human body parts and a kinematic model of the human body. The appearance of the body parts is represented by a set of local image descriptors. We employ a boosting classifier trained on the dataset of annotated human poses in order to learn which of these local features are informative for the presence of the body part at a given image location. Interpreting the output of each classifier as a local likelihood we infer the optimal configuration of the body parts using belief propagation.

3D human pose estimation and multi-person tracking

The second key component in our approach enables people tracking and 3D pose estimation. In contrast to prior work, we accumulate evidence in multiple stages thereby reducing the ambiguity of 3D pose estimation effectively in every step.

In particular we propose a novel multi-stage inference procedure for 3D pose estimation. Our approach goes beyond prior work in this area, in which pose likelihoods are often based on simple image features such as silhouettes and edge maps and which often assume conditional independence of evidence in adjacent frames of the sequence. In contrast to these approaches, our 3D pose likelihood is formulated in terms of estimates of the 2D body configurations and viewpoints obtained with a strong discriminative appearance model. In addition, we refine and improve these estimates by tracking them over time, which allows to detect occlusion events and group hypotheses corresponding to the same person. We demonstrate that the combination of these estimates significantly constrains 3D pose estimation and allows to prune many of the local minima that otherwise hinder successful optimization. ...



CONTACT

Mykhaylo Andriluka
 DEPT. 2 Computer Vision and Multimodal Computing
 Phone +49 681 9325-2119
 Email andriluk@mpi-inf.mpg.de



Bernt Schiele
 DEPT. 2 Computer Vision and Multimodal Computing
 Phone +49 681 9325-2000
 Email schiele@mpi-inf.mpg.de

Scalability of Object Class Recognition

Because of the omnipresence of digital imagery today, automatic techniques for the visual recognition of objects are becoming more and more important. The focus here lies on the recognition of entire classes of objects, such as cars, rather than specific instances, such as one particular red sports car. While current state-of-the-art systems achieve remarkable recognition performance for individual classes, the simultaneous recognition of multiple classes remains a major challenge: building reliable object class models requires a sufficiently high number of representative training examples, often in the form of manually annotated images. Since manual annotation is costly, our research aims at reducing the amount of required training examples for building object class models, thereby increasing scalability. We explore three different ways of achieving this goal.

Re-using object class model components

In this project, we are devising an object class model that represents objects as an assembly of spatially constrained partial shapes. Similar object classes, such as horses and giraffes [Figure 1], share similar representations. Components of one model, e.g., the expected variation based on the localisation of the horses' body parts, can be re-used in order to build another model, e.g., a giraffe model. Re-used model components do not have to be learned. Hence, the number of required training examples is reduced. The focus on object shape rather than on appearance allows model components to be re-used, even for relatively dissimilar classes, such as horses and swans.

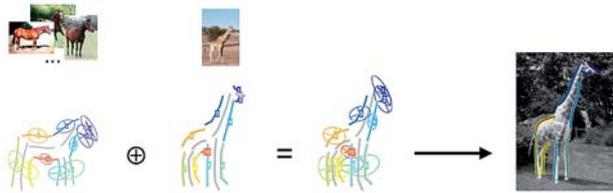


Figure 1: Re-using model components, example recognition output

Identifying re-usable components automatically

In order to apply the re-use of model components in practice, re-usable components have to be identified automatically. In this project, model components correspond either to attributes, i.e., distinct visual properties like color or texture, or to entire object class models. We propose to relate object classes and attributes by automatically mining linguistic knowledge bases. Examples of linguistic knowledge bases include, e.g., WordNet, Wikipedia, Flickr, and Yahoo web search. The mined relations can then be used to build object class models, even for classes for which no training examples are available, relying entirely on re-used components from other models. The recognition performance of the fully automatic system is demonstrated to be en par with providing the object class-attribute relations by human subjects.

Learning from 3D CAD data

In an attempt to abandon training images altogether, this project aims at learning object class models from 3D computer-aided design (CAD) data [Figure 2]. CAD models are typically

employed in product design, computer games, and movie making, and are readily available for many object classes. Since CAD models provide an accurate and detailed representation of object shape, they lend themselves to the learning of object class models. Since they are viewpoint independent, they allow for arbitrary view-points, that means artificial training examples, through rendering. The key challenge consists in bridging the gap between rendered training images and real-world test images, which we achieve by a shape-based abstraction of object appearance. Our experiments demonstrate state-of-the-art performance for the recognition of cars from multiple viewpoints, even compared to approaches using real-world training images. ...

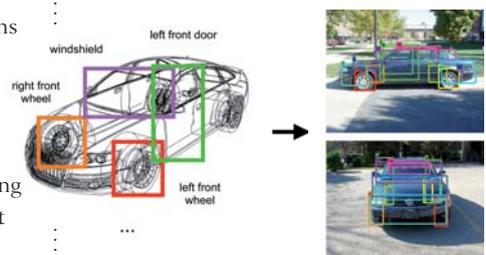


Figure 2: Learning from 3D CAD data, example recognition output



CONTACT

Marcus Rohrbach
 DEPT. 2 Computer Vision and Multimodal Computing
 Phone +49 681 9325-1206
 Email rohrbach@mpi-inf.mpg.de



Michael Stark
 DEPT. 2 Computer Vision and Multimodal Computing
 Phone +49 681 9325-2120
 Email stark@mpi-inf.mpg.de
 Internet <http://www.d2.mpi-inf.mpg.de/nlp4vision>

3D Scene Understanding from Monocular Cameras

Scene understanding – a longstanding goal in computer vision

Inspired by the human visual system, visual scene understanding has been advocated as the “holy grail” for machine vision since its beginning. In the early days this problem was addressed in a bottom-up fashion, starting from low-level features such as edges and their orientation as the key ingredients in the derivation of a complete scene description and scene understanding. Unfortunately, the reliable extraction of such low-level features proved very difficult due to their limited expressive power and their inherent ambiguities. As a result and despite enormous efforts, scene understanding remained an illusive goal even for relatively constrained and simple scenes. Disappointed by these early attempts, the research community has turned to easier sub-problems and has achieved remarkable results in areas such as camera geometry, image segmentation, object detection, and tracking. As the performance for these sub-tasks starts to achieve remarkable performance levels, we believe that computer vision should reinvestigate the problem to automatically understand 3D scenes from still images and video sequences.

Robotics and automotive applications

Robotics, such as mobile service robots, as well as automotive safety applications, such as pedestrian protection systems, are clearly of major scientific and commercial interest. Therefore, we use the recognition of pedestrians and cars from a moving camera mounted on a vehicle or robot as a test case and running example for our work. We can leverage domain knowledge for both application areas. The camera, for instance, can be calibrated relatively to its surrounding area, therefore, camera height above ground and the orientation, that is relatively to its surrounding area, are constrained and approximately known. Moreover, it can be assumed that the ground is flat in a local neighborhood and that all objects are supported by it. Additionally, we include a Gaussian prior

for pedestrian heights as well as for the heights of cars. These assumptions pose additional constraints to eliminate ambiguities and allow us to solve the scene understanding problem for such application domains more easily. See figure 1 for a sample scene.

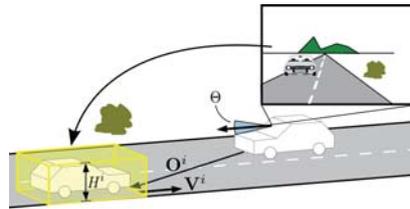


Figure 1: Our system aims to infer a 3D scene based on a combination of prior domain knowledge with videos of a monocular camera.

Integrated 3D scene model

The system that we have developed combines the above prior domain knowledge with powerful state-of-the-art *object class detectors*, *semantic scene labeling* and the notion of *tracklets*. Object class detectors determine 2D object locations in the image, and semantic scene labeling infers semantic classes such as roads, the sky, or objects for each pixel, whereas tracklets, due to geometric and dynamic consistency over a longer period of time, accumulate image information over several frames to achieve a more robust reasoning (see figure 2 for sample detections, segmentation, and system results).

By employing 3D reasoning, our model is able to represent complex interactions like inter-object occlusion, physical exclusion between objects, and geometric context. The results are encouraging and show that the combination of the individual components not only allows us to infer a 3D world representation but also improves the baseline given by the object class detectors. Throughout this research project we have used monocular cameras and are therefore not able to directly extract depth from the image disparity. Nonetheless, our approach is able to outperform similar systems that use a stereo camera setup.

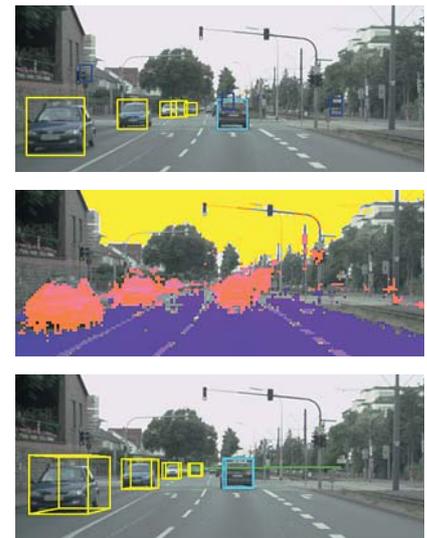


Figure 2: 2D detections, semantic scene labeling and results of our 3D scene estimation



CONTACT

Bert Schiele
 DEPT. 2 Computer Vision and Multimodal Computing
 Phone +49 681 9325-2000
 Email schiele@mpi-inf.mpg.de



Christian Wojek
 DEPT. 2 Computer Vision and Multimodal Computing
 Phone +49 681 9325-2000
 Email cwojek@mpi-inf.mpg.de
 Internet <http://www.d2.mpi-inf.mpg.de/monocular-3d-scene>

Image-based 3D Scene Analysis

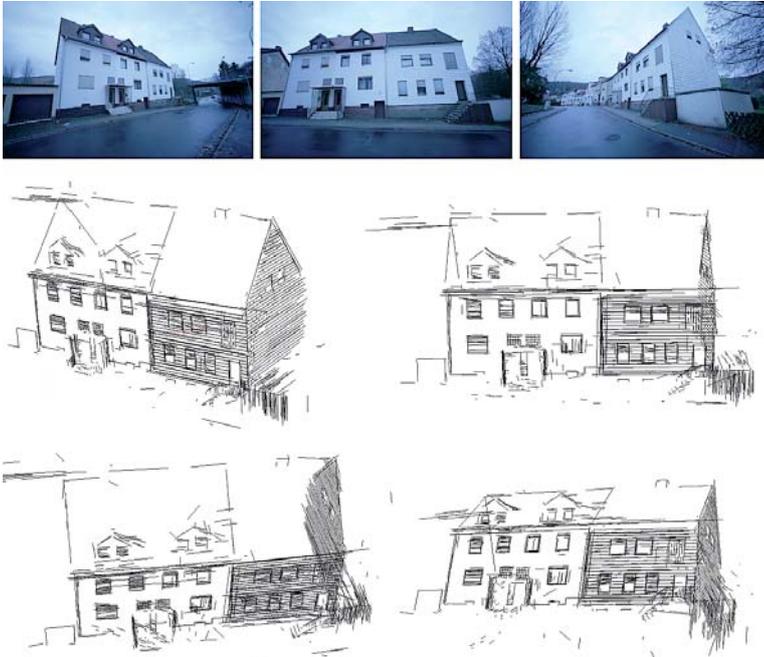


Figure 1: *Top: Input images, Bottom: line-based 3D reconstruction*

Line-based 3D reconstruction

Automatic approaches for 3D reconstruction from image sequences usually produce, in a first step, merely a 3D point cloud. The next step is then to generate a surface reconstruction of the 3D objects. The reliable reconstruction of 3D lines is thereby an important interim step. We have developed an approach that improves the reconstruction of 3D lines from image sequences. The advantage of our approach is that not every 3D line is estimated individually, but rather the global connectivity constraints between neighboring lines can be exploited. The 3D positions of connections of lines

can be estimated more reliably compared to estimates for individual lines. Our approach optimizes the 3D position of the lines and their connections at the same time, achieving a better 3D construction.

Learning skeletons for the manipulation of body shape and pose

The project uses a large number of 3D scans of people with different body shapes and in different poses to automatically extract a skeleton. There are already existing approaches that allow automatic extraction of a skeleton,

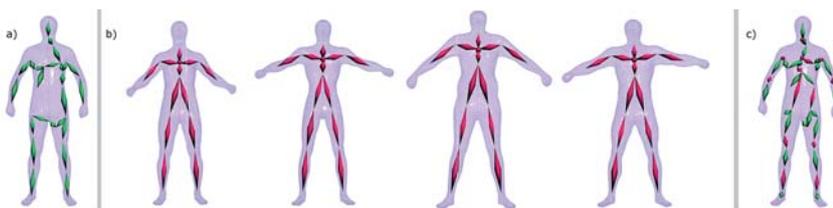


Figure 2:
 These skeletons were automatically generated from different 3D scans of people:
 a) a skeleton that changes body shape
 b) different body shapes with a skeleton that allows to change the person's pose
 c) a combined skeleton that allows to change body shape as well as the pose

but these skeletons are limited to manipulating the pose. Our approach produces a skeleton that can change the body shape; and, in addition, it produces a combined skeleton that can change both the shape and pose of the body [see figure 2]. This skeleton-based representation has the advantage that current 3D modeling packages or 3D games, which use skeletons to change poses, can now also change the body shape without adjusting their internal representation.

Production of artificial camera shake

Animated camera movements in virtual 3D scenes are frequently criticized because they look too perfect. They are missing the slight camera shake which is observed in true camera recording. In our project, we have developed a tool that analyzes the camera shake of real cameras and transfers this onto a virtual camera's path. A database of various real camera recordings was established, and camera movements were automatically estimated from the image sequences. In addition, the style of the observed camera shake was also stored in the database, for example, whether the camera movement took place in a helicopter, on a bicycle, on a skateboard, etc. A 3D animator can now produce a virtual camera path and then create camera shake in the desired style. With our approach only the spectrum of the camera shake is transferred and not the actual camera shake itself. This prevents unwanted repetitions in the shaking motion of the camera. ...



CONTACT

Thorsten Thormählen
 DEPT. 4 Computer Graphics
 Phone +49 681 9325-4017
 Email thormae@mpi-inf.mpg.de

Markerless Reconstruction of Dynamic Scenes

Reconstruction of detailed animation models from multi-video data

Computer-generated people, so-called avatars, have become an important component of visual media, such as, for example, in computer animations, films or networked virtual environments. In order to be able to present virtual humans in a convincing way, the individual characteristics of the person, such as, for example, his movement, geometry and surface texture, must be extremely realistically modeled. One way to achieve this is defining every partial aspect of the overall appearance manually in an animation program. This is a very time-intensive and complex process. The geometry of the person must be constructed in precise detail work, and each nuance of movement must be finely specified. It is therefore easy to imagine that completely manually-produced animations, especially with respect to the quality of movement animation, do not achieve the level of detail of a true human.

The alternative to manual modeling consists of measuring some aspects of animation from real people. Motion capture systems make it possible to reconstruct a skeletal model of a person using video image streams. Unfortunately, this method of movement capture is a very complex process, and the person who is measured must often wear a special skin-tight suit with optical markers. In addition, with such systems, neither the time changing geometry nor the texture of a person is captured.

In our research we have developed a new type of *Performance Capture Algorithms*. For the first time it was possible to reconstruct detailed movement, dynamic geometry and dynamic texture of a person in complex clothing, such as, for example, a dress or a ball gown, solely from multi-video streams. Our process requires no optical markers in the scene.



Figure 1: *Left:* An image from a multi-video sequence which was recorded with eight cameras. *Right:* The reconstructed geometric, skeletal and clothing model.

An important step forward was the new development of a process to estimate the movements of a person in arbitrary clothing, and to learn an animation model that enables us to change the movement of the reconstructed person arbitrarily [Figure 1]. A core component of this is a new algorithm which estimates the physical material and simulation parameters of clothing. A further milestone was the development of a combined segmentation and reconstruction process, with whose assistance, movement and geometry of several interacting people can be captured.

Real-time motion capture with depth sensors

The Performance Capture Methods for exact detailed reconstruction, which were described in the previous section, require very complex calculations and are therefore not real-time. In addition, these methods require several video cameras. Motion measurement from a single camera perspective is an extremely complex and highly underconstrained problem.

New types of depth cameras, such as, for example, so-called *Time-of-Flight cameras* measure 2.5D scene geometry in real time. By using both depth and video data together it is easier to reconstruct body movements than from a single camera perspective. Unfortunately, depth camera data is very noisy, has low resolution and exhibits a systematic measurement bias. We have therefore developed methods in order to calibrate depth sensors which eliminate noise, and increasing camera resolution through calculations on the original data (*super-resolution process*).

With improved depth data and through a newly-developed process from our research, all body movements of a person can be measured from one single camera perspective. The real time algorithm to measure movement combines a depth-based pose optimization procedure with a procedure for fast finding of similar poses from a large database of movement sequences [Figure 2].



Figure 2: The movement of the person (*bottom left as an intensity image*) is reconstructed from the *Time-of-Flight depth image (center)* in real time (*right*).



CONTACT

Christian Theobalt

DEPT. 4 Computer Graphics

Phone +49 681 9325-4028

Email theobalt@mpi-inf.mpg.de

Internet <http://www.mpi-inf.mpg.de/resources/perfcap/>

BIOINFORMATICS

Bioinformatics is a key discipline for progress in biological sciences, such as biotechnology, pharmacy and medicine. With computational methods Bioinformatics deepens and accelerates planning highly complex biological experiments as well as interpreting the resulting very large amounts of data.

For about 20 years, bioinformatics has significantly contributed to progress in the life sciences. Therefore, the field is a key ingredient of the revolution of biology. Bioinformatics methods and software support researchers in planning experiments, store data which come from all areas of biology, and assess this data with computer support. With the help of Bioinformatics, scientists can elucidate the molecular processes in cells – the basic units of living organisms – which form a complex system for processing matter, energy and information. The genome harbors the construction plan for the cell and the working plan of its metabolic and regulatory processes in a complex, encoded form. In order to facilitate these cellular processes, parts of the genome must be “read off”. This includes the genes, which amount to the blueprints of proteins, the molecular machinery of the cell. In turn, reading off the genes is controlled via complex molecular networks. For the synthesis of proteins, and also for their removal, there are specific molecular complexes which themselves are again subject to exquisite molecular control. The cells convert energy, they communicate with cells in their neighborhood, they take on different structures and forms, and they move. They react to changes in their environment, for example, pertaining to light, temperature and pH, and they mount defenses against invaders. Dysregulation of such processes constitutes the molecular basis of disease. Drug therapy aims to restore a tolerable molecular balance.

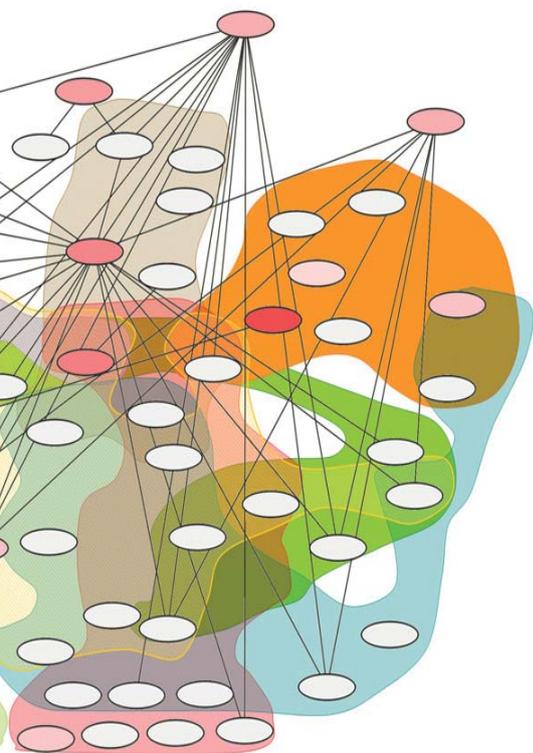
For the past two decades, classical biological research, which used to be concentrated on cellular sub-systems with limited complexity, has been complemented by high-throughput experiments. The respective screening methods

afford cell-wide capture of molecular data, e.g., by measuring the frequency of all read off genes (transcription), by capturing the population of the proteins used by the cell (proteome), and of their interactions (interactome). Generating new insight into the biology of the cell as well as the basis of diseases and providing new approaches to therapy involves highly complex information processing. Bioinformatics addresses this challenge. The Max Planck Institute for Informatics performs research on many of the challenges mentioned above.

Bioinformatics has the hybrid character of a basic science which poses and follows clear application perspectives. This unique quality is highlighted by a significant number of spin-offs of bioinformatics research groups. For example, Professor Lengauer co-founded the company BioSolveIT GmbH, which develops and distributes software for drug design. Pharmaceutical companies all over the world are among the users of this software.

The Bioinformatics Center Saar, whose speaker is Professor Lengauer, received top marks for its research in the final evaluation in 2007 of the five bioinformatics centers that have been funded by the German Research Foundation (DFG) in the past decade. ...

CONTRIBUTIONS



Structure Function Relationships of Proteins 38

Fighting the Hepatitis C Virus 39

Functional Analysis of Medically Relevant Proteins 40

Prediction of HIV Coreceptor Usage 41

The Genetic Foundation of Cancer 42

Regulatory Feedback Networks – How Little Do We Actually Know? 43

UNDERSTANDING
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

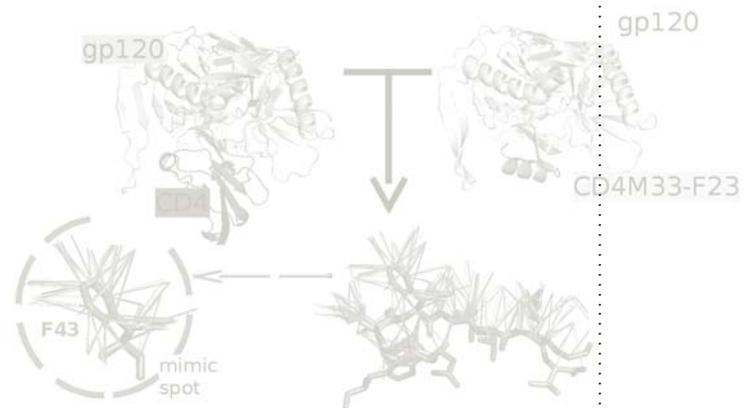
INFORMATION SEARCH
& DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION



Structure Function Relationships of Proteins

Proteins are the essential building blocks of life

Proteins carry out many tasks in cells. Most cellular processes are based on the interaction of several proteins. However, for many of the approximately 25,000 protein-coding genes in humans, nothing is known yet about the function of the respective proteins. A protein is a long chain of amino acids, which folds into a characteristic three-dimensional structure. The protein's 3D structure directly affects its function. For example, the structure of the protein surface determines whether and how well a protein can bind to another specific protein with a complementary surface. In addition to the interaction with other proteins, the structure also plays a role in interaction with DNA, RNA, or small molecules, such as drugs. The molecular structure therefore relates to questions such as: Which proteins bind together? Which protein regions bind together? How strong or lasting is such a binding? And what is the role of these individual functions in the larger context of cellular processes?

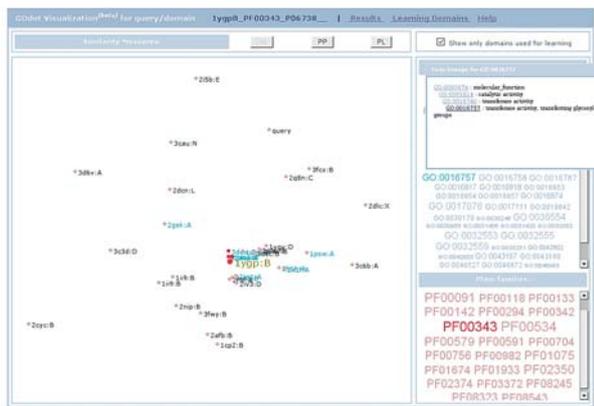
Methods

Structure-function relationships of proteins are one research focus in Department 3 of the Max Planck Institute for Informatics. We highlight two methods here: (1) The prediction of molecular function when protein structures are given. (2) The comparison of binding sites and interaction patterns at the interfaces of proteins to reveal recurring patterns.

Function prediction: GOdot

The GOdot system assesses the level of conservation of function in regions of protein structure space. Protein structure space is defined in terms of measuring all-against-all structural similarities of reference proteins. New proteins can be embedded into that space by relating them to the reference domains. The molecular functions annotated to these proteins are analyzed. Various regions in protein structure space exhibit considerable difference in the local conservation of function with respect to structure is modeled and based on these models. So the GOdot method predicts molecular function of a query protein with known structure but unknown function. The GOdot method is available online at <http://godot.bioinf.mpi-inf.mpg.de>.

In addition to prediction of function, the web-server visualizes target proteins embedded into the surrounding protein space; each protein is plotted as a point, and similar proteins are drawn close to each other. The proteins can then be colored according to molecular functions and family affiliations (in Figure 1, the local neighborhood of protein 1gyp:B is visualized, whereby proteins with the function 'GO:0006757 = transferase activity, transferring glycosyl groups' are marked in blue and members of the PFAM family 'PF00343 = phosphorylase' in red).



Fighting the Hepatitis C Virus

Infectious disease with world-wide spread

More than two hundred million people worldwide are affected by the infectious disease hepatitis C – most of them without even knowing it. The disease is triggered by the hepatitis C virus (HCV), which was only discovered in 1989. At first the disease seems to be inconspicuous, but without medication it progressively damages the liver, eventually provoking liver cancer and oftentimes leading to death. Like HIV and many other viruses, HCV is extremely adaptable and rapidly disguises its appearance by changing its genome sequence. As a result, antiviral medicine becomes ineffective, and to date there is still no vaccine against HCV. At the moment, currently approved drugs do not work effectively for fifty percent of all HCV infected patients in Europe. Therefore, the bioinformatics research at the Max Planck Institute for Informatics supports leading clinical research groups in the search for better drugs against HCV. To fight the virus, the scientists use two complementary approaches.

Analysis of viral sequence changes

HCV reacts to drugs that block its replication cycle by changing its genome. For this reason, a first approach for the bioinformatics support of HCV therapy is the computational analysis of the viral genomes found in patients. In this analysis, sequence changes that allow the virus to resist the latest antiviral drugs are identified. In cooperation with the University Hospital of Frankfurt, the sequences of viruses in patients are determined at different time points during drug therapy. Since the next-generation sequencing technology produces huge volumes of sequence data with high accuracy, millions of sequence fragments have to be mapped to the corresponding positions in the viral genome [Figure 1]. This enormous task can only be solved by powerful computers and requires the development and application of sophisticated algorithms. Subsequently, statistical methods are used to identify changes in the viral genomes that render the hepatitis C virus insensitive to antiviral drugs.

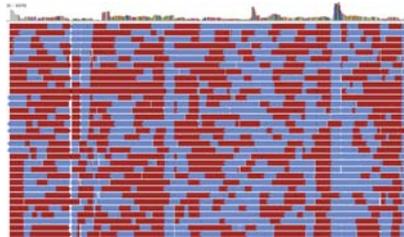


Figure 1: Viral sequence fragments (in red and blue) and their mapping to the genome of the hepatitis C virus (on the top)

This information gives doctors important clues about the course of the viral infection and also aids pharmaceutical research in developing personalized therapy and more effective drugs.

Analysis of human host factors

An alternative approach to fighting HCV consists of drugs that do not target viral molecules, but the molecular factors of the human host that are essential for the virus. The virus needs such host factors in order to replicate in human cells. Thus, these factors are promising targets for novel drugs as the virus cannot escape the drug treatment by changing its genome. To determine the host factors for HCV, researchers at the Max Planck Institute help virologists at the

University Hospital of Heidelberg since the experimental detection methods used for this purpose require special bioinformatics methods to evaluate and interpret the extensive measurement results [Figure 2]. In this way, medical bioinformatics contributes to the discovery of potential host factors and their molecular interactions in the human cell as well as to the identification of new drug targets. ...

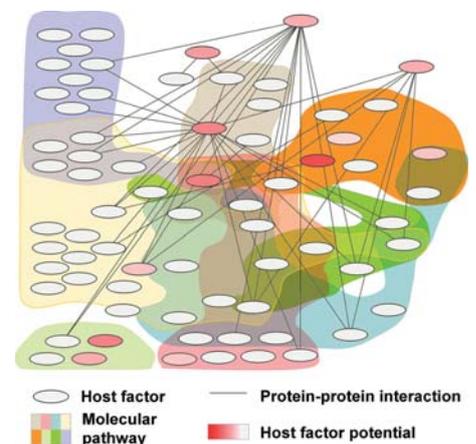


Figure 2: Schematic drawing of human host factors for the hepatitis C virus and their interactions in different molecular processes of the human cell

CONTACT

Mario Albrecht
 DEPT. 3 Computational Biology and Applied Algorithmics
 Phone +49 681 9325-3000
 Email mario.albrecht@mpi-inf.mpg.de



Hagen Blankenburg
 DEPT. 3 Computational Biology and Applied Algorithmics
 Phone +49 681 9325-3000
 Email hagen.blankenburg@mpi-inf.mpg.de



Sven-Eric Schelhorn
 DEPT. 3 Computational Biology and Applied Algorithmics
 Phone +49 681 9325-3000
 Email sven@mpi-inf.mpg.de
 Internet <http://medbioinf.mpi-inf.mpg.de>



Functional Analysis of Medically Relevant Proteins

Genetic variation with DNA

People differ from each other through individual genetic changes. The most frequent changes are *single nucleotide polymorphisms* (SNPs). They are point mutations in the DNA, that is, changes of single letters in the genetic code that occur with greater probability than other sequence variations. Population studies reveal that certain SNPs appear more often in particular subgroups and are responsible for the susceptibility to diseases and the severity of their course. Additionally, SNPs can influence the tolerability and effect of drugs. Medical research is thus particularly interested in the molecular changes that are caused by disease-relevant SNPs.

Protein structure and function

The DNA sequence of a gene is the blueprint of a protein as it determines the order of the amino acids, the molecular building blocks of proteins. Proteins have a characteristic three-dimensional structure, which is formed by atomic interactions within the protein structure. Most molecular processes occur due to interactions between proteins, and changes of the protein structure can disturb these interactions. Such modifications can be caused, for example, by point mutations like SNPs. A point mutation can lead to the exchange of an amino acid in the protein and result in significant changes of the intra- and intermolecular interactions of proteins. The thereby impaired function of the protein and the involved biological process can even provoke diseases. Well-known examples are sickle cell anemia and cystic fibrosis.

SNPs of inflammatory bowel disease

Clinical association studies, such as those carried out by our medical partners at the University Hospital of Kiel, systematically search the DNA of chronically ill persons for genetic commonalities. By comparing patients with a healthy control group, disease-associated SNPs are discovered in the DNA of the human genome. Bioinformaticians then analyze the data with computational methods to improve the understanding of the exact



The figure shows part of the very similar protein sequences of the human IL17REL protein and closely related proteins in human and other mammals. One of the two disease-associated SNPs is found at the end of the protein sequence of IL17REL at position 333. The exchange of the amino acid leucine (L) with proline (P) at this position (brown) can impact the function of IL17REL strongly.

function of relevant human proteins and their SNPs in the development and course of chronic diseases. One example of such a disease is ulcerative colitis, a chronic bowel inflammation with familial accumulation, whose molecular causes are still largely unknown. It is suspected that the barrier function of the immune system is disturbed in the bowel and the intestinal wall reacts oversensitively to the naturally occurring bacteria. Recently, a link between the disease and variations in the gene *IL17REL* could be established in an association study. The protein IL17REL is a yet little-explored cell receptor [see Figure], onto which certain chemical messengers of the immune system bind and thereby enable important signaling pathways of inflammations.

Computer analysis of the protein IL17REL

Possible effects of two SNPs on the structure and function of the IL17REL protein were more closely investigated by using bioinformatics methods. It was found that both point mutations can effect the binding sites of the cell receptor [see Figure]. Thus, this suggests

that these sequence variations disturb the interaction between IL17REL and the binding chemical messengers. Consequently, inflammatory signaling pathways can be affected, which promotes the emergence of chronic diseases. For this reason, the therapeutic influence of IL17REL is now considered as a possible treatment of patients with ulcerative colitis.

Medical bioinformatics

In this way, medical bioinformatics supports clinical research and helps to interpret and better understand genetic findings on a molecular level. Computational methods accelerate the discovery of molecular disease causes and the development of drugs to combat them. ...

CONTACT

Mario Albrecht

DEPT. 3 Computational Biology and Applied Algorithmics

Phone +49 681 9325-3000

Email mario.albrecht@mpi-inf.mpg.de

Internet <http://medbioinf.mpi-inf.mpg.de>

Gabriele Mayr

DEPT. 3 Computational Biology and Applied Algorithmics

Phone +49 681 9325-3000



The Genetic Foundation of Cancer

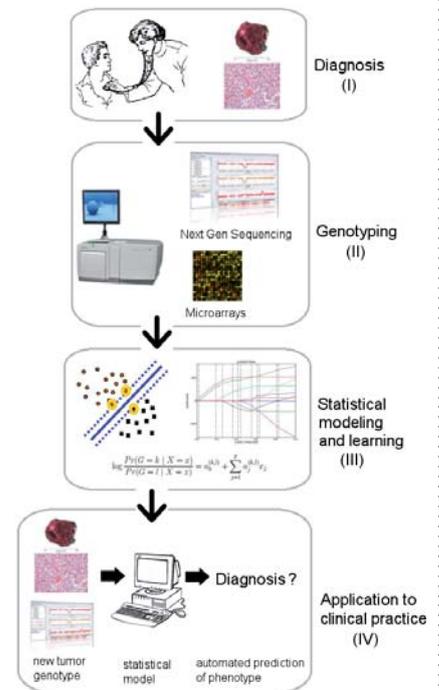
In the course of progression towards a cancerous state, the genome of a cell accumulates many mutations and structural alterations, which promote cell replication and prevent programmed cellular death (apoptosis). In our research, we study a specific type of structural alteration, namely *copy number aberrations*. While in healthy somatic cells most of the DNA is present in exactly two copies, some segments of the DNA of cancerous cells may be present in higher copy number (ranging from three to tens of copies) or in smaller copy number (one or zero copies). These segments are called copy number aberrations (CNAs). Most cancerous cells accumulate tens of CNAs in their genome during disease progression. As a consequence, the function of genomic sequences located within CNAs (e.g., genes) is either inhibited or promoted, which leads to aberrant cell behavior, e.g., uncontrolled proliferation.

Due to the recent technological advances in molecular biology such as microarrays and high-throughput sequencing, large cohorts of genome-scale measurements of CNAs have become available for most types of cancer. In our work, we use statistical tools for inferring the role of CNAs in cancer progression from sets of observed tumors. More specifically, we develop models for automated prediction of cancer phenotypes from genome-scale CNAs. For example, we plan to use data on monitored patients and their tumors for estimating statistical models that establish a statistical relation between the survival of a patient and the copy number aberrations observed in the genome of his/her tumor cells. Similar models can be used to identify genomic patterns that are susceptible or resistant to certain treatments, or that discriminate between early and advanced tumors.

One challenge of our endeavor lies in dealing with the high dimensionality of the copy number measurements, which adds difficulty to the statistical data mining task. The most important requirement for any prediction model that is trained on high-throughput genomics data is robustness. The number of genomic features that could potentially play a significant role in disease progression is so large compared to the number of observations available, that there is a great chance that the statistical learning process only grasps the idiosyncrasies of the few available training examples at the cost of generalization power. This phenomenon is called overtraining. We tackle this problem by grouping the genomic features into clusters, such as to achieve dimension reduction with no significant loss of relevant information. We choose robust classification models which rely on a small subset of highly relevant CNAs for prediction. The selected CNAs can be ranked by the weights with which they enter the model and become promising candidates for diagnostic biomarkers.

In our applications we investigate two types of cancer: neuroblastoma and lung cancer. In the case of neuroblastoma, we are able to predict tumor stage and survival from DNA copy number aberrations, with relatively high accuracies (up to 85% prediction accuracy). We also

identify copy number aberration patterns that can discriminate between the tumors of very young patients and tumors of older patients. Based on repeated clinical observations, age has been known to be a very important prognosis factor in neuroblastoma. Our results therefore show that there exists an underlying genetic explanation of this empirical observation.



The Figure illustrates the role of bioinformatics in the process of cancer treatment (*step III of the process*).



CONTACT

Laura Tolosi
DEPT. 3 Computational Biology and Applied Algorithmics

Phone +49 681 9325-3026
 Email laura@mpi-inf.mpg.de

Regulatory Feedback Networks – How Little Do We Actually Know?

A sobering picture

We have analyzed the current state of knowledge in one of the most important areas of modern genetics: gene regulation. By using statistical models, we can demonstrate very impressively that this research area is still in its early stages. We know very little about the molecular mechanisms that trigger and fine-tune the ability of living cells to adjust their genetic programs to variable environmental conditions.

Our work deals with former optimistic estimations, which must be corrected, and an increased need for research in the area of gene regulatory networks. Our statistical simulations show that even for the best-researched model organisms at most about one third of all gene regulatory interactions are known; this is a quite astonishing result.

Gene regulatory network

Gene regulatory networks emerge from complex interaction patterns of proteins with the DNA. So-called transcription factor proteins bind to certain areas of the DNA and regulate the production of proteins in any living organism. They ensure, for example, that insulin is only produced in a cell if sugar is actually present. The understanding of these

feed-back regulation networks is inevitable for a deeper understanding of the behavior of all living organisms.

Despite well over 1,500 completely sequenced organisms, it is astounding that one of the most important follow-up questions, the question about the interplay between the genes, cannot be answered completely for a single species. We provide statistically significant estimates of how much we already know about these fundamental networks, or rather, how little.

How do we estimate how little we know?

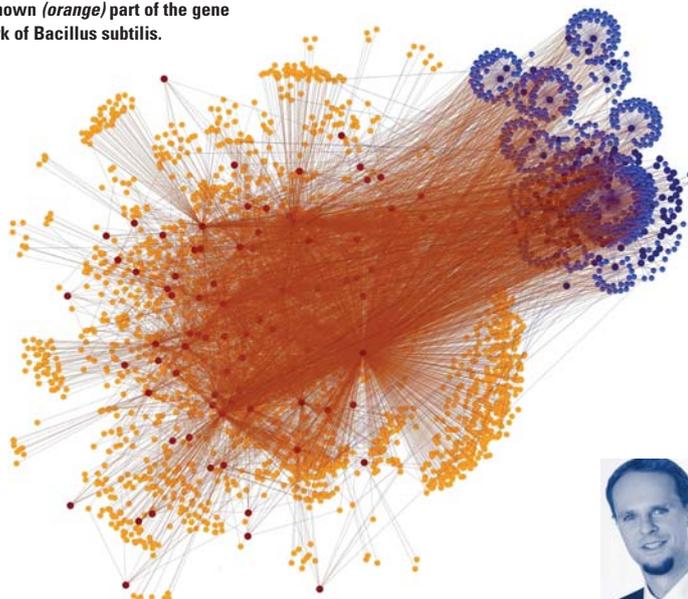
We utilized the known parts of different regulatory networks for the construction of statistical models, then we estimated the expected number of whole-organism gene regulatory interactions. The main problem was the validation of the results since there is not a single known, complete network to which one could compare the estimates today. Therefore, complex statistical methods were used to validate the developed model. In addition, confidence intervals were determined in order to provide scientifically sound and credible results.

Then, the model was applied to different organisms. As explained above, our results demonstrate our very limited knowledge base. The bacterium *Bacillus subtilis*, for example, is one of the best studied living organism and we know at most a quarter of its transcriptional gene regulatory relationships (see Figure 1 for an illustration). For complex mammals such as humans, the situation is even worse; here we know less than 1%.

What does this mean for us?

Our work demonstrates the importance of a constant development of further methods to determine gene regulatory relationships. The results also motivate consistent further research with established methods in order to understand the complex interplay of genes. These are not only issues that contribute to scientific knowledge, but also have urgent medical relevance. If we can show, that we know less than 1% of the human genetic interactions, this also means conversely, that we do not understand 99%. This is a great problem for personalized medicine as well as for the development of drugs. ...

An illustration of how little we know: The known (blue) vs. the unknown (orange) part of the gene regulatory network of *Bacillus subtilis*.



CONTACT

Jan Baumbach
DEPT. 3 Computational Biology and Applied Algorithmics
 Phone +49 681 302-70885
 Email jbaumbac@mpi-inf.mpg.de
 Internet <http://csb.mpi-inf.mpg.de>

GUARANTEES

Software must be reliable. The most important criterion for reliability is correctness. Nearly as important, however, is performance: A correct answer that is not received in time is not helpful. The search for correctness and performance guarantees is a cross-department research question at the institute.

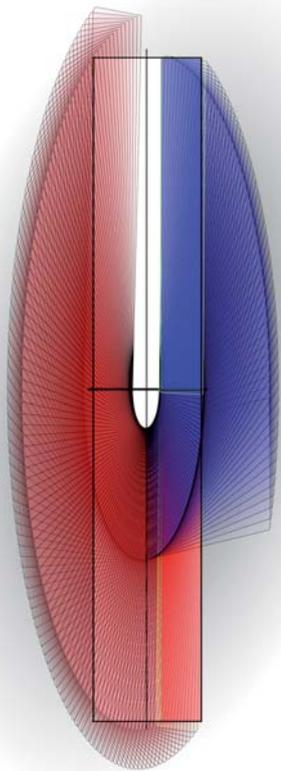
Today, computers, networks, and microprocessor controlled systems are a ubiquitous part of our lives. We constantly use them, sometimes knowingly, such as a computer on the desk, sometimes unknowingly, such as the electronic controls in the car, airplane or washing machine. The more our lives are dependent upon hardware and software, the more we pose the question whether the trust that we place in these products is justified. The answer to this question is especially difficult and complex as hardware and software are not robust. A single change in a program, a minor error in one line of a program with several hundred thousand lines of code, or a small change in the processor's calculations can cause the overall program to crash or to deliver wrong results. Or, the developers may not have considered an infrequent but possible border case, and software which has worked reliably for years can suddenly generate a wrong result. The granting of guarantees is even more complex when we examine not only a single program, but a system of programs in which parts work independently but also react to each other while pursuing different goals.

In the article *"Dealing with Selfishness in Optimization"*, the latter situation is investigated. An example of this are our current internet providers who have committed, on the one hand, to forwarding all Internet traffic, although they would ideally prioritize the data transfer of well-paying customers to optimize profit.

Guarantees of geometric calculations are especially important as today nearly every larger machine, such as a car or airplane, is designed on a computer. It would be undesirable if a calculation error caused, for example, the rudder of an airplane to become jammed during its service life. The articles *"Geometric Algorithms – Exact and Efficient"* and *"Controlled Perturbation – How to Avoid Difficult Special Situations"* show firstly, what is needed for accurate calculations in a geometric context, and secondly, how the substantial necessary effort required for this can be avoided in many situations.

In order to provide guarantees of programs in general, deductive methods must be used that enable the desired properties to be derived from the program in the form of proofs. Automatic verification procedures are required for this, which are discussed in the article *"Automated Deduction"*, as are techniques to properly divide and conquer resulting proof obligations. The complexity of verifications can thus be reduced. This aspect is discussed in the article *"Modular Proving in Complex Theories"*. Ultimately, it is currently not generally possible to give all desired guarantees of an arbitrary complex program within an acceptable period of time. In such a case, it makes sense to use procedures based on model assumptions, as is presented in the article *"Model Checking for Hybrid Systems"*. ...

CONTRIBUTIONS



Dealing with Selfishness in Optimization 46

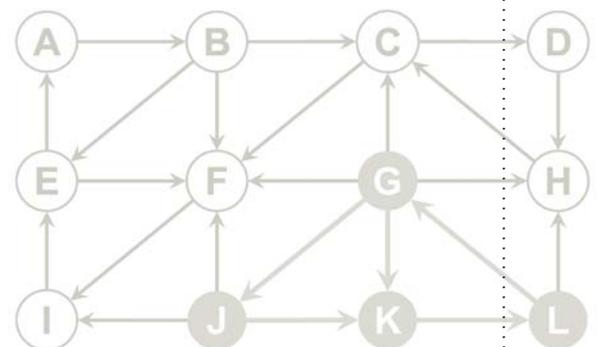
Geometric Algorithms – Exact and Efficient 47

**Controlled Perturbation – How to Avoid Difficult
Special Situations** 48

Automated Deduction 49

Model Checking for Hybrid Systems 50

Modular Proving in Complex Theories 51



UNDERSTANDING
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH
& DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION

Dealing with Selfishness in Optimization

In my research, I focus on optimization problems that involve selfishness in some way. For example, scheduling problems where the machines are controlled by selfish agents, and only each agent knows how fast his machine is. The agents are only interested in maximizing their own utility, which may be quite different from the overall optimization goal. Hence, we need to design mechanisms that ensure truth-telling is a dominant strategy for the agents.

A problem that I have recently studied in this area is maximizing the minimum load on selfish machines, i.e., optimizing the max-min fairness of the job allocation. This is a standard fairness measure in the area of fair division. Such an allocation is for instance relevant in the scenario where the jobs represent hard drives of various capacities, and we want to create a backup medium consisting of parallel storages of maximum possible capacity: this capacity is then bounded by the size of the smallest storage.

The utility of an agent is the payoff that it receives from the mechanism minus the cost that it has for processing the jobs assigned to it. This cost depends on its speed, which only each agent knows for itself. We want the agents to give their true speeds as input to the mechanism; otherwise we cannot optimize any objective function. It is known that if the allocation done by the mechanism is monotone, then telling the truth is a dominant strategy for the agents. Monotone means that if one agent increases its bid (i.e., its claimed cost for running one unit of work), the total size of the jobs assigned by the mechanism may not increase (the set of assigned jobs does not necessarily have to remain the same, however).

Together with Giorgos Christodoulou and Annamaria Kovacs, I have recently designed the first monotone approximation algorithm with a constant approximation ratio.

An area that I am just starting to work on is “selfish routing”, in particular on rings. On general graphs much is known about selfish routing, but the ring structure has not yet been studied in much detail (though some initial results have been obtained). This structure is very relevant, however, as the Internet in Europe in fact consists of interlocking rings (e.g., the German network is based on a ring).



CONTACT

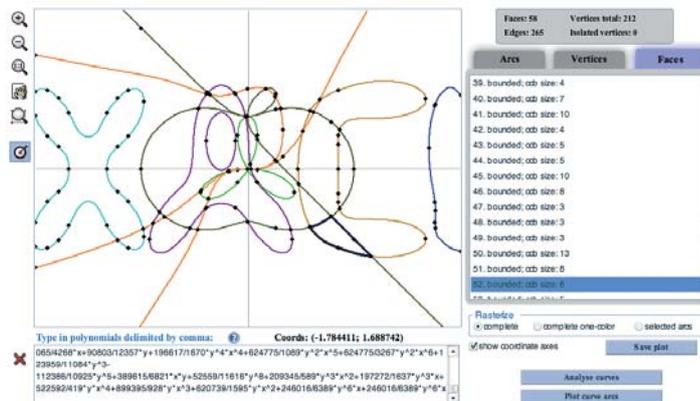
Rob van Stee
DEPT. 1 Algorithms and Complexity
 Phone +49 681 9325-1005
 Email vanstee@mpi-inf.mpg.de

Geometric Algorithms – Exact and Efficient

Geometric algorithms play an important role for the solution of many everyday problems. A classic example is the so-called “piano mover’s problem”, which involves, simply put, using calculations to decide whether or not the piano would fit through the narrow staircase the next time you move house. Another important application area is CAD (Computer Aided Design), where methods are developed to design and visualize of engineering components. Difficulties frequently occur when geometric algorithms are implemented, arising from wrong decision within the algorithm due to rounding errors. False results, crashes, or non-terminating programs can be the result.

The “Computer Algebra and Geometric Computing” group at the Max Planck Institute for Informatics aims to provide possible solutions for the above-mentioned problem. Here, we primarily deal with an algorithmic framework upon which respective specialized solution methods are developed. Such basic operations mainly include Boolean operations (intersection, union, etc.) of geometric objects such as straight lines, circles, ellipses, planes, spheres, ellipsoids, curves, and surfaces of higher degree. These operations are essential for the computation of so-called arrangements, that is, an explicit description of the decomposition of a plane or space implicitly defined by the given geometric objects.

The algorithms that we have developed are designed and implemented in such a way that all possible inputs can be exactly processed. Unfortunately, algorithms with this property often come along with a series of disadvantage. Firstly, the complexity of the algorithm increases and, what is much worse, one must trade off the completeness and exactness with a significant increase of computation times. The reason for this is that solutions to general polynomial systems must be found and further processed to handle curved geometric objects, and the corresponding operations are computationally intensive. For the



Arrangement calculation of plane algebraic curves

latter reason, commercial systems completely rely upon numerical methods. Such methods, however, only provide correct results if a “sufficiently benign” input is presupposed. Examples show, however, that there is certainly a practical justification in not seeing such a presumption as given.

In the first years of our research, we have been able to develop methods for the investigation of general curves and surfaces. We have combined symbolic methods from computer algebra with efficient (combinatorial) methods from algorithmic geometry. Measured against the best comparable methods, those which we have developed are among the fastest in the world. For the investigation of highly curved surfaces, the computational effort for the symbolic calculation steps is considerable, however, which severely restricts the practical application of our methods.

The focus of our latest research is therefore the development of more efficient methods for carrying out unavoidable symbolic calculations. Here, the calculations are broken down into many smaller units with significantly

lower complexity, and the individual results are ultimately reassembled to form a complete solution. Using massively parallel hardware architecture, as it is found, for example, on graphics cards, we have been able to accelerate important symbolic calculations (resultants, greatest common divisor, etc.) by a hundredfold. After we have adapted the more high-level algorithms based on it, they also fully benefit from the acceleration. When solving systems of polynomial equations, we have combined numeric methods with exact procedures for post-certification for the first time. Following this approach, we are able to find the solutions much faster than with previous methods and to finally ensure their correctness.

We are also interested in the theoretical complexity of our methods. This means that we want to specify an upper bound for the number of operations that is required for the processing of any input case of a certain size. We have managed to match or even improve existing bounds, both for basic operations such as the finding and approximation of zeroes and for more complex algorithms such as the arrangement computation of curves. ...

CONTACT

Michael Sagraloff
DEPT. 1 Algorithms and Complexity
 Phone +49 681 9325-1006
 Email msagraloff@mpi-inf.mpg.de
 Internet <http://www.mpi-inf.mpg.de/departments/d1/areas/GCCA.html>

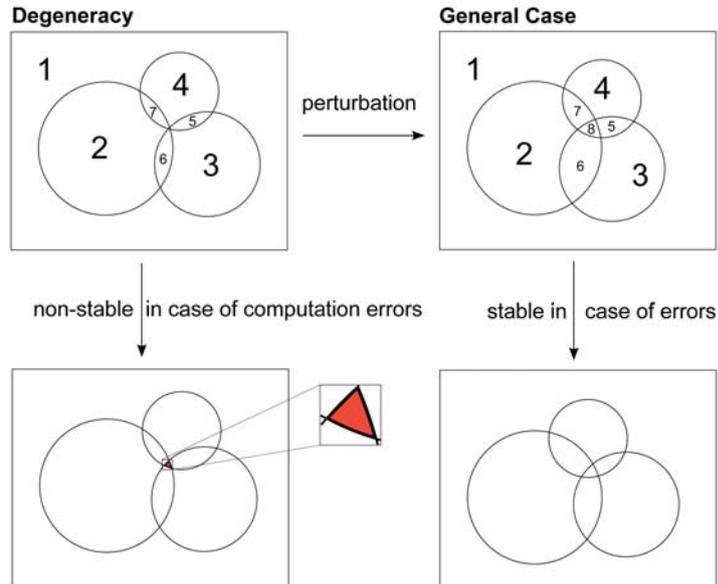


Controlled Perturbation – How to Avoid Difficult Special Situations

In the article on exact and efficient geometric algorithms “*Geometric algorithms – exact and efficient*”, we already address the need for exact and reliable methods for the treatment of geometric problems. The solution presented there is to formulate and implement all basic operations within the process in a way such that all occurring decisions can be made under guarantee. Such a decision is normally equivalent to determining the sign of a so-called predicate function (e.g. polynomial).

We would like to make this clear using the following example: Given three circles in a plane, compute the number of cells into which the plane is decomposed. The solution of this problem requires at a certain step in the algorithm that you can decide whether the three circles meet in exactly one point. The special case that the circles have one common intersection point, is thereby the most difficult. It may in fact be treated by an exact algebraic computation. In addition, it must also combinatorially be regarded as a special case, which complicates the overall algorithm. If the circles however do not share a common point, the solution is much easier. Namely, in the latter case, small errors are allowed in the calculation, and eventually lead to the same result. In more geometric terms, this is due to the fact that a slight shift or scaling of the circles does not change the number of cells.

Many idealized algorithms are not able to deal with special cases like the aforementioned one where the circles intersect in one point. Either they give false statements in the sub-steps (“The circles do not intersect” even though they do actually intersect) due to rounding errors in the approximate arithmetic, or they cannot handle the special case combinatorially. Both can lead to false end results, crashes, or to a non-terminating program.



A special case which can be generically made through perturbation. In a generic case, small errors lead to the same result, while the special case requires exact calculation.

Controlled Perturbation (CP) describes a method which uses an idealized algorithm to make all inputs run. The trick is to slightly perturb the inputs, consisting of a number of geometric objects, so that difficult special cases are excluded. In our example above, this means that the circles are first shifted slightly, which means that, with high probability they do not intersect in one point. Within the algorithm it is now tried to make all resulting decisions using inexact but very fast floating point arithmetic. In this case, possible rounding errors are taken into account. If this approach succeeds, the *algorithm gives an exact output for a slightly perturbed input* (but not necessarily for the exact input) without the need to handle special cases. In the event of failure (a decision cannot be made with approximate arithmetic), one perturbs the input again and increases the precision of the floating point calculations.

CP was introduced at the end of the 1990's by scientists at the University of Tel Aviv and, since then, has been successfully used with a number of methods for the treatment of lines, segments and circles. Within the “Computer Algebra and Geometric Computing” group, we addressed the question in which way CP is useful for algorithms in general. In addition, we are interested in the relationship between CP parameters (e.g., precision of floating point arithmetic, size of the perturbation) and the probability of success of the process. We managed to prove the feasibility of CP for a variety of algorithms and to provide a general analysis for the treatment of polynomial predicates. Working together with the University of Tel Aviv, we aim to extend CP for the first time to the practical study of nonlinear problems. ...

CONTACT

Michael Sagraloff
DEPT. 1 Algorithms and Complexity
 Phone +49 681 9325-1006
 Email msagraloff@mpi-inf.mpg.de
 Internet <http://www.mpi-inf.mpg.de/departments/d1/areas/GCCA.html>



Automated Deduction

In order to guarantee that a piece of hardware or software is working correctly, it must be verified – i.e., its correctness must be proven formally. Proving correctness means that one has to check whether certain properties of the system follow from other properties of the system that are already known. The question how to use computer programs for such proof tasks has been an important research topic for many decades. Ever since the fundamental theoretic results of Gödel and Turing at the beginning of the twentieth century, it is known that not everything that is *true* in a mathematical sense is actually *provable*, and that not everything that is *provable* is *automatically provable*. Correspondingly, deduction systems significantly differ in their expressiveness and properties. For example, decision procedures are specialized for a certain type of data (e.g., real numbers) and are guaranteed to detect the correctness or incorrectness of a statement from this area. Automatic provers for so-called first-order logic can handle any data types defined in a program. For these provers one can only guarantee, though, that they eventually find a proof if a proof exists; if none exists, the prover may continue searching without success, possibly forever. Even more complicated problems can be handled using interactive provers; these provers, however, only work with user assistance and without any guarantee of completeness.

How does a theorem prover for first-order logic work? It is not difficult to write a program that correctly derives new formulas from given formulas. A logically correct derivation is, however, not necessarily a useful derivation. For example, if we convert $2 \cdot a + 3 \cdot a$ to $2 \cdot a + 3 \cdot a + 0$ and then to $2 \cdot a + 3 \cdot a + 0 + 0$, we do not make any computation

errors, but we are not getting any nearer to our goal either. The actual challenge is thus to find the few *useful* derivations among infinitely many *correct* derivations. One quickly comes to the conclusion that it is useful to apply equations in such a way that the result is simplified, e.g., “ $x + 0 = x$ ” only from left to right and not the other way round.

$$\begin{aligned}
 x + 0 &= x \\
 x + (-x) &= 0 \\
 \frac{x \cdot z}{y \cdot z} &= \frac{x}{y}
 \end{aligned}$$

complicated → simple

Application of equations

However, this approach is not always sufficient. A typical example is fractional arithmetic: We all know that it is occasionally necessary to expand a fraction before we can continue calculating with it. Expansion, however, causes exactly what we are actually trying to avoid: The equation “ $(x \cdot z) / (y \cdot z) = x / y$ ” is applied from right to left – a simple expression is converted into a more complicated one. The superposition calculus developed by Bachmair and Ganzinger in 1990 offers a way out of this dilemma. On the one hand, it performs calculations in a forward direction; on the other hand, it systematically identifies and repairs the possible problematic cases in a set of formulas, where backward computation could be inevitable. Superposition is thus the foundation of almost all theorem provers for first-order logic with equality, including the prover SPASS that we have developed at the institute.

In our research group, we currently focus on refining the general superposition method for special applications. For instance, we are developing techniques for combining the capabilities of various proof procedures (for instance, superposition and arithmetic decision procedures). We are addressing the question of how superposition can be used to process even very large quantities of data, such as those that occur in the analysis of ontologies (knowledge bases). We also use superposition to verify network protocols and to analyze probabilistic systems (i.e., systems whose behavior is partially dependent on random decisions).



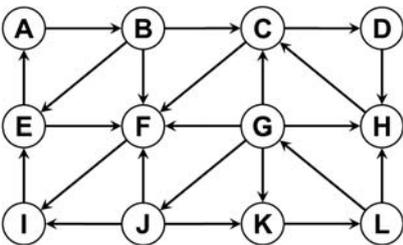
CONTACT

Uwe Waldmann
RG. 1 Automation of Logic
 Phone +49 681 9325-2905
 Email uwe@mpi-inf.mpg.de

Model Checking for Hybrid Systems

The behavior of many technical devices can be described by specifying their states and state transitions. A washing machine can, for example, be in the starting state “door open, drum empty, machine turned off” or the state “washing.” State transitions are changes from one state to another. These can take place automatically, such as the change between the washing and drainage, or through other effects, for example when a button is pressed. There are, however, possible states that are obviously undesirable, for example: “water inlet open, door open.” If we want to verify that the machine is *safe*, then we must show that there are no state transitions leading from the starting state to such an *unsafe* state.

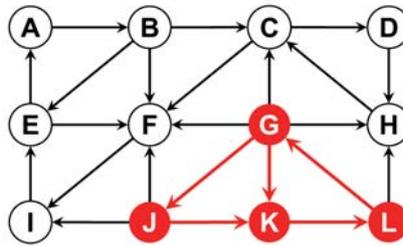
How can this problem be approached? We can describe states and state transitions using *graphs*. To this end, we symbolize each state by a circle and each possible state transition by an arrow. The result may look as follows:



Example

Let us assume that A is the starting state and L is unsafe. What are the predecessor states of L , i.e., the states from which an arrow points to L ? Obviously, there is only one such state: K . K is thus

also unsafe, since if K can be reached, then L can be reached from there. K 's predecessor states, namely G and J , are equally unsafe. If we take a look at the predecessor states of G and J (namely L and G), then we notice that these are already known to be unsafe. This means that there is no way of reaching the state set $\{G, J, K, L\}$ from outside.



Non-safe states

It is therefore clear that $\{G, J, K, L\}$ are the only states that are unsafe per se, or from which an unsafe state can be reached. All other states, in particular the starting state A , are safe. This technique to verify the safety of a system is known as *model checking*.

Unfortunately, the number of states can quickly lead to problems in practice. Imagine that the behavior of an electronic device depends on 60 bits (b_1, \dots, b_{60}), i.e., on 60 components that can each have the value of true or false. Now, 60 bits are not many, but if all the possible combinations are taken into account, we arrive at 2^{60} states, more than one quintillion. *Symbolic model checking* offers a way out here: We no longer enumerate the unsafe states and permissible state transitions explicitly, but we represent them

symbolically using logical formulas. For example, the formula

$$b_5 \text{ and not } b_{32}$$

stands for all states for which b_5 is true and b_{32} is false – this is a quarter of a quintillion states, which would otherwise have to be enumerated explicitly.

A new problem emerges if we try to verify the safety of a vehicle control system in this way. Such a control system is a *hybrid* system. This means that one is no longer solely dealing with *discrete* state changes (e.g., acceleration or braking), but also with variables such as speed, which can change continuously. The state formulae must therefore also include numerical variables, for example:

$$b_5 \text{ and not } b_{32} \text{ and } (x_2 > 50).$$

Using such formulas, we can deal with the discrete state transitions as before. What about the continuous state transitions? As long as the dynamic behavior is simple from a mathematical point of view (e.g., if the speed changes uniformly), one can use a method known as *quantifier elimination*. Certain technical tricks are needed, however, to prevent arriving at too large a formula too quickly. In our work group, we are currently focusing on developing methods that can be applied even to complicated dynamic behaviors. ...



CONTACT

Uwe Waldmann

RG. 1 Automation of Logic

Phone +49 681 9325-2905

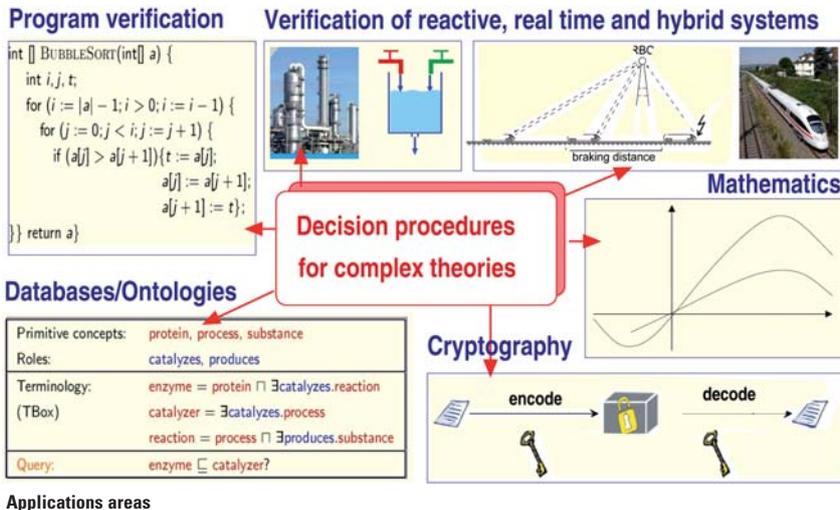
Email uwe@mpi-inf.mpg.de

Modular Proving in Complex Theories

The enormous advances in the development of information technology have led to the widespread use of complex, computer-controlled systems: in household appliances, cars, trains, airplanes or power plants. Especially in the last-mentioned safety-critical application domains, errors can have catastrophic consequences. It is therefore extremely important to guarantee the proper functioning of such systems and to mathematically prove their safety. In fact, it would be desirable to carry out such correctness proofs fully automatically by computer. However, fundamental theoretical results from Gödel, Church and Turing show that this is not possible. Nevertheless, effective automatic verification systems do exist for concrete application areas.

Our goal is to identify conditions under which efficient verification procedures for complex systems exist. In the formal description of a complex system, we need a diversity of data types (for example, numbers, data structures such as sets, lists or arrays, or functions over numerical domains). Efficient reasoning in such complex theories is an important goal of our research. We are interested in developing proof procedures which use the modular structure of complex theories and allow us to use specialized provers for reasoning in the component theories. Such a modular approach is particularly flexible and efficient as well as applicable in many areas (such as mathematics, verification, and knowledge representation [see Figure]).

The simplest form of complex theories we consider are extensions of a given theory (which we refer to as base theory) with additional functions. Such theory extensions can, for example, be used in parametric approaches to the verification of reactive or hybrid systems: e.g., in the verification of controllers regulating systems of trains in which either the number of trains or certain variables (e.g., their speed or their acceleration) and their changes can be considered to be parametric.



In general, it is difficult to deal with such extensions even when effective procedures for the base theory exist. We approach this problem by identifying a class of theory extensions in which it is possible to reduce the original problem to a problem in the base theory, which can then be solved with a specialized proof procedure for this theory.

In addition, we are developing methods for modular reasoning in combinations of theories. If, for example, we want to prove a statement in program verification which depends on both lists of numbers and arrays of numbers, then we can decompose this problem into two – one concerning only lists of numbers and another dealing with arrays only. For each of these two problems we can use existing proof methods for lists resp. arrays. If both methods exchange sufficient information about the shared data

(i.e., the numbers), it is guaranteed that in the end, the two partial solutions can be combined into one solution to the original problem.

Our theoretical contributions are the basis for the development of practically usable tools for the verification of safety-critical systems, particularly in the context of the SFB Transregio Project AVACS (*Automatic Verification and Analysis of Complex Systems*).

We use our method both for program verification and for verifying various control systems, such as a radio controller for train systems, as well as a (parametric) hybrid control system for chemical plant which supervise the relative concentration of various substances. Furthermore, we have also applied our approaches to the verification of distributed databases, in security, and for mathematical problems. ...



CONTACT

Viorica Sofronie-Stokkermans
RG. 1 Automation of Logic
 Phone +49 681 9325-2900
 Email sofronie@mpi-inf.mpg.de

INFORMATION SEARCH & DIGITAL KNOWLEDGE

Digital information has changed our society, economy, work in science, and everyday life.

Modern search engines deliver useful information for nearly every question, and the

Internet has the potential to be the world's most comprehensive knowledge base.

However, knowledge structures in the Internet are amorphous, and search engines rarely

have precise answers to expert questions for which users would consult reference books

and expert literature. A great challenge and research opportunity lies in making the leap

from processing raw information to computer-aided intelligent use of digital knowledge.

In parallel to this aim for a quantum leap, we observe a complexity explosion in digital information along several dimensions: quantity, structural variety, multimodality, digital history, and distribution.

- In addition to more than 20 billion web pages, important information sources are: online news streams, blogs, tweets and social networks with several hundred million users, Web2.0 communities for photos, music, books, scientific topics, and last but not least the encyclopedia Wikipedia. The total volume of this data is in the order of exabytes: 10^{18} of bytes – more than one million terabyte disks.

- There is an increasing amount of more expressive data representations: XML documents, RSS feeds, semantically connected RDF graphs, and much more. The richer structure and heterogeneity of the data increases the complexity of mastering this digital variety.

- In addition to text-oriented and structured data, we are also experiencing an explosion of multimodal information: billions of people are becoming data producers in the web by sharing images, videos, and sound recordings with the rest of the world. This is often accompanied by interpersonal contacts that arise from the Internet, and large online networks are organized.

- The history of digital information – for example earlier versions of our institute's web pages, which are partly conserved in the Internet Archive – is a potential gold mine for in-depth analyses along the time dimension. Sociologists and political scientists, but also media and market analysts as well as experts in intellectual property, can profit from this.

- The quantity and variety of the available information in the Internet has grown to such a scale that search engines are no longer able to maintain all relevant references in a centralized index. Therefore, global information search necessitates distributed algorithms, in which, for example, many local search engines are dynamically coupled for specific tasks. Here, not only local computing and search speed play an important role, but also communication efficiency in the network of networks, the Internet, and in embedded peer-to-peer networks.

This global issue is examined from various perspectives at the Max Planck Institute for Informatics. These include work on efficient search in semistructured XML documents, which play a big role in digital libraries and e-science data. Another research area is the scalable management of graph-structured RDF data which arises in the context of the semantic web but is also important as data representation in computational biology. In other projects, the user-centric view of Web 2.0 communities and multimodal data are the primary issues. The great vision of making the quantum leap towards knowledge search and discovery is pursued in work on automatic knowledge extraction from web sources such as Wikipedia. The Max Planck Institute for Informatics is a worldwide leader in this important research direction. ...



Random Telephone Chains – Efficient Communication in Data Networks	54
---	----

Search and Mining in Web Archives	55
--	----

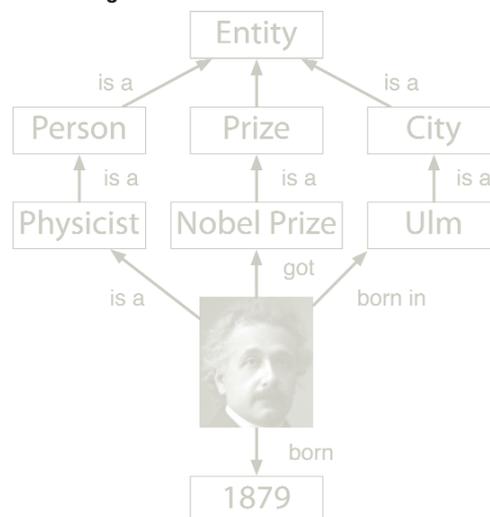
Information Search in Social Networks	56
--	----

YAGO – a Collection of Digital Knowledge	57
---	----

URDF – Efficient Reasoning in Uncertain RDF Knowledge Bases	58
--	----

EnBlogue – What is New and Interesting in Web 2.0?	59
---	----

Decision Procedures for Ontologies	60
---	----



UNDERSTANDING
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH
& DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION

Random Telephone Chains – Efficient Communication in Data Networks

How do you send an urgent message quickly and efficiently to a larger group of people? A classic procedure is the telephone chain. In the simplest case, the sequential telephone chain is based on an organized list of group members. Here, spreading news goes to the first person on this list. Then he calls the second person, who then calls the third, and so forth.

Similar task descriptions appear in informatics if information in computer centers must be distributed. A company with affiliates at several sites can store its customer data locally at each branch office. This has the benefit that in each affiliate one can quickly access this data with little complication, and that the data is also available if the central database is not reachable for a short period of time. When using such replicated databases, there must naturally be changes given in the data inventory of an affiliate, which is slowly communicated to all others.

In two central aspects, this telephone chain problem is different from informatics in the everyday world. First, such database systems can be very large: In a network with several thousand nodes, as in this affiliates example, the information distribution lasts much too long if the nodes are given the news in a fixed series, one after another.

A second aspect is the robustness of the procedure. Robustness means that a process can still function properly even if some partial steps are not error-free. In a telephone chain, the main problem might be, that a participant cannot be reached, or that he is reachable, but for some reason, he does not forward the news. The classic sequential telephone chain appears not to be very robust. If a participant does not pass on the news, all the following participants in the chain remain uninformed.

For this reason, processes based on classic telephone chains are inappropriate for application in informatics. However, there are telephone chain problems in informatics with a surprisingly simple solution. With the *random telephone chain* each caller who has received the news, calls random other participants. To simplify the presentation, one can accept that all calls last the same amount of time. This leads to the information exchange taking place in rounds. In each round, each informed participant calls another randomly chosen participant. At this latest point in time, this chosen participant is also informed.

The random telephone chain is amazingly efficient and robust. If we take an example of a data network with 1,024 nodes which can all communicate with one another, then the random telephone chain needs on average only 18.09 rounds in order to inform all nodes. The same is true for the robustness; even if we accept that every 10th participant is never reachable, it takes only 19.49 rounds on average to inform the remaining members. Therefore, it is not important which of the members fall within this 10%.

The good performance of the randomised telephone chain has stimulated other research efforts. A highly interesting question is: What is the optimal amount of randomness? The first results demonstrate that one sole random decision per participant gives results that are just as good as a completely random telephone chain. Actually, the optimum lies between the two extremes. The following very simple protocol demonstrates this.

We assume that all participants have one (e.g., alphabetically) organized list of all participants. If a participant is informed, he choose another random position on the list and calls, beginning at this position, all the other participants. If he reaches someone who is already informed, he dials a new, random starting position and continues to call in order. This simple protocol informs 1,024 nodes in only 14.11 rounds on average. Even if (randomly dialed) 10% of the participants are never reachable, only 15.09 rounds are needed on average.

A further benefit of this procedure is that we can give a termination criterion. We can establish, for example, that participants leave the procedure if they have called an informed node a fourth time. This is the first time that a termination criterion has been given for a randomized telephone chain, and it is guaranteed thereafter that all participants will be reached. Furthermore, it provides a very efficient procedure – on average, a participant makes only 5 calls. The recommended termination criterion leads in any case to a minimal increase in running time. In this case, for example, we need only 14.14 rounds in order to inform 1,024 nodes. Apparently, there is a trade-off situation: The earlier the participants stop, the longer the procedure takes. Even if the participants leave immediately after they have called an informed node, the new procedure is slightly better than the fully random telephone chain and needs 17.63 rounds instead of 18.09. ...



CONTACT

Benjamin Doerr
DEPT. 1 Algorithms and Complexity
 Phone +49 681 9325-1004
 Email doerr@mpi-inf.mpg.de

Search and Mining in Web Archives

The World Wide Web evolves constantly, and every day contents are added and removed. Some of the new contents are published first and exclusively on the Web and reflect current events. In recent years, there has been a growing awareness that such born-digital contents are part of our cultural heritage and therefore worth preserving. National libraries and organizations such as the Internet Archive [<http://www.archive.org>] have taken over this task. Other contents were published a long time ago and are now, thanks to improved digitization techniques, for the first time available to a wide public via the Web. Consider, as one concrete example, the archive of the British newspaper *The Times* that contains articles published from as early as 1785.

Our current research focuses on scalable search and mining techniques for such web archives. Improved search techniques, on the one hand, make it easier for users to access web archives. Mining techniques, on the other hand, help to gain insights about the evolution of language or popular topics. In the following, we describe three aspects of our current work.

Time travel in web archives

Existing search techniques ignore the time dimension inherent to web archives. For instance, it is not possible to restrict a search, so that only documents are retrieved that existed at a specified time in the past.

In our work, we consider time-travel keyword queries that combine a keyword query (e.g., "*bundestag election projection*") with a temporal context such as September 2009. For this specific query, only relevant documents that discuss election projections and which existed back in September 2009 should be retrieved as results.

Our approach builds on an inverted index that keeps a list of occurrences for every word. Depending on the type of query that has to be supported, the inverted index remembers an identifier, how often the word occurs, or the exact positions at which the word can be found in the document for every document in which a word occurs. We extend this information by a valid-time interval to also keep track of when a word was contained in a document and thus to enable time-travel queries.

Consecutive versions of the same document tend to differ only slightly. We exploit this observation to reduce index size drastically. To process time-travel queries more efficiently, we keep multiple lists for every word in the inverted index, each of which is responsible for a specific time interval. This introduces redundancy to the index, increases its size, and thus leads to a trade-off between index size and query-processing performance. Our approach casts this trade-off into optimization problems that can be solved efficiently and determine the lists to be kept for every word in the inverted index.

Temporal information needs

Information needs often have a temporal dimension, as expressed by a temporal phrase contained in the user's query and are best satisfied by documents that refer to a particular time. Existing retrieval models fail for such temporal information needs. For the query "*german painters 15th century*", a document with detailed information about the life and work of Albrecht Durer (e.g., mentioning 1471 as his year of birth) would not necessarily be considered relevant. This is because existing methods are unaware of the semantics inherent to temporal expressions and thus do not know that 1471 refers to a year in the 15th century.

To capture their semantics, we formally represent temporal expressions as time intervals. We then integrate them into a retrieval approach based on statistical language models that has been shown to improve result quality for temporal information needs.

Mining of characteristic phrases

Mining of web archives is another aspect of our current work. More precisely, we are interested in insights about ad-hoc subsets of the web archive, for instance, all documents that deal with Olympics. Given such an ad-hoc subset, we can identify persons, locations, or in general, phrases that are characteristic for documents published in a particular year. In our Olympics example, these could include Michael Phelps, Beijing and "*bird's nest*" for documents published in 2008.

To identify such characteristic phrases efficiently, one needs frequency statistics for so-called *n-grams* (i.e., sequences of one or more words). We develop efficient and scalable techniques to compute these *n-grams* statistics in a distributed environment. One design objective here is to allow for easy scale-out in order to keep up with the growth of web archives in the future. ...



CONTACT

Klaus Berberich
DEPT. 5 Databases and Information Systems
 Phone +49 681 9325-5018
 Email kberberi@mpi-inf.mpg.de

Information Search in Social Networks

The advent of Web 2.0 has revolutionized the way users interact with the Web. Instead of simply consuming information as before, now anyone can simply generate and publish contents. Online services such as del.icio.us, Flickr, LibraryThing, YouTube, MySpace, or Facebook offer a simple way to publish pictures, videos, bookmarks and similar data on the Web. In addition, they allow the user to interact with other users, enabling them, for example, to share content. Data published by other users can be explored, commented, assessed, and annotated with so-called tags. These tags are often very good content descriptions, because they can be freely chosen without having to stick to a predefined set. Various users frequently annotate the same image or the same video with different tags, reflecting their different interests. Most services offer easy to use and intuitive interfaces in order to find contents based on annotations, for example, the so-called “tag clouds”.

In addition, users can maintain explicit lists of their friends, and the system often provides additional value to such followers. It is easily possible to look for newly added content, new assessments or, especially lately, to inform friends about your current location. In this way, a social network is built, a dense network of relationships, in which the number

of friends is often taken as an indicator of one’s reputation on the network. While these lists of friends are initially filled with friends and other people known from “real” life using the same service, it grows over time to include previously unknown users with similar interests.

This dense user network, in combination with the user-defined annotations, allows to find valuable content recommended by friends (either directly or through transitive friends of friends), exploiting the “wisdom of the crowds”. Such a recommendation can either be explicit (e.g., through judgments and comments) or implicit (e.g., through intensive use of annotations). When searching in such systems, relations and the distance between users should therefore be taken into consideration, as close friends usually are more trusted than remote acquaintances. Some existing social networks already offer such functions, but only in a very limited fashion (for example, only for direct, but not for transitive friends); others do not support neighborhood search at all.

We have developed a search engine coined *SENSE* (for “Socially ENhanced Search and Exploration”) to close this gap. It provides tag-based search for content among the friends of the querying user, weighting annotations based on the

strength of relationship to the friend.

In addition to the distance in the social network, similarity in interests (derived from similarity of annotations) has influence in weights. In comparison to most existing systems that consider plain tag frequency, *SENSE* retrieves results with significantly better quality for such personalized searches. Our system uses highly efficient and scalable search algorithms to cope with the massive amount of data present in large social networks. The fast growth of these services, and especially the very high rate with which new contents and annotations are generated, requires optimized indexing strategies in order to keep the precomputed transitive user relationships always up to date. *SENSE* defers updates as long as possible to the point when up-to-date information is needed for query processing; then, updates are performed in a pay-as-you-go manner while processing the query. ...



CONTACT

Tom Crecelius
DEPT. 5 Databases and Information Systems

Phone +49 681 9325-5006
Email tcrecel@mpi-inf.mpg.de



Ralf Schenkel
DEPT. 5 Databases and Information Systems

Phone +49 681 302-70798
Email schenkel@mpi-inf.mpg.de

YAGO – a Collection of Digital Knowledge

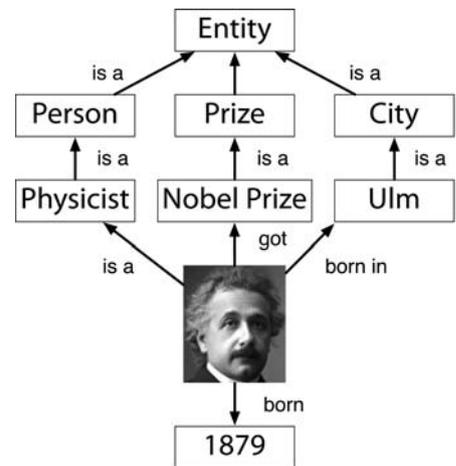
In recent years, the Internet has developed into the most significant source of information we have available today. Train schedules, news, and even entire encyclopedias are available online. Using search engines, we can very efficiently query this information. But, despite efficiency, search engines face limitations in terms of effectiveness, in particular for complicated queries. Assume we would like to know which scientists are also active in politics. This question can hardly be formulated in a way that it can be answered by Google. Queries like “politicians scientists” only return results about opinions on political events. The problem here is that the computers we use today can store a tremendous amount of data, but are not able to relate this data to a given context or, for that matter, even understand it. If it would be possible to make computers understand data as knowledge, this knowledge could be helpful not only for Internet search, but also for many other tasks, such as understanding spoken language or the automatic translation of a text into multiple languages. This is the goal of the “YAGO-NAGA” project at the Max Planck Institute for Informatics.

Before a computer can process knowledge, it must be stored in a structured fashion. Such a structured knowledge collection is called an ontology. The major building blocks of an ontology are *entities* and *relations* among entities. An entity is every type of concrete or abstract object: the physicist Albert Einstein, the year 1879, or the Nobel Prize. Entities are connected by relations, for example, Albert Einstein is connected to the year 1879 by the relation “born” (see graph). We have developed an approach to automatically create such an ontology using the online encyclopedia Wikipedia. Wikipedia contains articles about thousands of personalities, products and organizations. Each of these articles becomes an entity in our ontology.

There is for example an article about Albert Einstein, so the physicist can be recognized as an entity for the ontology. Each article in Wikipedia is classified into specific categories, the article about Einstein, for example, into the category “born in 1879”. The keyword “born” allows the computer to store the fact that Einstein was born in 1879. Using this approach, we get a very large ontology, in which all of the entities known to Wikipedia have their place. This ontology is called YAGO (*Yet Another Great Ontology*, <http://www.mpi-inf.mpg.de/yago-naga/yago/>). At the moment, YAGO contains nearly 10 million entities and about 80 million facts.

YAGO2, a recently created extension of the original knowledge base, pays particular attention to the organization of entities and facts in *space* and *time* – two dimensions that are highly useful when searching in a knowledge base. In fact, the great majority of the approximately 900,000 person-entities in YAGO2 are anchored in time by their birth and death dates, allowing us to position them in their historical context. For example, one can ask questions about important historical events during the lifetime of a specific president, emperor or pope, or also ask the question of when the person of interest was actually president.

Most of the approximately 7 million locations in YAGO2 have geographic coordinates which place them on the earth’s surface. Thus, spatial proximity between two locations can be used as search criterion. An exemplary search using the space and time criteria could be: Which 20th century scientists were awarded a Nobel Prize and were born in the vicinity of Stuttgart. In YAGO2, one finds, among others, Albert Einstein, since both his lifetime (1879–1955) and his birthplace Ulm (70 kilometers from Stuttgart) are stored in YAGO2.



In our next article about URDF, we describe how existing knowledge in YAGO2 can also be used to deduce new knowledge and then how to reason with this knowledge under different uncertainty models. ...



CONTACT

Johannes Hoffart
DEPT. 5 Databases and Information Systems
 Phone +49 681 9325-5028
 Email jhoffart@mpi-inf.mpg.de



Fabian Suchanek
DEPT. 5 Databases and Information Systems
 Phone +49 681 9325-5000
 Email suchanek@mpi-inf.mpg.de



Gerhard Weikum
DEPT. 5 Databases and Information Systems
 Phone +49 681 9325-5000
 Email weikum@mpi-inf.mpg.de
 Internet <http://www.mpi-inf.mpg.de/yago-naga/yago/>

URDF – Efficient Reasoning in Uncertain RDF Knowledge Bases

Reasoning in uncertain knowledge bases with soft and hard rules

Despite the vast improvements in the field of information extraction in recent years, the very nature of the applied extraction techniques entails that the resulting knowledge bases may exhibit a high degree of uncertainty or even inconsistency. Often, these inconsistencies are not obvious even for a human and can only be uncovered through intricate and potentially expensive inference steps. In large RDF knowledge bases with millions of facts, an automated search for inconsistencies can be accomplished only through complicated, logic-based inference techniques, and an overly eager removal of inconsistencies can even lead to the loss of correct data. Automatic reasoning in large knowledge bases therefore requires a high robustness with respect to incomplete, imprecise, or even inconsistent data, which we will here summarize under the term *uncertain data*. While traditional query processing techniques for RDF (based on the SPARQL query language) are limited to deterministic data, the research in our latest project, *URDF*, particularly focuses on efficient, rule-based, and statistical inference techniques for uncertain RDF knowledge bases.

In URDF, we can express, for example, that many people live in the same place as their spouses. Being a rather “soft” form of an inference rule, this rule will also be violated by a number of instances (in this case, people) in the real world, which means that any knowledge derived from such a soft rule will also be uncertain. On the other hand, we can surely exclude that a person could have been born in different geographic locations. This represents a strict constraint, i.e., a “hard” rule, which may not be violated by any instance in our knowledge base. In a logic-based representation, these rules are often formulated as implications, so-called *Horn clauses*, where a conjunctive condition of facts implies a new fact:

$$\begin{aligned} & \text{marriedTo}(\text{person}_1, \text{person}_2) \wedge \\ & \text{livesIn}(\text{person}_2, \text{location}_1) \\ & \rightarrow \text{livesIn}(\text{person}_1, \text{location}_1). \end{aligned}$$

The above rule is formulated in first-order predicate logic. More generally, our URDF framework allows for identifying common patterns between the instances in the knowledge base, and we can generalize (i.e., learn) first-order rules from these patterns in an inductive way. Conversely, we can then deductively reason about individual instances in our knowledge base by applying these first-order rules. In this case, we can, for example, apply the above first-order rule to the persons “Angela Merkel”, “Joachim Sauer”, and the place “Berlin-Mitte”. This form of uncertain reasoning allows us to infer that “Joachim Sauer” might live in “Berlin-Mitte”, given that we know that his wife, “Angela Merkel”, also lives in “Berlin-Mitte”:

$$\begin{aligned} & \text{marriedTo}(\text{Joachim_Sauer}, \text{Angela_Merkel}) \\ & \wedge \text{livesIn}(\text{Angela_Merkel}, \text{Berlin-Mitte}) \\ & \rightarrow \text{livesIn}(\text{Joachim_Sauer}, \text{Berlin-Mitte}) \end{aligned}$$

Moreover, a special type of Horn clause can be formulated by grouping purely negated facts into a disjunction, as in the following example:

$$\begin{aligned} & \neg \text{bornIn}(\text{Angela_Merkel}, \text{Hamburg}) \vee \\ & \neg \text{bornIn}(\text{Angela_Merkel}, \text{Bremen}) \end{aligned}$$

With this Horn clause, we can express that “Angela Merkel” cannot have been born both in “Hamburg” and also in “Bremen”. By using negations as in the above example, we now have a formal means to identify inconsistencies. If, in the extraction step, both facts $\text{bornIn}(\text{Angela_Merkel}, \text{Hamburg})$ and $\text{bornIn}(\text{Angela_Merkel}, \text{Bremen})$ have been inserted erroneously into the knowledge base, we can formulate that only one of the two facts may be true by using negations of the above form.

Efficient evaluation strategies

URDF supports both logic-based (propositional) and probabilistic evaluation strategies to answer user queries. In propositional reasoning, the user is guaranteed to obtain a consistent overview of a potentially inconsistent knowledge base in response to a query. This problem is reduced to the maximum satisfiability problem (Max-Sat), a classic problem in propositional logics. For URDF, we have developed a highly efficient method for solving a generalization of the Max-Sat problem, which is specifically tailored to the combination of soft and hard rules described above. In probabilistic reasoning, on the other hand, URDF does not only assign binary true/false values to facts, but also confidence weights, which correspond to a probabilistic interpretation of the derivation of the facts via the rules.

This enormous expressivity of rule-based and probabilistic inference techniques certainly poses major challenges to the development of new and efficient evaluation strategies for user queries. Both logic-based and probabilistic evaluation strategies underlie a much higher combinatorial complexity than traditional query evaluation strategies in relational databases. Therefore, exact evaluation algorithms cannot be applied to large volumes of data, which we obtain through information extraction from the World Wide Web. Our research concentrates on efficient approximation algorithms with good approximation guarantees, which are specifically tailored to these evaluation strategies. ...



CONTACT

Martin Theobald
DEPT. 5 Databases and Information Systems
 Phone +49 681 9325-5007
 Email martin.theobald@mpi-inf.mpg.de

EnBlogue – What is New and Interesting in Web 2.0?

The success of the Web 2.0 has drastically changed the way information is generated on the Web, with millions of users actively contributing content in form of blog entries or micro-news (e.g., Twitter). In order to maintain an overview of this continuously flowing stream of information, we need methods that extract the essence of current events and display them in a suitable way to users.

With enBlogue, the Max Planck Institute has developed an approach which continuously scans Web 2.0 streams for interesting events, which have received attention in blogs, Twitter, or other media. The name enBlogue is supposed to sound like “in vogue”, meaning it is popular.

EnBlogue processes web contents, annotated with the time of creation and semantic notes to detect surprising changes in the behavior of topics, which can be interpreted as interesting events. Not only individual semantic annotations (so-called *tags*) are assessed for their popularity, but all correlations between different tags are considered. Two tags are strongly correlated if the relative number of documents that report about two tags at the same time is high. Therefore, the dynamic increase in correlations is particularly interesting.

For example, normally the tags “Iceland” and “flight operation” are not very strongly correlated. This changed abruptly with the volcanic eruption of Eyjafjallajökull in the spring of 2010.



The enBlogue website

The choice of tags to be considered is a priori unlimited but is primarily based on annotations generated by users, for example, when commenting on a blog entry or a news paper article. Methods can also be used for classifying information into topic areas as well as for discovering names of persons or places. The enBlogue system then also enables a mixture of peoples, places, and regular topics.

A model that takes into account the quantity of change as well as time-lines is used to extract interesting emergent topics and to present them in a clear and user-friendly manner.

EnBlogue is the starting point to current information on the web and it helps users being linked to actual contents. Personalization, i.e., re-ranking

of topics depending on user preferences and interests, and additionally also considering the locations of events and users (if this information is available), is currently being researched. ...



Events identified by enBlogue on July 8, 1996 (with the assistance of the New York Times Archives)

CONTACT

Sebastian Michel
 DEPT. 5 Databases and Information Systems
 Phone +49 681 302-70803
 Email smichel@mpi-inf.mpg.de
 Internet <http://blogue.mmpi.uni-saarland.de>

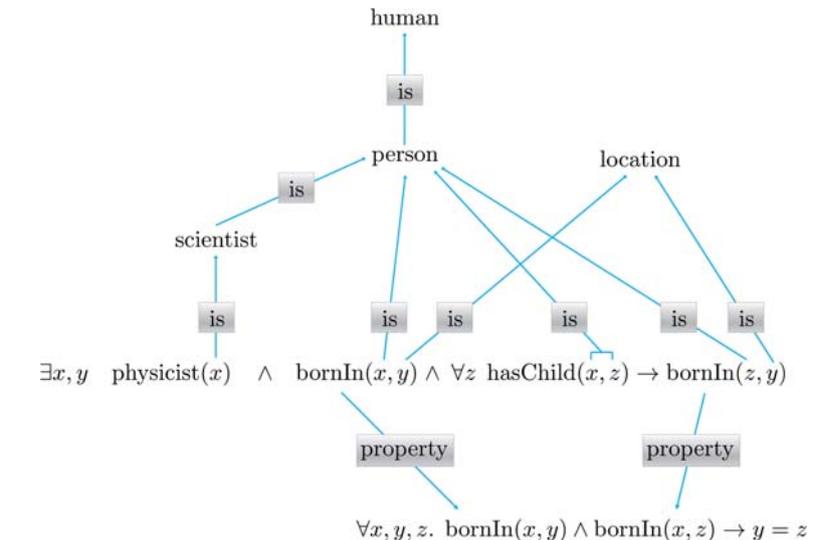


Decision Procedures for Ontologies

Today, we can already find answers to many questions in the Internet with the assistance of search engines. However, these search engines perform only a purely syntactic search based on keywords. For the query “Which physicists were born in the same place as all their children?” a search engine extracts the keywords “physicists”, “children”, “all”, and “born” and returns documents where these keywords occur exactly as a string. For the above query, a search engine delivers a vast number of documents that contain the keywords but typically not the answer to the question. If the query is reformulated, the search engine may return a document containing the answer. The problem is that the search engine is unaware of the meaning of the words in both the query and the documents. It only performs a search for the words, in particular the character strings, on a purely syntactic level.

In order to enable a computer to understand the meaning of words and sentences, certain knowledge is required, which must be explicitly available and processable by a computer. For example, the above query involves the knowledge that physicists and children are people, and that people have only one birth place. The aggregation of this kind of general knowledge is called an ontology. Ontologies can even be extracted automatically from different sources of knowledge, for example the YAGO ontology (see “YAGO – a Collection of Digital Knowledge”, page 57) consisting of several million entries.

One of today’s research challenges is to quickly answer the questions that can be answered by means of the knowledge comprised in an ontology. To do this, we first translate the ontology into a logical representation. With the help of the appropriate methods, a computer is able to fully automatically deduce everything that is implied by the logical representation of an ontology. In order to efficiently answer queries pertaining to an ontology, the logical representation of the ontology is compiled. More precisely, it is transferred into a compact logical representation containing all its ne-



cessary logical consequences and does not contain a contradiction. The computer uses this representation to efficiently find answers to questions that are logical consequences resulting from the knowledge in the ontology.

During the compilation process, intensive logical calculations are performed. This becomes an issue when ontologies with several tens of millions entries are considered because then existing procedures do not successfully compile ontologies of this size in reasonable time. Therefore, we are adjusting the logical deduction procedures such that they suit the particular structures of ontologies. This results in procedures that are several orders of magnitudes faster.

Our procedures currently compile ontologies like the YAGO ontology, with up to 10 million entries in about one hour and can answer queries within seconds.

The above-mentioned query contains a so called quantifier alternation: “there are (physicists) – for all (children)”. The answers for queries with quantifier alternations are particularly hard to compute. In principle, our approach supports queries with arbitrary many quantifier alternations. The query “Does every German physicist have at least one ancestor who was also a scientist?” is, for example, a “for all – there are” query. We are already able to answer these kind of queries from the compiled YAGO ontology in less than one second. ...



CONTACT

Christoph Weidenbach

RG. 1 Automation of Logic

Phone +49 681 9325-2900

Email weidenbach@mpi-inf.mpg.de



Patrick Wischniewski

RG. 1 Automation of Logic

Phone +49 681 9325-2918

Email wischnew@mpi-inf.mpg.de

MULTIMODAL INFORMATION

The volume, but also the importance of multimodal data has virtually exploded in recent years. Research on handling such data is thus becoming ever more important. This extends from the modeling and indexing to the understanding of multimodal sensor data. The Max Planck Institute for Informatics accommodates this development and is pursuing some of the most important issues arising from this area.

Sensors such as cameras, microphones, GPS and accelerometers are increasingly embedded in equipment and environments and are already useful in a variety of ways. The computer-controlled processing of sensor information has made enormous progress but is generally limited to simple matters. This means, in particular, that devices and computers with access to this sensor information do not fully interpret it and thus cannot truly understand their environment.

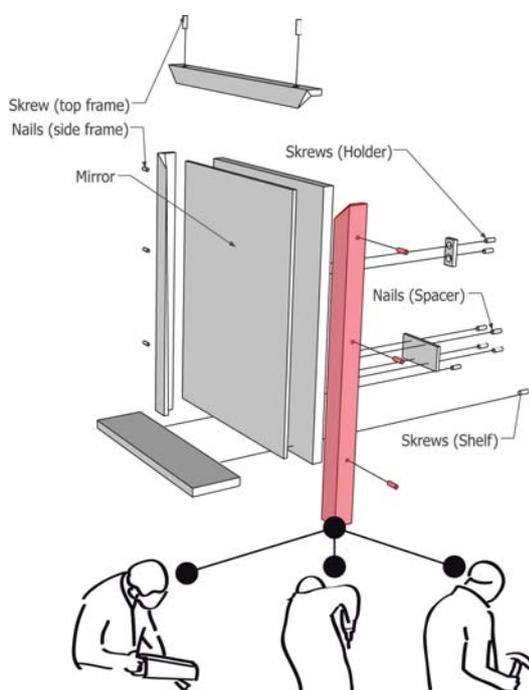
Two work areas will be presented below that address different problems involved in multimodal sensor processing. The first area is concerned with the automatic processing of music data. A central topic here is the efficient finding of individual audio segments or other music fragments in large music data collections. The processes being researched are not only useful for processing search queries, but also for content-based music analysis. This involves, among other things, the recognition and linking of semantic relationships in various versions of different modalities of a piece of music.

In a second area, we are researching the use of portable sensors that are already embedded in cell phones, or even in clothing. In this case, the focus is on capturing and understanding the context of a person playing an important role in man-computer interaction. Context awareness may enable natural communication, for example, with robots that understand the users' goals and offer support at the right time. The work described here concentrates on a specific type of context: the recognition of human activity. ...

CONTRIBUTIONS

Recognizing Human Activity 64

Automated Music Processing 65



UNDERSTANDING
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH
& DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION

Recognizing Human Activity

Recognizing human activity from wearable sensors

Sensing and understanding the context of the user plays an essential role in human-computer-interaction. It may enable natural communication, for instance, with robots that understand the user's goals and offer support at the right time. Our work focuses on a particular type of context, namely human activity recognition.

While impressive progress has been made in recognizing short and atomic activities (such as shaking hands and walking), research on more complex human activity lasting for minutes or hours (such as a morning routine or an assembly task) is far less mature. We thus focus on several important aspects of human activity recognition related to complex and long-term activities.



Figure 1: Wrist Sensor, Embedded Sensor Systems TU Darmstadt

Collecting human activity data

Wearable sensors attached to the body [Figure 1] have great potential for sensing what the user is doing, at any given time or place, from a first-person perspective. Given the advances in micro-technology, inexpensive sensors are already today becoming widely available in watches, cell phones, and even clothing. Motion data can then be collected and analyzed for activity understanding, using machine learning techniques.

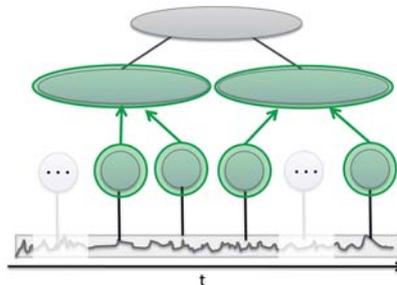


Figure 2: Relevant parts for composite activity

Identifying and combining relevant activity events

It is not usually the sequence of atomic activities that is interesting, but rather the higher level goal at which these activities are directed. There are several ways to infer the higher level goal from observing atomic activities. Since composite activities can contain large amounts of unrelated activity, using the *complete* observation can be suboptimal and therefore confuse the recognition. For many composite activities, it is sufficient to spot only a few underlying activity events to allow their recognition [Figure 2]. For example *having lunch* can be characterized by *walking at a certain time of day*, without even observing the actual *eating* activity. In a discriminative analysis, we observe that a surprisingly small fraction of relevant parts can be sufficient to recognize the higher level composite activity and allow efficient recognition algorithms.

Hierarchical model for composite activities

Preserving the structure of a hierarchical activity offers several benefits. Considering a construction manual for a

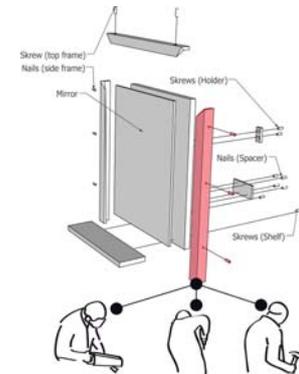


Figure 3: Composite activity for construction task

mirror in figure 3, one of several tasks is to fix the frame to the panel. This seemingly simple task consists of various steps, and it becomes obvious that composite activities add significant variation. Composite activities can be interrupted, the duration can vary strongly across different users, or the underlying activities can happen in different order. Using the same algorithms recognizing atomic activities can be suboptimal, as these require prohibitive amounts of training data. Therefore, we propose a hierarchical model that observes *relevant* activity events and combines them to recognize composite activities, similar to the way in which letters create words. Experiments show indeed superior performance compared to the standard approaches usually used for activity recognition.

Transferring knowledge

Parts that are similar in different composite activities can be shared, much like vocabulary. Instead of re-learning composite activities from scratch, transferring shared parts reduces the training effort for new composite activities significantly. ...



CONTACT

Ulf Blanke
DEPT. 2 Computer Vision and Multimodal Computing
Phone +49 681 9325-2000
Email blanke@mpi-inf.mpg.de



Bernt Schiele
DEPT. 2 Computer Vision and Multimodal Computing
Phone +49 681 9325-2000
Email schiele@mpi-inf.mpg.de

Automated Music Processing

The digital revolution has brought a massive increase in the availability and distribution of music-related documents. Therefore, the development of techniques and tools for organizing, structuring, retrieving, navigating, and presenting music-related data has become a major strand of research – the field is often referred to as *music information retrieval* (MIR). Major challenges arise because of the richness and diversity of music in form and content, leading to novel and exciting research problems. Exemplarily, we discuss three current research tasks in the context of content-based audio retrieval.

The task of *audio identification* (also called *audio fingerprinting*) is to identify a particular audio recording within a given music collection using a small audio fragment as query input. Even for large scale music collections and in the presence of signal distortions such as background noise, MP3 compression artifacts, and uniform temporal distortions, recent algorithms for audio identification yield good recognition rates and are used in commercial products such as Shazam. However, existing identification algorithms cannot deal with strong non-linear temporal distortions or with other musically motivated variations that concern, for example, the articulation or instrumentation.

The task of *audio matching* can be seen as an extension of audio identification. Here, giving a short query audio fragment, the goal is to automatically retrieve all musically related fragments contained in the documents (e. g., audio recordings, video clips) of a given music collection. Here, as opposed to traditional audio identification, one allows semantically motivated variations as they typically occur in different performances and arrangements of a piece of music.

For example, two performances may exhibit significant non-linear global and local differences in tempo, articulation, and phrasing as well as variations in executing ritardandi, accelerandi, fermatas, or ornamentations. Furthermore, one has to deal with considerable dynamical and spectral deviations, which are due to differences in instrumentation, loudness, tone color, accentuation, and so on. Recent matching procedures, which can deal with some of these variations, are based on so-called *chroma features*. Such features closely correlate to the musical aspect of harmony and have turned out to be a powerful mid-level representation applicable to a variety of multimodal retrieval scenarios. As illustration for such a scenario, the figure below shows an interface for simultaneous presentation of visual data (sheet music) and acoustic data (audio recording). The first measures of the third movement (Rondo) of Beethoven's Piano Sonata Op. 13 (Pathétique) are shown. Using a visual query (measures marked in green; theme of the Rondo), all audio documents that contain some matches are retrieved. Here, one audio recording may contain several matches (green rectangles; the theme occurs four times in the Rondo).



User interface for multimodal music retrieval.

Audio identification and audio matching are instances of *fragment-level* retrieval scenarios, where time-sensitive similarity measures are needed to locally compare the query with subsections of a document. In contrast, in *document-level* retrieval, a single similarity measure is considered to globally compare entire documents. One recently studied instance of document-level retrieval is referred to as *cover song identification*, where the goal is to identify different versions of the same piece of music within a database (including cover, remake, and remix versions). Also in this context, chroma-based audio features in combination with local alignment techniques (e.g. Smith-Waterman algorithm) have been applied successfully. By using so-called *shingling techniques* in combination with *locality sensitive hashing* (LSH), the document search can be significantly accelerated. ...



CONTACT

Meinard Müller

DEPT. 4 Computer Graphics

Phone +49 681 9325-4005

Email meinard@mpi-inf.mpg.de

OPTIMIZATION

Optimization procedures are today of central significance for companies' effectiveness. They are employed, for example, to reduce the need for expensive resources such as work or raw materials. The challenge to science is to develop efficient procedures for solving optimization problems. These procedures should quickly find an optimal solution, or at least a solution that is close to the optimum.

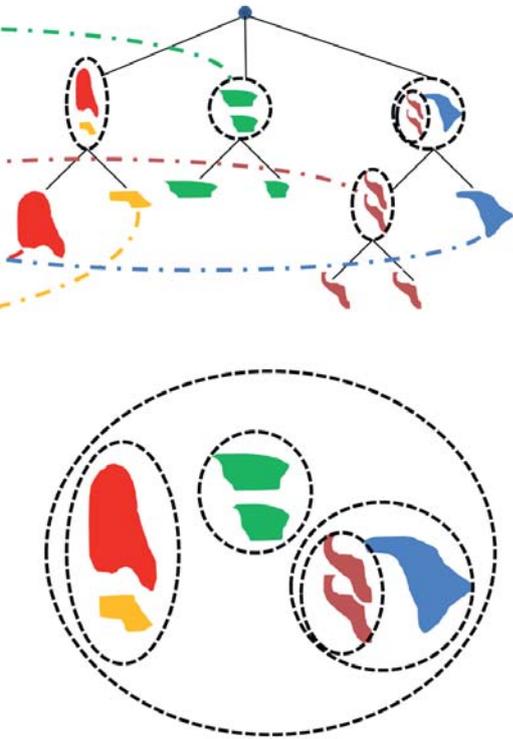
Efficient optimization procedures are of core importance in various areas. For large companies, they have a decisive influence on their competitiveness. With careful planning, a large amount of resources in industrial projects can be saved, leading to lower costs. However, such planning problems are usually very complex, and different requirements have to be taken into account. This makes it difficult for computers to find optimal, or at least very good, solutions.

At the Max Planck Institute for Informatics, we are working on difficult optimization problems arising in very different applications such as industrial optimization or medicine. On the one hand, we are developing elaborate procedures that find optimal solutions efficiently. On the other hand, if the underlying problem is too difficult to find an

optimal solution quickly, we develop procedures that can at least find a solution which is close to the optimum. In addition, we are investigating how we can employ random decisions to obtain more efficient and simpler optimization procedures. In this context, we are also studying procedures that are inspired by optimization processes in nature. Such methods often allow good solutions to be found for a given problem without spending too much effort on designing a specific optimization procedure.

Since optimization plays a significant role in many different areas, scientists in all research areas of the institute are working on optimization problems. Optimization is nowadays a crucial technique for the efficient design of planning processes. Its importance will continue to grow in the future. ...

CONTRIBUTIONS



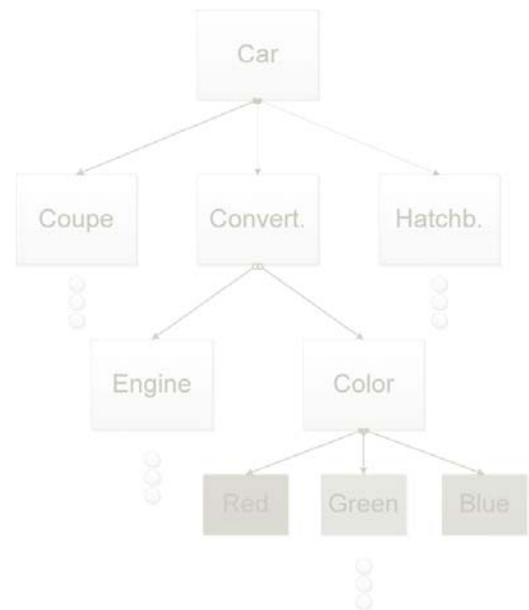
The Theory of Bio-inspired Algorithms 68

Approximation Algorithms for Hard Optimization Problems 69

Placement of Wind Turbines 70

Combinatorial Optimization in Information Science 71

Model-Based Product Configuration 72



UNDERSTANDING
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

INFORMATION SEARCH
& DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION

OPTIMIZATION

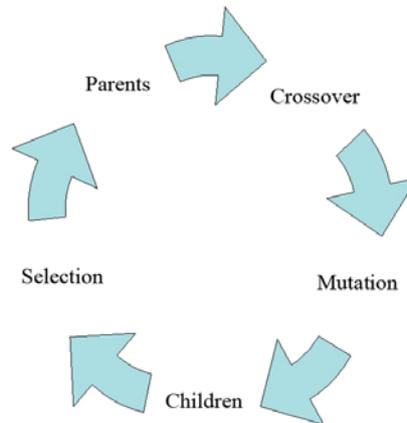
SOFTWARE

VISUALIZATION

The Theory of Bio-inspired Algorithms

Bio-inspired algorithms are general search algorithms, that are used in many ways in engineering disciplines and in the area of combinatorial optimization. An example of such algorithms are evolutionary algorithms (EA). This classical solution process follows from the evolutionary model and the Darwinian principle of “*survival of the fittest*.” Depending upon this natural evolution principle, a special solution candidate is called *individual* and a number of such candidates is called *population*. A so-called fitness function, which depends upon given problems, assesses the solution candidates. The offspring population can be produced from a parent population according to the biological principle. This occurs through so-called variation operators, in which the parents’ genetic material is inherited by the children. The most important operators are, in this case, *crossover* and *mutation*. Crossover usually produces a child from two adults, while mutation provides the child with new properties. Beginning with a starting population, the goal for a given problem is to retain as many good solution candidates as possible. After having produced the offspring a new parent population is created by selecting individuals from the parents and offspring. This process is repeated as long as necessary to find a good solution [see Figure].

Another class of bio-inspired algorithms is ant colony optimization, which uses the image of pathfinding and communication mechanisms of ants. These algorithms construct many possible solutions for a given problem from various components; components that participate in “good” solutions are marked with a chemical messenger, so-called pheromones, and are used in later solutions with higher probability. In doing so, we hope to develop an optimal solution as a combination of many good components.



Flowchart of an evolutionary algorithm

Bio-inspired algorithms are especially employed if, for a given (new) problem, no good, problem-specific algorithm is available. One cannot therefore expect that they can surpass a specially-designed solution process. It is therefore not the goal of our research to demonstrate that bio-inspired algorithms surpass problem-specific algorithms. Rather, the focus is to understand the working methods of this new process.

Research focus

While bio-inspired processes have already been successfully employed in many cases, the theoretical understanding of these algorithms in comparison to classic algorithms is still at an early stage. We are examining if bio-inspired search processes are capable of solving specific problems, and what structures are good or bad in their use. The main focus is on how much time the algorithms need to generate an optimal solution for a given problem. Because bio-inspired algorithms are a special class of randomized algorithms, one can use a large number of classic analytical methods. Furthermore,

new analytical methods are developed in order to respond to the specific characteristics of bio-inspired methods.

We have shown that bio-inspired methods often find good solutions for well-known problems without using problem specific knowledge. They are able to mimic problem-specific algorithms for a wide class of combinatorial problems. For example, it has been formally demonstrated that bio-inspired algorithms can efficiently calculate the shortest path between all nodes in a given graph. Another example is that EAs with recombination and mutation show demonstrable benefits against EAs that only use mutation as a variation operator.

Other studies show that approaches for multi-criteria optimization of bio-inspired algorithms bring additional benefits to an efficient search. Many optimization problems are given through a target function, which must be optimized through many side conditions. These side conditions are perceived as additional equal target functions in multi-criteria approaches. This gives the search for evolutionary algorithms additional search directions, which lead to more efficient processes with different combinatorial optimization problems. ...



CONTACT

Benjamin Doerr
DEPT. 1 Algorithms and Complexity
 Phone +49 681 9325-1004
 Email doerr@mpi-inf.mpg.de



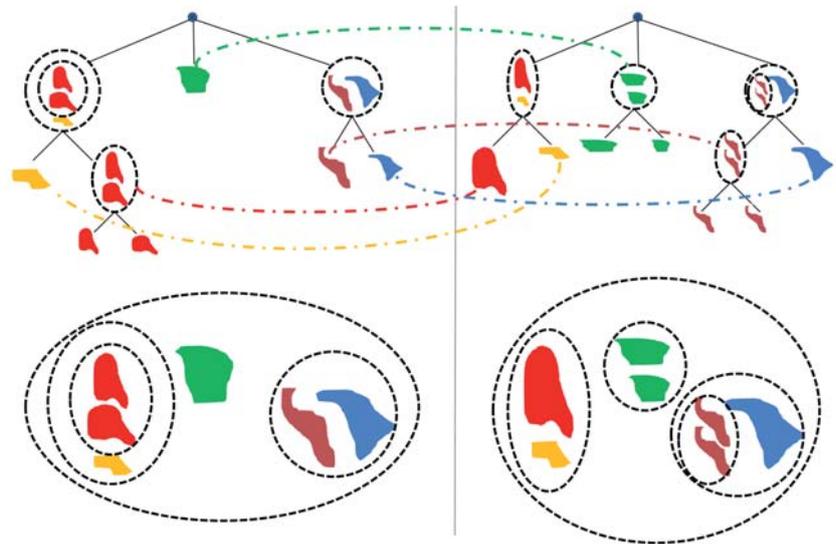
Frank Neumann
DEPT. 1 Algorithms and Complexity
 Phone +49 681 9325-1000
 Email fne@mpi-inf.mpg.de

Approximation Algorithms for Hard Optimization Problems

One of the most basic methods for dealing with NP-hard optimization problems is to design polynomial-time algorithms that find a provably near-optimal solution. Formally, an algorithm is said to be an x -approximation algorithm for a maximization (or minimization) problem if it runs in polynomial time in the size of the input instance and produces a solution whose objective value is within a factor x of the optimum solution. Under this general framework, we have been working on problems from different areas, including graph problems, computational economics, computational geometry and mathematical programming. We give an example below.

Matching of live cell video frames

In joint work with Stefan Canzar, Gunnar W. Klau and Julian Mestre, we have considered a problem that arises naturally in the computational analysis of live cell video data. Studying cell motility using live cell video data helps understand important biological processes such as tissue repair, the analysis of drug performance and immune system responses. Segmentation based methods for cell tracking typically follow a two stage approach: The goal of the first *detection* step is to identify individual cells in each frame of the video independently. In a second step, the linkage of consecutive frames, and thus the *tracking* of a cell, is achieved by assigning cells identified in one frame to cells identified in the next frame. However, limited contrast and noise in the video sequence often leads to *over-segmentation* in the first stage: a single cell is comprised of several segments. A major challenge in this application domain is therefore the ability to distinguish biological cell division from over-segmentation.



Example of a feasible matching between two segment sets in two frames

One approach to solve this problem is to match *sets of segments* between neighboring frames, where the segment sets are obtained by a certain *hierarchical* clustering procedure. See the example, which illustrates the hierarchical clustering of two consecutive frames. Matching here means identifying one set of segments at a certain level in the hierarchy in one frame with one set of segments in the other frame. Since segment sets representing different cells in the same frame must be disjoint, no two nodes on any root-to-leaf path in any of the two hierarchies can be matched at the same time. In the figure, we show a feasible matching using dashed lines. Typically, a non-negative weight is assigned to each pair of segment sets (one from the first frame and the other from the second), indicating how likely this pair is to represent the same cell. The objective is to find a feasible matching between the segment sets in the two frames, with maximum total weight.

While the problem was originally formulated in the Computational Biology community, our contribution was to study it from an approximation-algorithms point of view. We first showed that solving the problem exactly is NP-hard; in fact, we showed that there is a constant value $x > 1$ such that it is even NP-hard to obtain an x -approximation. On the other hand, we developed a 2-approximation algorithm for that problem. ...



CONTACT

Khaled Elbassioni
DEPT. 1 Algorithms and Complexity
 Phone +49 681 9325-1007
 Email elbassio@mpi-inf.mpg.de

Placement of Wind Turbines

Renewable energy plays an increasing role in the supply of energy worldwide and plays a major role in tackling climate change. Since 2005, the cumulative installed capacity of wind energy within the EU has almost doubled (from 40000 MW to 74000 MW). In 2009, 39% of all new energy capacity installed was based on wind. Currently, there are roughly 8800 wind turbines in Europe, which helped to save 180 Mio tonnes of CO₂ since the beginning of 2009. Recently, the installation and start-up of large offshore wind farms (for example the Thanet Wind Farm with a capacity of 300 MW) has gained increasing media attention.

To further increase the productivity of wind farms (onshore and offshore), it is necessary to exploit methods that help to optimize the performance of wind farms. Building wind farms and using wind as a resource to satisfy energy demands involves a lot of complex tasks that determine the success of its usage. Currently, there are only initial attempts that exploit the use of optimization methods in the context of wind energy. In this project, we consider how the placement of wind turbines on large wind farms can be optimized such that they achieve an optimal energy output.

We develop and evaluate new approaches for optimizing the micro-siting of turbines for large wind farms. Layout siting is too complex to fit into the class of optimization problems, which can be solved by traditional methods. Its non-linearity and search space size imply that exhaustive sampling is prohibitive and, instead, some kind of stochastic sampling is required. For example, the “optimizer” component of OpenWind conducts two strategies for sampling but has different weaknesses, as pointed out in their website FAQ.

The optimization task involves a model to predict the energy output for a certain siting. The cost of such an evaluation depends on the accuracy of the model used to predict the energy output. Accurate models may take into account wake effects, temperature, pressure, wind speed and directions as well as other conditions of the area where the wind turbines should be placed. In the case of our problem, the quality of a certain solution can only be evaluated by considering a complex model for the energy output. We tackle the optimization task with evolutionary algorithms (EAs). This kind of algorithms is in particular suitable to tackle highly complex non-linear optimization problems.

EAs work at each time step with a set of solutions called a population. These solutions can be evaluated in parallel in order to speed up the computation. However, current approaches can only deal with a small number of turbines, simple and imprecise models, and show rather weak performance. We demonstrate an accurate, efficient, and parallelizable optimization algorithm for the layout of hundreds, eventually 1000, turbines. It is modular and therefore allows different wake effect models to be incorporated. Its computational cost depends on how many candidate layouts it investigates and on the complexity of its wake loss calculation. We demonstrate how well it maximizes energy yield and show how it allows one to examine how wake loss scales with energy capture and number of turbines. ...



CONTACT

Frank Neumann
DEPT. 1 Algorithms and Complexity
 Phone +49 681 9325-1000
 Email fne@mpi-inf.mpg.de

Combinatorial Optimization in Information Science

We now have the ability to gather enormous amounts of data, and there are numerous challenges in organizing this data into usable information: Researchers are working on developing better web search engines. In biology, microarray experiments supply information about expression levels of thousands of genes, and this data needs to be organized and analyzed. Internet resellers would like to give customized recommendations to users based on known preferences, user input, and previous purchases.

One of the questions that arises when trying to solve these problems is the question of how to combine information from different sources: for example, how do we rank a list of candidates based on the different preference orders given by the members of a search committee? What is a good clustering of a set of genes if different microarray experiments all suggest different clusterings? Or, how do we cluster webpages based on similarity scores?

Rank Aggregation

Some of these questions have a long history, such as the question of how to rank candidates based on the preferences of different voters. The famous Arrow's impossibility theorem states that no voting system can simultaneously achieve three natural criteria: non-dictatorship (we do not simply want to output the preference order of a particular voter), independence of irrelevant alternatives (if the output order ranks Anna before Bob, then this should also be the case if one or more voters change their minds about the qualifications of Charlie), and finally, Pareto efficiency or unanimity (if all voters prefer alternative Anna to alternative Bob, then so does the output ordering).

Despite this negative result, we would still like to find a way to rank candidates based on the preferences of different voters. The computer science community became especially interested in this problem because of its application in *meta search* engines for web search. Different search engines use different algorithms, and they may perform better in certain aspects and worse in other aspects. In addition, algorithms are susceptible to spamming when websites attempt to manipulate the algorithm to achieve a higher position in the search results. Finally, search engines themselves may accept payment to display certain websites in a prominent position. The goal of a meta search engine is therefore to combine the ranked search results from different search engines (the "voters") into a new ranked list of results which minimizes the number of *disagreements* with the input lists, i.e., the number of pairs of results that are in reverse order in the new list compared to the original lists from the search engines.

Of course, other objectives than minimizing the number of disagreements are possible, too, and – by Arrow's impossibility result – no objective is good in all situations. We may want our ranking to be *fair*, for example, by ensuring that we do not contradict any of the input preference orders by too much. Or, we may want to guard against selfish voters and try to ensure that any majority of the voters cannot secede and obtain a better outcome.

Many of these objectives give rise to *NP-hard* problems; this means that the running time of the only algorithms that we know for optimally solving these problems grows exponentially with the size of the problem (and it is widely believed that this behavior is inherent to

the problems). We therefore try to find *approximation algorithms*. An approximation algorithm is an efficient algorithm for solving such a problem, i.e., an algorithm for which the running time grows polynomially rather than exponentially. The solution returned by an approximation algorithm is not necessarily optimal; however, it is guaranteed not to be too far from optimal.

In the past few years, algorithms have been developed for the meta search engine problem. These not only work well in theory, but computational studies have shown that they provide optimal solutions in many practical instances. In our current research we are investigating whether such good results can also be obtained for other objectives. ...



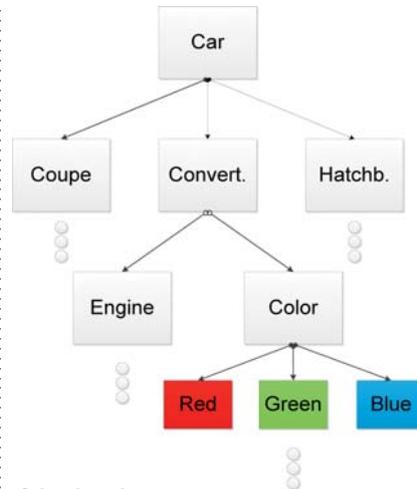
CONTACT

Anke van Zuylen
DEPT. 1 Algorithms and Complexity
 Phone +49 681 9325-1000
 Email anke@mpi-inf.mpg.de

Model-Based Product Configuration

Today, many automotive companies offer configurators for their products on the Internet. Typically, customers put together their own personal vehicle by selecting a basic model and then choosing various options. For some options, selection is obligatory, such as the choice of car color or the engine. Others are optional, such as a tow bar or a telephone module. There are dependencies between the options: a stronger engine must have a specific wheel combination due to the possible higher speeds, while choosing the tow-bar may exclude choosing lower-power engines. Our group researches generic methods for modeling products composed from individual parts, such as cars, to show their complete structure, and for answering queries on their configuration automatically.

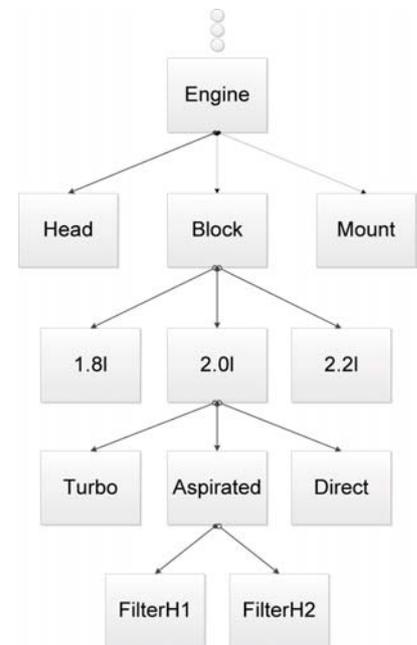
It should be borne in mind that the manufacturer's view of a car is much more complex than the customer's view, as the manufacturer's view not only includes options and parts that are relevant for sale, but ultimately the entire car down to its individual screws. Furthermore, manufacturers differentiate the design view (*Which parts compose the product?*) from the manufacturer's view (*How do I build the product?*) and the financial view (*What does the product cost? What profit margin does it provide?*). In addition, manufacturers no longer work on the basis of individual products, but rather they have so-called 'product lines', which are used to produce different products on the basis of the highest possible fraction of common parts. For example, the Audi A3 and the VW Golf belong to the same product line of the VW group. Today, the above-mentioned views are not typically generated together, but separately. Our goal is to make it possible to have an integrated view of a product and to automatically compute relevant properties.



Sales view of a car

Examples of such properties are the number of possible products from one product line with respect to a set of subassemblies to be used, or the consistency of subassemblies in relation to all products, i.e., whether these contribute to a product at all. Further possibilities could be an enquiry regarding a product with minimum manufacturing costs or the aforementioned sales-oriented configuration view of the product. All of these properties are in general "hard", i.e., there is currently no procedure that can compute the above properties for any desired product built from many parts within an acceptable period of time. For the car example, we can currently process product models up to a size of

about 6,000 parts in a reasonable period of time. The model then potentially represents 2^{6000} different products, much more than the estimated number 2^{1300} atoms in the universe. Such large models can only be processed through automatic and intelligent exploration of the underlying dependency structures. A complete product-line in automobile construction needs about 50,000 parts. Algorithms for problems of this size are the object of our current research. ...



Design view of an engine



CONTACT

Christoph Weidenbach

RG. 1 Automation of Logic

Phone +49 681 9325-2900

Email weidenbach@mpi-inf.mpg.de

SOFTWARE

Informatics is firstly a discipline of basic research that deals with universal computation and problem-solving methods and investigates fundamental properties such as correctness and complexity. Secondly, it also resembles an engineering science that supports a great variety of different applications. The development of software fulfills many functions in this context. It is the subject of basic research, the outcome of the implementation of new results from basic research, as well as the product out of an engineering effort for the solution of a concrete application problem. Software is thus an inherent part and a connecting link in informatics research.

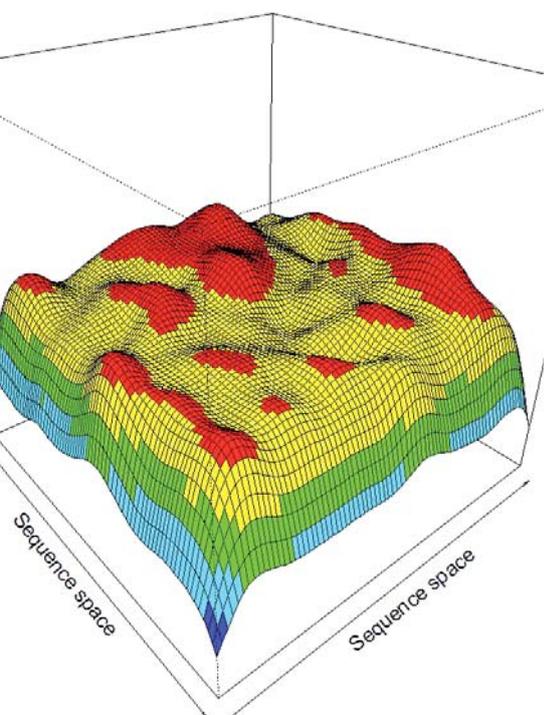
At the Max Planck Institute for Informatics this philosophy has been implemented with great success since the founding of the institute. All departments and groups work to implement the results of their basic research efforts into practical software systems and make them available to science and industry. There is a considerable number of software systems, developed in all departments and groups at the institute, that have found their way into the scientific community and that are used in research in many places around the world. Most of our software is distributed as free open source software. In some cases, however, start-up companies have been founded to further develop, distribute, and sell the software.

The clever implementation of mathematical models and algorithms into executable software is, as mentioned above, an important research subject. Algorithms that show, in general, nice mathematically analyzable run time and storage space properties (so-called asymptotic complexity dimensions) are not necessarily efficient when implemented on modern computers, particularly on current distributed computer hardware. Properties of communication, processors, storage and magnetic disks as well as the characteristics of real data must be appropriately considered in engineering if software systems are to be successfully applied

in practice. When we succeed in building a software system for a newly developed method, the system in turn provides valuable information on relevant special cases or meaningful generalizations of the problem investigated or the methods used.

The articles on the research focus of "software" are a cross-section through our activities. A current research topic is the enhancement of syntactic search possibilities in the World Wide Web (WWW) with semantic information. The articles "*TopX 2.0 – Efficient Search in Digital Libraries*" and "*RDF-3X – Fast Searches for Semantic Data*" describe two software solutions related to this topic. In bioinformatics, we work on software to automatically generate drug therapies. In the articles "*Analysis of HIV Drug Resistance*" and "*Analysis of HBV Resistance*", two software systems are presented for these problems. It is a great challenge to develop algorithms that increase performance on parallel hardware architectures which have meanwhile become standard. The article "*GBACE: Parallel Processing of Adaptive Data*" describes software that forms a layer between a parallel architecture and a software, aimed at solving a specific problem eventually running on this architecture. The additional layer eases the use of parallelism for the software developer. ...

CONTRIBUTIONS



Analysis of HBV Resistance 76

Analysis of HIV Drug Resistance 77

GBACE: Parallel Processing of Adaptive Data 78

TopX 2.0 – Efficient Search in Digital Libraries 79

RDF-3X – Fast Searches for Semantic Data 80

UNDERSTANDING
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

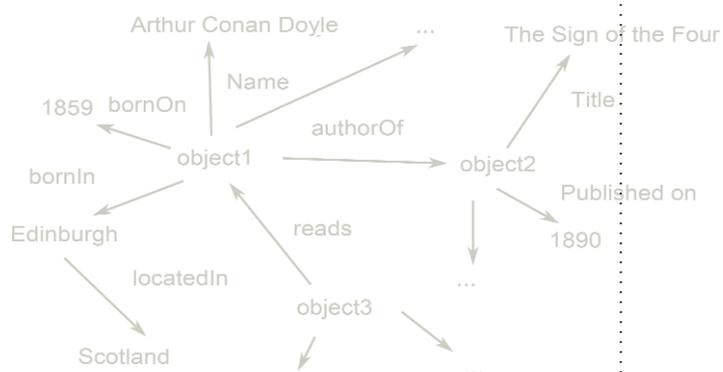
INFORMATION SEARCH
& DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION

OPTIMIZATION

SOFTWARE

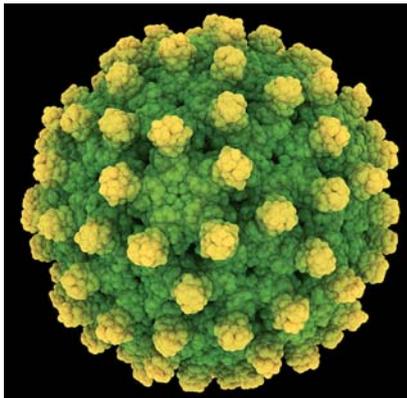
VISUALIZATION



Analysis of HBV Resistance

Hepatitis B Virus

The hepatitis B virus (*HBV*) infects human liver cells, which can result in an acute or chronic liver inflammation (*hepatitis*). The World Health Organization (*WHO*) estimates that 300 to 420 million people worldwide suffer from chronic hepatitis B and about one million people per year die as a consequence of this disease. In order to prevent the later sequelae such as liver cirrhosis and liver cancer, antiviral therapy must begin early. The goal of treatment is the permanent prevention of viral replication in liver cells. Several drugs are available for this purpose.



Hepatitis B virus particles
(Image source: Antonio S iber)

Reverse transcriptase inhibitors

Of great medical importance are the reverse transcriptase inhibitors. These prevent the reproduction of viral DNA through a simple, but elegant trick. The building blocks of DNA, the nucleotides, are normally connected in a long chain. The relevant chemical binding sites are found on both sides of the nucleotide. Reverse transcriptase inhibitors used in the treatment of hepatitis B have the property of their chemistry modeled after normal nucleotides; therefore, they are

incorporated into the DNA chain by the viral polymerase. However, one of the chemical groups that causes chain lengthening is missing so that a chain break is caused and the viral DNA cannot be completed. This therapy slows down the reproduction of the virus in a decisive way, thus preventing bound consequential diseases of the liver, but only as long as the virus has not adjusted to the treatment and developed resistance. The development of resistance is a major difficulty in the long-term treatment of hepatitis B.

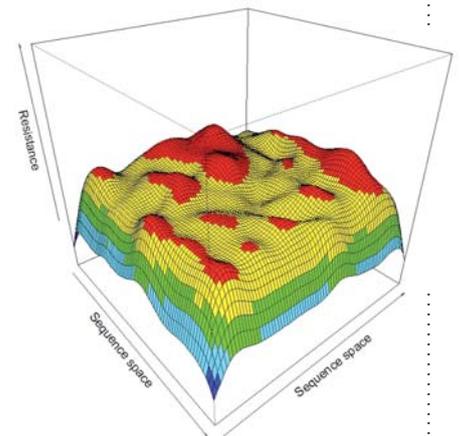
Resistance factor

In order to optimize therapy decisions, we have set a high goal, together with our collaborators in medicine and biology. Our task is to completely measure the hepatitis B virus in terms of resistance. In extensive laboratory tests, one can determine a virus' resistance in the form of an easily interpretable numerical value. This resistance factor must ideally be estimated before each therapeutic decision is made regarding a patient sample and any available drugs. Unfortunately, this is not possible for a variety of reasons. On the one hand, laboratory testing is expensive, and on the other hand, it takes several weeks to establish the results. We would therefore like to describe a method that we successfully tested with the human immunodeficiency virus (HIV).

Resistance landscape

The basis of our research is that the characteristics of the virus are encrypted in its genome. If we find a way to determine the resistance factor of the viral genome, we can always come back to this model. This requires several hun-

dred measurements of resistance factors for different viral genomes. The clever choice of these measurements is crucial for our project. First, we can only reflect contexts that we have previously seen in our measurements, and second, we have to infer the genotypic causes to derive the relevant conclusions. We have now developed a system which chooses several hundred viral prototypes from several thousand hepatitis B viruses from all of Europe. These prototypes reflect the diversity of the virus, and make it possible for us to develop an algorithm which can in the future support the choice of the most effective individual therapy. ...



Hepatitis B resistance landscape
(schematic presentation)

CONTACT

Bastian Beggel
DEPT. 3 Computational Biology and
Applied Algorithmics
Phone +49 681 9325-3016
Email beggel@mpi-inf.mpg.de
Internet <http://www.geno2pheno.org>



GBACE: Parallel Processing of Adaptive Data

The parallel revolution

Around 2004, a radical change occurred concerning the fabrication of microprocessors. Due to physical limitations, the leading development of larger and faster processor cores was given up and traded in against an exponentially growing hardware parallelization in many cores within the same chip. This basic change in the design of microprocessors poses an extreme challenge for software, which now requires explicit parallelization to benefit from proceeding hardware improvements. Without this effort, software cannot become any faster on new high-performance hardware.

Parallel co-processors

In 2004 the first double-core processors were released, however, these did not yet present the same requirements on software, which have come up with the exponentially increasing number of cores over the years. At that time, we found ourselves at the beginning of the multi-core era. In general, one differentiates between multi-core architectures with 2 to 8 cores, many-core architectures with 16 to 64 cores, and massively parallel architectures with 100 and more cores. While parallelization on multi-core architectures can be approached with traditional methods, the main challenge for software is in scaling to many-core and highly-parallel architectures such as graphics processors and future many-core CPUs.

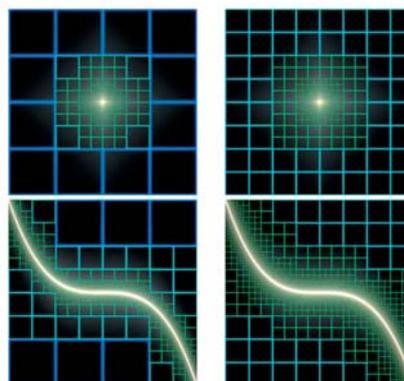
Parallelism on many levels

The use of many-core co-processors allows a significant acceleration for many applications. For larger scientific problems, one would like to use several co-processors

within a computer, so that the problem can be distributed across the memory of several co-processor cards. This presents the first parallelism level. A second parallelism level exists due to many cores within each co-processor. Finally, there are several computing units in each core which form a third parallelism level. For an efficient software execution, it is very important that all of these parallelism levels are used efficiently.

Adaptive data

Adaptive data presentation is often used in order to approximate continuous functions up to a given precision. The main idea is the principal of using a fine resolution in areas where it is difficult to represent the functions, and the use of a rough resolution in places where the representation is easier. In the left column of the illustration we show an adaptive grid representation of two functions. Depending on a precision threshold, the function is represented with different resolutions in different regions of the domain. We use varying colors and block sizes for the different resolution levels in the illustration: The larger the block, the rougher the resolution.



Adaptive grid generation and refinement using parallel recursion on graphics processors

Many applications which work on such adaptive grids allow only one level transitions between the resolutions of neighboring blocks. The realization of this condition leads to a recursive refinement of many blocks in the left column of the illustration. The appropriately refined grids can be found in the right column of the illustration. We see a softer transition between the resolutions as compared to the left column.

Parallel and adaptive

In order to solve scientific problems efficiently, one would like to combine the adaptive data representation with the parallelism on all three levels. Ideally, this would be achieved independently of the specific many-core co-processor, and could therefore run with different types of hardware. It is, however, already difficult to use the many levels of parallelism on problems that utilize the same resolution everywhere. Software complexity increases greatly for adaptive data representations and the hardware-independent coding poses an additional challenge. Due to this complexity, the potential of many-core co-processor, especially for larger software projects, has been little used up to now. The framework GBACE (*General Block Adaptive Concurrent Environment*) abstracts the parallel processing of adaptive data and incorporates many of the hardware requirements automatically, offering a simpler access to the great potential of many-core processors. In this way, the benefits of the parallel revolution will finally also have a positive impact on the performance of complex software projects. ...

CONTACT

Robert Strzodka

DEPT. 4 Computer Graphics

Phone +49 681 9325-4157

Email strzodka@mpi-inf.mpg.de

Internet <http://www.mpi-inf.mpg.de/~strzodka/software/GBACE>



TopX 2.0 – Efficient Search in Digital Libraries

Full text search in semistructured documents

TopX is a search engine for efficient full-text search in semi-structured documents, a data format which is frequently used, for example, in digital libraries. These documents are called semi-structured as they combine rich semantic annotations and structural markup with extensive text passages in a unified format, the popular XML (Extensible Markup Language). In particular semantic annotations are of interest for query evaluation because they can contribute to disambiguating the textual contents and thus to better understanding the textual information. Let us consider, for example, the keyword query “*Who were doctoral students of Max Planck, who themselves became scientists?*” Assuming that appropriate annotations exist in the collection, this keyword query can be translated into a very specific path query over a corresponding document structure:

```
//article[.//person ftcontains „Max Planck“]
//doctoral_students//scientist
```

(according to XPath 2.0 Full-Text Standard)

The results of this query over a Wikipedia-based, annotated XML corpus, will then be returned by the engine as direct references to “*Gustav Ludwig Hertz*” or “*Erich Kretschmann*” as compact links, for example, to their own encyclopedia entries.

Exploratory search

The prerequisite for a successful application of this query language is however the user’s detailed knowledge of the document structure, the so-called XML schema. In order to support the user in formulating structured queries over an initially unknown document schema, TopX also allows for an exploratory form of search. Starting with simple keyword queries, the user can employ the search engine to gradually learn more about the document’s schema and then refine the queries step by step via structural conditions. The expressiveness of

the query language (which supports the international XPath 2.0 and XQuery 1.0 Full-Text standards defined by the W3C) allows for a gradual clarification of the requests.

Following the above example, the user may start searching the annotated Wikipedia corpus with a pure keyword query for “*Max Planck*”. This query would provide (amongst others) links to the *person* Max Planck as well as to the various Max Planck {institutes} among its best matches. Taking advantage of the document structure, a simple refinement of the query might then look like the following:

```
//article[.//person ftcontains „Max Planck“],
```

which would tell the search engine that articles about the *person* “*Max Planck*” would be preferred.

Moreover, different forms of query relaxations and ontological query expansions are supported by TopX. On the one hand, overly specific queries that would not yield any exact matches can be dynamically relaxed at query time by automatically switching from a conjunctive query evaluation mode into a disjunctive evaluation strategy. On the other hand, through ontological query expansions, TopX can also find matches to semantically similar concepts.

Efficiency and scalability

Much of the implementation of TopX has been continuously improved and in part completely redesigned over

the past years. In particular, the internal index structure used by TopX was transformed from a conventional database system into a completely new, highly optimized index structure, which was specifically designed for TopX. This new index structure supports highly efficient forms of index compression and the distribution of indexes across multiple computers (i.e., compute cluster), which contributes significantly to the scalability and efficiency of the search engine. Our experimental results and the regular, very successful participation in international benchmark competitions support the globally strong position of TopX for information retrieval over semi-structured data that significantly outperforms commercial products.

Applications and open-source availability

TopX targets both the direct user application as well as the API-based integration with other systems. With the international XQuery 2.0 Full-Text standard, TopX also offers improved support for imperative programming language concepts, such as loops and conditional statements, which increases TopX’s usability with automatically generated queries.

The original prototype of TopX was already made available as an open-source package in 2006. A respective open-source distribution of the current TopX 2.0 version with support for distributed indexing, incremental index updates, and the XQuery 2.0 Full-Text standard is planned for 2011. ...



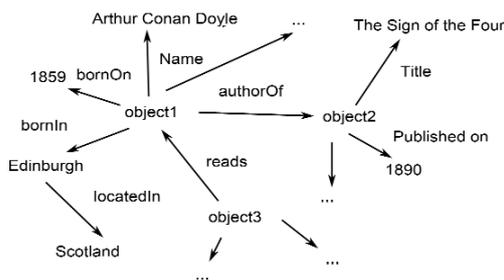
CONTACT

Martin Theobald
DEPT. 5 Databases and Information Systems
 Phone +49 681 9325-5007
 Email martin.theobald@mpi-inf.mpg.de

RDF-3X – Fast Searches for Semantic Data

Semantic Data

RDF-3X is a database system to manage and search for semantic data, that is, data which contains information about relationships between objects. Such relationships are ubiquitous. For example, a book has one or more authors, a protein takes part in specific reactions and a user of web 2.0 web pages has connections to his friends. These relationships between things, or, more abstractly: Entities, together with the entities themselves, form a network or graph structure. Using formal semantics associated with it, the data graphs form semantic networks, i.e., networks which allow for automatic reasoning.



Example RDF graph

A data format, which was specially designed for graph-structured data, is the *Resource Description Framework (RDF)*. An RDF data collection consists of triplets, each of which corresponds to an edge and an associated node pair in the data graph. A triplet consists in RDF writing of a subject, a predicate and an object. In the graph, this corresponds with the “predicate” labeled edge from subject to object. There are edges in the example $\langle id1, bornOn, 1859 \rangle$, which express that the subject “id1” is associated with the name Arthur Conan Doyle, born in 1859. Further edges, such as $\langle id1, authorOf, id2 \rangle$, $\langle id1, bornIn, Edinburgh \rangle$ and $\langle Edinburgh, is\ in, Scotland \rangle$ demonstrate that the object “id2” is written by id1, that id1 is born in Edinburgh, and that Edinburgh is in Scotland.

Search in RDF Graphs

Also, complex relationships can be formulated relatively simply with this kind of notation. The simplicity and flexibility, which is achieved through the graph structure of the data, leads to the fact that RDF searches are relatively expensive. Search queries are normally formulated in a query language such as SPARQL, for example, and describe a pattern that must be searched for in the data. If one searches, for example, for the title of books by Scottish authors, one can describe this with the following triple pattern:

```

?author <bornIn> ?city
?city <locatedIn> Scotland
?author <authorOf> ?book
?book <title> ?titel.
  
```

The parts beginning with question marks are individual variables whose values must be determined by database systems. The portion in brackets is the value given as search conditions by the user. In order to answer this query, the database system must find all triples that fulfill the patterns of the query. Because the data graphs are very large, there can be many millions of candidate triples, making efficient searches very difficult.

The database system we have developed, RDF-3X (RDF Triples Express), approaches these problems on several levels. First, the data itself is stored appropriately, so that individual triple patterns can be efficiently evaluated. We have compressed the data and indexed it using search trees, so that any triple pattern can be evaluated very quickly. But this is not enough, however; the users are mostly interested in larger

interrelationships and thus pose queries with linked triple patterns. Therefore, we map such interrelated queries into execution strategies using algebraic operators and choose the ones that can be expected to have the lowest execution time from the different execution alternatives. The different alternatives are assessed, and then the most efficient alternative is chosen with the assistance of statistical models. For the example query, one must estimate whether there are more authors or more cities in the data set, and how this affects the execution time. A careful choice of the execution strategy can often accelerate the query processing by a factor of 10 or more.

Efficiency and Scalability

Searching in semantic data is already difficult for even relatively small graphs because the graphs frequently show no known structure (i.e., they follow no schema). Despite this, it was important for us to use RDF-3X not only for small data graphs, but also to make sure that it scales efficiently to very large data sets with billions of edges. One big challenge is updating such large data sets, i.e., inserting new data and, if necessary, removing old data, all without having to stop the whole database. Over time, we have therefore incorporated many techniques, such as *Sideways Information Passing* and *Triple Versioning* in RDF-3X in order to scale to such data sizes. In experimental comparison with other systems, RDF-3X usually performs extremely well. Even for complex queries with many triple patterns, RDF-3X is frequently much faster than other systems. ...



CONTACT

Thomas Neumann
DEPT. 5 Databases and Information Systems
 Phone +49 681 9325-5007
 Email martin.theobald@mpi-inf.mpg.de

VISUALIZATION

Pictures are the fastest access to human consciousness. Algorithms for appropriate visualization of digital information therefore play a significant role in informatics. The requirements for visualization algorithms have dramatically increased: Artificial and natural worlds must be presented in an ever more realistic and fast manner – in flight simulators, surgical operation planning systems, computer games or with the illustration of large amounts of data in natural and engineering sciences.

The basis of qualitatively high-value, computer-generated images are accurate scene models. At the Max Planck Institute for Informatics, automated methods to digitize real objects are being researched, which next to their own scene geometry, also exactly measure lighting and reflection properties. Thus, the reconstruction of dynamic scenes gains more meaning and finds many applications in computer animation as well as in the areas of 3D TV and 3D video (see also here research reports from the area “*Understanding Images and Videos*”).

Another object of our research is the reconstruction of statistical models of specific object categories such as faces or human bodies from 3D scans. These models make it possible for us to simulate a movement of the face or the body, and also to simulate the changing face or body form (for example, after gaining weight). The solutions of many problems in image processing, computer animation, and 3D movement measurement are dramatically simplified with these models. One can discover underlying similarities and symmetries between different 3D models as well as within individual 3D models using statistical analysis of 3D geometry. In this way, the basic building blocks of 3D shapes can be seen, and rules can be automatically derived as new versions of these forms are created.

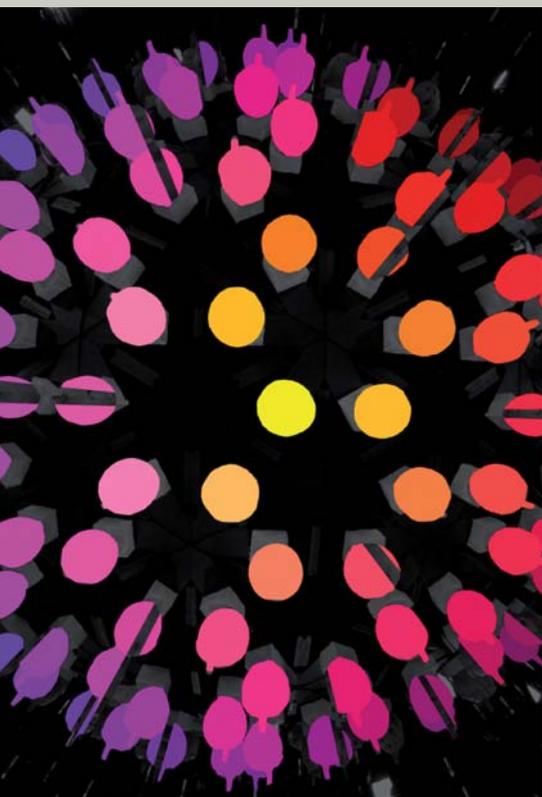
In order to show lifelike models of virtual worlds we have also developed new simulation methods for light propagation in scenes, so-called global lighting. In global illumination, we simulate (in addition to incident light, which is reflected directly from the objects) especially the indirect light, which bounces back and forth over different surfaces. For this, we focus on the development of real-time algorithms.

The images created through image synthesis or image construction typically have a brightness area (dynamic range) which is similar to the real world. We research algorithms for image processing of these so-called High Dynamic Range images (HDR) as well as methods for their portrayal on standard displays.

Photographs and videos also represent forms of visual media. These are both aesthetically appealing and important tools for data analysis. We are therefore working on new optical systems for cameras and turn the classical camera into a calculation instrument, which can extract much more information from individual images than pure light intensity, such as 3D geometry, for example. ...

CONTRIBUTIONS

HDR – Images and Video with Enhanced Contrast	84
Advanced Real-Time Rendering	85
Computational Photography	86
Statistical Models of People	87
Correspondences, Symmetries and Automatic Modeling	88



UNDERSTANDING
IMAGES & VIDEOS

BIOINFORMATICS

GUARANTEES

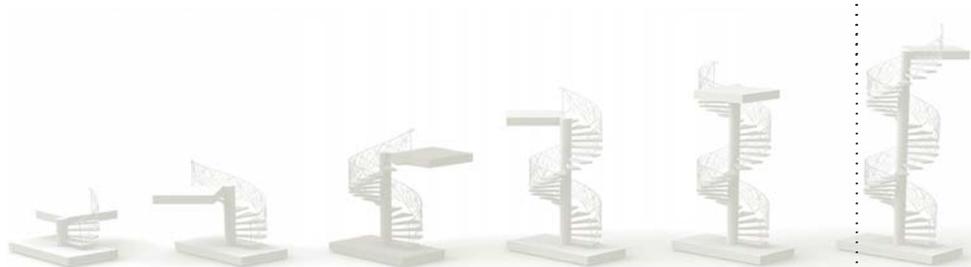
INFORMATION SEARCH
& DIGITAL KNOWLEDGE

MULTIMODAL INFORMATION

OPTIMIZATION

SOFTWARE

VISUALIZATION



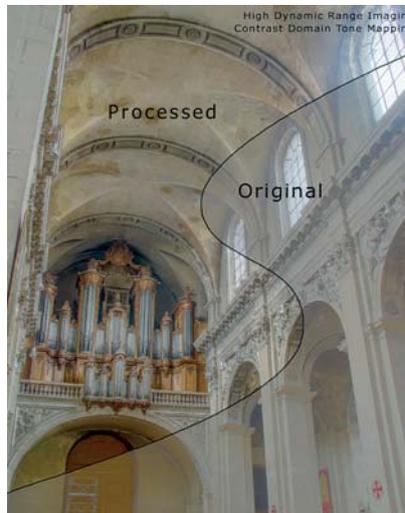
HDR – Images and Video with Enhanced Contrast

The vast majority of digital images and video material used today can capture only a small portion of colors and brightness levels visible to the human eye and does not even offer sufficient quality to reproduce them on the new generation of display devices. For instance, the popular JPEG format was designed, for efficiency reasons, to store as much information as could be displayed on the majority of displays, which were Cathode Ray Tube (CRT) monitors or TV sets at the time when the JPEG compression was developed. This assumption is no longer valid, as the new generations of LCD and Plasma displays can visualize much broader color gamut and dynamic range than their CRT ancestors.



Tone mapping reduces dynamic range in an HDR image (the right inset) by separately processing regions of different luminance levels (marked in blue and red in the left inset), which models the behavior of the human visual system as predicted by the anchoring theory of lightness perception.

The High Dynamic Range Imaging (HDRI) overcomes the limitation of traditional imaging by performing operations on color with much higher precision, which provides accuracy that exceeds the capabilities of human perception.



Rendering HDR images for the devices of a limited dynamic range

HDR images can faithfully represent real scenes with all perceivable colors and with no under- or over-exposed regions. HDRI does not only provide higher precision, but also enables to synthesize, store, and visualize a range of perceptual cues, which are not achievable with traditional imaging. Unlike traditional images, HDR imagery can represent such common visual phenomena as self-luminous surfaces (sun, shining lamps), bright highlights, dark shadows as well as vivid, strongly saturated colors.

HDRI has already gained vast popularity in photography, image editing, and realistic image rendering. It delivers more realistic virtual-world experience in video games and astounding special effects in movie productions. A consumer level HDR display and HDR cameras are going to be available on the market soon.

At the Max Planck Institute for Informatics, we have designed new video and image formats that are capable of encoding real-world scenes with exceptionally high fidelity. The formats can not only store HDR imagery, but also achieve good compression performance, necessary to reduce data to manageable sizes. We have also developed a suite of software for image processing on HDR images, intended mostly for research projects but also used by photographers [<http://www.mpi-sb.mpg.de/resources/pfstools/>]. In our work, we try to realize the concept of an imaging framework that would not be restricted by any particular imaging technology and, if storage efficiency is required, be limited only by the capabilities of the human visual perception. However, redesigning existing imaging software and hardware to work with HDR content requires a lot of effort and definition of new imaging standards. Our mission is to popularize the concept of HDR imaging, develop standard tools and algorithms for processing HDR content and research the aspects of human perception that greatly influence digital imaging.

Recently, our ideas have fallen on fertile ground in the EU Commission; our proposal to launch the so-called COST Action (European Cooperation in the field of Scientific and Technical Research) in cooperation with many partners from the European industry and academia has been approved. ...

CONTACT

Karol Myszkowski

DEPT. 4 Computer Graphics

Phone +49 681 9325-4029

Email karol@mpi-inf.mpg.de

Internet <http://www.mpi-inf.mpg.de/resources/hdr/>



Advanced Real-Time Rendering

Rendering is one of the fundamental research areas in computer graphics, which is striving to bring photorealistic imagery to digital media. Despite the recent strides made in graphics algorithms and hardware, real-time simulation of light propagation remains challenging. This article presents our attempts to realize photorealistic effects like the ones in normal photographs.

Image-based ray tracing

A typical approach to trace rays in a scene relies on the 3D mesh representation. Hierarchical data structures often help to accelerate the intersection test. Nonetheless, the typical performance of such approaches has been inappropriate for real-time applications. Our approach based on image-space scene representation precludes invisible data and thereby relieves costly computation. Moreover, its spatio-temporal coherency is appropriate to take advantage of parallel GPUs. Our rays are defined in a 2D image domain. The complexity of depth dimension is reduced with the layered data structure.

Simulation of optical lens effects

Optical simulation of a lens system is an example of efficient application for the image-space ray tracing. Unlike the pinhole lens model common in graphics, real cameras have a finite size of aperture. The traditional approach, which was the only accurate approach before ours, yet very costly, was to distribute multiple rays sampled in a thin lens. The results rendered with the image-based ray tracing [Figure 1] are hardly distinguishable from the traditional rendering, yet the speedup of rendering performance reaches up to an order of two magnitudes.



Figure 1: The depth-of-field effects rendered with a thin lens model (the upper-left half) greatly improves photorealism, in comparison to the image rendered with a standard pinhole lens model (the lower-left half).

A more advanced simulation can be handled with a spherical lens. One example is aberration, resulting from the deviation of rays at the focusing point. Although it is considered as an artifact in reality, the presence of aberration adds dramatic realism to a static scene. Figure 2 demonstrates the curvature of a field, one of the common aberration effects.



Figure 2: Curvature of a field makes an illusion of still velocity.

Server-based rendering

Server-client platforms that enable remote 3D graphics interaction are of high importance in the age of mobility and constant hardware change. By relegating rendering to powerful servers and streaming the resulting frames, the requirements on the client side can be reduced to standard web browsing and

video playing abilities. Such server-client graphics platforms have many other advantages, e.g., potentially huge data sets remain on the server, which simplifies maintenance, improves security, and enables consistent updates in collaborative work scenarios. In our research, we focus on reducing the server load by perception-based adjustment of the rendering quality to the video streaming bandwidth. This way, image details, which would otherwise be removed due to the compression, are not rendered on the server.

In another approach, to reduce server workload, we capitalize on the often significant computation capabilities of the client, who, for the pure-video-streaming scenario, stays mostly idle. To transfer high-precision auxiliary depth data to the client, our solution uses an efficient edge encoding and a new fast diffusion process. The availability of such additional information makes it possible to employ many applications, such as spatio-temporal upsampling (superresolution). Consequently, costly per-pixel shaders are only applied to low resolution frames on the server, which are streamed to the client where the corresponding high-resolution frames are reconstructed, hereby lowering server workload and bandwidth. ...



CONTACT

Sungkil Lee
DEPT. 4 Computer Graphics
Phone +49 681 9325-4000
Email slee@mpi-inf.mpg.de



Karol Myszkowski
DEPT. 4 Computer Graphics
Phone +49 681 9325-4029
Email karol@mpi-inf.mpg.de

Computational Photography

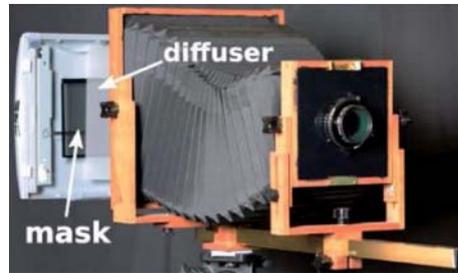
Traditional photography was replaced by its digital counterpart years ago. While the traditional process operates on a chemical basis, digital cameras convert incident light directly into numerical values. This conversion is performed by a computer unit located on the camera. The newly emerging field of computational photography deals with the question of how conventional cameras can be improved by computational methods. Our group investigates innovative algorithms in order to correct common weak points in ordinary cameras and develop new applications, such as 3D reconstruction.

A camera with high dynamic range (HDR)

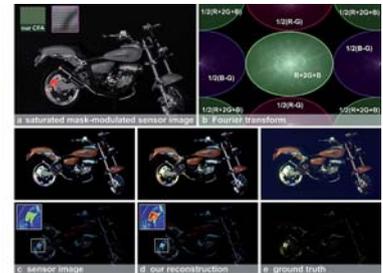
Although digital photography has mostly replaced analog recording processes outside the film industry, digital cameras are not superior to film-based cameras in every way. The most significant limitation of the digital cameras being used today, as compared to film-based approaches, pertains to the dynamic range, i.e., the contrast relationship between the brightest and darkest perceivable objects in the image. While the best digital cameras have a dynamic range of up to only 10 f-stops, film-based cameras can accurately depict up to 12-14 f-stops – the difference of one f-stop being double the amount of light. Our group has developed a camera prototype which has a 1.5 f-stop higher dynamic range than standard digital cameras. In comparison, professional lenses gain a maximum of 1 f-stop over less expensive models.

Theoretical analysis of multiplex camera system

Cameras are often used in multiplex mode, i.e., various forms of information, such as color, are digitized by suitable optical filters in different pixels of a single sensor. This process allows for less expensive production in comparison to the use of separate sensor chips for each information channel – for the colors red, green and blue, for example. These filters are known as color filter arrays (CFAs).



The camera prototype (left) uses specially designed filters, which code the dynamic range in the frequency range of the image (top right). Computational optimization methods can restore the lost dynamic range, utilizing the prior encoding (bottom right).

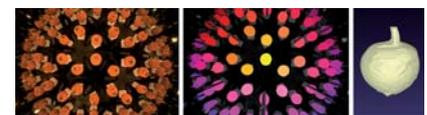


Recent analysis shows that, in addition to color information, the directional distribution of light, the so-called light field, can be recorded in different channels. The recording of light fields makes it possible, among other things, to change the focus of an image *after* the recording. Our group has, in the course of our theoretical work, established that many of the types of new cameras studied in recent years are special variations of one single overarching model. This knowledge can be used to create improved reconstruction algorithms for color and directional image reproduction. Moreover, the application of the new model enables to compare the different camera prototypes in terms of their noise patterns.

Kaleidoscopic cameras for generating hemispherical viewpoint distributions

The three-dimensional digitization of objects is another important application area for new camera technologies. Digitization, also called 3D scanning, is used more and more commonly in measurement, quality control, in the digital preservation of cultural artifacts, but also in the entertainment industry, above all in film production. A fundamental problem with 3D scanning systems

is the inability of conventional cameras to have an all-around view; an object can only be viewed from one perspective at a time. This implies that all-round digitization of *dynamic* objects can only be achieved at considerable expense: by employing multi-camera systems. In addition, there is a direct correlation between the number of cameras and the quality of the digitization. Just like in computed tomography technology, a greater number of measurement positions for an object provides more robust results and yields an increased reconstruction quality. We have developed a method in our group that uses a mirror system to generate a high number of virtual viewpoints (>200). These can be recorded *simultaneously* with a *single* camera. ...



One camera in a kaleidoscopic mirror system observes manifold views of the same object (left). The difficulty of using these images for 3D reconstruction consists in the mutual masking of reflections. With an algorithm invented by our group, it is possible to divide the image into its individual view-points (center, different colors indicate different viewpoints). An interim step in our procedure delivers a rough 3D model of the scene (right).



CONTACT

Ivo Ihrke
 DEPT. 4 Computer Graphics
 Phone +49 681 302-70756
 Email ihrke@mpi-inf.mpg.de
 Internet <http://www.mpi-inf.mpg.de/departments/d4/areas/giana/>

Statistical Models of People

Statistical face models

By using a so-called Lightstage (a sphere covered with light sources, in the middle of which is a subject's face), the properties of a face can be captured much better than in the past. Now, it is possible to precisely capture the 3D geometry of a face, in which each individual pore is visible. Further measurable properties are, for example, the diffuse and specular components of the reflected light, the shininess of the skin, or the strength of the scattering of light into the skin. This improved measurement device enables us, for the first time, to capture the influence of makeup on the properties of the skin.

We have developed a software based on such a Lightstage to recommend an appropriate makeup for a subject and present her with a convincing 3D preview. Up to now, literature only reports on approaches which enable transferring makeup from a photo of a person to the photo of another person. For our new system, we established a database with 56 women, which captures their faces with and without makeup. If a suitable makeup is needed for a woman who is not yet in the database, her skin characteristics and typical facial features are compared with women in the database so that an appropriate makeup match can be found.

Statistical body models

A detailed model of the human body is highly important in many applications for image recognition and computer graphics, such as, for example, in marker-free motion analysis, man-machine interaction, or computer animation. Previously used models, as a rule, could only show the geometry and movement of a specific person with great accuracy. The need to produce a person-specific model prevented, however, applying these image processing methods to any person. We have therefore developed a new body model of people, which is described by two low-dimensional parameter spaces. The first parameter space defines the joint angles



Figure 1: From left to right: Input image, recommended makeup, non-recommended makeup, detail enlargements

of the human skeleton, and therefore the movement of the body. The main innovation is the second space, which allows modifying the body's shape and constitution through a few intuitive parameters. Examples of such parameters are: Weight, body size, leg length, waist girth, muscularity, etc. The model was generated by applying a machine learning method to laser scans of over 120 real people of various ages and both sexes. It thereby reflects not only the range of motions a person can do, but also the wide span of variability in body shape among the population.

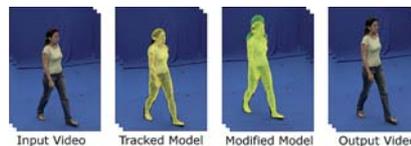


Figure 2: MovieReshape: The movement of the actress in a video can be captured with a statistical body model. Her body size can then be modified in the entire film segment.

This new model can be adapted to image video data and thus measures the movement and the constitution of a certain person in a far more accurate way. Using another newly-developed approach – MovieReshape – we can also measure human motion in 2D video streams, that are, normal films and videos, using one single camera perspective. After capturing the movement of a person, the body shape parameters can be modified throughout the entire video sequence [Figure 2]. So actors can, for example, be made taller, smaller or more muscular without needing to reshoot the scene [Figure 3].



Figure 3: MovieReshape: The muscularity of the actor from the above original sequence was increased in the lower sequence.



CONTACT

Christian Theobalt

DEPT. 4 Computer Graphics

Phone +49 681 9325-4028

Email theobalt@mpi-inf.mpg.de



Thorsten Thormählen

DEPT. 4 Computer Graphics

Phone +49 681 9325-4017

Email thormae@mpi-inf.mpg.de

Correspondences, Symmetries and Automatic Modeling

As humans, we can recognize the structures in objects intuitively. For example, when we look at the model of a parking garage, as shown in Figure 1, we can immediately tell that it consists of various floors that are themselves made up of repeated segments, together with entrance and exit ramps. An analysis of this type is considerably more difficult for the computer. Nevertheless, it would be extremely useful if structures of 3D models could be automatically recognized to a certain degree. This would, for example, enable new floors to be added to the parking garage without any additional modeling work.



Figure 1: Three examples of 3D models with a striking structure immediately recognized by the human eye. Can we recognize the structure of geometric models using the computer?

What is structure?

Two things are needed to solve such a problem: Firstly, we need a criterion for what the structure of an object is made up of. Secondly, we need an algorithm that automatically recognizes structures in line with this model. For the model, we must naturally lower our sights compared with the complex cognitive processes in the human brain; the trick is to capture important aspects with simple formal assumptions. Here, we use a model that has its origins in image analysis: An object is structurally similar to another if it looks the same on a local level. If we only ever look at a small, locally restricted part of an object, no difference between objects of the same structure should be noticeable. In particular, this means that a model made up of different building blocks is equivalent to all other models made up of the same building blocks, but perhaps put together in a different way.

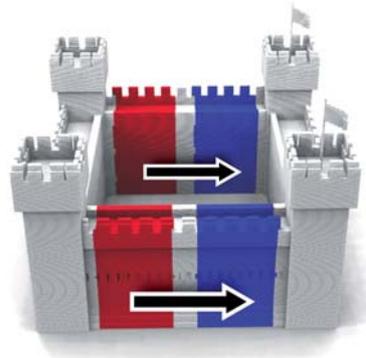


Figure 2: A partial symmetry: In the illustrated shift, the input object is partially projected onto itself (red on blue).

Correspondences and symmetries

We calculate such building blocks by using a correspondence analysis: When we find two pieces of geometry that are identical except for a rotation or a shift, we have found *correspondences* between these pieces under a rigid body transformation. We then apply this correspondence analysis to an individual object: We calculate all partial correspondences between object parts [an example is given in Figure 2]. Such correspondences within an object are referred to as *partial symmetries*. We then find our building blocks by cutting the model apart within the symmetrical areas: A cut through the red area in Figure 2 generates a symmetrical cut in the blue area. The three pieces thus formed can then be assembled in various combinations.

Automatic modeling

Our algorithm then defines all cuts of this kind and uses them to build up a kit of building blocks that can be used to create new models. These models are guaranteed to be similar to the input model from which the building blocks were calculated. Figure 3 gives some examples for how the models shown in Figure 1 (red) can be automatically altered. The newly generated models (gray) display a similar structure to the input examples. With this process, we have taught the computer to understand a little about the structure of geometric models and to use this knowledge to build similar models independently. In practice, this is an important tool for supporting 3D modeling: The user's work is facilitated by semi-automatic functions in the processing and modification of 3D models. ...

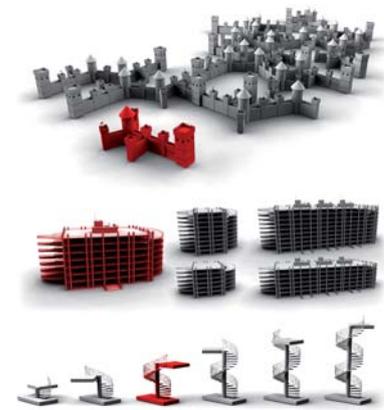


Figure 3: Results of the automatic modeling. The example objects marked in red are analyzed and similar objects are calculated automatically.



CONTACT

Martin Bokeloh

DEPT. 4 Computer Graphics

Phone +49 681 9325-4008

Email mbokeloh@mpi-inf.mpg.de



Michael Wand

DEPT. 4 Computer Graphics

Phone +49 681 9325-4008

Email mwand@mpi-inf.mpg.de

Internet <http://www.mpi-inf.mpg.de/departments/d4/areas/statgeoprocessing>

The International Max Planck Research School for Computer Science (IMPRS-CS)

The training of young scientists is fundamental to the future of science, research, and innovation in Germany. The Max Planck Society, in cooperation with German universities, has launched an initiative to promote young scientists: the International Max Planck Research Schools (IMPRS). They offer especially gifted German and foreign students the possibility to earn a doctorate within a structured program that provides excellent conditions for research. The aim is to strengthen the recruitment and training of young scientists.



IMPRS - CS

Promotion of young scientists

The IMPRS-CS is an opportunity for young scientists and scholars who are between the bachelor's or master's degree and the PhD. This includes first-class training programs, academic specialization, often with thematic linking of individual doctorates, and close collaboration between doctoral students and their academic advisors.

One focus is on international cooperation: The IMPRS-CS strives especially to attract foreign applicants to doing doctoral studies in Germany, familiarize them with the research institutions, and arouse their interest in future work at or in cooperation with German research institutions. Over 50 percent of our doctoral candidates come from abroad, with the largest contingent coming from Bulgaria, China, India and Poland.

IMPRS-CS programs

Together with Saarland University and the Saarbrücken Graduate School of Computer Science, the IMPRS-CS offers programs to establish graduate level competency and also to achieve a doctoral degree. All graduate programs are offered in close cooperation with the Max Planck Institute for Informatics, the Max Planck Institute for Software Systems and the Department of Computer Science at Saarland University. The projects are jointly supervised by the scientists of the Max Planck Institutes and their colleagues in the Department of Computer Science at Saarland University. Outstanding knowledge of English is required for all candidates.

Financial support

IMPRS-CS students receive a scholarship which covers fees, living expenses, and health insurance coverage for themselves as well as for their spouses and children who accompany them. In addition, we assist our students in finding accommodation and with administrative matters of all kinds. We offer English and German classes at several levels, joint activities, and excursions. :::

CONTACT

Jennifer Gerling
IMPRS-CS

Phone +49 681 9325-1801

Email jgerling@mpi-inf.mpg.de

Internet <http://www.imprs-cs.de>



CeBIT 2010

Participation of the Max Planck Institute for Informatics at the CeBIT fair

The CeBIT in Hanover is traditionally one of the world's most important events for the IT industry. It combines a trade fair with conferences, keynote addresses and corporate events, and provides a major forum for economic contacts. In addition, it is a well-used platform for presenting innovation in IT to the public.

In March 2010, the Max Planck Institute for Informatics was again present at the Saarland research booth, in cooperation with the Computer Science Department of Saarland University, the



Minister Hartmann (middle right) at the presentation of an MPI project

Cluster of Excellence “*Multimodal Computing and Interaction*”, and the Competence Center Computer Science Saarland (KIS).

One of the projects shown at the Saarland research stand is a technique developed by Dr. Thorsten Thormählen [<http://www.mpi-inf.mpg.de/resources/ReflectionEditing>], which enables one to divert reflections in computer-generated 3D scenes in an interactive way, giving great artistic freedom in the composition of a scene. Dr. Meinard Müller's research project is a Multimodal Music Player which searches notes, audio and text from pieces of music, and compares and synchronizes them (“*Automated Music Processing*”, page 65). Prof. Dr. Christoph Weidenbach presented “*Version Management in the Future*” as part of the “*Future Talk*” series.

On March 4, a Saarland delegation of representatives from politics and business visited the research stand together with the Minister for Economics and Science, Dr. Christoph Hartmann. ...



Saarland research stand

Images from Science

The calendar of the Max Planck Society

The Max Planck Society annually produces a year planner with aesthetic, surprising, and unusual images from current research to accompany every week of the year. In order to select the images, a competition among researchers from the 80 Max Planck Institutes is organized.

As in the previous year, two contributions were chosen from the Max Planck Institute for Informatics and can be found in the 2011 calendar. ...

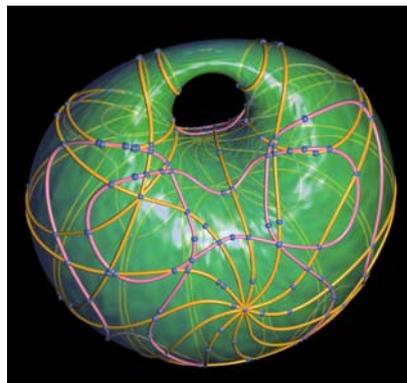


Illustration of a ring cyclide whose section curves have two further surfaces by Eric Berberich, Pavel Emeliyanenko and Michael Kerber



Combination of multiple recordings of a sunken ship motor by Martin Cadik and Glenn Lawyer

Business Location Saarland

Saarland – a top location for national and international investors

On August 18, 2010, the media project “Business Location Saarland” was presented in the auditorium of the Saarbrücken city hall. For this trilingual location marketing project that aims to highlight the benefits of the business location Saarland, the European Economic Publishing House presented a book, a film, and a website created in collaboration with companies and institutes from the region [<http://ebn24.com/index.php?id=34983>]. The Max Planck Institute for Informatics as well as the Cluster of Excellence “Multimodal Computing and Interaction” are representatives of the quality and attractiveness of this scientific location and its world-wide reputation as a center for cutting-edge research.

Prof. Dr. Hans-Peter Seidel was invited as a keynote speaker. Other prominent guests included Prime Minister



Mayor Ralf Latz, Director Prof. Hans-Peter Seidel, Prime Minister Peter Müller and Publisher Christian Kirk (from left to right)

Peter Müller, Mayor Ralf Latz and publisher Christian Kirk. In his speech, Prof. Seidel emphasized the importance of informatics for Saarland and pointed out an impending paradigm shift in information technology. The collection, preparation and processing of multimodal content will become one of the main topics of informatics in the future.

Following Prof. Seidel’s speech, the audience, consisting of more than 150 representatives from politics, business and science, as well as many representatives of the press, took part in the movie premiere of the promotional film [<http://ebn24.com/index.php?id=35888>]. At the concluding get-together, the book “Business Location Saarland”, still hot of the press, was viewed and discussed.

Produced with the participation of renowned companies and authors from the worlds of politics, business, science, art, culture and sports, and released in German, English and French, the book and film together with the free website aim to attract the interest of national and international investors by emphasizing Saarland’s strengths. ...

Intel Visual Computing Institute (Intel VCI)

A new research center on the campus of Saarland University

In May 2009, the new Intel Visual Computing Institute (Intel VCI) was opened in the presence of Saarland Prime Minister Peter Müller, University President Prof. Volker Linneweber, and Intel’s CTO Justin Rattner. Intel VCI is located on the campus of Saarland University and is carried jointly by Intel, Saarland University, the German Research Centre for Artificial Intelligence (DFKI), the Max Planck Institute for Software Systems and the Max Planck Institute for Informatics. Intel Corporation has invested \$12 Mio to create the new research center.

Visual computing is the analysis, enhancement and display of visual information to create life-like, real-time

experiences and more natural ways for people to interact with computers and other devices. Intel’s vision on visual computing is to create computer applications that look real, act real and feel real.

The Intel VCI will be led by Professors Thorsten Herfet (Chair for Telecommunications) and Philipp Slusallek (Chair for Computer Graphics) at Saarland University. The MPG representative on the Governance Board of the Intel VCI is Hans-Peter Seidel. Christian Theobalt is a member of the Steering Committee. The Intel VCI is a part of Intel Labs Europe and, according to Intel, the largest research cooperation by Intel in Europe. ...



Signing the Intel VCI founding contract: Justin Rattner (CTO Intel), Prof. Peter Druschel (MPI-SWS), Prof. Wolfgang Wahlster (DFKI), University President Prof. Volker Linneweber (University of Saarland), Prof. Hans-Peter Seidel (MPI-INF) (from left to right)

The Cluster of Excellence on Multimodal Computing and Interaction

Successful Scientific Advisory Board review

The past three decades have brought dramatic changes in the way we live and work. This phenomenon is widely characterized as the advent of the Information Society. Ten years ago, most digital content was textual. Today, it has expanded to include audio, video, and graphical data. The challenge is now to interface, organize, understand, and search this multimodal information in a robust, efficient and intelligent way, and to create dependable systems that allow natural and intuitive multimodal interaction. The term multimodal describes the different kinds of information such as text, speech, images, video, and graphics, and the way it is perceived and communicated, particularly through vision, hearing, and human expression.

The Cluster of Excellence on Multimodal Computing and Interaction, established by the German Research Foundation (DFG) within the framework of the German Excellence Initiative, addresses this challenge. More than 20 Junior Research Groups with almost 100 researchers in total have been newly established within the Cluster. Four institute directors are among the 13 PIs. Scientific coordinator of the Cluster of Excellence is Prof. Hans-Peter Seidel.

The Cluster of Excellence is advised by an international Scientific Advisory Board that has been appointed by the President of the Saarland University for a period of five years. The Advisory Board meets every two years for an extended on-site visit to get an overview of the ongoing research and to discuss and evaluate progress and possible research directions.



Participants in the Scientific Advisory Board Meeting

Following the kick-off workshop in November 2008, the second regular Advisory Board review of the Cluster of Excellence was held in November 2010. The Chairman of the Committee, Prof. Thomas Gross (ETH Zürich), particularly highlighted the interdisciplinary nature of the research effort: *“The scientists work together very closely within the framework of this cluster. This allows them to achieve much more than isolated individual projects.”*

In its written report, the Advisory Board stated: *“The primary purpose of the Cluster is to perform excellent research in multiple media. The Advisory Board finds impressive progress towards this purpose. The diversity and depth of research being performed is commendable. Especially noteworthy is the integration across groups and across areas, and the growing and increasingly deep multidisciplinary of both the work and the people in the Cluster. This is working well and stands as*

a model to other centers [...]

An impressive achievement of the Cluster is the establishment of the model of Junior Research Group Leaders for the development of the leading researchers of tomorrow. The Cluster has been successful in attracting young researchers of the highest quality while simultaneously maintaining a diversity of topics in the research groups [...] The success of this model in developing human potential is evident in many of the researchers in the Cluster already successfully finding academic positions.” ...

Promotion of Talent

Attracting more pupils to informatics

In 2010, the Max Planck Institute for Informatics, as a partner of the Competence Center Computer Science Saarland (KIS), further extended its activities for promoting young talent. The competence center unites the skills of about 460 scientists in informatics and related faculties of Saarland University and Saarland University of Applied Sciences (HTW) with two Max Planck research institutes and the German research center for artificial intelligence. One of the goals of this unique group is to create an environment where highly qualified young scientists can further develop their skills. Workshops for university freshmen and high-school students, information congresses and school cooperations are commonly used instruments to recognize and support young talents as early as possible. A special focus here is the promotion of girls, who, as experience in Germany has shown, despite sufficient qualifications rarely choose a career in technical subjects such as informatics.



Booth of the Competence Center Computer Science Saarland

The collaboration between the KIS and the Max Planck Institute for Informatics' public relations department is particularly close in this respect. The center's participation at the "First Girl's Technology Congress" in the Pirmasens Dynamikum museum, at the information fair "Graduation – what's next?" in the Saarbrücken Congress Hall, particularly dedicated to local high-school students,

as well as at the "Apprentice and Student Days" in Frankfurt, were organized by KIS and supported by staff, demonstrations and informational material from the institute.

Even closer is the cooperation at workshops and lectures, particularly in the framework of long-term programs, since they can benefit greatly from the excellent technical and staff infrastructure of the Max Planck Institute. The following events were supported by the institute:

April 22, 2010: Girls' Day 2010

[<http://www.girls-day.de>]

At Girls' Day, female high-school students meet professionals from areas in engineering, science, as well as women in leading positions in business and politics. Girls get an insight into these often unknown areas and can make first contact with professionals. Workshops, focusing on the girls' own activities complete the program of this day.

June 18, 2010: MINToring 2010

[<http://www.sdw.org/schuelerakademie/mintoring/>]

MINToring = MINT (mathematics, informatics, science and engineering) + mentoring. The training program "MINToring" is a joint initiative of the Foundation of German Business (sdw), the Federal Ministry of Education and Research (BMBF) as well as regional sponsors.

August 4, 2010: UniCamp 2010

[<http://www.uni-saarland.de/verwalt/beauftr/frauen/unicamp2010>]

In addition to the Germany wide organized Girls' Day, the UniCamp has a similar purpose with a more regional focus. It is organized by Saarland University and the Ministry of Education, Culture and Science of Saarland. The goal is to encourage female high-school students age 13–15 to discover their science/technical talents.



A variety of activities was offered for these events: In one workshop, for example, the procedure, the opportunities and the limits of laser scanning were explained and then tested in practice. In another workshop, pupils could learn to build and program robots and test their new knowledge using an operational robot. In addition, the basic concepts of informatics were communicated through exercises from the "CS Unplugged Activities" program.

In 2010, the Max Planck Institute for Informatics also participated in compiling and organizing new concepts for teaching informatics to pupils and high-school students: In May, the institute was visited by the youth project Klick-Clic [<http://www.klick-clic.eu>], where children were presented with informatics problems in the form of games. In September and November, two events, lectures, experiments, and workshops have taken place within a seminar cooperation with a local high-school. This cooperation has been planned for an entire school year and aims to use the concept of the seminar subject to bring participating high-school students closer to scientific work and the field of informatics. ...

Indo-German Max Planck Center for Computer Science (IMPECS)

Fundamental research in collaboration between Indian and German researchers

The “Indo-German Max Planck Center for Computer Science” (IMPECS) was opened by the President of Germany Dr. Horst Köhler on February 3, 2010 at IIT Delhi. The goal of the centre is to support first-class fundamental research in informatics in close cooperation between Indian and German scientists. This cooperation includes joint research in up to 10 projects, the exchange of doctoral candidates and postdocs, a Max Planck guest professor, and workshops and schoolings.

IMPECS will be financed for at least the first five years by the Max Planck Society, the Federal Ministry of Education and Research (BMBF) and the Indian Ministry for Science and Technology (DST). Kurt Mehlhorn (MPI-INF) and Rupak Majumdar (MPI-SWS) from

the Max Planck Society side, and Naveen Garg (IIT Delhi) and Manindra Agrawal (IIT Kanpur) from the Indian side, are the project leaders.

IMPECS finances up to 10 German-Indian research groups. Any university/institute can participate from the Indian side; from the German side, one partner must be from the Max Planck Institute for Informatics or the Max Planck Institute for Software Systems. Eight groups are now in operation.

There is exchange of postdocs and students on both sides. Prof. Pankaj Agarwal from Duke University spent several months at IIT Dehli as a Max Planck visiting professor. A kick-off event was held in April 2010 in Delhi in form of a joint workshop.



Federal President Horst Köhler at the opening ceremony

Furthermore, IMPECS co-financed a “School on Parameterized and Exact Computation” at IISc Chennai (December 2010), a school (October 2010) and a workshop (November 2010) on “Geometric Computing” at IIT Delhi and, together with Microsoft Research, a “School on Approximability” at IISc Bangalore (January 2011).

Cooperation with Siemens AG

Cooperation in the form of doctoral scholarships



Prof. Dr. Reinhold Achatz and Prof. Dr. Kurt Mehlhorn

Prof. Dr. Reinhold Achatz, Director of Siemens Corporate Research and Technologies and Prof. Dr. Kurt Mehlhorn, current Managing Director of the Max Planck Institute for Informatics, signed

an agreement for collaboration in the form of doctoral scholarships in Saarbrücken on May 10, 2010. One of the main factors that brought this agreement about was the commitment of Prof. Dr. Hartmut Raffer, a longtime member of the institute’s Curatorship Board. The agreement will facilitate the implementation of joint research projects by means of doctoral scholarship holders and deepen the cooperation between our institute and Siemens AG.

This cooperation can already look back on a long tradition, particularly between the Algorithm and Complexity Department headed by Kurt Mehlhorn and Christoph Moll from Siemens AG, Technology Field Modeling, Simulation and Optimization. Further areas of co-

operation beyond the topics of algorithms and optimization were defined as part of a joint workshop on June 21, 2010 in Munich. The area of information systems, represented by Gerhard Weikum from the Max Planck Institute for Informatics, cooperates with the Knowledge Management area, represented by Hermann Friedrich from Siemens AG. The areas of logic automation and constraint-based configuration, each represented by Christoph Weidenbach from the Max Planck Institute for Informatics and Herwig Schreiner from Siemens AG, will also share research efforts.

We expect the first joint doctoral candidates in 2011.

Alumni

Informatics Saarbrücken - hotbed for outstanding researchers

One of the goals of the Max Planck Institute for Informatics is to promote junior researchers who will eventually take up leading positions in academia or industry. Since our founding in 1990, 403 scientists have been doing research at our institute and then moving on to continue their professional development. These young scientists typically come to us either after completing their diploma or master's degree or following their doctorate. We have so far hosted 231 doctoral students, of which 153 moved to other research institutes after completing their PhD at Saarland University. The remaining 78 took employment in industry. Of the 172 postdoc researchers, 131 went on to other research institutes after spending two to five years with us, and 41 went to industry. Since the time the institute was founded, a total of 84 of its researchers have gone on to professorial positions.

An excellent example for the fact that the institute, and thus also the overall Saarbrücken location for informatics,

has created outstanding researchers for computer science talent is Prof. Dr. Susanne Albers. Susanne Albers works on the design and analysis of algorithms, in particular online- and approximation algorithms. She is also interested in algorithmic game theory and algorithm engineering. After completing her studies in mathematics, informatics and business administration in her home town of Osna-brück, Susanne Albers earned her doctorate in the first DFG graduate college for informatics in 1993 under the supervision of Kurt Mehlhorn at Saarland University. She then worked at the Max Planck Institute for Informatics until 1999 and made several research stays to the US, Japan and around Europe.

At the age of just 33, Susanne Albers qualified as a professor in 1999 and immediately received three offers for professorships; she ultimately decided on the one in Dortmund. In 2001, she moved to take the professorship for algorithms and complexity at the University of Freiburg. Since 2009, she has been a

professor at the Humboldt University in Berlin. Susanne Albers has received numerous awards: In 1993, she was awarded with the Max Planck Society's Otto Hahn medal for her outstanding doctorate; in 2008, she received the Gottfried Wilhelm Leibniz Prize of the German Research Foundations (DFG) for her contributions on efficient algorithms. With its 2.5 million Euros, the Leibniz prize is Germany's most prestigious scientific award. In 2010, Susanne Albers was accepted into the national academy of sciences, the Leopoldina. ...



Lise Meitner Award for Outstanding Female Computer Scientists

Award winner Dr. Anke van Zuylen

Female computer scientists are rare. Less than 20 percent of the freshmen in informatics are women. At higher career levels, the percentage of women is even lower. Therefore, the promotion of women in informatics is important. Saarbrücken has already achieved that the percentage of women over the different qualification levels (bachelor, master, doctorate, professorship) is constant. However, constant at a much too low level.

In 2010, the institute introduced the Lise Meitner Award for the promotion of outstanding female postdoctoral students. The prize consists of a two-year

scholarship and a small research budget. The first winner is Dr. Anke van Zuylen. Anke van Zuylen received her doctorate in 2008 at Cornell University in the area of Operations Research. Afterwards, she was a postdoc with Andy Yao at the Institute for Theoretical Computer Science in Beijing, China. In September 2010, she joined the Max Planck Institute for Informatics.

She primarily works on the design and analysis of algorithms for NP-hard optimization problems, especially for network design. The aim is to design networks that fulfill specific requirements with respect to robustness, bandwidth,



Dr. Anke van Zuylen

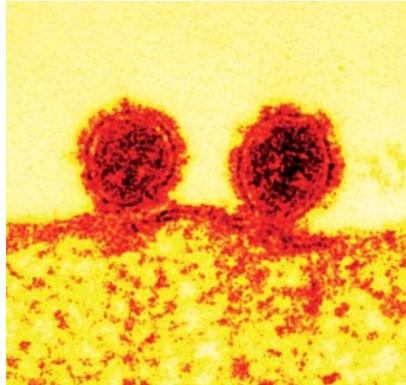
costs, and other factors. Anke van Zuylen has developed approximation algorithms for a wide variety of network design problems. ...

GISAID – Global Initiative on Sharing All Influenza Data

A database and communication platform for influenza research

Infectious diseases continue to be a great threat to mankind. Pathogens change constantly thereby finding new ways of evading the human immune system as well as drugs applied for therapy. So, the threat posed by pathogens continuously renews itself. Preventive care against infections, e.g., via vaccination, as well as the development of effective drugs, is based on research holding pace with the continued evolution of the pathogens. The basis for this research is the continued measurement of the genomic changes of the pathogen and the collection of sufficient volumes of such data in databases, which are as comprehensive as possible. For about ten years, this process has been taking place regarding HIV (see “*Analysis of HIV Drug Resistance*”, page 77, and “*Prediction of HIV Coreceptor Usage*”, page 41). For Hepatitis B (see “*Analysis of HBV Resistance*”, page 76) and Hepatitis C (see “*Fighting the Hepatitis C Virus*”, page 39) such databases are currently being initiated.

A special problem has occurred in recent years regarding influenza. Here many viral variants with a high risk potential are emerging, not in humans and in the Western world, but in animals such as birds and pigs in the Third World (Asia and Africa). In recent years, countries in these geographic regions have mounted their hesitation to provide data on the relevant viral genome sequences, because research will then be performed on the basis of these data in the Western world, whereas the providers of the data will not benefit appropriately from the



Influenza virus (Image: Volker Brinkmann, Max Planck Institute for Infection Biology, Berlin)

results. This bottleneck in the availability of influenza data became particularly evident with the emergence of the H5N1 bird flu virus in 2005 in Southeast Asia.

In response to this situation, the international initiative GISAID (Global Initiative on Sharing All Influenza Data) was formed in 2006. It joins scientists from all regions of the world and from biological as well as methodical disciplines. On the internet page, www.gisaid.org, the initiative freely offers the EpiFlu™ database, which contains the largest collection of influenza virus sequences in the world, including important sequences from the developing and emerging countries. These data could be made accessible through a new type of Data Access Agreement, the signing of which obliges scientists to fair and sustainable use of the data. Thus, the providers of the data have secured their share in the added value of the results of research based on their data. The Max Planck Institute for

Informatics develops the software for the GISAID portal, and especially for the EpiFlu™ database. Since the middle of 2008, this database has also been used by all WHO collaborating laboratories in order to perform the semiannual selection of viral strains for the development of the seasonal vaccine against influenza.

The German Federal Ministry of Nutrition, Agriculture and Consumer Protection has become the official host of the GISAID portal as of the beginning of the year 2011, based on a cooperation agreement with the GISAID Foundation. GISAID promises to become an essential resource for international influenza research. ...

Promoting Young Computer Science Talent in Germany: BWINF

Supporting the German computer science high-school competition

Since April 2010, the Max Planck Institute for Informatics, in conjunction with the Fraunhofer Group for IUK Technology and the German Computer Science Association GI, has been the fourth sponsor of the BWINF initiative. The program is given basic funding by the German Federal Ministry for Education and Research and includes the Informatics Beaver [<http://www.informatik-biber.de/>], the Federal Contest for Informatics [<http://www.bundeswettbewerb-informatik.de/>], the training for the International Olympiad in Informatics [<http://www.informatik-olympiade.de/>], and the Informatics Entry Portal [<http://www.einstieg-informatik.de/>]. The idea is to attract young people to informatics right after elementary school. The first encounter is then the Informatics Beaver Competition. In the Federal Informatics Competition, further talents are identified and promoted, and the best of them are accepted and trained on the Olympiad team for Informatics. These activities are accompanied by a portal that offers in-depth support.

The Max Planck Institute for Informatics is particularly committed, together with the computer science department of Saarland University, to the



Federal Competition for Informatics and the training of the Olympiad team. Every year, we invite the successful participants from the second competition round of the federal contest to a research day in Saarbrücken. The motto is “*teach and challenge*”. The program consists of work in small groups on current informatics questions, lectures, discussions and a framework program where participants can get to know one another.

There is a development program for the best students from each year that provides the first step towards participation in European competitions; former

participants assist in the workshops that prepare the students for the federal competition and the Olympiad. In a second step, four participants from this group are chosen for the Olympiad team. These students are then given final training. In 2010, this training was carried out by the Max Planck Institute for Informatics and the computer science department, by professors Markus Bläser, Sebastian Hack, and Christoph Weidenbach in Saarbrücken. And with great success! Two of the four students won gold medals at the 2010 Olympiad in Waterloo, Canada. ...

Open House

„Informatics Saarland“

At the open house of Saarland University, the Max Planck Institute for Informatics, along with the Max Planck Institute for Software Systems, the German Research Center for Artificial Intelligence (DFKI), the computer science department of Saarland University, and the student council presented their research projects and results to interested visitors. Scientists and students gave lectures and demonstrations of current scientific achievements from running projects, providing insight into the current state of informatics research. ...



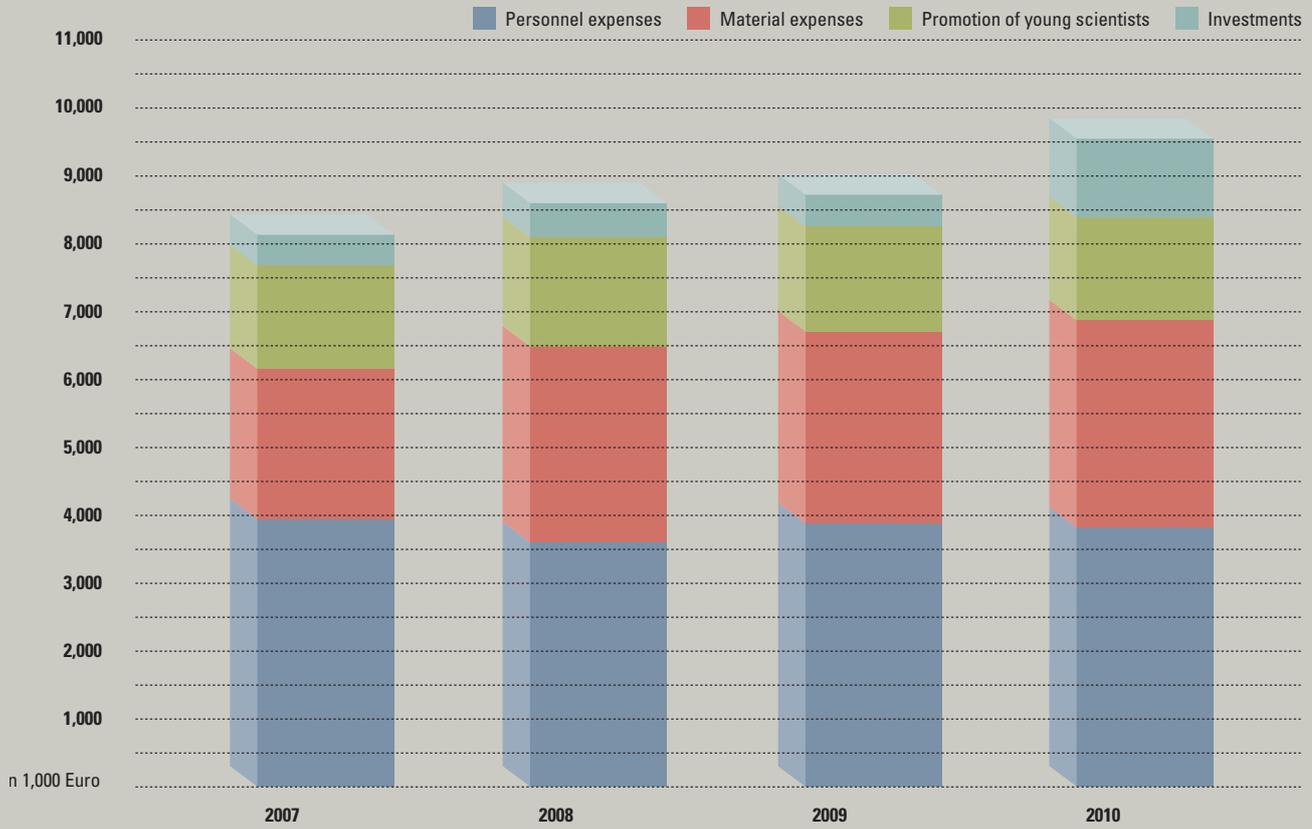
STEP 2010

Introduction phase for students

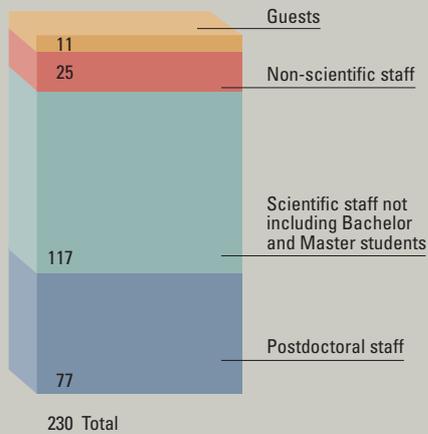
One week before the beginning of the lecture period of the computer science winter term at Saarland University, the student council of the computer science department organized a student introduction phase (called “STEP”) [<http://cs.fs.uni-saarland.de/step>]. In three days, about 130 freshmen from the departments of computer science, bioinformatics, and computing and digital media were given all the information they need for a successful start at the university. The Max Planck Institute for Informatics supported the event with informative literature and talks. ...

The Institute in Figures

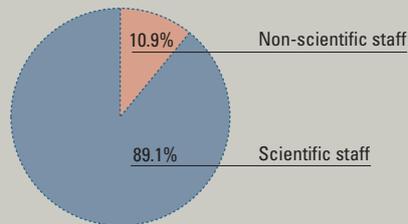
Budget without third-party funds from 2007 to 2010



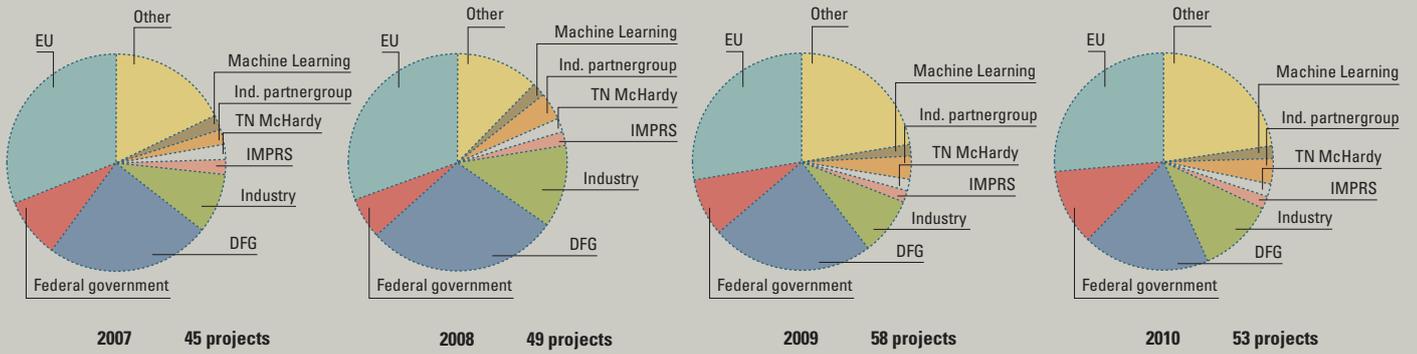
Staff Current as of 1/1/2011



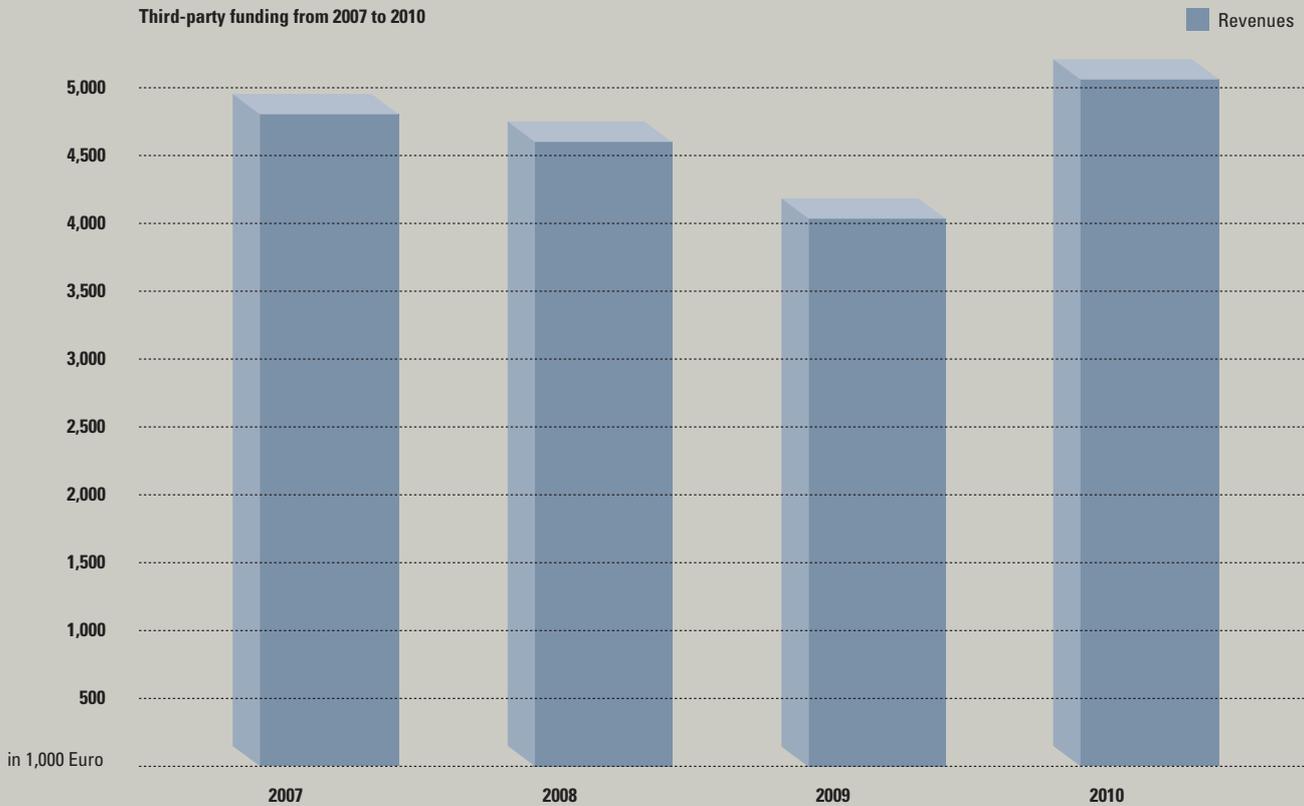
Ratio of scientific to non-scientific staff members Current as of 1/1/2011



Third-party funding from 2007 to 2010 Numbers and distribution



Third-party funding from 2007 to 2010



CONTACT

Volker Maria Geiß
Joint Administration
 Phone +49 681 9325-5700
 Email geiss@mpi-inf.mpg.de

Information Services and Technology

Unhindered global cooperation and communication in a motivating environment form the basis for an institute with the aim of producing first-class research. The flexibility, quality and reliability of the equipment as well as its usability make a decisive contribution to this research.



This aim can be transferred to our IT infrastructure: We operate a versatile system that can adapt to rapidly evolving requirements and provide the user with consistency and reliability. It does not neglect security despite the openness necessary to support international cooperation.

Multiplicity of tools

We deploy systems from various manufacturers, depending on the projects' requirements. Linux and Windows dominate in both notebook and workstation use. MacOS is gaining favor. Most servers run on Linux. If the applications or services require it, however, Windows or Solaris are used. We use Solaris in file servers due to its stability and other performance characteristics, for example.

Dynamics and innovation

Research on the front line also means constantly using innovative technologies. Sometimes, we even have to deploy and support prototypes.

Extensively uniform user experience

The research projects are often cross-platform, as they are either intended for a heterogeneous composite, or the benefits of various computer and operating system architectures must be used.

The typical workstations are completely exchangeable, because the user data and most important software packages are available independently of equipment and platforms. This homogeneity makes it easier to work with the entire system and promotes its acceptance.

Reliability of the installations

Regular updates and upgrades with software and hardware for all supported platforms pose high demands on the reliability of the installation and administration.

Cooperation and communication

Our network is divided into various areas according to organizational and security-relevant points. The end-user equipment in network sections in individual areas is powered via its floor switch with Gigabit Ethernet. This equipment is connected in a fail-safe manner with 10 GB Ethernet at the central station (two redundant backbone switches). From there, the central servers are supplied with this bandwidth as their services require it. Server farms and computer clusters are also connected over redundant switches in a fail-safe manner.

The external area includes the necessary 1GB and 10GB connections to different facilities on the campus of Saarland University in Saarbrücken and on the university campus in Kaiserslautern. Internet connectivity is realized with a 2.6 GB XWIN DFN association connection that is used jointly with the university. Our external connections are also delivered in pairs, as a rule, in order to implement the required fail-safe reliability.

Externally accessible, but anonymous services (DNS (Internet address book), WWW, FTP (data transfer), SMTP (e-mail) etc.) are combined at the institute's firewall in several demilitarized

zones (DMZs), which are differentiated according to their meaning and their hazard potential.

The network is structured in a way that promotes the integration of guest scientists and students; they can connect notebooks they have brought with them by wire or via WiFi, all without additional software installations. The network infrastructure automatically treats them as external machines.

International cooperation demands external access possibilities to internal resources in the infrastructure (Intranet). Here, we offer, among other things, secure login and access to e-mail and other important databases and services. Cooperation in software development is supported by protected access to a revision control database (software repository: Subversion).



Security and protection

Fixed mechanisms which protect against sabotage and espionage are not possible in open systems. They limit usability to too great a degree. The security guidelines can therefore be nothing more than a compromise that flexibly follows the requirements.

Some direct hazards can be deflected through the structure of the network, the firewall, the encryption for external access, or virus scanners in the mail server. Indirect hazards, such as the connection to virus-infected computers in the Intranet or faults in externally acting software systems, must be fought against through current software statuses and continuously updated local virus scanners.

Automation as guarantee

Continually developing the infrastructure and adjusting to the newest demands of the research projects and security concept, while retaining homogeneity and reliability, are our central tasks.

The mere variety of hardware and software components forbids an ad-hoc approach when supporting each device. Instead, we use extensive, advanced, automated processes that are designed to work with individual needs. Only this can allow us to resolve the necessary work within a reasonable amount of time with the required reliability and precision for the systems involved.

Our environment is too dynamic and heterogeneous for simple installations based on disk images.

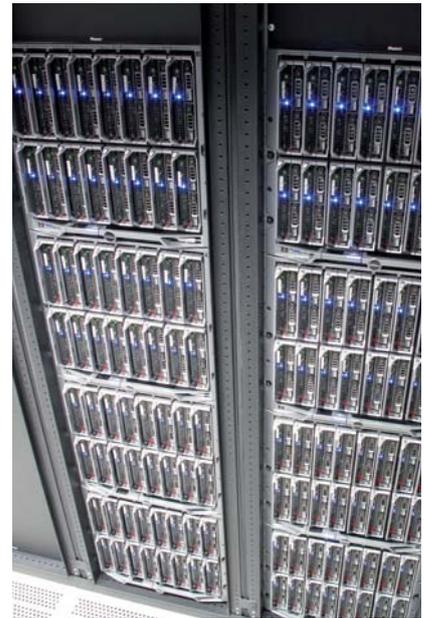
We therefore favor operating systems with package-oriented installation mechanisms. Solaris, but above all Debian in the Linux environment, is by default based on package systems that also consider dependencies between the installed packages. For Windows platforms, we have introduced such a package system with the assistance of additional software (netinstall).

Following the example of Solaris, we have implemented an installation and administration system for the current operating system versions of Debian and Windows. Results that have been achieved can be repeated as required and used very quickly across the whole infrastructure. This implementation work costs effort and, in some cases, slows down reaction times. The advantages for operating security, total expenses and time until larger changes are ready, however, clearly outweigh these disadvantages.

The system is flexibly structured, so it can be customized to suit changing requirements (such as hardware requirements) quickly through specific extensions of the installation system. However, it is only used for repetitive tasks. The integration of short-lived special installations is not practical due to the high expense associated with its automation.

Compute Service

We operate several larger systems that work with up to 64 closely coupled processors and have up to one terabyte of main memory. These machines are used for scientific research and applications which require high parallelism and uniform access to a large main memory.



Our largest cluster to date, with 128 systems with 8 cores and 48 GB main memory apiece, was brought into operation in mid-2010. It is, like its predecessor, operated under the Sun Grid Engine (SGE). The automatic distribution of processes on the cluster's individual computers allows us to reach high utilization on the whole system. The homogeneity of the operating system installations allows us to flexibly integrate the above-mentioned larger systems into this system. Jobs are categorized according to requirements before their release for the cluster, so that the SGE software can distribute them optimally on the available hardware. The prioritization of specific processes helps with the appropriate and fair distribution of resources across all the waiting jobs.

File service and data security

The data of the institute is made available over the current 20 file servers via NFS and CIFS (SMB). The now roughly 147 TB of central data is distributed across 52 RAID systems with 2843 disks, which are connected to a redundant Storage Area Network (SAN). All RAID systems are operated in a paired mirroring architecture, which is also protected against the failure of an entire RAID system. This mirroring by the file system software utilizes test sums to protect against creeping data changes (ZFS). Any detected data errors on a mirrored half are automatically repaired during running operation, whereby the false data is overwritten with the correct data. The two mirrored halves are accommodated in at least two different fire compartments, so that a fire would have little chance of causing data loss, even at the deepest level. Three file servers apiece share a view of the disk status and can substitute for one another within a very short period of time using virtual network addresses and SAN technology. This also allows for server updates without any perceivable interruptions to service.

Our data security is based on two different systems: a conventional tape backup, which secures the data directly using a Tivoli Storage Manager TSM on a tape robot (StorageTek), and an online disk backup system that minimizes space requirements by using data comparison and always keeps the data online (open

source system BackupPC). The numbers for this online data security are impressive. Because no data in this system is actually held with unnecessary redundancy, about 615 TB of gross data can be reduced to about 36 TB in the different backup runs. In order to combine the benefits of disk and tape technologies, we will combine this system with the tape robot in the future.

The tape robot currently has the ability to access 700 tapes with a total capacity of 700 TB (uncompressed). The robot is located in a specially-prepared room to maximize the protection of its high-value data. In the next construction phase, the system will be extended to a second robot that can then operate at a remote site to protect the data against disaster situations. This is at least currently achieved offline through the storage of copies in a special fire-safe data safe or at a second site.



Special systems

For special research tasks, especially in the area of computer graphics, diverse special systems are required. Available systems include a video editing system, several 3D scanners, multivideo recording systems and 3D projection systems. External communication and collaboration is supported through the operation of several videoconference installations.

Operational reliability

A monitoring system based on Nagios (open source) provides information on critical states of the server systems and the network, and on malfunctions of complex processes via e-mail and text messaging.

Power supply and cooling are designed in a way that server operation can be maintained, even during a power outage. As the last link in the chain, a generator with a power output of roughly one megawatt supplies the core infrastructure and the cooling systems with sufficient power to guarantee uninterrupted operation. Our new computer room, which was realized as part of the renovation and enlarging of infrastructure, is now operating. It hosts, among other things, the aforementioned redundant systems, which could thus finally be installed in two different locations.

Responsibilities

Procurement, installation, administration, operation, applications support and continuation of the described systems and techniques are the duty of IST (Information Services and Technology), the joint IT department for the Institutes for Software Systems and Informatics.

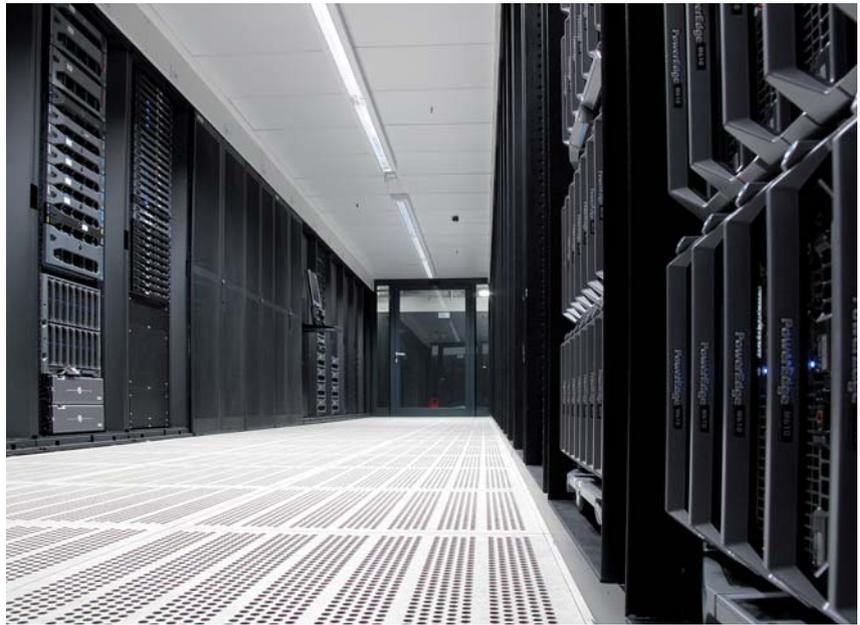
As a result of the work of the two institutes in cooperation with departments and institutes of the University, IST is also responsible for the entire campus library and the Cluster of Excellence “Multimodal Computing and Interaction (MMCI)”.

Flexible support of scientific projects

The described services, servers and computing clusters are deployed for many scientific projects in very different scenarios, sometimes even in international cooperations such as the “CADE ATP System Competition,” [<http://www.cs.miami.edu/~tptp/CASC/22/>]. In order to meet the various resulting requirements, IST offers specifically tailored support that extends from pure server hosting in individual cases to applications support. The division of tasks is in most cases dictated by the intersection of project requirements and the IST portfolio. For problems with self-administered systems, however, IST only provides consultation.

Staff structure

In addition to management and procurement (two positions), the Max Planck Institute for Informatics provides six scientific staff members and one technician. IST is supplemented by five scientific staff members from Max Planck Institute for Software Systems. There is also one employee apiece for additional administrative tasks to support the joint library and MMCI.



IST is divided into one core and several front groups. Presented in a simplified fashion, the core group is responsible for services that are identical for both institutes or are even operated together. The front groups accordingly cover the specific needs of the institutes. The members of the front groups are therefore active at multiple sites.

The front groups are supported by a 10-member team of students for their helpdesk. They form, together with two one-year student apprentices, the communications interface for users, and are reachable either by e-mail or on a web interface, or in person during business hours. In addition to processing questions concerning the use of the infrastructure, this group also maintains information systems, including inter alia, a summary of interesting questions and answers (FAQ), and a bulletin board. ...

CONTACT



Jörg Herrmann
IT Department

Phone +49 681 9325-5800

Email jh@mpi-inf.mpg.de

Internet <http://www.mpi-inf.mpg.de/services/ist/>

Selected Cooperations

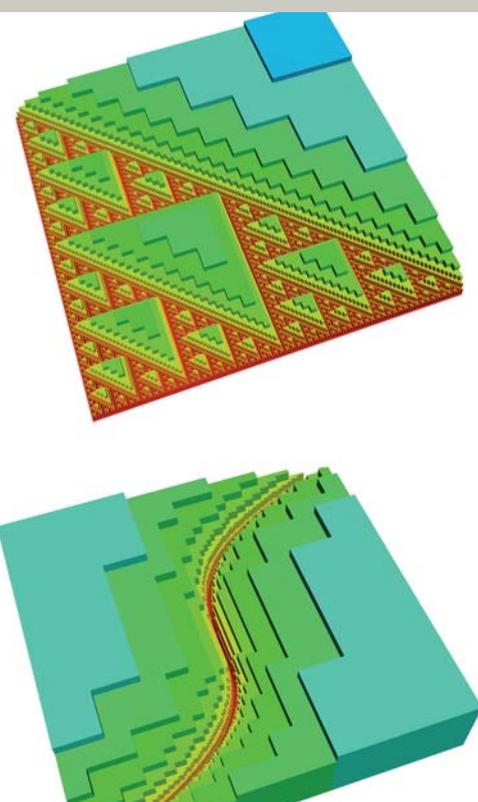
UNDERSTANDING IMAGES & VIDEOS

- ::: Bosch, *Hildesheim, Germany*
- ::: California Institute of Technology, *Pasadena, USA*
- ::: Indian Institute of Technology Delhi, *Delhi, India*
- ::: Intel Corporation, *Santa Clara, USA*
- ::: Intel Visual Computing Institute, *Saarbrücken, Germany*
- ::: Leibniz University Hanover, *Hanover, Germany*
- ::: Microsoft Research, *Cambridge, UK*
- ::: Stanford University, *Stanford, USA*
- ::: Swiss Federal Institute of Technology, *Zurich, Switzerland*
- ::: Technical University of Darmstadt, *Darmstadt, Germany*
- ::: Toyota, *Brussels, Belgium*
- ::: University College London, *London, UK*
- ::: University of Adelaide, *Adelaide, Australia*

BIOINFORMATICS

- ::: BioSolveIT GmbH, *Sankt Augustin, Germany*
- ::: Centrum Wiskunde & Informatica, *Amsterdam, Netherlands*
- ::: Christian-Albrechts-University, *Kiel, Germany*
- ::: CSIRO Livestock Industries, *St. Lucia, Queensland, Australia*
- ::: European Bioinformatics Institute, *Hinxton, UK*
- ::: Federal Office for Agriculture and Food, *Bonn, Germany*
- ::: Friedrich Loeffler Institute, *Insel Riems, Germany*
- ::: GISAID Foundation, *Washington DC, USA*

- ::: Gladstone Institute of Cardiovascular Disease, *San Francisco, USA*
- ::: Goethe University Frankfurt am Main, *Frankfurt, Germany*
- ::: Harvard University and Broad Institute, *Cambridge, USA*
- ::: IBM Haifa, *Haifa, Israel*
- ::: Informa SRL, *Rome, Italy*
- ::: Karolinska Institute, *Stockholm, Sweden*
- ::: Laboratory Dr. Thiele, *Kaiserslautern, Germany*
- ::: Max Planck Institute for Biophysical Chemistry, *Göttingen, Germany*
- ::: Max Planck Institute for Molecular Physiology, *Dortmund, Germany*
- ::: Max Planck Institute for Neurological Research, *Cologne, Germany*
- ::: Saarland University, *Saarbrücken, Germany*
- ::: Stanford University, *Stanford, USA*
- ::: Roche Diagnostics, *Penzberg, Germany*
- ::: Ruprecht-Karls-University, *Heidelberg, Germany*
- ::: Technical University Dortmund, *Dortmund, Germany*
- ::: University of Barcelona, *Barcelona, Spain*
- ::: University of British Columbia, *Vancouver, Canada*
- ::: University of California, Santa Cruz, *California, USA*
- ::: University of Düsseldorf, *Düsseldorf, Germany*
- ::: University of Freiburg, *Freiburg, Germany*
- ::: University of Heidelberg, *Heidelberg, Germany*
- ::: University of Cologne, *Cologne, Germany*
- ::: University of Nijmegen, *Nijmegen, Netherlands*



::: University of Padua, *Padua, Italy*
 ::: University of Salzburg,
Salzburg, Austria
 ::: University of Siena, *Siena, Italy*
 ::: University Tor Vergata, *Rome, Italy*
 ::: Wellcome Trust Sanger Institute,
Hinxton, UK
 ::: WHO Collaborating Centers for
 Reference & Research on Influenza,
*Melbourne, Australia, Peking, China,
 Tokyo, Japan, Atlanta, USA, London,
 UK*

GUARANTEES

::: Centrum Wiskunde & Informatica,
Amsterdam, Netherlands
 ::: Efi Arazi School of Computer
 Science and Engineering,
Herzliya, Israel
 ::: Entrepreneurial University of
 Munich, *Munich, Germany*
 ::: Indian Institute of Science
 Bangalore, *Bangalore, India*
 ::: Indian Institute of Technology Delhi,
Delhi, India
 ::: International Computer Science
 Institute, *Berkeley, USA*
 ::: Japan Advanced Institute of Science
 and Technology, *Ishikawa, Japan*
 ::: Laboratoire Lorrain de Recherche
 en Informatique et ses Applications,
Nancy, France
 ::: Max Planck Institute for Software
 Systems, *Saarbrücken, Germany*
 ::: National Information and
 Communications Technology
 Australia, *Canberra, Australia*
 ::: Politecnico di Milan, *Milan, Italy*
 ::: Purdue University,
West Lafayette, USA
 ::: Rutgers Center for Operations
 Research, *New Jersey, USA*
 ::: School of IT - The University of
 Sydney, *Sydney, Australia*

::: Tel Aviv University, *Tel Aviv, Israel*
 ::: Università degli Studi di Udine,
Udine, Italy
 ::: University of Freiburg,
Freiburg, Germany
 ::: University of Haifa, *Haifa, Israel*
 ::: University of Liverpool,
Liverpool, UK
 ::: University of Milan, *Milan, Italy*
 ::: University of Oldenburg,
Oldenburg, Germany
 ::: University of Twente,
Twente, Netherlands

INFORMATION SEARCH & DIGITAL KNOWLEDGE

::: Centrum Wiskunde & Informatica,
Amsterdam, Netherlands
 ::: Computer and Automation Research
 Institute, Hungarian Academy of
 Sciences, *Budapest, Ungarn*
 ::: Consorzio Nazionale Interuniversitario
 per le Telecomunicazioni, *Parma,
 Italy*
 ::: Documentation Research Training
 Center, Indian Statistical Institute,
Bangalore, India
 ::: Ecole Polytechnique Fédérale de
 Lausanne, *Lausanne, Switzerland*
 ::: Emory University Atlanta,
Atlanta, USA
 ::: European Archive Foundation,
Amsterdam, Netherlands
 ::: Federal Library of Niedersachsen,
Göttingen, Germany
 ::: German-Israeli Foundation for
 Scientific Research and
 Development, *Israel*
 ::: Google, *Zurich, Switzerland*
 ::: Leibniz University Hanover,
Hanover, Germany
 ::: Hanzo Archives Limited, *London,
 UK*

::: Hebrew University, *Jerusalem, Israel*
 ::: Indian Institute of Technology
 Bombay, *Bombay, India*
 ::: International Computer Science
 Institute, *Berkeley, USA*
 ::: Japan Advanced Institute of Science
 and Technology, *Ishikawa, Japan*
 ::: Jülich Supercomputing Center,
Jülich, Germany
 ::: Karlsruher Institute for Technology
 (KIT), *Karlsruhe, Germany*
 ::: Magyar Tudományok Akadémia
 Számítástechnikai Es Automatizálási
 Kutató Intézet, *Hungary*
 ::: Microsoft Research, *Cambridge, UK*
 ::: Microsoft Research, *Redmond, USA*
 ::: National Library of the Czech
 Republic, *Prague, Czech Republic*
 ::: New York University, *New York, USA*
 ::: Polish Academy of Science,
Warsaw, Poland
 ::: Rutgers Center for Operations
 Research, *New Jersey, USA*
 ::: RWTH Aachen, *Aachen, Germany*
 ::: School of IT - The University of
 Sydney, *Sydney, Australia*
 ::: Siemens AG, *Munich, Germany*
 ::: Simon Fraser University, *Burnaby,
 Canada*
 ::: SORA Institute for Social Research
 and Analysis Ogris & Hofinger
 GmbH, *Vienna, Austria*
 ::: Stanford University, *Stanford, USA*
 ::: Stichting European Archive
Amsterdam, Netherlands
 ::: Technical University of Dresden,
Dresden, Germany
 ::: Tel Aviv University, *Tel Aviv, Israel*
 ::: T-Systems Solutions for Research
 GmbH, *Germany*
 ::: Università degli Studi di Pavia,
Pavia, Italy

::: University of Patras,
Patras, Greece

::: University of Rome „La Sapienza“,
Rome, Italy

::: University of Southampton,
Southampton, UK

::: University of Trento, *Trento, Italy*

::: University of Tübingen,
Tübingen, Germany

::: University of Twente,
Twente, Netherlands

::: University of Würzburg,
Würzburg, Germany

::: Yahoo! Research,
Barcelona, Spain

MULTIMODAL INFORMATION

::: Saarland University, Excellence
Cluster Multimodal Computing and
Interaction, *Saarbrücken, Germany*

::: Saarland University of Music,
Saarbrücken, Germany

::: Technical University of Darmstadt,
Darmstadt, Germany

::: University of Bonn,
Bonn, Germany

OPTIMIZATION

::: Aberystwyth University,
Aberystwyth, UK

::: Cornell University, *Ithaca, USA*

::: Future SOC Lab of the Hasso-
Plassner-Institute, *Potsdam, Germany*

::: Indian Institute of Technology Dehli,
Dehli, India

::: Indian Institute of Technology
Kanpur, *Kanpur, India*

::: INRIA Saclay Île-de-France,
Orsay, France

::: Josip Juraj Strossmayer, University
of Osijek, *Osijek, Croatia*

::: Maastricht University,
Maastricht, Netherlands

::: MIT Computer Science and
Artificial Intelligence Lab,
Cambridge, USA

::: New York University,
New York, USA

::: PROSTEP AG, Darmstadt
Darmstadt, Germany

::: Swiss Federal Institute of Technology
Zurich, *Zurich, Switzerland*

::: Technical University of Berlin,
Berlin, Germany

::: Technical University of Denmark,
Copenhagen, Denmark

::: Technical University of Dortmund,
Dortmund, Germany

::: Technische Universität Wien,
Vienna, Austria

::: The Commonwealth Scientific and
Industrial Research Organisation,
Melbourne, Australia

::: Universidad de Chile,
Santiago, Chile

::: Università di Roma „La Sapienza“,
Rome, Italy

::: University of Adelaide,
Adelaide, Australia

::: University of Birmingham,
Birmingham, UK

::: University of Freiburg,
Freiburg, Germany

::: University of Iowa, *Iowa, USA*

::: University of Kiel, *Kiel, Germany*

::: University of Tokyo, *Tokyo, Japan*

::: Vrije Universiteit Amsterdam,
Amsterdam, Netherlands

::: Warwick Mathematics Institute,
Warwick, UK

SOFTWARE

::: Fraunhofer Institutes, *Wachtberg,
Germany*

::: Linköping University, *Linköping,
Sweden*

::: Los Alamos National Laboratory,
Los Alamos, USA

::: Technical University of Dortmund,
Dortmund, Germany

::: University of Bonn, *Bonn, Germany*

::: University of California, *Davis, USA*

::: University of Hanover,
Hanover, Germany

::: University of Siegen, *Siegen,
Germany*

VISUALIZATION

::: Indian Institute of Technology Delhi,
Delhi, India

::: International Digital Laboratory,
University of Warwick,
Warwick, UK

::: Saarland University, Excellence
Cluster Multimodal Computing and
Interaction, *Saarbrücken, Germany*

::: Stanford University, *Stanford, USA*

::: Technical University of Denmark,
Copenhagen, Denmark

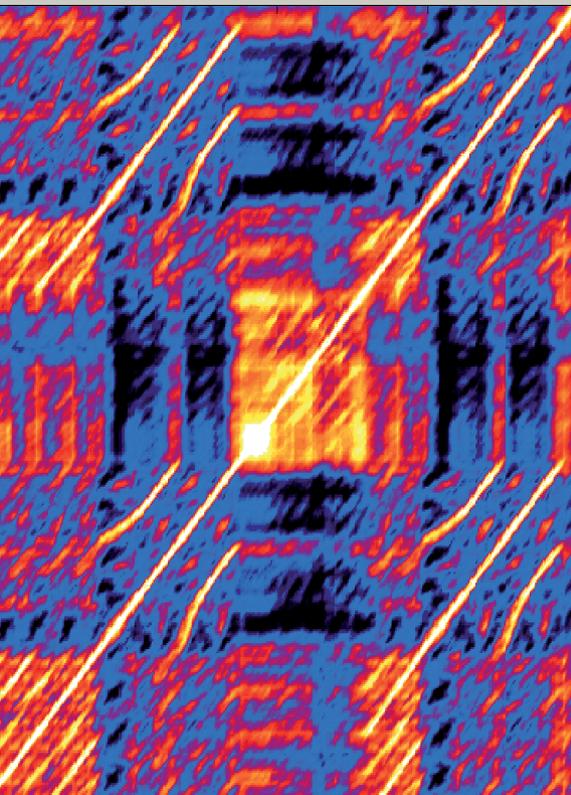
::: Telecom ParisTech, *Paris, France*

::: University College London,
London, UK

::: University of Bristol, *Bristol, UK*

::: University of British Columbia,
Vancouver, Canada

Selected Publications



- [1] E. ALTHAUS, E. KRUGLOV AND C. WEIDENBACH. Superposition modulo linear arithmetic $\text{sup}(la)$. In S. Ghilardi and R. Sebastiani, eds., *Frontiers of Combining Systems, 7th International Symposium, FroCoS 2009*, Trento, Italy, September 2009, LNAI 5749, pp. 84–99. Springer.
- [2] A. ALTMANN, M. DÄUMER, N. BEERENWINKEL, Y. PERES, E. SCHÜLTER, J. BÜCH, S.-Y. RHEE, A. SÖNNERBORG, W. J. FESSEL, R. W. SHAFER, M. ZAZI, R. KAISER AND T. LENGAUER. Predicting the response to combination anti-retroviral therapy: Retrospective validation of geno2pheno-THEO on a large clinical database. *The Journal of Infectious Diseases*, 199(7):999–1006, April 2009.
- [3] A. ALTMANN, T. SING, H. VERMEIREN, B. WINTERS, E. VAN CRAENENBROECK, K. VAN DER BORGH, S.-Y. RHEE, R. W. SHAFER, E. SCHÜLTER, R. KAISER, Y. PERES, A. SÖNNERBORG, W. J. FESSEL, F. INCARDONA, M. ZAZI, L. BACHELER, H. VAN VLIJMEN AND T. LENGAUER. Advantages of predicted phenotypes and statistical learning models in inferring virological response to antiretroviral therapy from HIV genotype. *Antiviral Therapy*, 14(2):273–283, 2009.
- [4] A. ALTMANN, L. TOLOSI, O. SANDER AND T. LENGAUER. Permutation importance: an unbiased feature importance measure. *Bioinformatics*, 26(10):1340–1347, 2010.
- [5] O. AMINI AND M. MANJUNATH. Riemann-roch for sublattices of the root lattice a_n . *Electronic Journal of Combinatorics*, 17(1):R124, 1–R124, 50, September 2010.
- [6] M. ANDRILUKA, S. ROTH AND B. SCHIELE. Monocular 3D pose estimation and tracking by detection. In *2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. – Pt. 1, San Francisco, USA, June 2010, pp. 623–630. IEEE.
- [7] R. AWADALLAH, M. RAMANATH AND G. WEIKUM. Language-model-based pro/con classification of political text. In F. Crestani, S. Marchand-Maillet, H.-H. Chen, E. N. Efthimiadis and J. Savoy, eds., *SIGIR, 2010*, pp. 747–748. ACM.
- [8] T. O. AYDIN, M. CADIK, K. MYSKOWSKI AND H.-P. SEIDEL. Visually significant edges. *ACM Transactions on Applied Perception*, 7(4):27:1–27: 15, July 2010.
- [9] N. BANSAL, A. GUPTA, J. LI, J. MESTRE, V. NAGARAJAN AND A. RUDRA. When LP is the cure for your matching woes: Improved bounds for stochastic matchings (extended abstract). In M. de Berg and U. Meyer, eds., *Algorithms – ESA 2010, 18th Annual European Symposium. – Pt. II*, Liverpool, UK, 2010, LNCS 6347, pp. 218–229. Springer.
- [10] R. BAR-YEHUDA, D. HERMELIN AND D. RAWITZ. Minimum vertex cover in rectangle graphs. In M. de Berg and U. Meyer, eds., *Algorithms – ESA 2010, 18th Annual European Symposium. – Pt. I*, Liverpool, England, 2010, LNCS 6347, pp. 255–266. Springer.
- [11] S. BASWANA, S. BISWAS, B. DOERR, T. FRIEDRICH, P. P. KURUR AND F. NEUMANN. Computing single source shortest paths using single-objective fitness functions. In T. Jansen, I. Garibay, R. Wiegand and A. S. Wu, eds., *FOGA'09, revised selected papers from ACM SIGEVO Foundations of Genetic Algorithms X*, Orlando, USA, 2009, pp. 59–66. ACM.
- [12] P. BAUMGARTNER AND U. WALDMANN. Superposition and model evolution combined. In R. A. Schmidt, ed., *Automated Deduction – CADE-22, 22nd International Conference on Automated Deduction*, Montreal, Canada, August 2009, LNAI 5663, pp. 17–34. Springer.
- [13] S. J. BEDATHUR, K. BERBERICH, J. DITTRICH, N. MAMOULIS AND G. WEIKUM. Interesting-phrase mining for ad-hoc text analytics. *PVLDB*, 3(1):1348–1357, 2010.
- [14] E. BERBERICH, E. FOGEL, D. HALPERIN, K. MEHLHORN AND R. WEIN. Arrangements on parametric surfaces I: General framework and infrastructure. *Mathematics in Computer Science*, 4(1):45–66, 2010.
- [15] K. BERBERICH, S. BEDATHUR, O. ALONSO AND G. WEIKUM. A language modeling approach for temporal information needs. In C. Gurrin and U. Kruschwitz, eds., *Proceedings of the 32nd European Conference on Information Retrieval (ECIR 2010)*, Milton Keynes, UK, 2010, pp. 13–25. Springer.
- [16] U. BHASKAR, L. FLEISCHER, D. HOY AND C.-C. HUANG. Equilibria of atomic ow games are not unique. In C. Mathieu, ed., *Proceedings of the Twentieth Annual ACM-SIAM Symposium*

- on *Discrete Algorithms (SODA 2009)*, New York, U.S.A., January 2009, pp. 748–757. ACM.
- [17] U. BLANKE AND B. SCHIELE. Remember and transfer what you have learned – recognizing composite activities based on activity spotting. In H.-J. Yoo and K. Van Laerhoven, eds., *International Symposium on Wearable Computers 2010, the 14th Annual ISWC*, Seoul, Korea, 2010, pp. 1–8. IEEE.
- [18] U. BLANKE AND B. SCHIELE. Towards human motion capturing using gyroscopeless orientation estimation. In H.-J. Yoo and K. Van Laerhoven, eds., *International Symposium on Wearable Computers 2010, the 14th Annual ISWC*, Seoul, Korea, 2010, pp. 1–2. IEEE.
- [19] H. BLANKENBURG, R. D. FINN, A. PRLIĆ, A. M. JENKINSON, F. RAMÍREZ, D. EMIG, S.-E. SCHELHORN, J. BÜCH, T. LENGAUER AND M. ALBRECHT. Dasm: exchanging, annotating and assessing molecular interaction data. *Bioinformatics*, 25(10):1321–1328, 2009.
- [20] C. BOCK. Epigenetic biomarker development. *Epigenomics*, 1(1):99–110, 2009.
- [21] C. BOCK, G. VON KUSTER, K. HALACHEV, J. TAYLOR, A. NEKRUTENKO AND T. LENGAUER. Web-based analysis of (epi-) genome data using EpiGRAPH and Galaxy. *Methods Mol Biol*, 628:275–296, 2010.
- [22] J. BOGOJESKA, S. BICKEL, A. ALTMANN AND T. LENGAUER. Dealing with sparse data in predicting outcomes of hiv combination therapies. *Bioinformatics*, 26(17):2085–2092, July 2010.
- [23] M. BOKELOH, M. WAND AND H.-P. SEIDEL. A connection between partial symmetry and inverse procedural modeling. *ACM Transactions on Graphics (Proc. Siggraph)*, 29(3), 2010.
- [24] V. BONIFACI, H.-L. CHAN, A. MARCHETTI-SPACCAMELA AND N. MEGOW. Algorithms and complexity for periodic real-time scheduling. In M. Charikar, ed., *21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2010)*, Austin (TX), USA, 2010, pp. 1350–1359. SIAM.
- [25] V. BONIFACI AND A. MARCHETTI-SPACCAMELA. Feasibility analysis of sporadic real-time multiprocessor task systems. In M. de Berg and U. Meyer, eds., *Algorithms – ESA 2010, 18th Annual European Symposium. – Pt. II*, Liverpool, United Kingdom, November 2010, LNCS 6347, pp. 230–241. Springer.
- [26] E. BOROS, K. ELBASSIONI AND K. MAKINO. Left-to-right multiplication for monotone Boolean dualization. *SIAM Journal on Computing*, 39(7):3424–3439, 2010.
- [27] K. BOZEK AND T. LENGAUER. Positive selection of hiv host factors and the evolution of lentivirus genes. *BMC Evolutionary Biology*, 10(1):186, 1–186, 15, 2010.
- [28] A. BUTMAN, D. HERMELIN, M. LEWENSTEIN AND D. RAWITZ. Optimization problems in multiple interval graphs. *ACM Transactions on Algorithms*, 6(2):40, 1–40, 18, 2010.
- [29] J. CASE AND T. KÖTZING. Strongly non- u -shaped learning results by general techniques. In A. T. Kalai and M. Mohri, eds., *COLT 2010, The 23rd Conference on Learning Theory*, Haifa, Israel, 2010, pp. 181–193. COLT.
- [30] S.-H. CHAN, T.-W. LAM AND L.-K. LEE. Non-clairvoyant speed scaling for weighted ow time. In M. de Berg and U. Meyer, eds., *Algorithms – ESA 2010, 18th Annual European Symposium. – Pt. I*, Liverpool, United Kingdom, 2010, LNCS 6346, pp. 23–35. Springer.
- [31] G. CHRISTODOULOU, C. CHUNG, K. LIGETT, E. PYRGA AND R. VAN STEE. On the price of stability for undirected network design. In E. Bampis and K. Jansen, eds., *Approximation and Online Algorithms, 7th International Workshop, WAOA 2009*, Copenhagen, Denmark, June 2010, LNCS 5893, pp. 86–97. Springer.
- [32] G. CHRISTODOULOU, E. KOUTSOPIAS AND A. VIDALI. A lower bound for scheduling mechanisms. *Algorithmica*, 85(4):729–740, 2009.
- [33] G. CHRISTODOULOU AND A. KOVÁCS. A deterministic truthful PTAS for scheduling related machines. In M. Charikar, ed., *21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2010)*, Austin, Tx., USA, 2010, pp. 1005–1016. SIAM.
- [34] Y. CUI, S. SCHUON, D. CHAN, S. THRUN AND C. THEOBALT. 3D shape scanning with a time-of-flight camera. In *Proc. of IEEE Intl. Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2010, pp. 1173 – 1180.
- [35] W. DAMM, H. DIERKS, S. DISCH, W. HAGEMANN, F. PIGORSCH, C. SCHOLL, U. WALDMANN AND B. WIRTZ. Exact and fully symbolic verification of linear hybrid automata with large discrete state spaces. *Science of Computer Programming*, to appear, 2011.
- [36] P. DIDYK, E. EISEMANN, T. RITSCHEL, K. MYSZKOWSKI AND H.-P. SEIDEL. Apparent display resolution enhancement for moving images. *ACM Transactions on Graphics*, 29(4):113:1–113:8, July 2010.
- [37] P. DIDYK, E. EISEMANN, T. RITSCHEL, K. MYSZKOWSKI AND H.-P. SEIDEL. Perceptually-motivated real-time temporal upsampling of 3D content for high-refresh-rate displays. *Computer Graphics Forum (Proceedings Eurographics 2010, Norrköpping, Sweden)*, 29(2):713–722, May 2010.
- [38] B. DOERR, T. FRIEDRICH AND T. SAUERWALD. Quasirandom rumor spreading: Expanders, push vs. pull, and robustness. In S. Albers, A. Marchetti-Spaccamela, Y. Matias, S. Nikolettseas and W. Thomas, eds., *Automata, Languages and Programming, 36th International Colloquium, ICALP 2009*, Rhodes, Greece, 2009, LNCS 5555, pp. 366–377. Springer.
- [39] B. DOERR AND L. A. GOLDBERG. Drift analysis with tail bounds. In R. Schaefer, C. Cotta, J. Kolodziej and G. Rudolph, eds., *Parallel Problem Solving from Nature – PPSN X: 11th International Conference*, Krakow, Poland, 2010, LNCS 6238, pp. 174–183. Springer. Best Paper Award.
- [40] B. DOERR, D. JOHANNSEN, T. KÖTZING, F. NEUMANN AND M. THEILE. More effective crossover operators for the all-pairs shortest path problem. In R. Schaefer, C. Cotta, J. Kolodziej and G. Rudolph, eds., *Parallel Problem Solving from Nature – PPSN XI, 11th International Conference*, Krakow, Poland, 2010, LNCS 6238, pp. 184–193. Springer.
- [41] B. DOERR, D. JOHANNSEN AND C. WINZEN. Multiplicative drift analysis. In M. Pelikan and J. Branke, eds., *Proceedings of 12th Annual Conference on Genetic and Evolutionary Computation (GECCO-2010)*, Portland, USA, 2010, pp. 1449–1456. ACM. Best Paper Award.

- [42] D. DUMITRIU, S. FUNKE, M. KUTZ AND N. MILOSAVLJEVIC. How much geometry it takes to reconstruct a 2-manifold in \mathbb{R}^3 . *ACM Journal of Experimental Algorithms*, 14:2.2:1–2.2:17, May 2009.
- [43] S. EBERT, D. LARLUS AND B. SCHIELE. Extracting structures in image collections for object recognition. In K. Daniilidis, P. Maragos and N. Paragios, eds., *Computer Vision – ECCV 2010, 11th European Conference on Computer Vision*, Crete, Greece, 2010, LNCS 6311, pp. 720–733. Springer.
- [44] K. ELBASSIONI AND H. CHAN. A QPTAS for TSP with fat weakly disjoint neighborhoods in doubling metrics. In *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2010)*, Austin, USA, January 2010, pp. 256–267. SIAM.
- [45] K. ELBASSIONI, R. RAMAN, S. RAY AND R. SITTERS. On the approximability of the maximum feasible subsystem problem with 0/1-coefficients. In C. Mathieu, ed., *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2009)*, New York, NY, USA, 2009, pp. 1210–1219. ACM.
- [46] S. ELBASSUONI, K. HOSE, S. METZGER AND R. SCHENKEL. ROXXI: Reviving witness documents to explore extracted information. *PVLDB*, 3(2):1589–1592, 2010.
- [47] S. ELBASSUONI, M. RAMANATH, R. SCHENKEL AND G. WEIKUM. Searching RDF graphs with SPARQL and keywords. *IEEE Data Eng. Bull.*, 33(1):16–24, 2010.
- [48] E. ELLINGHAUS, D. ELLINGHAUS, P. E. STUART, R. P. NAIR, S. DEBRUS, J. V. RAEALSON, M. BELOUCHI, H. FOURNIER, C. REINHARD, J. DING, Y. LI, T. TEJASVI, J. GUDJONSSON, S. W. STOLL, J. J. VOORHEES, S. LAMBERT, S. WEIDINGER, B. EBERLEIN, M. KUNZ, P. RAHMAN, D. D. GLADMAN, C. GIEGER, H. E. WICHMANN, T. H. KARLSEN, G. MAYR, M. ALBRECHT, D. KABELITZ, U. MROWIETZ, G. R. ABECASIS, J. T. ELDER, S. SCHREIBER, M. WEICHENTHAL AND A. FRANKE. Genome-wide association study identifies a psoriasis susceptibility locus at TRAF3IP2. *Nature Genetics*, 42(11):991–995, 2010.
- [49] A. K. ELMAGARMID AND D. AGRAWAL, eds. *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD 2010)*, 2010. ACM.
- [50] D. EMIG, N. SALOMONIS, J. BAUMBACH, T. LENGAUER, B. R. CONKLIN AND M. ALBRECHT. AltAnalyze and DomainGraph: analyzing and visualizing exon expression data. *Nucleic Acids Research*, 38(Suppl. 2):W755–W762, 2010.
- [51] L. EPSTEIN, A. LEVIN AND R. VAN STEE. Max-min online allocations with a reordering buffer. In S. Abramsky, C. Gavoille, C. Kirchner, F. Meyer auf der Heide and P. G. Spirakis, eds., *Automata, Languages and Programming, 37th International Colloquium, ICALP 2010*, Bordeaux, France, July 2010, LNCS 6198, pp. 336–347. Springer.
- [52] J. FABER, C. IHLEMANN, S. JACOBS AND V. SOFRONIE-STOKKERMANS. Automatic verification of parametric specifications with complex topologies. In D. Mery and S. Merz, eds., *Integrated Formal Methods, 8th International Conference, IFM 2010*, Nancy, France, 2010, LNCS 6396, pp. 152–167. Springer.
- [53] A. FARZAN, T. GAGIE AND G. NAVARRO. Entropy-bounded representation of point grids. In O. Cheong, K.-Y. Chwa and K. Park, eds., *Algorithms and Computation, 21st International Symposium, ISAAC 2010. - Pt. II*, Jeju Island, Korea, 2010, LNCS 6507, pp. 327–338. Springer.
- [54] A. FIETZKE, H. HERMANN AND C. WEIDENBACH. Superposition-based analysis of first-order probabilistic timed automata. In C. G. Fermüller and A. Voronkov, eds., *Logic for Programming, Artificial Intelligence, and Reasoning, 17th International Conference, LPAR-17*, Yogyakarta, Indonesia, October 2010, LNCS 6397, pp. 302–316. Springer.
- [55] N. FOUNTOLAKIS, A. HUBER AND K. PANAGIOTOU. Reliable broadcasting in random networks and the effect of density. In *2010 Proceedings IEEE INFOCOM*, San Diego, USA, 2010, pp. 2552–2560. IEEE.
- [56] N. FOUNTOLAKIS AND K. PANAGIOTOU. Orientability of random hypergraphs and the power of multiple choices. In S. Abramsky, C. Gavoille, C. Kirchner, F. Meyer auf der Heide and P. G. Spirakis, eds., *Automata, Languages and Programming, 37th International Colloquium, ICALP 2010. - Pt. 1*, Bordeaux, France, 2010, LNCS 6198, pp. 348–359. Springer.
- [57] A. FRANKE, T. BALSCHUN, C. SINA, D. ELLINGHAUS, R. HÄSLER, G. MAYR, M. ALBRECHT, M. WITTIG, E. BUCHERT, S. NIKOLAUS, C. GIEGER, H. E. WICHMANN, J. SVENTORAITYTE, L. KUPCINSKAS, C. M. ONNIE, M. GAZOULI, N. P. ANAGNOU, D. STRACHAN, W. L. McARDLE, C. G. MATHEW, P. RUTGEERTS, S. VERMEIRE, M. H. VATN, M. KRAWCZAK, P. ROSENSTIEL, T. H. KARLSEN AND S. SCHREIBER. Genomewide association study for ulcerative colitis identifies risk loci at 7q22 and 22q13 (IL17REL). *Nature Genetics*, 42(4):292–294, 2010.
- [58] T. FRIEDRICH, M. GAIRING AND T. SAUERWALD. Quasirandom load balancing. In M. Charikar, ed., *21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2010)*, Austin, Texas, 2010, pp. 1620–1629. SIAM.
- [59] M. FRITZ, G.-J. M. KRUIJFF AND B. SCHIELE. Tutor-based learning of visual categories using different levels of supervision. *Computer Vision and Image Understanding*, 114(5):564–573, 2010.
- [60] M. FRITZ, A. MYKHAYLO, S. FIDLER, M. STARK, A. LEONARDIS AND B. SCHIELE. Categorical perception. In H. I. Christensen, G.-J. M. Kruijff and J. L. Wyatt, eds., *Cognitive Systems, Cognitive Systems Monographs*, vol. 8, ch. 4, pp. 131–164. Springer, Berlin, 2010.
- [61] N. GARG, T. KAVITHA, A. KUMAR, K. MEHLHORN AND J. MESTRE. Assigning papers to referees. *Algorithmica*, 58(1):119 – 136, 2010.
- [62] D. GÖDDEKE AND R. STRZODKA. Mixed precision GPU-multigrid solvers with strong smoothers. In J. J. Dongarra, D. A. Bader and J. Kurzak, eds., *Scientific Computing with Multi-core and Accelerators*. CRC Press, Dec. 2010.
- [63] A. GUBICHEV, S. BEDATHUR, S. SEUFERT AND G. WEIKUM. Fast and accurate estimation of shortest paths in large networks. In *Proceedings of the 19th ACM International Conference on Information and Knowledge Management (CIKM 2010)*, Toronto, Canada, October 2010, pp. 499–508. ACM.

- [64] N. HASLER, T. THORMÄHLEN AND H.-P. SEIDEL. Multilinear pose and body shape estimation of dressed subjects from image sets. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2010)*, 2010, pp. 1823–1830. IEEE.
- [65] M. HORBACH AND C. WEIDENBACH. Superposition for fixed domains. *ACM Transactions on Computational Logic*, 11(4):27,1–27,35, November 2010.
- [66] M. B. HULLIN, J. HANIKA, B. AJDIN, H.-P. SEIDEL, J. KAUTZ AND H. P. A. LENSCH. Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions. *ACM Trans. Graph. (Proc. SIGGRAPH 2010)*, 29(4):97:1–97:7, 2010.
- [67] C. IHLEMANN AND V. SOFRONIE-STOKKERMANS. On hierarchical reasoning in combinations of theories. In J. Giesl and R. Hähnle, eds., *Automated Reasoning, 5th International Joint Conference, IJCAR 2010*, Edinburgh, 2010, LNAI 6173, p. 30–45. Springer.
- [68] I. IHRKE, K. N. KUTULAKOS, H. P. A. LENSCH, M. MAGNOR AND W. HEIDRICH. Transparent and specular object reconstruction. *Computer Graphics Forum*, 29(8):2400–2426, 2010.
- [69] I. IHRKE, G. WETZSTEIN AND W. HEIDRICH. A theory of plenoptic multiplexing. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010, pp. 1–8.
- [70] S. JACOBS. Incremental instance generation in local reasoning. In A. Bouajjani and O. Maler, eds., *Computer Aided Verification, 21st International Conference, CAV 2009*, Grenoble, France, June 2009, LNCS 5643, pp. 368–382. Springer.
- [71] A. JAIN, C. KURZ, T. THORMÄHLEN AND H.-P. SEIDEL. Exploiting global connectivity constraints for reconstruction of 3D line segment from images. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2010)*, San Francisco, USA, June 2010, pp. 1586–1593. IEEE.
- [72] A. JAIN, T. THORMÄHLEN, H.-P. SEIDEL AND C. THEOBALT. Moviereshape: Tracking and reshaping of humans in videos. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia)*, 29(6):148–157, 2010.
- [73] S. JAIN, S. SEUFERT AND S. J. BEDATHUR. Antourage: mining distance-constrained trips from Flickr. In M. Rappa, P. Jones, J. Freire and S. Chakrabarti, eds., *WWW*, 2010, pp. 1121–1122. ACM.
- [74] O. JURCA, S. MICHEL, A. HERRMANN AND K. ABERER. Continuous query evaluation over distributed sensor networks. In F. Li, M. M. Moro, S. Ghandeharizadeh, J. R. Haritsa, G. Weikum, M. J. Carey, F. Casati, E. Y. Chang, I. Manolescu, S. Mehrotra, U. Dayal and V. J. Tsotras, eds., *ICDE*, 2010, pp. 912–923. IEEE.
- [75] S. KRATSCH AND M. WAHLSTRÖM. Preprocessing of min ones problems: A dichotomy. In S. Abramsky, C. Gavoille, C. Kirchner, F. Meyer auf der Heide and P. G. Spirakis, eds., *Automata, Languages and Programming, 37th International Colloquium, ICALP 2010*, Bordeaux, France, 2010, LNCS 6198, pp. 653–665. Springer.
- [76] A. KUROISHI, K. BOZEK, T. SHIODA AND E. E. NAKAYAMA. A single amino acid substitution of the human immunodeficiency virus type 1 capsid protein affects viral sensitivity to TRIM5 alpha. *Retrovirology*, 7:58,1–58,10, 2010.
- [77] C. M. LANGE, K. ROOMP, A. DRAGAN, J. NATTERMANN, M. MICHALK, U. SPENGLER, V. WEICH, T. LENGAUER, S. ZEUZEM, T. BERG AND C. SARRAZIN. HLA class I allele associations with HCV genetic variants in patients with chronic HCV genotypes 1a or 1b infection. *Journal of Hepatology*, 53(6):1022–1028, December 2010.
- [78] S. LEE, E. EISEMANN AND H.-P. SEIDEL. Real-Time Lens Blur Effects and Focus Control. *ACM Transactions on Graphics (Proc. ACM SIGGRAPH'10)*, 29(4):65:1–7, 2010.
- [79] T. Lengauer, A. Altmann, A. Thielen and R. Kaiser. Chasing the aids virus. *Communications of the ACM*, 53(3):66–74, 2010.
- [80] T. LENGAUER AND R. KAISER. Computerjagd auf das AIDS-VIRUS. *Spektrum der Wissenschaft*, 08:62–67, August 2009.
- [81] T. LU, S. MERZ AND C. WEIDENBACH. Model checking the pastry routing protocol. In J. Bendisposto, M. Leuschel and M. Roggenbach, eds., *10th Intl. Workshop Automatic Verification of Critical Systems (AVOCS 2010)*, Düsseldorf, Germany, September 2010, vol. 10, pp. 19–21. Universität Düsseldorf. Short contribution.
- [82] Y. MASS, M. RAMANATH, Y. SAGIV AND G. WEIKUM. IQ: The case for iterative querying for knowledge. In *Proceedings of the 5th Biennial Conference on Innovative Data Systems Research (CIDR 2011)*, Asilomar, USA, 2011, p. 38–44.
- [83] R. A. MCGOVERN, A. THIELEN, T. MO, W. DONG, C. K. WOODS, D. CHAPMAN, M. LEWIS, I. JAMES, J. HEERA, H. VALDEZ AND P. R. HARRIGAN. Population-based v3 genotypic tropism assay: a retrospective analysis using screening samples from the a4001029 and motivate studies. *AIDS*, 24(16):2517–2525, October 2010.
- [84] A. C. MCHARDY AND B. ADAMS. The role of genomics in tracking the evolution of influenza a virus. *PLoS Pathog.*, 5(10):e1000566, October 2009.
- [85] K. MEHLHORN AND M. SAGRALOFF. A deterministic descartes algorithm for real polynomials. Accepted for ISSAC 2009, January 2010.
- [86] G. DE MELO AND G. WEIKUM. MENTA: Inducing multilingual taxonomies from Wikipedia. In J. Huang, N. Koudas, G. Jones, X. Wu, K. Collins-Thompson and A. An, eds., *Proceedings of the 19th ACM Conference on Information and Knowledge Management (CIKM 2010)*, Toronto, Canada, 2010, pp. 1099–1108. ACM.
- [87] J. MEYER, P. SCHNITZSPAN, S. KOHLBRECHER, K. PETERSEN, O. SCHWAHN, M. ANDRILUKA, U. KLINGAUF, S. ROTH, B. SCHIELE AND O. VON STRYK. A semantic world model for urban search and rescue based on heterogeneous sensors. In *RoboCup 2010: Robot Soccer World Cup XIV*, Singapore, Singapore, 2011 to appear, pp. 1–12. Springer.
- [88] M. MÜLLER, M. CLAUSEN, V. KONZ, S. EWERT AND C. FREMERY. A multimodal way of experiencing and exploring music. *Interdisciplinary Science Reviews (ISR)*, 35(2):138–153, 2010.
- [89] M. MÜLLER AND S. EWERT. Towards timbre-invariant audio features for harmony-based music. *IEEE Transactions on Audio, Speech, and Language Processing (TASLP)*, 18(3):649–662, 2010.

- [90] T. MÜTZE, T. RAST AND R. SPÖHEL. Coloring random graphs online without creating monochromatic subgraphs. In D. Randall, ed., *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 11)*, San Francisco, USA, January 2011, pp. 145–158. Society for Industrial and Applied Mathematics.
- [91] N. NAKASHOLE, M. THEOBALD AND G. WEIKUM. Scalable knowledge harvesting with high precision and high recall. In *4th International Conference on Web Search and Data Mining (WSDM 2011)*, Hong Kong, 2011, ACM. To appear.
- [92] G. PONS-MOLL, A. BAAK, T. HELTEN, M. MÜLLER, H.-P. SEIDEL AND B. ROSENHAHN. Multisensorfusion for 3D full-body human motion capture. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, San Francisco, CA, USA, June 2010, pp. 663–670.
- [93] N. PREDÁ, G. KASNECI, F. M. SUCHANEK, T. NEUMANN, W. YUAN AND G. WEIKUM. Active knowledge: dynamically enriching RDF knowledge bases by web services. In *Elmagarmid and Agrawal* [49], pp. 399–410.
- [94] F. RAMÍREZ AND M. ALBRECHT. Finding scaffold proteins in interactomes. *Trends in Cell Biology*, 20(1):2–4, 2010.
- [95] E. REINHARD, G. WARD, S. PATTANAIK, P. DEBEVEC, W. HEIDRICH AND K. MYSKOWSKI. High Dynamic Range Imaging: Acquisition, Display, and Image-based Lighting. Elsevier (Morgan Kaufmann), Burlington, MA, 2010.
- [96] S.-Y. RHEE, S. MARGERIDON-THERMET, M. NGUYEN, T. LIU, R. KAGAN, B. BEGGEL, J. VERHEYEN, R. KAISER AND R. SHAFER. Hepatitis b virus reverse transcriptase sequence variant database for sequence analysis and mutation discovery. *Antiviral Research*, 88(3):269–275, December 2010.
- [97] T. RITSCHEL, T. THORMÄHLEN AND H.-P. SEIDEL. Interactive on-surface signal deformation. *ACM Transactions on Graphics (TOG)*, 29(4):1–8, July 2010.
- [98] M. ROHRBACH, M. STARK, G. SZARVAS, I. GUREVYCH AND B. SCHIELE. What helps Where - and Why? Semantic Relatedness for Knowledge Transfer. In *2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. – Pt. 2, San Francisco, USA, June 2010, pp. 910–917. IEEE.
- [99] A. RYBALCHENKO AND V. SOFRONIE-STOKKERMANS. Constraint solving for interpolation. *Journal of Symbolic Computation*, 45(11):1212–1233, 2010. Accepted for publication.
- [100] A. SCHLICKER AND M. ALBRECHT. FunSimMat update: new features for exploring functional similarity. *Nucleic Acids Research*, 38(Database issue):D244–D248, 2010.
- [101] A. SCHLICKER, T. LENGAUER AND M. ALBRECHT. Improving disease gene prioritization using the semantic similarity of Gene Ontology terms. *Bioinformatics*, 26(18):i561–i567, 2010.
- [102] P. SCHNITZSPAN, S. ROTH AND B. SCHIELE. Automatic discovery of meaningful object parts with latent CRFs. In *2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. – Pt. 1, San Francisco, USA, June 2010, *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, vol. 23, pp. 121–128. IEEE.
- [103] V. SETTY, S. J. BEDATHUR, K. BERBERICH AND G. WEIKUM. InZeit: Efficiently identifying insightful time points. *PVLDB*, 3(2):1605–1608, 2010.
- [104] S. SEUFERT, S. BEDATHUR, J. MESTRE AND G. WEIKUM. Bonsai: Growing interesting small trees. In *10th IEEE International Conference on Data Mining (ICDM 2010)*, Sydney, Australia, December 2010, pp. 1013–1018. IEEE Computer Society.
- [105] M. SHMUELI-SCHUEER, C. LI, Y. MASS, H. ROITMAN, R. SCHENKEL AND G. WEIKUM. Best-effort top-k query processing under budgetary constraints. In *ICDE*, 2009, pp. 928–939. IEEE.
- [106] D. SKOČAJ, K. MATEJ, A. VRECKO, A. LEONARDIS, M. FRITZ, M. STARK, B. SCHIELE, S. HONGENG AND J. L. WYATT. Multi-modal learning. In H. I. Christensen, G.-J. M. Kruijff and J. L. Wyatt, eds., *Cognitive Systems, Cognitive Systems Monographs*, vol. 8, ch. 7, pp. 265–309. Springer, Berlin, 2010.
- [107] V. SOFRONIE-STOKKERMANS. Locality results for certain extensions of theories with bridging functions. In R. A. Schmidt, ed., *Automated Deduction – CADE-22, 22nd International Conference on Automated Deduction*, Montreal, Canada, 2009, *LNAI 5663*, p. 67–83. Springer.
- [108] M. SOZIO AND A. GIONIS. The community-search problem and how to plan a successful cocktail party. In B. Rao, B. Krishnapuram, A. Tomkins and Q. Yang, eds., *KDD*, 2010, pp. 939–948. ACM.
- [109] M. STARK, M. GOESELE AND B. SCHIELE. Back to the future: Learning shape models from 3D CAD data. In F. Labrosse, R. Zwiggielaar, Y. Liu and B. Tiddeman, eds., *21st British Machine Vision Conference (BMVC)*, Aberystwyth, Wales, September 2010, *Proceedings of the British Machine Vision Conference*, vol. 21, pp. 106,1–106,11. BMVA.
- [110] U. STEINHOFF AND B. SCHIELE. Dead reckoning from the pocket – an experimental study. In *IEEE 2010 International Conference on Pervasive Computing and Communications (PerCom)*, Mannheim, Germany, 2010, pp. 162–170. IEEE.
- [111] C. STOLL, J. GALL, E. DE AGUIAR, S. THRUN AND C. THEOBALT. Video-based reconstruction of animatable human characters. *ACM Transactions on Graphics (Proc. of SIGGRAPH Asia)*, 29(6):139–149, 2010.
- [112] R. STRZODKA, M. SHAHEEN, D. PAJAK AND H.-P. SEIDEL. Cache oblivious parallelograms in iterative stencil computations. In *ICS '10: Proceedings of the 24th ACM International Conference on Supercomputing*, 2010, pp. 49–59. ACM.
- [113] M. SUDA, C. WEIDENBACH AND P. WISCHNEWSKI. On the saturation of YAGO. In J. Giesl and R. Hähnle, eds., *Automated Reasoning, 5th International Joint Conference, IJCAR 2010*, Edinburgh, UK, 2010, *LNAI 6173*, pp. 441–456. Springer.
- [114] C. THEOBALT, E. DE AGUIAR, C. STOLL, H.-P. SEIDEL AND S. THRUN. Image and Geometry Processing for 3D-Cinematography, ch. *Performance Capture from Multi-view Video*. Springer, 2010.

- [115] A. THIELEN, N. SICHTIG, R. KAISER, J. LAM, P. R. HARRIGAN AND T. LENGAUER. Improved prediction of HIV-1 coreceptor usage with sequence information from the second hypervariable loop of gp120. *Journal of Infectious Diseases*, 202(9):1435–1443, November 2010.
- [116] L. H. U, N. MAMOULIS, K. BERBERICH AND S. J. BEDATHUR. Durable top-k search in document archives. In *Elmagarmid and Agrawal* [49], pp. 555–566.
- [117] L. VALGAERTS, A. BRUHN, H. ZIMMER, J. WEICKERT, C. STOLL AND C. THEOBALD. Joint estimation of motion, structure and geometry from stereo sequences. In *Proc. of European Conference on Computer Vision (ECCV)*, 2010, pp. 1–12.
- [118] S. WALK, N. MAJER, K. SCHINDLER AND B. SCHIELE. New features and insights for pedestrian detection. In *2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. – Pt. 2, San Francisco, USA, 2010, pp. 1030–1037. IEEE.
- [119] S. WALK, K. SCHINDLER AND B. SCHIELE. Disparity statistics for pedestrian detection: Combining appearance, motion and stereo. In K. Daniilidis, P. Maragos and N. Paragios, eds., *Computer Vision – ECCV 2010, 11th European Conference on Computer Vision*. – Pt. VI, Heraklion, Crete, Greece, September 2010, LNCS 6341, pp. 182–195. Springer.
- [120] M. WAND, B. ADAMS, M. OVSJANIKOV, A. BERNER, M. BOKELOH, P. JENKE, L. GUIBAS, H.-P. SEIDEL AND A. SCHILLING. Efficient reconstruction of non-rigid shape and motion from real-time 3D scanner data. *ACM Transactions on Graphics*, 28(2), 2009.
- [121] G. WEIKUM, N. NTAMOS, M. SPANIOL, P. TRIANTAFILLOU, A. BENCZÚR, S. KIRKPATRICK, P. RIGAUX AND M. WILLIAMSON. Longitudinal analytics on web archive data: It’s about time! In *Proceedings of the 5th Biennial Conference on Innovative Data Systems Research (CIDR 2011)*, Asilomar, USA, 2011, pp. 199–202.
- [122] G. WEIKUM AND M. THEOBALD. From information to knowledge: harvesting entities and relationships from web sources. In J. Paredaens and D. V. Gucht, eds., *PODS*, 2010, pp. 65–76. ACM.
- [123] G. WETZSTEIN, I. IHRKE AND W. HEIDRICH. Sensor saturation in Fourier multiplexed imaging. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010, pp. 1–8.
- [124] S. E. WHANG, D. MENESTRINA, G. KOUTRIKA, M. THEOBALD AND H. GARCIA-MOLINA. Entity resolution with iterative blocking. In U. Çetintemel, S. B. Zdonik, D. Kossmann and N. Tatbul, eds., *SIGMOD Conference*, 2009, pp. 219–232. ACM.
- [125] T. WITTKOP, D. EMIG, S. LANGE, S. RAHMANN, M. ALBRECHT, J. H. MORRIS, S. BÖCKER, J. STOYE AND J. BAUMBACH. Partitioning biological data with transitivity clustering. *Nature Methods*, 7(6):419–420, 2010.
- [126] I. WÖHLERS, F. S. DOMINGUES AND G. W. KLAU. Towards optimal alignment of protein structure distance matrices. *Bioinformatics*, 26(18):2273–2280, September 2010.
- [127] C. WOJEK, S. ROTH, K. SCHINDLER AND B. SCHIELE. Monocular 3D scene modeling and inference: Understanding multi-object traffic scenes. In K. Daniilidis, P. Maragos and N. Paragios, eds., *Computer Vision – ECCV 2010, 11th European Conference on Computer Vision*. – Pt. IV, Heraklion, Crete, Greece, September 2010, LNCS 6341, pp. 467–481. Springer.

Directions

The Max Planck Institute for Informatics (Building E14) is located on the campus of Saarland University, approximately 5 km north-east of downtown Saarbrücken, in a wooded area near Dudweiler.

Saarbrücken has its own airport (Saarbrücken-Ensheim) and is well-connected to the Frankfurt and Luxemburg airports by shuttle bus, train, or car. There are good hourly train connections from Saarbrücken to cities all over Germany, and trains to Metz, Nancy and Paris leave at regular intervals. The express highways – or Autobahns – provide fast routes to Mannheim, Frankfurt, Luxemburg, Trier, Köln, and Straßburg as well as Metz, Nancy, and Paris.

The campus is accessible...

from the airport at Saarbrücken-Ensheim

by taxi in approximately 20 minutes

from Saarbrücken central train station

by taxi in approximately 15 minutes

by bus in approximately 20 minutes

direction „Dudweiler-Dudoplatz“ or „Universität Campus“

exit „Universität Mensa“, or

alternatively exit „Universität Campus“

from Frankfurt or Mannheim on Autobahn A6

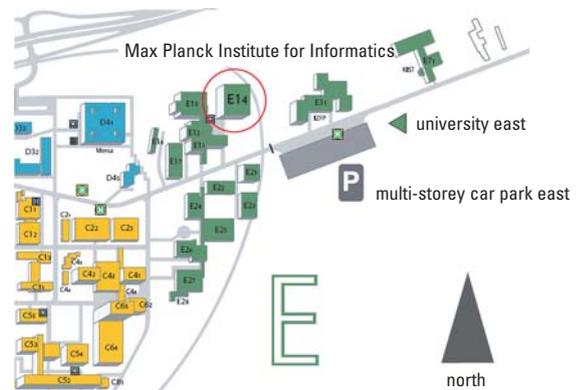
exit „St.Ingbert-West“

follow the white signs for „Universität“, leading to the campus

from Paris on Autobahn A4

exit „St.Ingbert-West“

follow the white signs for „Universität“, leading to the campus



: IMPRINT

: **Publisher**

: Max Planck Institute for Informatics
: Campus E1 4
: D-66123 Saarbrücken

: **Editors & Coordination**

: Volker Maria Geiß
: Samir Hammann
: Manuel Lamotte-Schubert
: Thomas Lengauer
: Kurt Mehlhorn
: Jennifer Müller
: Marcus Rohrbach
: Thomas Sauerwald
: Bernt Schiele
: Hans-Peter Seidel
: Ingolf Sommer
: Martin Theobald
: Christian Theobalt
: Uwe Waldmann
: Christoph Weidenbach
: Gerhard Weikum

: **Proof-reading**

: Krista Ames
: Kamila Kolesniczenko
: Ebba Schröder
: Nathalie Strassel

: **Photos & Photo Editing**

: Martin Cadik
: Ivo Ihrke
: Manuel Lamotte-Schubert
: Kristina Scherbaum

: **Contact**

: Max Planck Institute for Informatics
: phone +49 681 9325-0
: fax +49 681 9325-5719
: email info@mpi-inf.mpg.de
: internet http://www.mpi-inf.mpg.de

: **Report Period**

: 1st of January, 2009 to 31st of December, 2010

: **Design**

: Behr Design | Saarbrücken

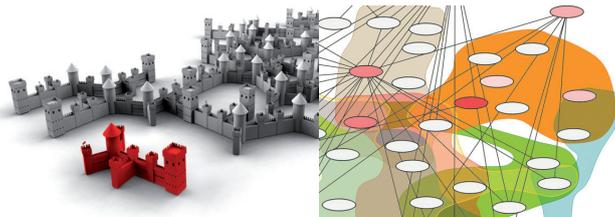
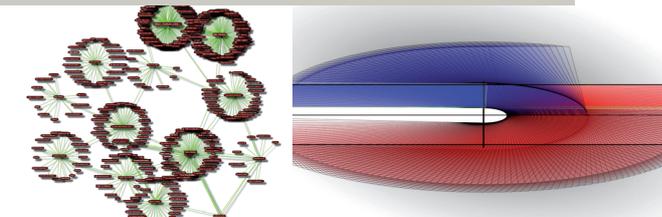
: **Printing**

: Kubbli GmbH





max planck institut
informatik



Max Planck Institute for Informatics
Campus E1 4
D-66123 Saarbrücken

phone +49 681 9325-0
fax +49 681 9325-5719
email info@mpi-inf.mpg.de
internet <http://www.mpi-inf.mpg.de>



max planck institut
informatik