

Kranc: a Mathematica package to generate numerical codes for tensorial evolution equations [☆]

Sascha Husa ^{a,*}, Ian Hinder ^b, Christiane Lechner ^a

^a *Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, Am Mühlenberg 1, D-14476, Potsdam, Germany*

^b *School of Mathematics, University of Southampton, Southampton, SO17 1BJ, United Kingdom*

Received 12 July 2005; received in revised form 12 February 2006; accepted 24 February 2006

Available online 17 April 2006

Abstract

We present a suite of Mathematica-based computer-algebra packages, termed “Kranc”, which comprise a toolbox to convert certain (tensorial) systems of partial differential evolution equations to parallelized C or Fortran code for solving initial boundary value problems. Kranc can be used as a “rapid prototyping” system for physicists or mathematicians handling very complicated systems of partial differential equations, but through integration into the Cactus computational toolkit we can also produce efficient parallelized production codes. Our work is motivated by the field of numerical relativity, where Kranc is used as a research tool by the authors. In this paper we describe the design and implementation of both the Mathematica packages and the resulting code, we discuss some example applications, and provide results on the performance of an example numerical code for the Einstein equations.

Program summary

Title of program: Kranc

Catalogue identifier: ADXS_v1_0

Program summary URL: http://cpc.cs.qub.ac.uk/summaries/ADXS_v1_0

Program obtainable from: CPC Program Library, Queen’s University of Belfast, N. Ireland

Distribution format: tar.gz

Computer for which the program is designed and others on which it has been tested: General computers which run Mathematica (for code generation) and Cactus (for numerical simulations), tested under Linux

Programming language used: Mathematica, C, Fortran 90

Memory required to execute with typical data: This depends on the number of variables and gridsize, the included ADM example requires 4308 KB

Has the code been vectorized or parallelized: The code is parallelized based on the Cactus framework.

Number of bytes in distributed program, including test data, etc.: 1 578 142

Number of lines in distributed program, including test data, etc.: 11 711

Nature of physical problem: Solution of partial differential equations in three space dimensions, which are formulated as an initial value problem. In particular, the program is geared towards handling very complex tensorial equations as they appear, e.g., in numerical relativity. The worked out examples comprise the Klein–Gordon equations, the Maxwell equations, and the ADM formulation of the Einstein equations.

Method of solution: The method of numerical solution is finite differencing and method of lines time integration, the numerical code is generated through a high level Mathematica interface.

Restrictions on the complexity of the program: Typical numerical relativity applications will contain up to several dozen evolution variables and thousands of source terms, Cactus applications have shown scaling up to several thousand processors and grid sizes exceeding 500^3 .

[☆] This paper and its associated computer program are available via the Computer Physics Communications homepage on ScienceDirect (<http://www.sciencedirect.com/science/journal/00104655>).

* Corresponding author.

E-mail address: sascha.husa@uni-jena.de (S. Husa).

Typical running time: This depends on the number of variables and the grid size: the included ADM example takes approximately 100 seconds on a 1600 MHz Intel Pentium M processor.

Unusual features of the program: based on Mathematica and Cactus

© 2006 Elsevier B.V. All rights reserved.

PACS: 04.25.Dm; 2.70.Bf; 2.60.Cb

Keywords: Numerical relativity; Computer algebra; Code generation; Partial differential equations; Finite differencing

1. Introduction

The numerical solution of partial differential equations (PDEs) is a common and often quite challenging problem in physics, applied mathematics and other fields. The Kranc software presented here is designed as a tool for initial boundary value problems where the complexity of the PDE systems is a principal problem, and manual coding of equations and general equation-dependent code is too error prone or simply not feasible. Our software is released under the GPL (GNU General Public License), and is available at the URL <http://numrel.aei.mpg.de/Research/Kranc>. Currently, Kranc only directly supports the numerical solution of initial boundary value problems with simplistic boundary conditions. Extension to fully user-specified boundary conditions is planned for the future, though it is possible to link the automatically generated code with hand-written code for the boundary treatments, e.g., following standard Cactus examples.

Our work is motivated by the field of numerical relativity, i.e. the numerical solution of the Einstein equations [1]. An important open research problem is to find an evolution system for the Einstein equations which is well-suited for numerical simulations. The presented software drastically reduces the complexity of the task of comparing different formulations of the Einstein equations; these are complicated nonlinear systems which can have dozens of evolution variables and thousands of source terms.

We aim to address the problems of handling complexity and of dealing with tensorial (or other multi-component) quantities in a transparent way. To this end, we present a suite of Mathematica packages, termed “Kranc”, to convert (optionally tensorial) systems of partial differential evolution equations to parallelized C or Fortran code. Our approach is to use Mathematica to perform high level tensor operations and convert the equations into component form. This form is then converted to C or Fortran 90. See Section 2.1 for the rationale behind this choice. For evolution problems in three spatial dimensions performance is critical. The design decisions that we have made (see Sections 2, 3) allow us to automatically generate codes with efficiencies comparable to hand-written code. Kranc provides a simple programming model for code generation. Modifying Kranc to generate more efficient code by—for example—re-arranging loops, is straightforward.

In numerical relativity, the use of computer algebra (CA) methods based on symbolic tensor calculus cannot only save scientists valuable time, but can also help to substantially reduce errors. It can be used not only for manipulating the evolution equations, but also for deriving constraint propagation systems and perturbation formalisms. Having a code-generation system directly compatible with the CA system is obviously desirable. Generally, we feel that our use of CA makes it easier to focus on algorithms, detached from a particular system of equations. Stressing a more abstract point of view is not only mathematically more appealing but also increases flexibility, which benefits scientific productivity.

Our paper is organized as follows: In Section 2 we describe the computational infrastructure that we use. Since our software only targets the issue of complexity of the equation systems and of processing tensorial quantities, we rely on external software to provide a general computer algebra environment and a framework for parallelization. The design of the generated codes is described in Section 3, while the design of the Kranc software itself is described in Section 4. In Section 5 we describe the use of Kranc to generate codes for the Klein–Gordon equation, the Maxwell equations, and an evolution system for full general relativity, the Arnowitt–Deser–Misner (ADM) equations. We consider our software both as a “rapid prototyping” system and as a system to create efficient production code, we therefore also include some results on performance.

2. Choice of infrastructure, and related design decisions

2.1. The role of computer algebra

When designing a numerical code to solve tensorial equations, two principal alternatives come to mind. A first possibility is to implement high-level tensor operations in a language like C++ or Fortran 90 using tensorial operations like contractions which operate on tensor objects. An interesting example of this approach has been implemented by Schnetter [2]. The second option is to use computer algebra to expand all tensorial expressions into components, and then use only standard arithmetics in the numerical code. Both approaches have their advantages. While the first approach avoids mixing of technologies, the second approach uses specialized technologies to deal with parts of the problem. For previous work on generating code from high-level problem descriptions see also Refs. [3–8].

We decided to adopt the second strategy for the following reasons. Most importantly, we needed a framework that extends beyond code generation, and can also be used as a tool for the analytical part of our work. This requirement naturally leads to the use of a computer algebra system. Including the transition from tensors to components and the process of code generation in the computer algebra system is relatively easy to implement, and greatly simplifies the structure of the resulting numerical code, which is written purely at the level of components. At least part of this task could easily have been implemented in a language like PERL, which would have decreased the dependence on a commercial computer algebra system (we chose Mathematica, see Section 2.2). However, this would have demanded knowledge of (at least) three different languages from many users and developers: Mathematica, PERL and C/Fortran, which would have significantly steepened the learning curve.

Our design, where the numerical code is not based on tensorial data structures or operations, simplifies performance optimization as the code is readable and transparent down to the level of standard arithmetics. The resulting code can in principle be implemented in any suitable language, and we provide the option of generating either Fortran 90 or C (all examples below will be C). Furthermore, in our approach the full power of the chosen computer algebra system can be utilized in the process of code generation, e.g., to perform optimizations which rely on the structure of the equations. Generally, manipulations on the computer algebra level make it much easier to perform experiments involving significant rewrites of the numerical code.

2.2. Computer algebra software

All of the tasks performed by our software which are related to the tensorial nature of equation systems can in principle be accomplished by component calculations. These can be carried out quite conveniently and efficiently by computer algebra systems like GRTensorII [9], which allows expressions to be entered in an abstract notation and yields results in component form. This is often what one wants, but when considering deriving and analyzing different systems of equations, as is often necessary in numerical relativity research, it is clear that more is needed. Apart from the fact that this approach may result in very large calculations requiring significant time and memory, this method is unwieldy and not very intuitive. Rather one would like to keep an abstract notation as long as possible, and in particular get results in this notation. We have decided to use abstract index tensor calculus; i.e. tensors are manipulated symbolically according to rules based on an index notation (see e.g. Wald [10], Section 2.4, and Penrose and Rindler [11]).

We chose Mathematica as our base CA system for several reasons. The authors were already familiar with its use, the language has proven remarkably stable, and Mathematica provides excellent support for pattern matching in the core language. Also, with the freely available Ricci [12], xTensor [13] and the commercial MathTensor [14], there are already three add-on packages for abstract index tensor manipulation available. xTensor has only been released very recently and was not available during the main development phase of the present project. Note that pattern matching seems quite essential for tensor manipulations, e.g., T_{ab} and T_{cd} are not the same expression but still are equivalent mathematical objects, which can easily be identified with pattern matching techniques.

Code generation with Kranc can be performed using either MathTensor or the TensorTools package, which was developed by Hinder during the course of this work as a non-commercial alternative. MathTensor has also been made the basis of our extensive analytical calculations, but since this article is restricted to the code generation part of our project, we will discuss only the usage with TensorTools, which is included in the distribution of Kranc. Both MathTensor and TensorTools represent tensors as plain functions of indices: $h[la, lb] \rightarrow h_{ab}$, $CD[Metricg[la, lb], lc] \rightarrow 0$. This straightforward syntax is less error-prone than Ricci's corresponding $h[L[a], L[b]]$ or $h[L[a], L[b]] [L[c]]$ for a covariant derivative.

TensorTools uses approximately the same syntax as MathTensor, and will accept reasonably straightforward MathTensor expressions involving partial derivatives and tensors. This means that tensorial analytical work can be done using MathTensor where TensorTools is inadequate, and the results can be fed into TensorTools and hence Kranc.

2.3. Parallel computing infrastructure

Aiming at “production quality” evolution codes in three spatial dimensions, we regarded parallelization as a vital feature. The primary requirements for the parallelization strategy were a short development time and portability—our code should run both on inexpensive small clusters, e.g., commodity Linux workstations, but should also scale up to very large runs performed at super-computer centres. We therefore did not want to rely on shared memory or proprietary solutions. The current de facto standard for distributed (scalable) computing is MPI (message passing interface) [15]. Several software packages are available which introduce a software layer between the application programmer and MPI, and thus significantly reduce the effort to write parallel applications. Two examples we considered as the basis for our work are Cactus [16] and PETSc [17]. PETSc is a general purpose tool developed at the Argonne National Laboratory as parallel framework for numerical computations, and Cactus is an open source problem solving environment originating in the numerical relativity community, originally developed at the Max Planck Institute for Gravitational Physics. While PETSc offers more support for numerical algorithms, in particular for parallel elliptic solvers, Cactus already contains some general numerical relativity functionality like horizon finders [18,19] and is geared toward flexibility and collaborative projects (it also allows the use of PETSc as a library). We finally decided to base our code on Cactus, also because two

of the authors currently work at the Max Planck Institute for Gravitational Physics, and thus have a chance of direct collaboration with the Cactus development team.

Cactus mainly targets the issues of parallelization, modularization and portability. The name Cactus derives from the design of a central core (or “flesh”) which connects to application modules (or “thorns”) through an extensible interface. Thorns can implement problem specific code on any level from concrete physics applications to low-level infrastructure. A set of thorns comprising the “Cactus computational toolkit” provide infrastructure such as parallel I/O, data distribution and checkpointing. Current data distribution implementations are based on the principles of domain decomposition and message passing using the MPI standard [15]. This choice yields an open and reasonably *documented* infrastructure, parallelization, a variety of I/O methods and allows easy interfacing with a growing community writing numerical relativity Cactus applications. Modifications to interface with other systems with capabilities similar to those of Cactus or with standalone codes should be straightforward.

We now describe those aspects of Cactus which are important for the remainder of this paper. Cactus does not have the structure of a library which provides a set of functions that can be called by the user application. Instead, all Cactus thorns are compiled into libraries, and management tasks such as code execution and allocation of distributed memory are handled by Cactus and steered via configuration files. The end user supplies a “parameter file” which specifies all user-controllable aspects of the run.

The basic module structure within Cactus is called a “thorn”. All user-supplied code is organized into thorns, which communicate with each other via calls to the flesh API (application programmer interface) or APIs of other thorns. The integration of a thorn into the flesh or with other thorns is specified in configuration files which are parsed at compile time. Thorns are logically grouped into *arrangements* for organizational purposes. The arrangements live in the `arrangements` subdirectory of the main Cactus directory. Inside an arrangement directory there are directories for each thorn belonging to the arrangement. One current shortcoming of Cactus is that this arrangement directory structure can only be one level deep. Arrangement and thorn names must be (case independently) unique, must start with a letter, and can only contain letters, numbers or underscores. In addition, thorn names must contain 27 characters or less.

A key concept for a thorn is the *implementation*, which defines a group of variables and parameters which are used to implement some well-defined functionality. Relationships among thorns are all based upon relationships among the implementations they provide. Different thorns providing the same implementation can be compiled into the same binary; the decision for which thorn is executed (“active”) can be made at run time.

A thorn consists of a subdirectory of an arrangement containing at least three administrative files, illustrated here by examples taken from a Kranc-generated arrangement to solve the Maxwell equations (see the example in Section 5.2). The complete code is provided with the Kranc distribution. These files are text files written in the *Cactus Configuration Language* (CCL). The CCL syntax is case independent, and the ‘#’ character indicates that the rest of the line is a comment.

- `interface.ccl`

This defines the implementation the thorn provides, and the variables the thorn needs, along with their visibility to other implementations.

```
# excerpt from Example code EM/EMBase/interface.ccl
implements: EMBase
inherits:   Grid
public:

CCTK_REAL E1 type=GF timelevels=2 # the electric field
{
  E11, E12, E13                # 3 components of the vector E1
} "E1"

CCTK_REAL Elrhs type=GF timelevels=1 # RHS for E1-evolution eq.
{
  E11rhs, E12rhs, E13rhs
} "Elrhs"

# analogous for the magnetic field B...
```

- `param.ccl`

This defines the thorn parameters along with their visibility to other implementations. The parameters are numbers, strings and switches which are read at runtime and control the thorn behavior.

```
# excerpt from Example code EM/EMsetID/param.ccl
restricted:
```

```
CCTK_REAL sigma "sigma"      # a floating point parameter called sigma
{
  *.* :: "no restrictions"    # all values are allowed
} 0                            # default value is 0
```

- `schedule.ccl`

By default no routine of a thorn will be run. Those routines which should be run by the Cactus infrastructure are registered in the file `schedule.ccl` along with directives for appropriate memory allocation and interprocessor communication for grid variables. Functions are scheduled in *scheduling bins*—slots in the main loop such as BASEGRID (for setting up coordinates), EVOL (the evolution step) or ANALYSIS (for analyzing data).

```
# excerpt from Example code EM/EMBase/schedule.ccl
schedule EMTTBase_RegisterSymmetries at BASEGRID # BASEGRID is a scheduling bin
{
  LANG: C
} "register symmetries"
```

3. Design of Kranc evolution codes

3.1. Method of lines

Aiming both at simplicity and flexibility, we decided to restrict time evolution algorithms to those compatible with the method of lines (MoL) e.g. [20]. Here, the discretization of a system of PDEs describing an evolution problem is split into a separate discretization of “space” and “time”. Discretizing only the spatial derivatives (e.g., by finite difference, spectral or finite element methods), while keeping the time derivatives continuous yields a *coupled* set of ordinary differential equations (ODEs) in time, with one ODE per spatial grid point. These ODEs are coupled via the spatial discretization. For a stable spatial finite differencing scheme, any stable ODE integrator can then be used to time-integrate this large system of ODEs. Further key advantages of the method of lines are that:

- the stability theory is well developed, and is simplified by splitting the analysis into the analysis of spatial and time discretizations;
- numerical methods for the method of lines are well-developed; and
- by decoupling the taking of spatial derivatives, e.g., spatial finite differences, from the time integration, the resulting numerical scheme and computer code are simplified and can easily be modularized.

In particular, the method of lines makes it straightforward to combine options for spatial discretization and for time integration. For the concrete implementation, we choose to base our code on the MoL [21] method of lines thorn within Cactus, developed by Hawke. This code provides a parallel ODE integrator, implementing generic Runge–Kutta and iterative Crank–Nicolson methods.

3.2. Spatial discretization

Within Kranc, spatial discretization is implemented in a centralized and flexible way. At the moment, only a small number of finite differencing methods are implemented for spatial discretization, but extension to other methods based on a regular grids is straightforward. All finite difference formulas are collected in a single header file, currently via C-style macros. Support for unstructured meshes and irregular boundary shapes, e.g., for finite element codes, depends on supporting such meshes in the Cactus infrastructure. This is planned for future Cactus releases, and can then be incorporated into Kranc in a relatively straightforward way.

For an evolution code in three spatial dimensions, which evolves equations as complicated as the Einstein equations, performance issues are critical. It is typical that a code spends most of its time evaluating the right-hand side expressions of the evolution equations. It is thus crucial to simplify these expressions as much as possible (this is done by the computer algebra part) and optimize the evaluation of the finite difference formulas. It turns out that significant performance gains can be obtained when finite difference expressions (and any sub-expressions that are very lengthy or re-used many times) are promoted to local scalar variables that are pre-computed once before being used many times. Apart from cutting down runtime, this method often drastically reduces compile time when heavy optimization is used. Pre-computing is the default behavior.

The user of the Kranc system will denote derivatives using an abstract notation when passing equations to the Kranc functions; e.g., a first derivative of the variable f in the x -direction would be denoted as $D1[f]$. The complete list of abstract operators for first and second derivatives is $\{D1, D2, D3, D11, D22, D33, D21, D31, D32\}$.

At compile time, the user has two choices to make:

- Should the derivatives be precomputed at the start of the grid loop?
- What sort of finite differencing should be performed?

This decision is currently made at compile time in order to maximize efficiency and may change in future releases of the software. In order to disable precomputation, the user should set the preprocessor macro `NOPRECOMPUTE`, typically through the compiler option `-DNOPRECOMPUTE`. The choice of finite differencing is controlled by defining macros: `FD_C2` for second order centred, `FD_C4` for fourth order centred, or `FD_C0` for zeroth order (a projection of all derivative terms to zero, useful for debugging and performance profiling). These macros may be defined by compiling with, e.g., `-DFD_C2`. This corresponds to the default setting.

The `GenericFD` thorn defines several preprocessor macros for performing finite differencing. The concrete finite differencing macro name is formed by appending the method to the abstract name; for example, `D1_c2[f]` denotes the first derivative of f in the x -direction, using centred second order differences. At the moment only second and fourth order centred methods are implemented.

In the C or Fortran code which is produced by Kranc, the abstract notation is replaced with a grid function style notation; e.g., `D1[f]` is replaced by `D1[f, i, j, k]`, where $\{i, j, k\}$ are the loop indices representing the grid points. These “derivative operators” are in fact preprocessor macros, which have different definitions depending on whether or not precomputation is enabled. When precomputing is switched off, macros such as `D1[f, i, j, k]` are defined as `D1gf[f, i, j, k]`. When precomputing is switched on, macros such as `D1[f, i, j, k]` are defined as `D1f`. Here `D1f` is used as a local variable which is pre-computed at the start of the loop over grid points; e.g.,

```
D1f = D1gf(f, i, j, k);
```

The operators with suffix `gf` are defined in terms of whichever concrete differencing macros have been selected.

3.3. Boundary conditions

Currently, the stable version of Kranc only directly supports generating code with periodic boundary conditions. This is achieved via the use of “ghost points” on the sides of the grid which Cactus automatically populates with data from the opposite side so that finite difference operators can be used right up to the last “real” grid point. Extension to fully user-specified boundary conditions is planned for future releases, also it is certainly possible to link the automatically generated code with hand-written code for the boundary treatments, e.g., following standard Cactus examples.

3.4. Code organization

On the user level, the Kranc code generation toolkit consists of Mathematica functions that generate an entire code module (*thorn* in the Cactus terminology). Within the Kranc system, thorns are classified into a small number of thorn types, and complex codes can be composed from a larger number of thorns which fall into one of these classes. The basic tasks of a thorn are:

- define Cactus grid functions;
- assign values to grid functions;
- set grid function attributes (e.g., symmetries or boundary conditions); and
- define Cactus parameters.

We define five types of thorn with specific characteristics regarding their functionality: the base, setter, translator, evaluator and MoL thorns. While the base thorn only defines parameters, grid functions and their properties, all the other thorns also perform one or more calculations, i.e. they assign new values to one or more grid functions. Here the translator, evaluator and MoL thorns are in a sense “added value” versions of the basic setter thorn.

A *calculation* is one of the essential concepts of Kranc. The end result of a calculation is to assign new values to a grid function in a loop over grid points. In order to simplify lengthy calculations, all calculations are carried out in terms of local scalar variables, and the assignment to grid functions is made at the end of the loop. This avoids repeated multiplications due to array accesses. Intermediate expressions, which are only needed as local scalars (i.e. of which no derivatives are required), are called *shorthands* in the Kranc context.

As an example of a calculation, we list a slightly abbreviated (to save space) version of the file `MKGMoL_CalcRHS.c` (in directory `Examples/KleinGordon/MKG/MKGMoL/src`), which assigns the right-hand sides for the MoL evolution of the Klein–Gordon equation, Eq. (1) in Section 5.

```
// several header files are included and some macros get defined
/* Define macros used in calculations */ // C++ style comments do not appear in original code
```

```

#define INITVALUE (42)
void MKGMoL_CalcRHS(CCTK_ARGUMENTS)
{
    DECLARE_CCTK_ARGUMENTS
    DECLARE_CCTK_PARAMETERS

    /* Declare the variables used for looping over grid points */
    int i = INITVALUE; // same for j, k, index, istart, jstart, kstart, iend, jend, kend

    /* Declare finite differencing variables */
    CCTK_REAL dx = INITVALUE; // same for dy, dz, dxi, dyi, dzi, hdx, hdy, hdz

    /* Declare shorthands */ // none in this example

    /* Declare local copies of grid functions */
    CCTK_REAL phiL = INITVALUE; // same for phitL, phirhsL, phitrhsL, D11phi, D22phi, D33phi

    /* Initialize finite differencing variables */
    dx = CCTK_DELTA_SPACE(0);
    dxi = 1 / dx;
    hdx = 0.5 * dxi; // analogously for x and y directions

    /* Set up variables used in the grid loop with stencils suitable for finite differencing */
    istart = cctk_nghostzones[0];
    iend = cctk_lsh[0] - cctk_nghostzones[0]; // analogously for jstart, kstart, jend, kend

    /* Loop over the grid points */
    for (k = kstart; k < kend; k++)
    {
        for (j = jstart; j < jend; j++)
        {
            for (i = istart; i < iend; i++)
            {
                index = CCTK_GFINDEX3D(cctkGH,i,j,k); // a Cactus macro to compute the 3D array index

                /* Assign local copies of grid functions */
                phiL = phi[index];
                phitL = phit[index];

                /* Precompute derivatives */
                D11phi = D11gf(phi,i,j,k); // analogous for D22phi and D33phi

                /* Calculate grid functions */
                phirhsL = phitL;

                phitrhsL = D11(phi,i,j,k) + D22(phi,i,j,k) + D33(phi,i,j,k) -
                    phiL*SQR(mass); // use macro SQR for mass * mass

                /* Copy local copies back to grid functions */
                phirhs[index] = phirhsL;
                phitrhs[index] = phitrhsL;
            }
        }
    }
}

```

A Cactus arrangement will typically contain one base thorn, which defines the basic quantities of the model to be studied, and an assortment of several of the other Kranc thorn types which implement the equations. Functionality which is required but not presently covered by the code that is generated by Kranc, e.g., more involved boundary conditions, can easily be added as hand coded thorns. There are currently several restrictions as to what Kranc is able to generate: only a few boundary conditions are implemented, and only periodic boundaries are fully tested and supported.

Note that all thorns can define their own parameters, which can then be shared with other thorns. Apart from defining parameters and displaying banner information, the Kranc thorns implement the following tasks:

- Base thorn :
 - Register grid functions with Cactus. Discriminate between grid functions that can be evolved and need two sets of storage (timelevels) and those that only need one time level allocated.
 - Register grid function symmetries. The transformation behavior under coordinate reflections $x^i \rightarrow -x^i$ needs to be defined in order to use Cactus infrastructure to run in bitant/quadrant/octant mode for data which are compatible with such symmetries. Via the *Cartoon method* [22] this infrastructure can also be used to handle axisymmetric situations in Cartesian coordinates while keeping manifest axisymmetry.
- Setter thorn:
 - Set grid functions to values, e.g., to set initial data or to update auxiliary grid functions. Functions are scheduled in the POSTINITIAL and EVOL time bins at appropriate times relative to the time stepping functions.
- MoL evolve thorn:
 - Register grid functions with the MoL thorn. During registration, the grid functions need to be designated as either “evolved” or “primitive”, with respect to this thorn. Evolved variables are those that *this thorn* will evolve using MoL. Primitive variables are those that are used by this thorn in the calculation of the right hand sides of the evolution equations, but are not evolved themselves.
 - Generate code to apply boundary conditions.
 - Set MoL right-hand side variables through a function scheduled in the EVOL bin.
- Translator thorn:
 - This is a specialized version of a setter thorn which translates between Kranc variables and some other set of variables. The idea is that initial data may be specified in one set of variables by an external thorn, and can be translated by a translator thorn into the variables you wish to evolve using Kranc. The evolved variables can then be translated back into the original set, so that external thorns can use them. These translations occur in the POSTINITIAL and POSTSTEP bins, respectively. This interface is essential for compatibility with many of Cactus’ existing numerical relativity thorns, which use the ADM variables via the ADMBase thorn.
- Evaluator thorn:
 - Define new grid functions.
 - Assign values to new grid functions in routines scheduled in the Cactus ANALYSIS bin.

When defining grid functions, tensor components are grouped together as Cactus groups. The reflection symmetries of grid functions are read off from their tensor properties automatically (through the function `calcSymmetry[gf]` in file `Thorn.m`). By default, grid functions are treated as scalars, but, e.g., a grid function called `h11` is interpreted as the 11-component of a 2-tensor, which defines its parity information.

4. Kranc design

Kranc is composed of several Mathematica packages (textual.m files); each one performs a distinct function. The user only needs to be concerned with calling functions from one package: `KrancThorns`. This package contains functions for creating the different types of Kranc thorn.

The diagram in Fig. 1 illustrates the relationships between the Kranc packages `KrancThorns`, `TensorTools`, `CodeGen`, `Thorn` and `MapLookup`, which are described in the following subsections. Structuring the code in this way, separating out the different logically independent functions, promotes code reuse. For example, none of the thorn generation packages need to know anything about tensors, and none of the packages other than `CodeGen` (in principle) need to know anything about the programming language that is being written (C or Fortran). We have chosen to define the concept of a “setter thorn”, and “evaluator thorn”, etc., but the mechanics of producing a thorn implemented in `Thorn` and `CodeGen` are completely independent of this decision.

4.1. *KrancThorns: Constructing the different Kranc thorns*

We have already discussed the different types of Cactus thorn used in a Kranc arrangement (MoL, setter, base, translator and evaluator). The `KrancThorns` package provides functions to create thorns of these types given high level descriptions. These `Create*Thorn` functions are the ones directly called by users; see Appendix B for detailed argument descriptions. Internally the `KrancThorns` package uses the `Thorn` package to create the Cactus thorns.

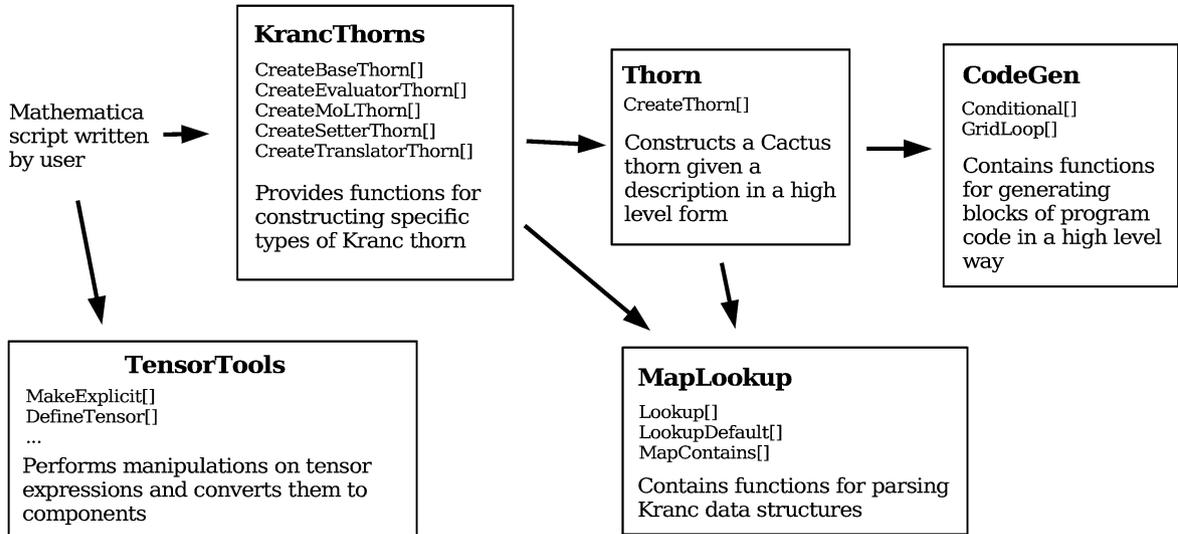


Fig. 1. Relationships between Kranc packages: Each block represents a package, with the main functions it provides indicated with square brackets. An arrow indicates that one package calls functions from another.

4.1.1. Types of arguments

Mathematica allows two types of arguments to be passed to a function: “positional arguments” and “named arguments” (referred to in the Mathematica book as “optional arguments”). It is possible for some named arguments to be omitted from a function call; in this case a suitable default will be chosen. Positional arguments are useful when there are few arguments to a function, and their meaning is clear in the calling context. Named arguments are preferred when there are many arguments, as the argument names are given explicitly in the calling context.

For each type of Kranc thorn, there is a function to create it (`Create*Thorn`). There is a certain set of named arguments (“Common named arguments”) which can be passed to any of these functions (e.g., the name of the Thorn to create, where to create it, etc.). Then, for each type of thorn, there is a specific set of named arguments specifically for that thorn type. All of the functions accept some positional arguments as well. See [Appendix B](#) for descriptions of the arguments that can be given to each of the `KrancThorns` functions.

4.1.2. Common data structures

Kranc is a relatively complex system. It consists of several packages which need to pass data between themselves in a structured way. Mathematica does not have the concept of a C++ “class” or a C “structure”, in which collections of named objects are grouped together for ease of manipulation. Instead, we have adopted the definition of a “Kranc structure” as a list of rules of the form *key* \rightarrow *value*. We have chosen to use the “rule” symbol “ \rightarrow ” for syntactic convenience. For example, one might describe a person using a “Person” structure as follows:

```
alice = {Name -> "Alice", Age -> 20, Gender -> Female}
```

Once a structure has been built up, it can be parsed with the `lookup` function in the `MapLookup` package. `lookup[structure, key]` returns the value in `structure` corresponding to `key`. For example, `lookup[alice, Age]` would return the number 20. This usage mirrors what is known as an association list (or alist) in LISP style languages. Based on this concept we have defined a number of data structures which will be used to describe the thorns to construct. Each of these data structures is introduced below, and full details are given in [Appendix A](#).

GroupDefinition

A `GroupDefinition` structure lists the grid functions that are members of a specific Cactus group. A list of such structures should be supplied to all the `KrancThorns` functions so that Kranc can determine which group each grid function belongs to.

Calculation

Calculation structures are the core of the Kranc system; a detailed specification is given in [Appendix A.1](#). The idea is that the user provides a list of equations of the form *variable* \rightarrow *expression*. When the calculation is performed, for each point in the grid, *expression* is evaluated and placed into the grid function *variable*. Here *expression* may contain partial derivatives of grid functions written using the following notation:

$$\frac{\partial g}{\partial x^i} \rightarrow \text{Di}[g], \quad \frac{\partial g}{\partial x^i \partial x^j} \rightarrow \text{Di}j[g]$$

where i and j must be given explicitly as integers 1, 2 or 3 representing x , y or z .

You may specify intermediate (non-grid) variables called “shorthands” which can be used as *variable* for precomputing quantities which you will use later in the calculation. To identify these variables as shorthands, they should be listed in a “Shorthands” entry of the Calculation.

The arrangement of the terms in the equations can have a marked effect on both compile time and run time. It is often helpful to tell Mathematica to collect the coefficients of certain types of term, rather than expanding out entire expressions. To this end, the user can include a “CollectList” entry in a calculation; this is a list of variables whose coefficients should be collected.

There is the facility for performing multiple loops in a single calculation structure; this can be used to set a grid function in one loop, then evaluate derivatives of it in a later loop. For this reason, the equations are given as a list of lists of equations. Explicit synchronization of ghost zones is performed after each loop for those groups whose grid functions have been set.

Note that the system is not designed to allow the same grid function to be set more than once in a single loop of a calculation.

GroupCalculation

A GroupCalculation structure associates a group name with a Calculation which is used to update the grid functions in that group. This is used when creating evaluator thorns, where the calculations are triggered by requests for output for specific groups.

4.2. TensorTools

4.2.1. Overview

The TensorTools package was written specifically for the Kranc system, though it is in no way tied to it. We needed to perform certain operations on tensorial quantities, and there was no free software available which met our needs. TensorTools has the following features:

- Expand covariant derivatives in terms of partial derivatives and Christoffel symbols (multiple covariant derivatives can be defined).
- Expand Lie derivatives in terms of partial derivatives.
- Automatic relabeling of dummy indices to avoid conflicts.
- Convert abstract tensor expressions into components.

Before using any TensorTools functions, the TensorTools package must be loaded. It is self-contained, with no dependencies on any other packages.

```
Get["TensorTools.m"];
```

4.2.2. Representation of tensor quantities

Tensorial expressions are entered in mostly the same syntax as is used by MathTensor. An abstract tensor consists of a *kernel* and an arbitrary number of abstract *indices*, each of which can be *upper* or *lower*. Abstract indices are alphabetical characters (a–z, A–Z) prefixed with either an “l” or a “u” depending on whether the index is considered to be lower or upper. The tensor is written using square brackets as

```
kernel [ indices separated by commas ]
```

For example, T_a^b would be written as `T[la, ub]`. There is no automatic index raising or lowering with any metric. Before using a tensor, it must be registered with the TensorTools package using the `DefineTensor` function:

```
DefineTensor[T]
```

Entering a tensorial expression causes it to be rendered in an easy to read form:

```
In := T[la, lb]
Out =  $T_{ab}$ 
```

Internally, tensors are represented as `Tensor[kernel, TensorIndex[label, type], ...]` where *label* is the alphabetical index, and *type* is either “u” or “l” depending on the position of the index. This representation helps in pattern matching, and allows TensorTools to identify whether a certain object is a tensor or not.

4.2.3. Expansion of tensorial expressions into components

The function `MakeExplicit` converts an expression containing abstract tensors into a list of component expressions:

```
In := MakeExplicit[T[la, lb]g[ub, uc]]
Out = {g11 T11 + g21 T12 + g31 T13, g12 T11 + g22 T12 + g32 T13,
       g13 T11 + g23 T12 + g33 T13, g11 T21 + g21 T22 + g31 T23,
       g12 T21 + g22 T22 + g32 T23, g13 T21 + g23 T22 + g33 T23,
       g11 T31 + g21 T32 + g31 T33, g12 T31 + g22 T32 + g32 T33,
       g13 T31 + g23 T32 + g33 T33}
```

Note here that upper and lower indices are not distinguished between in the component form. TensorTools was written mainly for automated code generation rather than symbolic manipulation; we suggest using different kernels for the different forms if this is a problem.

4.2.4. Covariant derivatives

TensorTools allows the user to define as many covariant derivatives (and associated Christoffel symbols) as are necessary. The following defines a covariant derivative operator CD with Christoffel symbol H:

```
DefineConnection[CD, H]
```

The function CDtoPD is used to replace covariant derivatives with partial derivatives in any expression:

```
In := CDtoPD[CD[V[ua], lb]]
Out = Va,b + HabcVc
```

The function MakeExplicit will automatically do this before converting expressions into components.

The fundamental operation is to convert a first covariant derivative of an abstract tensor into a partial derivative plus some extra terms involving Christoffel symbols. For convenience, TensorTools understands that covariant derivatives are linear and satisfy the Leibniz property. Also, higher order covariant derivatives are converted to repeated application of a first order derivative.

We define a number of rules to perform the operations in the reduction to the desired form. In the following, x and y represent expressions which may or may not contain tensorial indices.

- Replace any high order covariant derivatives with repeated application of a first order covariant derivative. This ensures that we only need to know how to evaluate a first derivative.

$$\nabla_d \nabla_a V^b \rightarrow \nabla_d (\nabla_a V^b).$$

- Replace the covariant derivative of a product using the Leibniz rule:

$$\nabla_a (xy) \rightarrow (\nabla_a x)y + x(\nabla_a y).$$

- Replace the covariant derivative of a sum using the linearity property:

$$\nabla_a (x + y) \rightarrow \nabla_a x + \nabla_a y.$$

- Replace the covariant derivative of an arbitrary expression containing tensorial indices with its expansion in terms of a partial derivative and Christoffel symbols, one for each index in the expression: e.g.

$$\nabla_a V^b \rightarrow \partial_a V^b + \Gamma^b_{ac} V^c.$$

4.2.5. Lie derivatives

The Lie derivative of an expression x with respect to a vector V is written

```
Lie[x, V]
```

where V has been registered using DefineTensor and is written *without* indices. The function LieToPD is used to replace Lie derivatives with partial derivatives:

```
In := LieToPD[Lie[T[ua, lb], V]]
Out = Tab,cVc + TacVc,b - TcbVa,c
```

Lie derivatives of products and sums are supported. The function MakeExplicit will automatically perform this replacement before converting expressions into components.

4.2.6. Automatic dummy index manipulation

When two expressions both containing a dummy index b are multiplied together, one dummy index is relabeled so as not to conflict with any other index in the resulting expression:

```
In := (T[la, lb]g[ub, uc])v[ub, ld, lb]
Out = TabgbcVede
```

This requires that every multiplication be checked for tensorial operands. This can be a performance problem, so the feature can be enabled and disabled with `SetEnhancedTimes[True]` and `SetEnhancedTimes[False]`. It is enabled by default.

4.3. CodeGen: High level code generation

During the development of the Kranc system, we explored two different approaches to generating Cactus files using Mathematica as a programming language. Initially, we used a very straightforward system whereby C statements were included almost verbatim in the Mathematica script and output directly to the thorn source file. This approach has two main deficiencies:

- The same block of text might be used in several places in the code. When a bug is fixed in one place, it must be fixed in all.
- It is not easy to alter the language that is produced—for example, it is difficult to output both C and Fortran.
- The syntax in the Mathematica source file is ugly, with lots of string concatenation, making it difficult to read and edit.

To address the first problem, we decided to create Mathematica functions to represent each block of text. This also allows the block to be customized by giving the function arguments. By making this abstraction, it became very easy to change between outputting C and Fortran, and the resulting Mathematica scripts are more elegant. We call this system “CodeGen”. The effect of migrating our old code [23] to the CodeGen system was that the Mathematica source files became dramatically shorter and easier to maintain.

Fundamental to the CodeGen system is the notion of a CodeGen “block”; this can be either a string or a list of CodeGen blocks (this definition is recursive). All the CodeGen functions return CodeGen blocks, and the lists are all flattened and the strings concatenated when the final source file is generated. This is because it is syntactically easier in the Mathematica source file to write a sequence of statements as a list than to concatenate strings.

Many programming constructs are naturally block-structured; for example, C `for` loops need braces after the block of code to loop over. For this reason, it was decided that CodeGen functions could take as arguments any blocks of code which needed to be inserted on the inside of such a structure.

By default, Kranc generates C code. To generate Fortran code the command `SetSourceLanguage"Fortran"` has to be issued after loading `KrancThorns.m`. Note that while Kranc-generated C-code should always compile (if different thorns have been created consistently), a Kranc generated Fortran code may not compile in certain cases. The reasons are that the Fortran standard limits names to 31 characters, and specifies a maximum number of 19 continuation lines per statement (typical compilers have much larger continuation line limits). Both limits can be compiler dependent, and are not checked by Kranc.

4.4. Thorn: Constructing the most general Cactus thorn

We have decided to define several different types of Kranc thorn (base, MoL, setter, evaluator, translator) and to provide convenient interfaces for generating each of them. However, there is a lot that is common to each of these, and it would be wasteful for the underlying implementation to have separate code for generating each type of thorn. The Kranc system is designed to be as modular as possible, with as much code re-use as possible. To this end, there is a package called “Thorn” which takes as input a high level description of a Cactus thorn, and which creates the necessary thorn files.

The `interface.ccl`, `schedule.ccl` and `param.ccl` files are first created by calling `CreateInterface`, `CreateParam` and `CreateSchedule` with high-level descriptions of these files, which are converted into CodeGen blocks representing the CCL files. These blocks are then passed, along with blocks for each source file, to `CreateThorn`, which creates the final thorn. The Thorn package provides every aspect of thorn generation that is not specific to the Kranc thorns we have defined, and could be used for creating any type of thorn.

5. Examples

In order to illustrate how the packages `TensorTools.m` and `KrancThorns.m` can be used, we give three examples. The first example is the massive Klein–Gordon equation which, although very simple, contains the essential steps for automatic thorn generation. It is described in some detail to serve as a reference for analogous steps in the other examples. As a second (still trivial) example, we consider the vacuum Maxwell equations. Here the evolution variables are components of tensors and therefore the concept of groups and a `GroupDefinition` structure becomes essential. Finally, as a non-trivial example, we consider the Einstein equations in the standard ADM form. The automatic generation of Cactus thorns for these equations involves (almost) the full functionality of the packages `TensorTools.m` and `KrancThorns.m`.

The following descriptions should be understood as accompanying explanations of the scripts `Examples/KleinGordon/MKGT.m`, `Examples/Maxwell/EMTT.m` and `Examples/ADM/KrancADMTT.m` which actually perform the automatic gen-

eration of code. The reader therefore is strongly advised to read the text along with the example scripts. Detailed instructions for using the examples are provided in the file README.

5.1. The massive Klein–Gordon field

As a simple example we consider the massive Klein–Gordon field equation

$$\square\phi - m^2\phi = 0, \quad (1)$$

where \square is the wave operator on flat space, $\square \equiv -\partial_t^2 + \delta^{ij}\partial_i\partial_j$ in standard Cartesian coordinates (t, x^i) .

5.1.1. Continuum formulation

Introducing $\Pi \equiv \partial_t\phi$ reduces Eq. (1) to a system of PDEs which are first order in time

$$\begin{aligned} \partial_t\phi &= \Pi, \\ \partial_t\Pi &= \delta^{ij}\partial_i\partial_j\phi - m^2\phi. \end{aligned} \quad (2)$$

We consider Eqs. (2) on a 3-torus \mathcal{T} (i.e. on a rectangular domain with periodic boundary conditions), with initial data

$$\phi(0, x^i) = \cos(2\pi x/d), \quad \Pi(0, x^i) = 0, \quad (3)$$

where d is the range of the coordinate x , $-d/2 \leq x \leq d/2$.

As an example of a quantity which is evaluated for analysis purposes, we introduce the total energy $E = \int_{\mathcal{T}} \rho d^3x$, with energy density ρ ,

$$\rho = \Pi^2 + \delta^{ij}\partial_i\phi\partial_j\phi + m^2\phi^2. \quad (4)$$

The energy is conserved in time and can be used to monitor the accuracy of the numerical solution.

5.1.2. Specifications for Kranc

The script `MKGTT.m` in the directory `Examples/KleinGordon` generates Cactus thorns to solve this time evolution problem. After loading the packages `TensorTools.m` and `KrancThorns.m` we specify Eqs. (2) and (4) as rules written in TensorTools syntax; i.e.

```
EvolEqs = {dot[phi] -> phit,
           dot[phit] -> PD[phi, la, lb] KD[ua, ub] - mass^2 phi},
```

where `KD` denotes TensorTools' Kronecker delta and the variable Π has been renamed to `phit` in order to avoid conflicts with the constant π in a Fortran code. The sums over dummy indices are expanded automatically when applying the function `MakeExplicit`,

```
{dot[phi] -> phit,
 dot[phit] -> D11[phi] + D22[phi] + D33[phi] - mass^2*phi}
```

The numerical implementation of the above problem is as follows. The evolution variables `phi`, `phit`, and the parameter `mass` are declared in a base thorn. Initially the grid functions `phi`, `phit` are set to the values (3) by a setter thorn. The right-hand sides of the evolution equations (2) are specified in the MoL thorn. Finally the energy density `rho` (4) is evaluated by an evaluator thorn.

The evolution variables `phi` and `phit` have to be given in form of groups

```
MKGvars = {"phi", {phi}}, {"phit", {phit}};
```

as arguments to the functions `CreateBaseThorn`, `CreateMoLThorn`, `CreateSetterThorn` and should also be given to `CreateEvaluatorThorn`. There is no need to introduce any primitives and therefore we set

```
MKGprimitives = {};
```

which is given as a positional argument to the function `CreateBaseThorn`.

The parameter `mass`

```
MKGRealBaseParameters = {mass};
```

is declared in the base thorn and is used in the MoL thorn and the evaluator thorn. It is listed in the `RealBaseParameters` argument to the functions `CreateBaseThorn`, `CreateMoLThorn` and `CreateEvaluatorThorn`.

Eqs. (2), (3) and (4) and the corresponding grid functions are specified in calculations. For the evolution equations we define

```
MoLCalculation = {Equations -> {{dot[phi] -> phit,
                                dot[phit] -> D11[phi] +...}}};
```

Note that the equations are given as a list of lists (consisting of one element here), each sublist resulting in a separate loop in the code.

The calculations `endensCalculation` for the energy density and `IDCalculation` for the initial data are set up in an analogous way. In addition we assign the name `Name -> "time_symmetricID"` to the function that sets the initial data. The function `"time_symmetricID"` in the Setter thorn should be called only at the beginning, which is achieved by `SetTime -> "initial_only"`.

The function `CreateEvaluatorThorn` takes two positional arguments, `EvaluationDefinitions` and a group structure which should contain the groups to be evaluated by the thorn (as well as all the groups whose variables are referred to). `EvaluationDefinitions` specifies the calculation `endensCalculation` which is responsible for updating the variables in the group `"endens"`,

```
EvaluationGroups = {"endens", {rho}}};
EvaluationDefinitions = {"endens", endensCalculation}};
```

We also specify a name for the arrangement,

```
MKGSystemName = "MKG";
```

Summarizing, the following information is given to the functions `Create*Thorn`:

```
baseThornInfo = CreateBaseThorn[Join[MKGvars, MKGprimitives],
  Map[First, MKGvars],
  Map[First, MKGprimitives],
  SystemName -> MKGSystemName,
  RealBaseParameters -> MKGRealParameters];

mainMoLThornInfo = CreateMoLThorn[MoLCalculation, MKGvars,
  RealBaseParameters -> MKGRealParameters,
  PrimitiveGroups -> Map[First, MKGprimitives],
  SystemName -> MKGSystemName,
  ThornName -> MKGSystemName <> "MoL",
  Debug -> True];

setIDThornInfo = CreateSetterThorn[IDCalculation, MKGvars,
  SystemName -> MKGSystemName,
  ThornName -> MKGSystemName <> "setID",
  SetTime -> "initial_only",
  Debug -> True];

evaluateThornInfo = CreateEvaluatorThorn[EvaluationDefinitions,
  Join[EvaluationGroups, MKGvars],
  RealBaseParameters -> MKGRealParameters,
  SystemName -> MKGSystemName,
  ThornName -> MKGSystemName <> "evalEnergydens",
  Debug -> True];
```

As a last step, a thorn list `MKG.th` is created by the function `CreateThornList`,

```
thorns = Join[baseThornInfo, mainMoLThornInfo, setIDThornInfo, evaluateThornInfo];
CreateThornList[thorns, SystemName -> MKGSystemName];
```

The parameter file for running the code has to activate all necessary thorns. In this case, all the thorns in the file `MKG.th` are necessary. The syntax is

```
ActiveThorns = "Boundary CoordBase..."
```

The parameter `mass` is specified as

```
MKGBase::mass = 1.0
```

We refer the reader to the Cactus thorn documentation for details of the many standard parameters which can be adjusted, for example, to set up the grid in a different way, or to change the MoL time integration scheme. Here we just mention the parameters necessary for changing the resolution and the Courant factor ($\Delta t/\Delta x$). The number of grid points in the x -direction is set by `driver::global_nx`. This number includes both the physical grid points plus the ghost points, which are specified by `pugh::ghost_size`. `grid::dxyz` gives the grid spacing (the same for x , y and z). `cactus::cctk_initial_time` and `cactus::cctk_itlast` set the initial time and the last time step. The Courant factor is specified by `time::dtfac`.

5.2. The Maxwell equations

As a second example we consider the vacuum Maxwell equations on a Minkowski background in standard Cartesian coordinates (t, x^i) ,

$$\partial_\mu F^{\mu\nu} = 0, \quad \partial_{[\mu} F_{\nu\sigma]} = 0, \quad (5)$$

where $F_{\mu\nu}$ is the electromagnetic field strength.

5.2.1. Continuum formulation

An observer at rest with respect to the above coordinates measures an electric and magnetic field

$$E_i = F_{i0}, \quad B_i = \epsilon_i^{kl} F_{kl}, \quad (6)$$

where i, k, l run from 1 to 3 and ϵ_{ijk} is the alternating symbol with $\epsilon_{123} = 1$. Eqs. (5) can be decomposed into the evolution equations

$$\partial_t E_i = \epsilon_i^{jk} \partial_j B_k, \quad \partial_t B_i = -\epsilon_i^{jk} \partial_j E_k \quad (7)$$

and the constraints

$$C_E \equiv \delta^{ij} \partial_i E_j = 0, \quad C_B \equiv \delta^{ij} \partial_i B_j = 0. \quad (8)$$

If the constraints (8) are satisfied initially the evolution equations (7) guarantee that they are satisfied for later times as well. It is therefore sufficient to solve (7) subject to initial data satisfying (8).

Again we take the spatial domain to be a 3-torus \mathcal{T} and consider the initial data

$$\begin{aligned} E_1(0, x^i) &= \sigma \cos\left(\frac{2\pi(x+y)}{d}\right), \\ E_2(0, x^i) &= (\sigma - 1) \cos\left(\frac{2\pi x}{d}\right) - \sigma \cos\left(\frac{2\pi(x+y)}{d}\right), \\ B_3(0, x^i) &= (1 - \sigma) \cos\left(\frac{2\pi x}{d}\right) + \sigma \cos\left(\frac{2\pi(x+y)}{d}\right), \\ E_3(0, x^i) &= B_1(0, x^i) = B_2(0, x^i) = 0, \end{aligned} \quad (9)$$

where $-d/2 \leq x, y \leq d/2$. For $\sigma = 0$ these data evolve as a wave traveling in x direction, $E_2(t, x^i) = -\cos(\frac{2\pi(x+t)}{d}) = -B_3(t, x^i)$, with all other components vanishing. For $\sigma = 1$ the wave evolves in a diagonal direction in the xy plane.

The energy density ρ is given by

$$\rho = \frac{1}{2} \delta^{ij} (E_i E_j + B_i B_j), \quad (10)$$

and gives rise to the total energy $E = \int_{\mathcal{T}} \rho d^3x$, which again is conserved in time.

5.2.2. Specifications for Kranc

The procedure for generating Cactus thorns is similar to the previous example, the major difference being that the variables are components of tensors and Eqs. (7) have to be split into equations for the individual components.

The automatic generation of code is performed by `Examples/Maxwell/EMTT.m`. We define the tensors `E1` and `B` with `TensorTools'` function `DefineTensor`, where the electric field has been renamed in order to avoid a name clash with the exponential constant in `Mathematica`. The evolution equations (7) in `TensorTools` syntax read

```
EvolEqs = {{dot[E1[1a]] -> (Eps[1a,1b,1c] KD[ub,ue] KD[uc,uf] PD[B[1f],1e]),
            dot[B[1a]] -> -(Eps[1a,1b,1c] KD[ub,ue] KD[uc,uf] PD[E1[1f],1e])}};
```

These equations are split into individual components by applying the function `MakeExplicit`.

As before, we generate a base thorn, a MoL thorn, a setter thorn for the initial data and an evaluator thorn for evaluating the constraints and energy density.

The evolution variables are the electric and magnetic field

```
EMvars = {"E1", {E11, E12, E13}}, {"B", {B1, B2, B3}};
```

As before, we do not introduce any primitives. The parameter `sigma` entering the initial data (9) only has to be known by the setter thorn. It is therefore given as the named argument `RealParameters` to the function `CreateSetterThorn`.

We define two groups of variables to be set by the evaluator thorn

```
EvaluationGroups = {"Constraints", {CE1, CB}},
{"endens", {rho}};
```

and correspondingly two calculations that are responsible for evaluating the constraints and the energy density.

5.3. The Einstein equations in ADM form

As a last example we consider the Einstein equations in ADM form, which provide a non-trivial example that illustrates the advantages of automatic code generation. A mere 200 lines of Mathematica input is needed to generate a whole evolution code including the evaluation of the constraints. Compared to the previous example, we make use of more features of `TensorTools`, for example defining a covariant derivative and dealing with tensors with symmetries. The code generation part illustrates the concept of primitive grid functions and the possibility of generating two different setter thorns to compute the same quantity in different ways. Further, it explains how grid functions are computed that depend on derivatives of other grid functions.

5.3.1. Continuum formulation

We decompose the four-dimensional vacuum Einstein equations $G_{\mu\nu} = 0$ into a set of evolution equations and constraints corresponding to the standard ADM formulation given by York [24]. We will not touch upon issues of well-posedness, as these are irrelevant for our current purpose. Space–time is foliated by a family of spacelike hypersurfaces Σ_t , labeled by a parameter t . Introducing coordinates (t, x^i) adapted to this foliation, the four-dimensional metric can be written as

$$ds^2 = -(\alpha^2 - \beta_i \beta^i) dt^2 - 2\beta_i dt dx^i + h_{ij} dx^i dx^j, \quad (11)$$

where h_{ij} is the induced metric on Σ_t , α is the lapse function and β^i the shift vector and where we set $\beta_i = h_{ij} \beta^j$. Lapse and shift describe the components of the time flow ∂_t orthogonal and parallel to the hypersurface Σ_t , $\partial_t = \alpha n + \beta^i \partial_i$. The rate of change of the induced metric h_{ij} along the normal direction n defines the extrinsic curvature $K_{ij} = \frac{1}{2} \mathcal{L}_n h_{ij}$.

Projections of the vacuum Einstein equations in the form ${}^{(4)}R_{\mu\nu} = 0$ into the hypersurface Σ_t , together with the definition of the extrinsic curvature, yield evolution equations for h_{ij} and K_{ij} ,

$$\partial_t h_{ij} = \mathcal{L}_\beta h_{ij} + 2\alpha K_{ij}, \quad (12)$$

$$\partial_t K_{ij} = \mathcal{L}_\beta K_{ij} + D_i D_j \alpha + 2\alpha h^{lm} K_{lj} K_{mi} - \alpha K K_{ij} - \alpha R_{ij}. \quad (13)$$

Projection of $G_{\mu\nu} = 0$ in the orthogonal direction yields the Hamiltonian and momentum constraints,

$$\begin{aligned} C_H &\equiv R - h^{lm} h^{pq} K_{lp} K_{mq} + K^2 = 0, \\ C_{Mi} &\equiv h^{lm} D_l K_{im} - D_i K = 0. \end{aligned} \quad (14)$$

\mathcal{L}_β denotes the Lie derivative in the direction of the shift vector β^i and D_i is the covariant derivative compatible with the 3-metric h_{ij} . The three-dimensional scalar curvature R , Christoffel symbols $\gamma^i{}_{jk}$ and the trace of the extrinsic curvature K have been introduced as abbreviations,

$$\begin{aligned} K &= h^{ij} K_{ij}, \\ \gamma^i{}_{jk} &= h^{il} (\partial_k h_{lj} + \partial_j h_{kl} - \partial_l h_{jk})/2, \\ R &= h^{ij} R_{ij}. \end{aligned} \quad (15)$$

Again the constraints (14) are propagated by the evolution equations (12) and it is sufficient to evolve freely (i.e. without imposing the constraints at each time step) initial data satisfying (12).

We express the Ricci tensor R_{ij} in two ways, one involving second derivatives of the metric

$$R_{ij} = h^{lm} (-\partial_i \partial_j h_{lm} + \partial_j \partial_m h_{li} + \partial_i \partial_l h_{mj} - \partial_l \partial_m h_{ij})/2 + h_{lm} h^{pq} (\gamma_{mi}^l \gamma_{qj}^p - \gamma_{ij}^l \gamma_{mq}^p), \quad (16)$$

and the other is the standard representation in terms of Christoffel symbols

$$R_{ij} = \partial_p \gamma_{ij}^p - \partial_j \gamma_{pi}^p + \gamma_{ij}^p \gamma_{pq}^q - \gamma_{qi}^p \gamma_{pj}^q. \quad (17)$$

Note that Eqs. (16) and (17) lead to different stencil sizes when discretized, which might change the behavior of the numerical evolution code.

For our choice of gauge, initial data and boundary conditions, we follow the “gauge wave testbed” described in [25]. The spatial domain is again taken to be a 3-torus \mathcal{T} and initial data are given by

$$\begin{aligned} h_{11}(0, x^i) &= 1 - A \sin\left(\frac{2\pi x}{d}\right), \\ K_{11}(0, x^i) &= \frac{\pi A}{d} \frac{\cos\left(\frac{2\pi x}{d}\right)}{\sqrt{1 - A \sin\left(\frac{2\pi x}{d}\right)}} \end{aligned} \quad (18)$$

with all other components of the extrinsic curvature and all off-diagonal 3-metric components vanishing initially. The remaining diagonal metric components are 1 initially. A is a constant amplitude and d again denotes the range of the coordinate x . The shift β^i is set to zero and the lapse α is given by the square root of the determinant of the 3-metric $\alpha = \sqrt{\det h}$.

The analytical solution to this initial value problem is a pure gauge wave traveling in the x direction at the speed of light, $h_{11}(t, x^i) = h_{11}(x - t)$, $K_{11}(t, x^i) = K_{11}(x - t)$, where all other evolution variables maintain their initial values.

5.3.2. Specifications for Kranc

The conversion of tensorial equations to equations for individual components using the package TensorTools clearly is more involved than in the previous examples. We start by defining the tensors appearing in Eqs. (12)–(15)

```
Map[DefineTensor, {h, hInv, K, beta, Ricci, gamma, CM}]
```

The inverse metric h^{ij} has been renamed to `hInv` as it is going to be a shorthand computed from the metric h_{ij} . `Ricci` and `gamma` denote the Ricci tensor and the Christoffel symbols respectively, and `beta` is the shift.

Note that since the tensors h_{ij} and K_{ij} are symmetric, we only evolve a subset of linearly independent components. To this effect, before applying `MakeExplicit` to expressions involving symmetric tensors, the symmetries are registered. Note that, in conformance with the conventions of `MathTensor`, in our example covariant and contravariant indices are treated differently with respect to the ordering of coordinate indices. The “12” component of a covariant symmetric 2-tensor automatically is converted to the “21” component, whereas for contravariant symmetric tensors the index ordering is the opposite. The corresponding commands for symmetric (co/contravariant) tensors are `AssertSymmetricDecreasing` and `AssertSymmetricIncreasing`,

```
AssertSymmetricIncreasing[hInv[ua, ub]];
Map[AssertSymmetricDecreasing, {h[la, lb], K[la, lb], Ricci[la, lb]}];
AssertSymmetricDecreasing[gamma[ua, lb, lc], lb, lc];
```

The symmetry handling in `TensorTools` is rudimentary, and currently only supports symmetries for pairs of indices in any particular tensor. When applying `MakeExplicit` to one of the above tensors all components are kept thus resulting in a list with duplicate entries. These have to be removed by a user defined function which is called `MakeExplicitSet` here.

To complete the description of the geometric structure for `TensorTools` the covariant derivative together with the corresponding Christoffel symbol is defined,

```
DefineConnection[CD, gamma];
```

Computing the determinant of the spatial metric h by the Mathematica component expression `hDet = Det[MatrixOfComponents[h[la, lb]]]` the inverse metric is given by a series of shorthands

```
deth      -> hDet,
invdeth   -> 1 / deth,
hInv[ua, ub] -> invdeth hDet MatrixInverse[h[ua, ub]]
```

Here, `deth` is computed as a shorthand at run time to avoid it being calculated for each component of `hInv[ua, ub]`. `hDet MatrixInverse[h[ua, ub]]` is the component inverse of h with the common determinant factor cancelled.

The quantities appearing in Eqs. (12) and (14) split in the following way into evolution variables, primitives and shorthands: Evolution variables are the 3-metric and the extrinsic curvature

```
ADMvars = {"h", {h11, h21, h31, h22, h32, h33}}, {"K", {K11, K21, K31, K22, K32, K33}}};
```

All other quantities appearing in (12) and (14) have to be either primitive grid functions, which are set in a setter thorn, or shorthands defined inside the loop in which they are used. The decision about which quantities are primitives is based on the following considerations. Quantities that appear in differentiated form in the evolution equations, e.g., lapse and shift, have to be grid functions. Quantities which can be computed in various different ways, e.g., the Ricci tensor, can be set in different setter thorns which then can be compiled into the same executable, where one of them is “activated” at run time. Quantities that are needed for both the evolution equations and the constraints, e.g., the Christoffel symbols, can be set in a setter thorn in order to reduce computational costs. Furthermore all quantities whose output is required, e.g., the trace of the extrinsic curvature as specified in the tests in [25], have to be grid functions. All of these quantities should be declared as primitives. The complete list of primitives for our example therefore is

```
ADMprimitives = {"alpha", {alpha}}, {"beta", {beta1, ...}}, {"Ricci", {Ricci11, ...}},
                {"gamma", {gamma111, ...}}, {"trK", {trK}}
```

The scalar curvature is computed as a shorthand from the Ricci tensor.

Eqs. (12)–(17) in `MathTensor` syntax are read from the file `ADMEqs.m` and expanded into components with `MakeExplicit`. In order to increase performance once the functionality of `TensorTools` is not needed anymore we switch off the special treatment of products by setting

```
SetEnhancedTimes[False]
```

As a last preparational step we define the components of the inverse metric and the lapse in a `CollectList` which is going to be used to simplify the equations at the level of components, essentially by applying the `Mathematica` function `Collect[rhs, CollectList, Simplify]` to the right-hand sides of assignments.

The generation of the base thorn, the MoL thorn, the various setter thorns and the evaluator thorn run along the same lines as in the previous examples with the following additional features. All primitive grid functions have to be set at every time step, so the argument `SetTime` for all setter thorns therefore has to be `"initial_and_poststep"`.

The Ricci tensor in the form Eq. (17) requires taking derivatives of the Christoffel symbols, which therefore have to be known on the whole grid before the Ricci tensor is computed. This is done in the code by using two loops, where the first loop is used to compute the Christoffel symbols. At the end of the loop the grid functions that store the Christoffel symbols are synchronized such that the ghost zones are filled in with data from the opposite side of the grid. The Ricci tensor can then be computed in the second loop. In order to generate two separate loops, the calculation passed as an argument to the function `CreateSetterThorn` has to specify the equations in two separate (sub)lists, i.e.

```
Equations -> {Join[InvMetricEqs, ChristoffelEqs], RicciStandardEqs};
```

Another major difference to the previous examples is the way the initial data is set. We use the existing thorn `Exact` in the arrangement `AEIThorns` to compute the metric and extrinsic curvature as given in Eq. (18). The only thing left to do is to convert the initial data from the variables used by the `Exact` thorn, `{gxx, ..., kxx, ...}`, defined by the `CactusEinstein` arrangement, to our evolution variables and correct for different sign conventions in the definition of the extrinsic curvature, i.e. `h11 == gxx, ..., K11 == - kxx, ...`. This is done in the `Translator` thorn, which implements functions to translate either way: from the `CactusEinstein` variables to the evolution variables at the `POSTINITIAL` `Cactus` time bin, and from the evolution variables to the `CactusEinstein` variables after every time step. The latter translation is performed to allow use of existing analysis tools, e.g., in `CactusEinstein`. The translator functions are defined by the Kranc calculations `TranslatorInCalculation` and `TranslatorOutCalculation`, respectively.

The thorn list is generated as in the other examples. This time, however, it has to be supplemented by the thorns `CactusEinstein/ADMCoupling`, `AEIThorns/Exact`, `CactusEinstein/CoordGauge` and `CactusEinstein/StaticConformal` which are needed by the `Exact` thorn to set up the initial data Eqs. (18).

All the thorns contained in the thorn list—including both Setter thorns for the Ricci tensor—are compiled into one executable. The parameter file specifies the way the Ricci tensor is computed by activating the relevant thorn. The specification of the parameters determining the initial data can be found in the documentation of the thorn `AEIThorns/Exact`.

5.3.3. Performance

We include some timing results for simulations to demonstrate that our computer-generated code can achieve a performance that is comparable to carefully optimized handwritten code, which again is based on the `Cactus` computational toolkit. We do not attempt a systematic study of performance issues, and restrict ourselves to a single example based on the “gauge wave” data as defined in Eq. (18). Our test run is a 3D version of the gauge wave spacetime defined with a grid size of 50^3 (plus one ghost zone); we evolve for 100 time steps. In our experience, such simple tests already give a reasonable estimate on the efficiency of a finite difference evolution code which is based on the method of lines as implemented within the `Cactus` framework. A detailed

Table 1

Timing results from a Dell D600 Laptop running Redhat Linux 9.0, with a 1.6 GHz Intel Pentium M CPU and 1 GB of RAM, comparing results obtained with the Intel 8 compilers vs. a combination of Pacific Sierra VAST F90 and gcc 3.2.2 compilers. Optimization options for the Intel compilers are `-O3 -xN -ip`, and `-O3` for VAST and gcc

Code version	Compile time	User runtime	Cactus timing
CactusEinstein ADM, Intel 8	184	145.2	145.6
CactusEinstein ADM, VAST f90 & gcc	154	155.0	155.2
Kranc ADM, C, Intel 8	119	154.9	155.4
Kranc ADM, C, VAST f90 & gcc	89	166.8	167.4
Kranc ADM, F90, Intel 8	129	159.7	159.5
Kranc ADM, F90, VAST f90 & gcc	92	168.3	168.8

Table 2

Timing results from a Dell Precision 450 workstation running Redhat Linux 7.2, with a 1.7 GHz Intel Pentium Xeon CPU and 1 GB of RAM, binaries have been built with the Intel 7.1 compilers with optimization options `-O3 -xW -ip`

Code version	Compile time	User runtime	Cactus timing
CactusEinstein ADM, Intel 7.1	303	188.1	189.5
Kranc ADM, C, Intel 7.1	186	131.3	131.1
Kranc ADM, F90, Intel 7.1	194	115.3	115.7

performance analysis, which takes into account various effects of memory handling and parallelization would go beyond the scope of this article, and work on optimizing codes within the Kranc approach will be reported elsewhere.

We compare the performance of the ADM code presented above to a code that has been developed as a production code by the AEI astrophysical numerical relativity group. This code is publicly available from the Cactus CVS server as the CactusEinstein arrangement [26]. The Kranc and CactusEinstein ADM codes use the same numerical methods (3 step Iterative Crank–Nicolson and second order centred spatial finite differencing). As a physics production code, CactusEinstein ADM is significantly more complex, but the evolution algorithm is essentially equivalent. The numbers given below should only be taken as a rough guide. In contrast to Kranc ADM, which uses computer generated code and a generic infrastructure for spatial differencing and time evolution, CactusEinstein ADM is hand coded and does not support general numerical schemes (although different time update schemes have been implemented as well).

In Tables 1 and 2, *compile time* and *user runtime* have been timed with Linux's time utility and show the total number of CPU-seconds that the process spent in user mode. The results of the *cactus timing* column have been obtained with the internal timing infrastructure of Cactus.

Tests have been performed on two machines, a Dell D600 Laptop running Redhat Linux 9.0, with a 1.6 GHz Intel Pentium M CPU and 1 GB of RAM; and a Dell precision 450 workstation running Redhat Linux 7.2, with a dual Intel Pentium Xeon CPU running at 1.7 GHz and again 1 GB of RAM. All runs have been performed as single processor runs.

Note that with the exception of the result for CactusEinstein ADM compiled with Intel 7.1, the Kranc generated code is around 6–10% slower. Possible explanations are that the CactusEinstein ADM code partially uses Fortran 90 array language which allows further optimizations (e.g. the vectorization feature of the Intel compiler), and the scheme for synchronizing grid functions within Kranc is not yet optimal (synchronizing is not necessary if no derivatives appear in assignments, as the ghost zones can be updated locally). We speculate that the CactusEinstein result for the Intel 7.1 compiler might be caused by the compiler reaching internal limits and giving up on certain kinds of optimization.

Acknowledgements

The authors thank the members of the AEI numerical relativity group and the Cactus developers for helpful discussions, in particular D. Pollney and I. Hawke for many stimulating discussions, I. Hawke for helping with MoL, E. Schnetter for resolving subtleties with preprocessing and compiling Kranc-generated code, J. Thornburg for careful reading of an earlier version of the manuscript, and K. Roszkowski and N. Dorband for feedback about using Kranc. I. Hinder thanks AEI for hospitality.

Appendix A. Data structure specifications

Here we describe in detail the data structures which are used when calling the KrancThorns functions.

A.1. Calculation

Key	Type	Description
Equations	list of lists	{loop1, loop2}—Each loop is a list of rules of the form <i>variable</i> → <i>expression</i> where <i>variable</i> is to be set from <i>expression</i>
Shorthands (optional)	list of symbols	Variables which are to be considered as ‘shorthands’ for the purposes of this calculation
Name (optional)	string	A name for the calculation
Before (optional)	list of strings	Function names before which the calculation should be scheduled
After (optional)	list of strings	Function names after which the calculation should be scheduled

A.2. GroupCalculation

A GroupCalculation structure is a list of two elements; the first is the name (a string) of a Cactus group and the second is the Calculation to update the variables in that group.

A.3. GroupDefinition

A GroupDefinition structure is a list of two elements. The first is the name (string) of a Cactus group and the second is the list of variables (symbols) belonging to that group. The group name can be prefixed with the name of the Cactus implementation that provides the group followed by two colons (e.g. “ADMBase::metric”). If this is not done, then the KrancThorns functions will attempt to guess the implementation name, usually using the name of the thorn being created.

Appendix B. KrancThorns function reference

Here we document the arguments which can be specified for the functions CreateBaseThorn, CreateMolThorn, CreateSetterThorn, CreateTranslatorThorn and CreateEvaluatorThorn.

Note that we use Mathematica syntax for function-specific section headers. Underscores denote function arguments, and Opt-Arguments stands for optional arguments, also referred to as named arguments below. These are given in the form myFunction[... , argumentName → argumentValue].

B.1. Common named arguments

The following named arguments can be used in any of the Create*Thorn functions:

Argument	Type	Description	Default
SystemName	string	A name for the evolution system implemented by this arrangement. This will be used for the name of the arrangement directory	“MyPDESystem”
SystemParentDirectory	string	The directory in which to create the arrangement directory	“.”
ThornName	string	The name to give this thorn	SystemName + thorn type
Implementation	string	The name of the Cactus implementation that this thorn defines	ThornName
SystemDescription	string	A short description of the system implemented by this arrangement	SystemName
InheritedImplementations	list of strings	A list of implementations to inherit from	{}
DeBug	Boolean	Whether or not to print debugging information during thorn generation	False

B.2. Arguments relating to parameters

The following table describes named arguments that can be specified for any of the thorns except CreateBaseThorn. CreateBaseThorn is special because it can be used to define parameters which are inherited by each thorn in the arrangement, so the arguments it can be given are slightly different.

Argument	Type	Description	Default
RealBaseParameters	list of strings	Real parameters used in this thorn but defined in the base thorn	{}
IntBaseParameters	list of strings	Integer parameters used in this thorn but defined in the base thorn	{}
RealParameters	list of strings	Real parameters defined in this thorn	{}
IntParameters	list of strings	Integer parameters defined in this thorn	{}

B.3. CreateBaseThorn[groups_, evolvedGroupNames_, primitiveGroupNames_, OptArguments___]

B.3.1. Positional arguments

Argument	Type	Description
groups	list of GroupDefinition structures	Definitions of any groups referred to in the other arguments. Can supply extra definitions for other groups which will be safely ignored
evolvedGroupNames	list of strings	Names of groups containing grid functions which will be evolved by MoL in any of the thorns in the arrangement
primitiveGroupNames	list of strings	Names of groups containing grid functions which will be referred to during calculation of the MoL right hand sides in any of the thorns in the arrangement

B.3.2. Named arguments

Argument	Type	Description	Default
RealBaseParameters	list of strings	Real parameters defined in this thorn and inherited by all the thorns in the arrangement	{}
IntBaseParameters	list of strings	Integer parameters defined in this thorn and inherited by all the thorns in the arrangement	{}

B.4. CreateEvaluatorThorn[groupCalculations_, groups_, OptArguments___]

B.4.1. Positional arguments

Argument	Type	Description
groupCalculations	list of GroupCalculation structures	The GroupCalculations to evaluate in order to set the variables in each group
groups	list of GroupDefinition structures	Definitions for each of the groups referred to in this thorn. Can supply extra definitions for other groups which will be safely ignored

B.5. CreateMoLThorn[calculation_, groups_, OptArguments___]

B.5.1. Positional arguments

Argument	Type	Description
calculation	Calculation	The calculation for setting the right hand side variables for MoL. The equations should be of the form $\text{dot}[gf] \rightarrow \text{expression}$ for evolution equations, and <i>shorthand</i> $\rightarrow \text{expression}$ for shorthand definitions, which can be freely mixed in to the list
groups	list of GroupDefinition structures	Definitions for each of the groups referred to in this thorn. Can supply extra definitions for other groups which will be safely ignored

B.5.2. Named arguments

Argument	Type	Description	Default
PrimitiveGroups	list of strings	These are the groups containing the grid functions which are referred to but not evolved by this evolution thorn	{}

B.6. CreateSetterThorn[calculation_, OptArguments___]

B.6.1. Positional arguments

Argument	Type	Description
calculation	Calculation	The calculation to be performed
groups	list of GroupDefinition structures	Definitions for each of the groups referred to in this thorn. Can supply extra definitions for other groups which will be safely ignored

B.6.2. Named arguments

Argument	Type	Description	Default
SetTime (optional)	string	“initial_and_poststep”, “initial_only” or “poststep_only”	“initial_and_poststep”

B.7. CreateTranslatorThorn[groups_, OptArguments_...]

B.7.1. Positional arguments

Argument	Type	Description
groups	list of GroupDefinition structures	Definitions for each of the groups referred to in this thorn. Can supply extra definitions for other groups which will be safely ignored

B.7.2. Named arguments

Argument	Type	Description
TranslatorInCalculation	Calculation	The calculation to set the evolved variables from some other source
TranslatorOutCalculation	Calculation	The calculation to convert the evolved variables back into some other set of variables

References

- [1] For an overview of current issues and PDE problems arising in numerical relativity, and for further references, see, e.g., L. Lehner, *Class. Quant. Grav.* 18 (2001) R25.
- [2] E. Schnetter, Gauge fixing for the simulation of black hole spacetimes, PhD thesis, Universität Tübingen, 2003, <http://w210.ub.uni-tuebingen.de/dbt/volltexte/2003/819/>.
- [3] G.O. Cook Jr., ALPAL, a program to generate physics simulation codes from natural descriptions, *Internat. J. Modern Phys. C* 1 (1990) 1.
- [4] G.O. Cook Jr., J.F. Painter, S.A. Brown, How symbolic computation boosts productivity in the simulation of partial differential equations, *J. Sci. Comput.* 6 (2) (1991) 193–209.
- [5] T. Ruppelt, G. Wirtz, Automatic transformation of high-level object-oriented specifications into parallel programs, *Parallel Comput.* 10 (1989) 15.
- [6] J. Thornburg, A PDE compiler for full-metric numerical relativity, in: *Frontiers in Numerical Relativity*, Cambridge University Press, 1989, p. 370.
- [7] M. Pasquali, L.E. Scriven, Free surface flows of polymer solutions with models based on the conformation tensor, *J. Non-Newtonian Fluid Mech.* 108 (2002) 363–409.
- [8] M. Pasquali, Polymer molecules in free surface coating flows, PhD Thesis, University of Minnesota, 2000.
- [9] P. Musgrave, D. Pollney, K. Lake, *Fields Institute Comm.* 15 (1996) 313, <http://grtensor.phy.queensu.ca>.
- [10] R. Wald, *General Relativity*, University of Chicago Press, 1984.
- [11] R. Penrose, W. Rindler, *Spinors and Space-Time: Two-Spinor Calculus and Relativistic Fields*, Cambridge University Press, 1984.
- [12] J. Lee, Ricci, a Mathematica package for doing tensor calculations in differential geometry, <http://www.math.washington.edu/~lee>.
- [13] J.M. Martín-García, xTensor, <http://metric.imaff.csic.es/Martin-Garcia/xAct/>.
- [14] L. Parker, S. Christensen, *MathTensor: A System for Doing Tensor Analysis by Computer*, Addison-Wesley, 1994.
- [15] The MPI Forum, <http://www-unix.mcs.anl.gov/mpi>.
- [16] T. Goodale, G. Allen, G. Lanfermann, J. Massó, T. Radke, E. Seidel, J. Shalf, The Cactus framework and toolkit: Design and applications, in: *Vector and Parallel Processing—VECPAR'2002*, 5th International Conference, in: *Lecture Notes in Computer Science*, Springer, Berlin, 2003; G. Allen, et al., *Cluster Comput.* 4 (2001) 179; B. Talbot, S. Zhou, G. Higgins, NASA report, http://ct.gsfc.nasa.gov/esmf_tasc/Files/Cactus_b.html; Cactus development team, <http://www.cactuscode.org>.
- [17] Portable, Extensible toolkit for scientific computation, <http://www-unix.mcs.anl.gov/petsc/petsc-2>.
- [18] J. Thornburg, *Class. Quant. Grav.* 21 (2) (2004) 743–766.
- [19] P. Diener, *Class. Quant. Grav.* 20 (2003) 4901–4918.
- [20] B. Gustafsson, H. Kreiss, J. Olinger, *Time Dependent Problems and Difference Methods*, John Wiley and Sons, 1995.
- [21] I. Hawke, for a short description and download instructions see <http://www.aei.mpg.de/~hawke/MoL/MoL.html>.
- [22] M. Alcubierre, S. Brandt, B. Bruegmann, D. Holz, E. Seidel, R. Takahashi, J. Thornburg, *Internat. J. Modern Phys. D* 10 (2001) 273–290.
- [23] S. Husa, C. Lechner, in: A. Lobo, et al. (Eds.), *Proceedings of the Spanish Relativity Meeting (ERE 2002)*, Maó, Menorca, Spain, 22–24 September 2002, Barcelona University Press, 2003, e-Print archive gr-qc/0301076.
- [24] J.W. York Jr., Kinematics and dynamics of general relativity, in: L. Smarr (Ed.), *Sources of Gravitational Radiation*, Cambridge University Press, 1979.
- [25] M. Alcubierre, et al., *Class. Quant. Grav.* 21 (2004) 589–613.
- [26] See <http://www.cactuscode.org/Community/Relativity.html>.