



Cactus



Einstein Toolkit

Beyond the Algorithm: Supporting Infrastructure for Large Scale Simulation Codes

Ian Hinder



Max Planck Institute for Gravitational Physics
(Albert Einstein Institute)
Potsdam, Germany

NCSA Software Seminar, Urbana-Champaign, July 2016



Overview

- The Einstein Toolkit and Cactus
- Automatic code generation
- Abstracting the machine
- Reproducibility
- Software quality control

The Einstein Toolkit and Cactus: Background


- Open code for **Numerical Relativity**
- einstein toolkit.org
- Origin: **Ed Seidel's** group at AEI: binary black hole problem, 1995–
- **100** registered users in **56** different groups worldwide
- Einstein equations:
 - complicated **partial differential equations** solved with **finite difference methods**
- Based on **Cactus** framework
- OpenMP/MPI **parallelism** from 100s to 1000s of cores
- **Framework** vs library
 - large number of **components** ("**thorns**") plugged together
 - communication via well-defined **simple interfaces**

Binary black hole merger and gravitational waves



Cactus framework

- Cactus modules called **thorns**, all talk to the **flesh**



cactuscode.org

- Each thorn has:
 - **Executable**
 - **Metadata** files (interfaces, parameters, scheduling)
 - **Source** files (C/C++/Fortran: physics equations, algorithms, infrastructure)
- The flesh:
 - Defines APIs for **communication** between thorns
- Intent: many groups can **independently develop** public and private codes which all **work together**

The Einstein Toolkit

- **Cactus:**
 - flesh
 - support thorns
- **Physics** thorns
- **Mesh refinement** thorns
- **Numerical methods** thorns (interpolator, time integrator, etc)
- **Infrastructure** thorns (3D output, base interfaces, termination management etc)
- **226** Cactus thorns in total
- **Kranc** automatic code generator: generates Cactus thorns
- **SimFactory**: manage simulations/compilation across diverse HPC machines

Automatic code generation



- Sascha Husa, IH, Christiane Lechner, 2004
- **High level** description of equations, including tensorial
- "Compiled" to **complete Cactus thorn**
- Application developer sees **equations**, not code
- Almost all Cactus **boilerplate** hidden
- Implemented in **Mathematica**
- kranccode.org

Equation
script

Kranc

Cactus
thorn

Compiler

Executable

Cluster

Results

Automatic code generation:

Example wave equation

```
initialSineCalc =  
{  
  Name -> "initial_sine",  
  Schedule -> {"AT INITIAL"},  
  Equations ->  
  {  
    phi -> Sin[2 Pi (x - t)],  
    pi -> -2 Pi Cos[2 Pi (x - t)]  
  }  
};
```

```
evolveCalc =  
{  
  Name -> "calc_rhs",  
  Schedule -> {"in MoL_CalcRHS"},  
  Equations ->  
  {  
    dot[phi] -> pi,  
    dot[pi] -> Euc[ui,uj] PD[phi,li,lj]  
  }  
};
```

```
CreateKrancThornTT[groups, ".", "SimpleWave",  
  Calculations -> {initialSineCalc, evolveCalc},  
  PartialDerivatives -> derivatives, DeclaredGroups -> {"evolved_group"}];
```


Automatic code generation

- Solve **time evolution** PDEs in 3D
- **High performance** parallel codes
- End users can treat as **black box**
- Arbitrary order finite differencing
- Existing codes benefit from **new Kranc features**



Automatic code generation: Features

- Arbitrary order finite differencing
- **Mesh refinement** and **multi-block** grids
- **OpenMP**
- High level optimisations:
 - Common **subexpression** elimination
 - **Loop** splitting and joining
- Floating point instruction **vectorisation**: generate compiler intrinsics for all operations
- Experimental support for **CUDA/OpenCL**, as well as **Xeon Phi**
- Can implement many of these algorithms at a very **high level** (in Mathematica)



Abstracting the machine: The Simulation Factory



- **Manage simulations:** uniform interface across supercomputers
- by **Erik Schnetter**
- Hide **low-level** cluster-specific details
- Machine database: many **XSEDE** and institution clusters
- Enforce/encourage best practices and avoid common **mistakes**
- **New version** under development (IH, Barry Wardell, Erik Schnetter)
 - **Code-agnostic**; not specific to Cactus
 - Want to try it out with **other codes**
 - **Mostly working**
 - <http://simfactory.org>

```
sim setup
```

```
sim build --thornlist  
thornlists/mythorns.th
```

```
sim submit mysim --  
parfile par/mysim.par  
--procs 128 --walltime  
12:00:00
```

```
sim list-simulations
```

```
sim stop mysim
```

Abstracting the machine: Source code and building



- Centralisation:
 - *"Which machine did I fix that bug on?"*
 - Keep source tree in single **central location**
 - **Sync** code to remote clusters (rsync)
- Building:
 - Database of Cactus "**optionlists**" for each machine
 - **Environment** setup commands: module etc

```
sim sync bluewaters
```

```
sim --remote bluewaters build
```

Abstracting the machine: Simulations



- Submitting a simulation to the queue:
 - **Submit scripts** for each machine in machine database
 - Specify parameter file, number of cores, walltime
 - Also: undersubscribing, OpenMP threads, more

```
sim --remote bluewaters submit mysim  
    --parfile par/mysim.par --procs 128 --walltime 12:00:00
```

- Simulation lifecycle management: simulation **states**:
active(running,queued)/inactive

Abstracting the machine: The Simulation Factory



- Long simulations split into **segments**
- One segment per **job walltime**, e.g. 24 h
- Best practice:
 - No job should **overwrite data** from a previous job
- Checkpoint files hardlinked between segments
- Simulations automatically **submit the next segment** (new version)
- **Termination conditions** defined by regexps
 - FinalTime, EndOfWalltime, DiskQuotaExceeded, UnknownError
- **Termination actions**: Continue, Error, Email
- sim pause, sim continue - request immediate **checkpoint**/restart

Reproducibility

- Copy of **source code** preserved in every simulation (tar.gz archive)
 - Always know **what code was run**, even 10 years later
 - But: difficult to relate to **version controlled** commits
- Multiple components: **multiple repositories**
 - Difficult to identify a single **revision** of "the code"
 - Experimental use of **git submodules** to pull everything together



Releases

- Every **6 months**
- **Run tests** on all supported HPC systems, track down and fix problems
- Commit to **backporting** serious issues to last release
- "**Known good**" version that people can use

Testing

- Cactus has standard mechanism for **test cases**
- Mostly **regression tests**: "does this parameter file lead to the same results as the reference data?"
- **Problem**: developers don't run them
- **Solution**: Tests run **after each commit** on a central server
- **Jenkins** web application to manage
- Distributed build nodes
- Integrating with HPC systems is a problem

Testing: Jenkins web application

The screenshot shows the Jenkins web application interface. The browser window title is "EinsteinToolkit [Jenkins]" and the address bar shows "https://build.barrywardell.net". The Jenkins logo and name are at the top left, with a search bar and "log in | sign up" links at the top right. Below the header, there are tabs for "All", "Chemora", "Dashboard", "EinsteinToolkit", "EinsteinToolkitMulti", and "EinsteinToolkitSet". The "EinsteinToolkit" tab is selected, showing a table of build jobs. The table has columns for "S", "W", "Name", "Last Success", "Last Failure", "Last Duration", and "Test Result". The jobs listed are "EinsteinToolkit", "EinsteinToolkitDoc", "EinsteinToolkitFull", and "EinsteinToolkitReleased". On the left side, there are links for "People", "Build History", "Claim Report", "Disk usage", and "Job Import Plugin". Below these, there is a "Build Queue" section showing "No builds in the queue." and a "Build Executor Status" section showing the status of various executors: "master" (1 idle), "AEI" (1 idle), "cct" (1 idle), "perimeter1" (offline), and "ucd" (offline). At the bottom right, there is a "Legend" section with links for "RSS for all", "RSS for failures", and "RSS for just latest builds".

People
Build History
Claim Report
Disk usage
Job Import Plugin

Build Queue
No builds in the queue.

Build Executor Status

- master
1 idle
- AEI
1 idle
- cct
1 idle
- perimeter1 (offline)
- ucd (offline)

S	W	Name	Last Success	Last Failure	Last Duration	Test Result
		EinsteinToolkit	19 hr - #828	4 days 1 hr - #825	1 hr 7 min	(no failures)
		EinsteinToolkitDoc	19 hr - #423	2 mo 8 days - #371	18 min	N/A
		EinsteinToolkitFull	19 hr - #279	4 days 1 hr - #275	31 min	(no failures)
		EinsteinToolkitReleased	1 mo 8 days - #101	1 mo 8 days - #99	1 hr 30 min	(no failures)

Icon: [S](#) [M](#) [L](#)

Legend [RSS for all](#) [RSS for failures](#) [RSS for just latest builds](#)

Tickets and review

- **Ticket system** (TRAC) where people can report problems and track discussion and patches
- Changes **discussed** in a ticket and fixes or enhancements **reviewed** by someone else
- 2nd pair of eyes
- Not always applicable

Thank you!

- Room 4018, leaving tomorrow
- ian.hinder@aei.mpg.de
- <https://members.aei.mpg.de/ianhin/>

einsteintoolkit.org

cactuscode.org

simfactory.org

kranccode.org

build.barrywardell.net