

THE SIMULATION FACTORY

Objectives

- **High-level control** over scientific runs; low-level, HW/SW details (compile options, submission scripts) hidden to user, only maintained by reduced number of people, easy to swap without larger user community's disruption.
- Consistent **tracking** of simulation data, status, chaining;
- Designed to **prevent traditional Cactus mishaps** (overwriting important data, recovering from wrong checkpoint files, et cetera);
- Get started on Erik's page:

<http://www.cct.lsu.edu/~eschnett/SimFactory/>

THE SIMULATION FACTORY

Functionality

\$ sim

Simulation Factory:

sim <command> <options>

Manage jobs on batch systems.

Available commands are:

help:	Print this list of commands
list-mach[ines]:	List all available machines
remote <machine_name> <command>:	Execute <command> remotely
login <machine_name>:	Open a remote shell
print-machine:	Output machine name
execute <command>:	Execute an arbitrary shell command
sync <machine_name(s)>:	Copy source tree to another machine
build <configuration_name(s)>:	Build a configuration
sync-remote-build <machine_name> <configuration_name(s)>	
remote-build <machine_name> <configuration_name(s)>	
list-conf[igurations]:	List all configurations
create <simulation_name>:	Create a job skeleton
list-sim[ulations]:	List all simulations
find-sim[ulation] <job id(s)>:	Find simulation with a given job id
submit <simulation_name>:	Submit a job
create-submit <simulation_name>	
cleanup-submit <simulation_name>	
status <simulation_name>:	Determine job status
show-output <simulation_name>:	Show the job's stdout and stderr
stop <simulation_name(s)>:	Stop (qdel, cancel) job
cleanup <simulation_name(s)>:	Clean up after a job has finished
stop-cleanup <simulation_name(s)>	
purge <simulation_name(s)>:	Completely remove a simulation
comment <simulation_name>:	Add a comment to a simulation

THE SIMULATION FACTORY

Developments

- Check out IDEAS/TODO files under `simfactory/`;
- Even higher-level control over runs:
 - Handle parameter sweeps;
 - Query for resource availability (`showstart` comparison);
- Explicit integration with archiving infrastructure (Formaline and NumRel portal);
- Explicit integration with visualization and debugging tools;
- GUI (especially useful for simulation monitoring and archival);

ALPACA Objectives

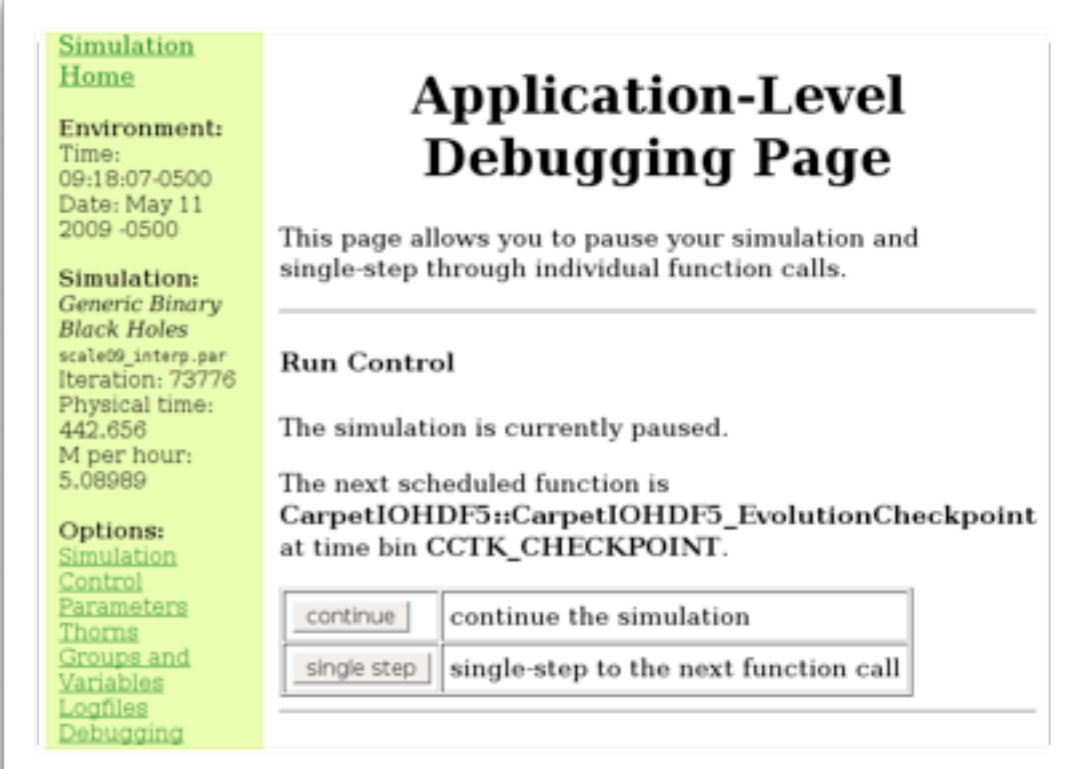
- Scale **ease of control** up to thousands of processes:
 - interactive visualization;
 - agile execution control;
 - rapid debugging.
- **Application-level** tools: smarter, Cactus-aware solutions.
- For further information, take a look at:

<http://www.cct.lsu.edu/~eschnett/Alpaca/>



ALPACA HTTPS

- SSL **webserver**, with X.509 certificate authorization.
- Connect to a running Cactus simulation, pause it, and steer the simulation parameters.
- Fine-grained execution control (single-step through individual thorn functions)
- Use on supercomputers with simple SSH **port forwards**.



The screenshot shows a web interface for debugging a simulation. On the left is a green sidebar with navigation links: Simulation Home, Environment, Simulation, Options, Simulation Control, Parameters, Thorns, Groups and Variables, Logfiles, and Debugging. The main content area has a title 'Application-Level Debugging Page' and a description: 'This page allows you to pause your simulation and single-step through individual function calls.' Below this is a 'Run Control' section stating 'The simulation is currently paused.' and 'The next scheduled function is CarpetIOHDF5::CarpetIOHDF5_EvolutionCheckpoint at time bin CCTK_CHECKPOINT.' At the bottom are two rows of buttons: 'continue' (continue the simulation) and 'single step' (single-step to the next function call).

Simulation Home

Environment:
Time: 09:18:07-0500
Date: May 11 2009 -0500

Simulation:
Generic Binary
Black Holes
scale09_interp.par
Iteration: 73776
Physical time: 442.656
M per hour: 5.08989

Options:
[Simulation Control](#)
[Parameters](#)
[Thorns](#)
[Groups and Variables](#)
[Logfiles](#)
[Debugging](#)

Application-Level Debugging Page

This page allows you to pause your simulation and single-step through individual function calls.

Run Control

The simulation is currently paused.

The next scheduled function is **CarpetIOHDF5::CarpetIOHDF5_EvolutionCheckpoint** at time bin CCTK_CHECKPOINT.

<input type="button" value="continue"/>	continue the simulation
<input type="button" value="single step"/>	single-step to the next function call

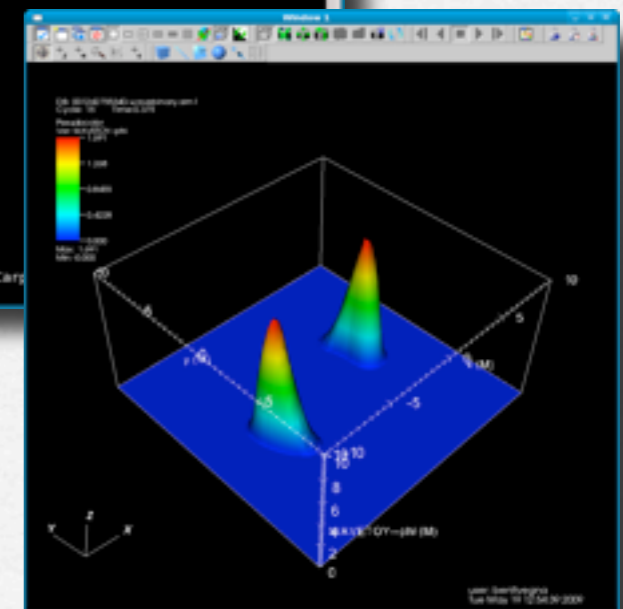
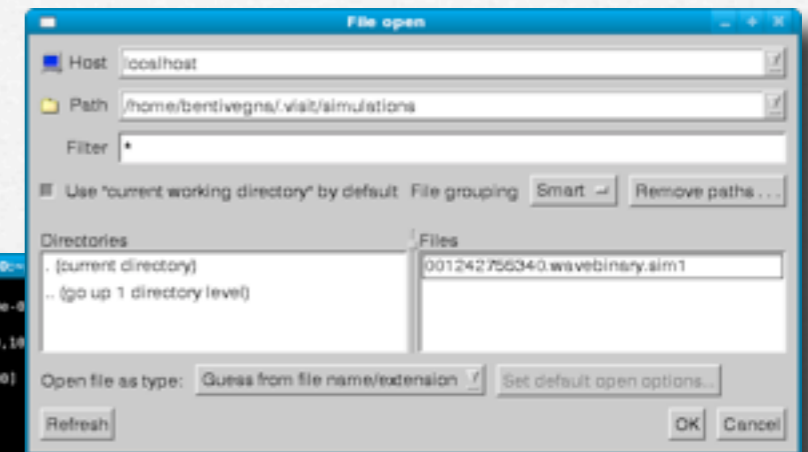
ALPACA VisitConnect

- Concept of **runtime visualization**: skip I/O step, take a look at meshes and arrays in memory.
- HDF5 streaming, ParaView's coprocessing;
- VisIt's **simulation library**, memory access plus basic execution control: watch 3D effect of single Cactus functions on variables.
- VisIt's design also oriented to remote visualization: can one connect to a remote simulation? Communication model currently prevents it, in development (port forwarding).

```
bentivegna@numrel30:~$
INFO (CartGrid3D): Grid Spacings:
INFO (CartGrid3D): dx=1.0000000e-01 dy=1.0000000e-01
INFO (CartGrid3D): Computational Coordinates:
INFO (CartGrid3D): x=[-10.000,10.000] y=[-10.000,10.000]
INFO (CartGrid3D): Indices of Physical Coordinates:
INFO (CartGrid3D): x=[0,200] y=[0,200] z=[0,200]
INFO (Time): Timestep set to 0.025 (courant_static)

Iteration  Time |           NAVETOY::phi
-----
0  0.000 | 0.0000000  0.0000000
1  0.025 | 0.0000000  0.0238732
2  0.050 | 0.0000000  0.0756197
3  0.075 | 0.0000000  0.1432394
4  0.100 | 0.0000000  0.2387324
5  0.125 | 0.0000000  0.3588906
6  0.150 | 0.0000000  0.5013381
7  0.175 | 0.0000000  0.6684505
8  0.200 | 0.0000000  0.8448124
9  0.225 | 0.0000000  1.0241810
10 0.250 | 0.0000000  1.1987415
11 0.275 | 0.0000000  1.3676590
12 0.300 | 0.0000000  1.5312526
13 0.325 | 0.0000000  1.6918430
14 0.350 | 0.0000000  1.8706160
15 0.375 | 0.0000000  2.0472401

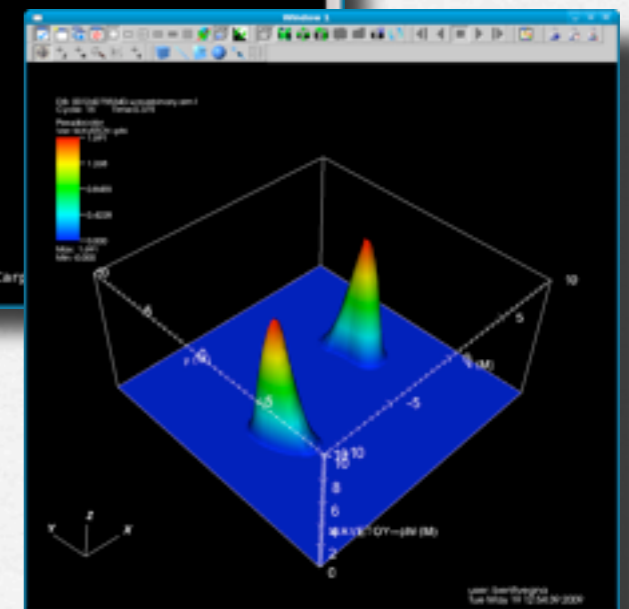
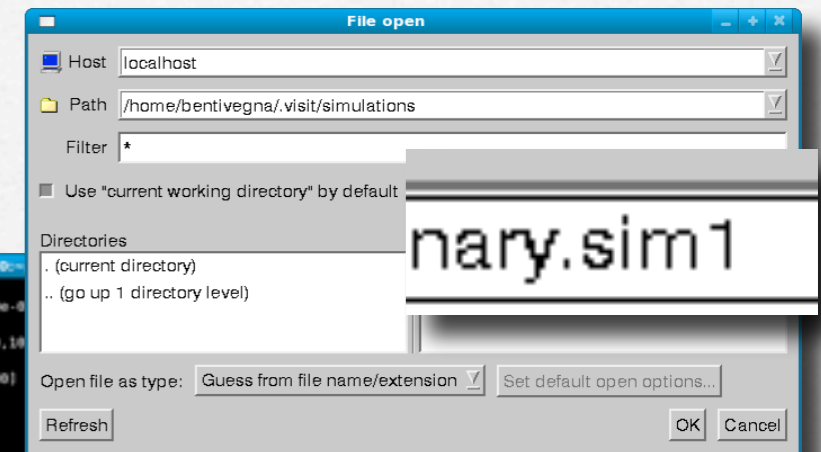
INFO (VisitConnect): VisIt connected.
INFO (VisitConnect): Next scheduled call at CCTK_ANALYSIS is Carr
```



ALPACA VisitConnect

- Concept of **runtime visualization**: skip I/O step, take a look at meshes and arrays in memory.
- HDF5 streaming, ParaView's coprocessing;
- VisIt's **simulation library**, memory access plus basic execution control: watch 3D effect of single Cactus functions on variables.
- VisIt's design also oriented to remote visualization: can one connect to a remote simulation? Communication model currently prevents it, in development (port forwarding).

```
bentivegna@numrel20:~$
INFO (CartGrid3D): Grid Spacings:
INFO (CartGrid3D): dx=1.000000e-01 dy=1.000000e-01
INFO (CartGrid3D): Computational Coordinates:
INFO (CartGrid3D): x=[-10.000,10.000] y=[-10.000,10.000]
INFO (CartGrid3D): Indices of Physical Coordinates:
INFO (CartGrid3D): x=[0,200] y=[0,200] z=[0,200]
INFO (Time): Timestep set to 0.025 (courant_static)
.....
Iteration  Time |           NAVETOY::phi
          | minimum      maximum
-----
0         0.000 | 0.0000000    0.0000000
1         0.025 | 0.0000000    0.0238732
2         0.050 | 0.0000000    0.0756197
3         0.075 | 0.0000000    0.1432394
4         0.100 | 0.0000000    0.2387324
5         0.125 | 0.0000000    0.3588906
6         0.150 | 0.0000000    0.5013381
7         0.175 | 0.0000000    0.6684505
8         0.200 | 0.0000000    0.8448124
9         0.225 | 0.0000000    1.0241810
10        0.250 | 0.0000000    1.1987415
11        0.275 | 0.0000000    1.3676590
12        0.300 | 0.0000000    1.5312526
13        0.325 | 0.0000000    1.6918430
14        0.350 | 0.0000000    1.8706160
15        0.375 | 0.0000000    2.0472401
INFO (VisitConnect): VisIt connected.
INFO (VisitConnect): Next scheduled call at CCTK_ANALYSIS is Carr
```





ALPACA VisitConnect

- Concept of **runtime visualization**: skip I/O step, take a look at meshes and arrays in memory.
- HDF5 streaming, ParaView's coprocessing;
- VisIt's **simulation library**, memory access plus basic execution control: watch 3D effect of single Cactus functions on variables.
- VisIt's design also oriented to remote visualization: can one connect to a remote simulation? Communication model currently prevents it, in development (port forwarding).

```
INFO (CartGrid3D): Grid Spacings:
INFO (CartGrid3D): dx=1.000000e-01 dy=1.000000e-01
INFO (CartGrid3D): Computational Coordinates:
INFO (CartGrid3D): x=[-10.000,10.000] y=[-10.000,10.000]
INFO (CartGrid3D): Indices of Physical Coordinates:
INFO (CartGrid3D): x=[0,200] y=[0,200] z=[0,200]
INFO (Time): Timestep set to 0.025 (courant_static)

Iteration  Time |           NAVETOY::phi
          |           minimum    maximum
-----
0         0.000 | 0.0000000  0.0000000
1         0.025 | 0.0000000  0.0238732
2         0.050 | 0.0000000  0.0756197
3         0.075 | 0.0000000  0.1432394
4         0.100 | 0.0000000  0.2387324
5         0.125 | 0.0000000  0.3588906
6         0.150 | 0.0000000  0.5013381
7         0.175 | 0.0000000  0.6684505
8         0.200 | 0.0000000  0.8448124
9         0.225 | 0.0000000  1.0241810
10        0.250 | 0.0000000  1.1987415
11        0.275 | 0.0000000  1.3676590
12        0.300 | 0.0000000  1.5312526
13        0.325 | 0.0000000  1.6918430
14        0.350 | 0.0000000  1.8706160
15        0.375 | 0.0000000  2.0472401

INFO (VisitConnect): VisIt connected.
INFO (VisitConnect): Next scheduled call at CCTK_ANALYSIS is Carr
```

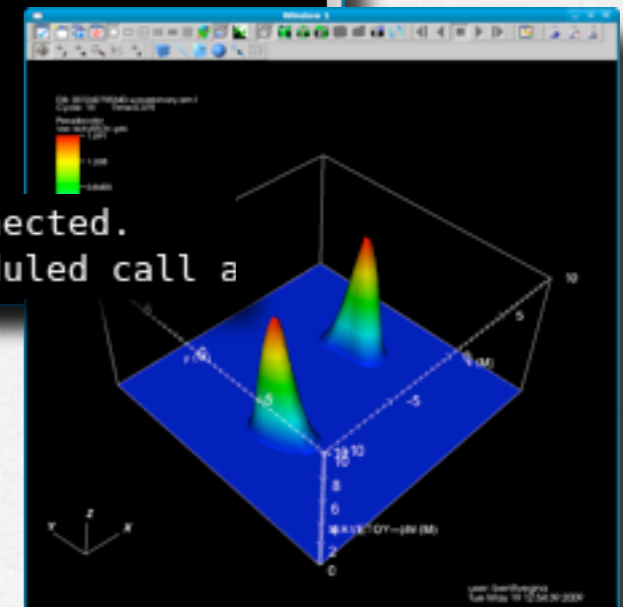
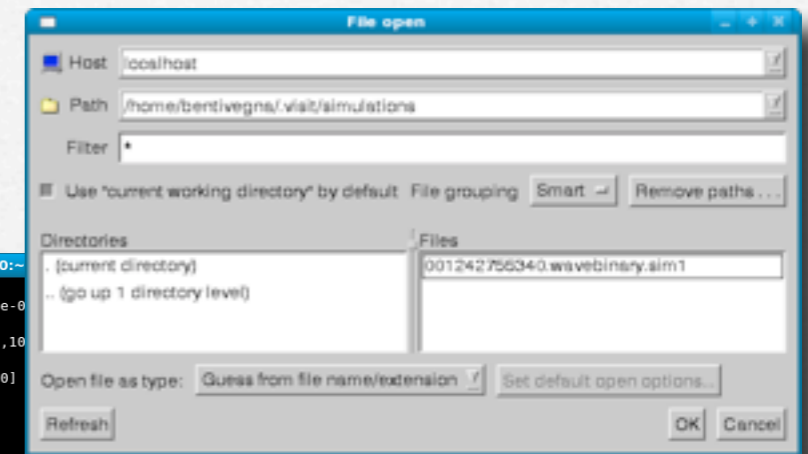


ALPACA VisitConnect

- Concept of **runtime visualization**: skip I/O step, take a look at meshes and arrays in memory.
- HDF5 streaming, ParaView's coprocessing;
- VisIt's **simulation library**, memory access plus basic execution control: watch 3D effect of single Cactus functions on variables.
- VisIt's design also oriented to remote visualization: can one connect to a remote simulation? Communication model currently prevents it, in development (port forwarding).

```
bentivegna@numrel10:~$  
INFO (CartGrid3D): Grid Spacings:  
INFO (CartGrid3D): dx->1.0000000e-01 dy->1.0000000e-01  
INFO (CartGrid3D): Computational Coordinates:  
INFO (CartGrid3D): x->[-10.000,10.000] y->[-10.000,10.000]  
INFO (CartGrid3D): Indices of Physical Coordinates:  
INFO (CartGrid3D): x->[0,200] y->[0,200] z->[0,200]  
INFO (Time): Timestep set to 0.025 (courant_static)  
-----  
Iteration  Time  |  minimum  maximum  
-----  
0  0.000  |  0.0000000  0.0000000  
1  0.025  |  0.0000000  0.0238732  
2  0.050  |  0.0000000  0.0716197  
3  0.075  |  0.0000000  0.1432394  
4  0.100  |  0.0000000  0.2387324  
5  0.125  |  0.0000000  0.3580986  
6  0.150  |  0.0000000  0.5013381  
7  0.175  |  0.0000000  0.6684505  
8  0.200  |  0.0000000  0.8440124  
9  0.225  |  0.0000000  1.0241010  
10 0.250  |  0.0000000  1.1987415  
11 0.275  |  0.0000000  1.3676590  
12 0.300  |  0.0000000  1.5312526
```

```
INFO (VisitConnect): VisIt connected.  
INFO (VisitConnect): Next scheduled call a
```





ALPACA VisitConnect

- Concept of **runtime visualization**: skip I/O step, take a look at meshes and arrays in memory.
- HDF5 streaming, ParaView's coprocessing;
- VisIt's **simulation library**, memory access plus basic execution control: watch 3D effect of single Cactus functions on variables.
- VisIt's design also oriented to remote visualization: can one connect to a remote simulation? Communication model currently prevents it, in development (port forwarding).

```
INFO (CartGrid3D): Grid Spacings:
INFO (CartGrid3D): dx=1.000000e-01 dy=1.000000e-01
INFO (CartGrid3D): Computational Coordinates:
INFO (CartGrid3D): x=[-10.000,10.000] y=[-10.000,10.000]
INFO (CartGrid3D): Indices of Physical Coordinates:
INFO (CartGrid3D): x=[0,200] y=[0,200] z=[0,200]
INFO (Time): Timestep set to 0.025 (courant_static)

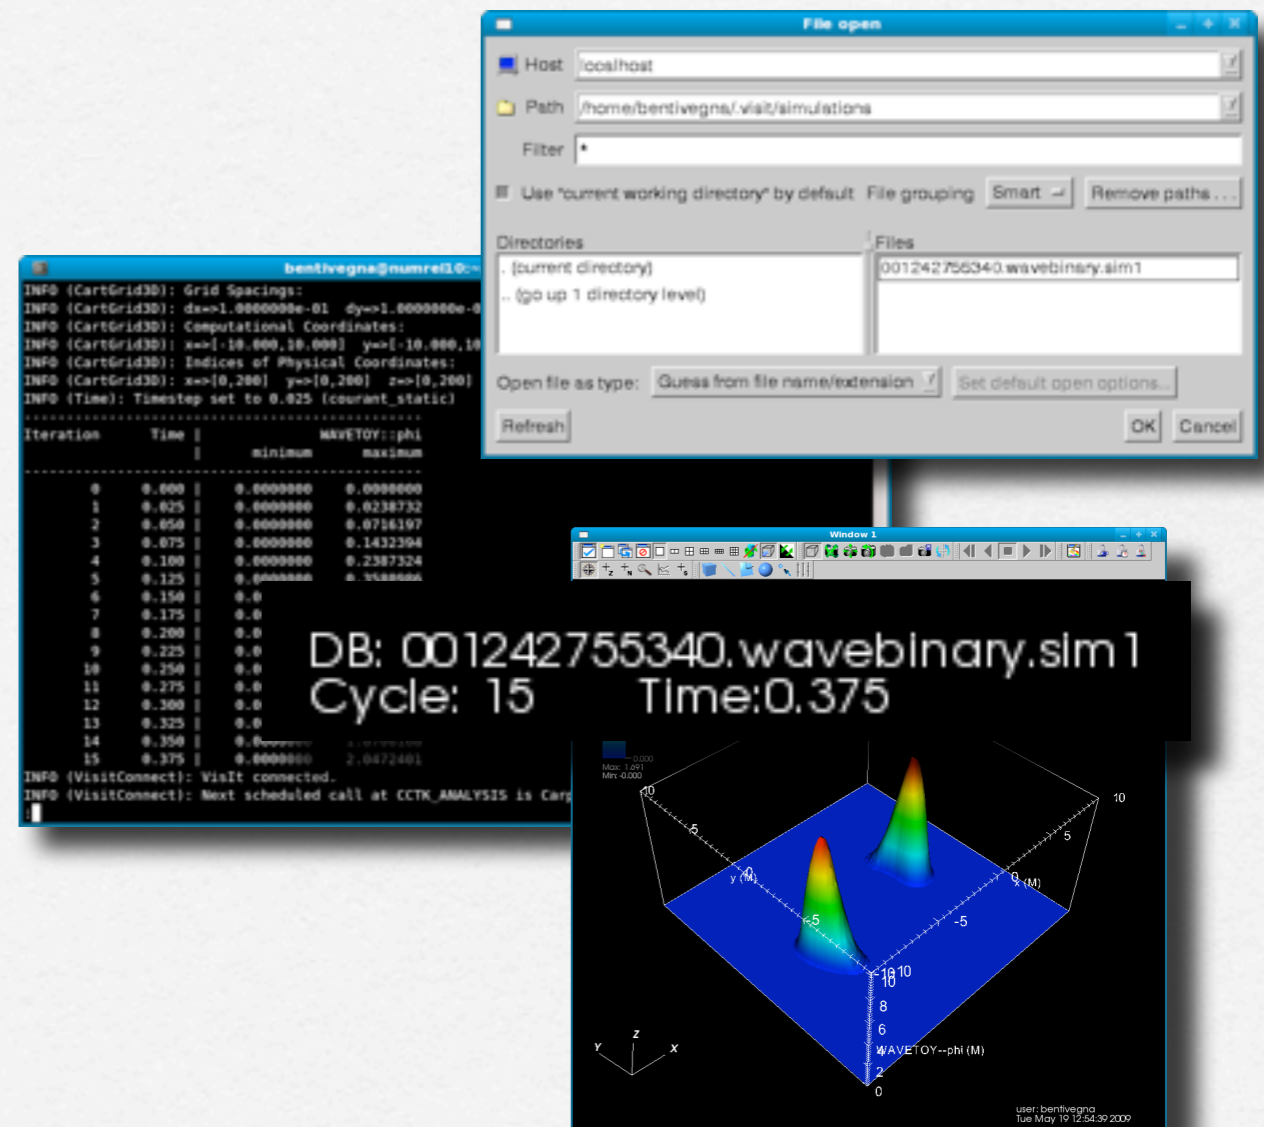
Iteration  Time |           NAVETOY::phi
-----
0  0.000 | 0.0000000  0.0000000
1  0.025 | 0.0000000  0.0238732
2  0.050 | 0.0000000  0.0756197
3  0.075 | 0.0000000  0.1432394
4  0.100 | 0.0000000  0.2387324
5  0.125 | 0.0000000  0.3588906
6  0.150 | 0.0000000  0.5013381
7  0.175 | 0.0000000  0.6684505
8  0.200 | 0.0000000  0.8448124
9  0.225 | 0.0000000  1.0241810
10 0.250 | 0.0000000  1.1987415
11 0.275 | 0.0000000  1.3676590
12 0.300 | 0.0000000  1.5312526
13 0.325 | 0.0000000  1.6918439
14 0.350 | 0.0000000  1.8796160
15 0.375 | 0.0000000  2.0472401

INFO (VisitConnect): VisIt connected.
INFO (VisitConnect): Next scheduled call at CCTK_ANALYSIS is Carr
```



ALPACA VisitConnect

- Concept of **runtime visualization**: skip I/O step, take a look at meshes and arrays in memory.
- HDF5 streaming, ParaView's coprocessing;
- VisIt's **simulation library**, memory access plus basic execution control: watch 3D effect of single Cactus functions on variables.
- VisIt's design also oriented to remote visualization: can one connect to a remote simulation? Communication model currently prevents it, in development (port forwarding).





ALPACA VisitConnect

- Concept of **runtime visualization**: skip I/O step, take a look at meshes and arrays in memory.
- HDF5 streaming, ParaView's coprocessing;
- VisIt's **simulation library**, memory access plus basic execution control: watch 3D effect of single Cactus functions on variables.
- VisIt's design also oriented to remote visualization: can one connect to a remote simulation? Communication model currently prevents it, in development (port forwarding).

```
INFO (CartGrid3D): Grid Spacings:
INFO (CartGrid3D): dx=1.000000e-01 dy=1.000000e-01
INFO (CartGrid3D): Computational Coordinates:
INFO (CartGrid3D): x=[-10.000,10.000] y=[-10.000,10.000]
INFO (CartGrid3D): Indices of Physical Coordinates:
INFO (CartGrid3D): x=[0,200] y=[0,200] z=[0,200]
INFO (Time): Timestep set to 0.025 (courant_static)

Iteration  Time |           NAVETOY::phi
          | minimum      maximum
-----
0         0.000 | 0.0000000  0.0000000
1         0.025 | 0.0000000  0.0238732
2         0.050 | 0.0000000  0.0756197
3         0.075 | 0.0000000  0.1432394
4         0.100 | 0.0000000  0.2387324
5         0.125 | 0.0000000  0.3588906
6         0.150 | 0.0000000  0.5013381
7         0.175 | 0.0000000  0.6684505
8         0.200 | 0.0000000  0.8448124
9         0.225 | 0.0000000  1.0241810
10        0.250 | 0.0000000  1.1987415
11        0.275 | 0.0000000  1.3676590
12        0.300 | 0.0000000  1.5312526
13        0.325 | 0.0000000  1.6918430
14        0.350 | 0.0000000  1.8706160
15        0.375 | 0.0000000  2.0472401

INFO (VisitConnect): VisIt connected.
INFO (VisitConnect): Next scheduled call at CCTK_ANALYSIS is Carr
```



ALPACA Developments

- Debugging tool prototypes ready: extend to **production machines?**
- **Performance analysis;**
- **Fault tolerance;**
- Grand Challenge: conceivable to hide simulation size from user?