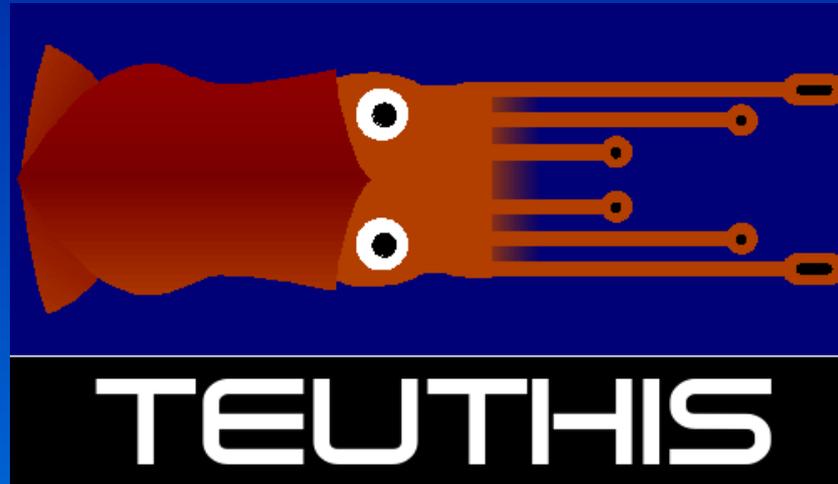


Teuthis: Simulation Management Cyberinfrastructure



Paul Ricker

*National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign*



DES Collaboration Meeting
December 12, 2006





What do simulators need to do?

- **Control simulation jobs**

- ✓ • Configure and build applications on remote systems
- ✓ • Submit and track remote jobs
- ✓ • Create parameter studies and restart jobs

- **Manage simulation data and tasks**

- ✓ • Stage and archive data
- ✓ • Keep track of where datasets are stored
- ✓ • Organize job metadata by purpose and disposition

- **Share data with collaborators**

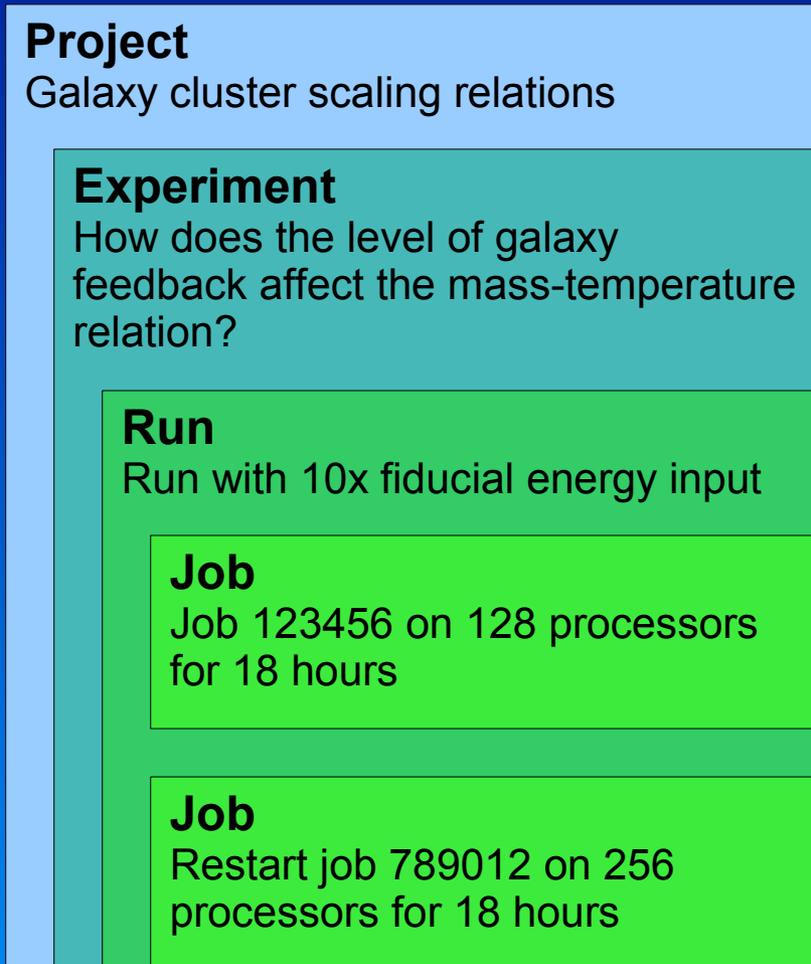
- Make data available
- Convert data among different formats

✓ Teuthis addresses these aspects of simulation tasks with a small, simple, cross-platform GUI (~ 7,000 lines of Python + PyGTK).



Objects manipulated by Teuthis

Workflow hierarchy



Resources

Application

- Accepts text parameter file
- Executes noninteractively
- May need to be compiled
- Produces log file, screen output, data files

Machine

- Login host
- Access method
- Queuing system
- Paths and commands



Project view

Simulation Manager 1.0			
File View Settings Help			
Name	Description	Status	Date last modified
▼ FLASH testing	Testing Simulation Manager using FLASH		Tue Oct 4 17:1
▼ Basic Sedov test (local)	Test of local jobs		Tue Oct 4 19:14:5
Run A			Tue Oct 4 19:15:0
▼ Basic Sedov test (cobalt)	Test of jobs on a machine with PBS queuing and using :		Wed Oct 5 01:43:
Run A	Single run using default parameters to test job submission	Complete	Wed Oct 5 01:17:
Job A0004	Original	28909 [21:26 10/04/2005] 1 CPU/00:10 (Complete) Successful completion	Wed Oct 5 01:50:2
Job A0004	Restart of 28909	28910 [21:27 10/04/2005] 1 CPU/00:10 (Complete) Successful completion	Wed Oct 5 01:16:
▼ Sedov scaling test (cobalt)	Test with varying number of processors		Wed Oct 5 01:43:
Run A1		Complete	Wed Oct 5 02:07:
Job A10001	Original	28924 [01:44 10/05/2005] 1 CPU/00:10 (Complete) Successful completion	Wed Oct 5 02:05:3
Run A2		Complete	Wed Oct 5 02:07:
Job A20001	Original	28925 [01:45 10/05/2005] 2 CPUs/00:10 (Complete) Successful completion	Wed Oct 5 02:05:2
Run A4		Complete	Wed Oct 5 02:07:
Job A40001	Original	28926 [01:45 10/05/2005] 4 CPUs/00:10 (Complete) Successful completion	Wed Oct 5 02:05:7
Run A8		Complete	Wed Oct 5 02:07:
Job A80001	Original	28927 [01:45 10/05/2005] 8 CPUs/00:10 (Complete) Successful completion	Wed Oct 5 02:05:0
▼ Sedov test with varying parameter (cobalt)	Test of jobs with a single varying parameter (Irefine_mi		Wed Oct 5 02:09:
Run A		Complete	Wed Oct 5 02:08:
Job A0001	Original	28928 [01:52 10/05/2005] 1 CPU/00:10 (Complete) Successful completion	Wed Oct 5 01:58:3
Job A0001 Copy	Original	29231 [15:01 10/05/2005] 1 CPU/00:10 (Complete) No data	Wed Oct 5 15:01:5
Run B		Complete	Wed Oct 5 02:08:
Job B0001	Original	28929 [01:52 10/05/2005] 1 CPU/00:10 (Complete) Successful completion	Wed Oct 5 01:58:4
Run C		Complete	Wed Oct 5 02:08:
Job C0001	Original	28930 [01:53 10/05/2005] 1 CPU/00:10 (Complete) Successful completion	Wed Oct 5 01:58:5
Run D		Complete	Wed Oct 5 02:08:
Job D0001	Original	28931 [01:53 10/05/2005] 1 CPU/00:20 (Complete) Successful completion	Wed Oct 5 01:59:0
Run E		In progress	Wed Oct 5 02:08:
Job E0001	Original	28932 [01:53 10/05/2005] 1 CPU/00:20 (Complete) Successful completion; k	Wed Oct 5 02:02:5
Run F		In progress	Wed Oct 5 02:09:
Job F0001	Original	28933 [01:53 10/05/2005] 1 CPU/00:30 (Complete) Exceeded MAXBLOCKS	Wed Oct 5 01:55:2



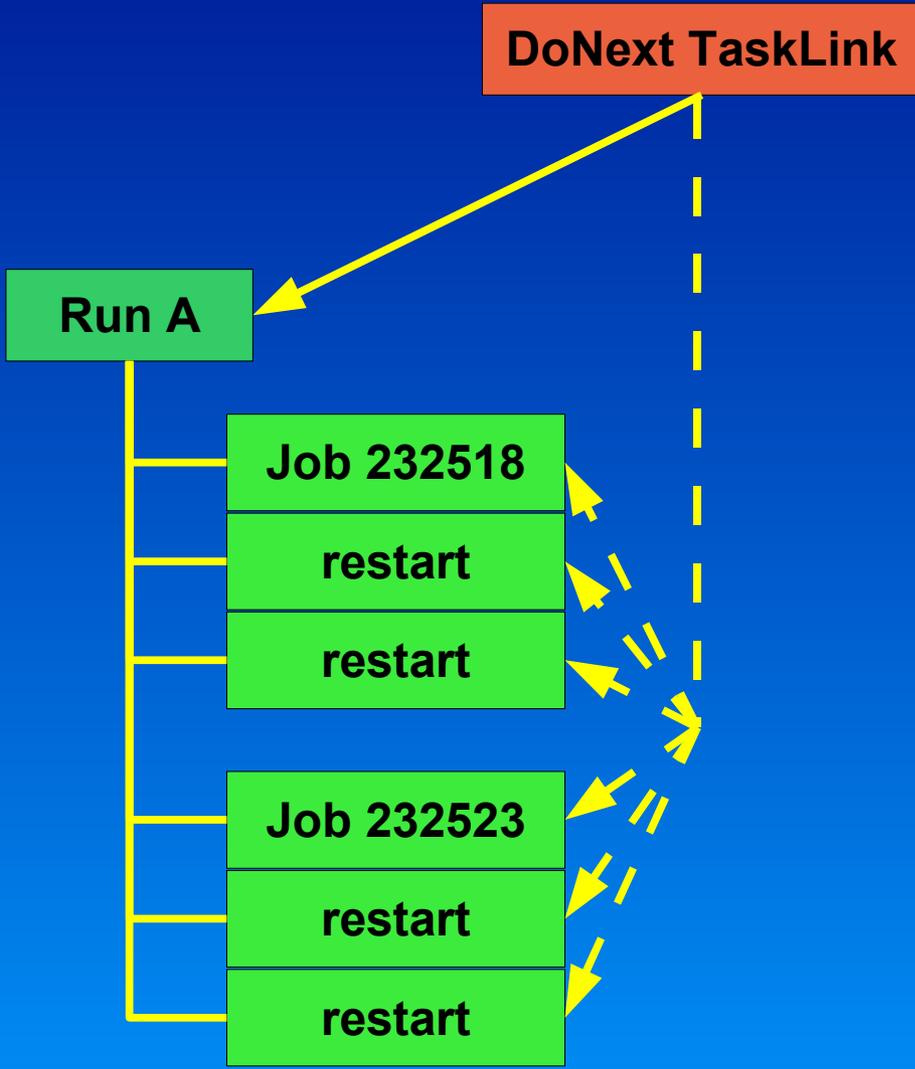
Teuthis workflow management

• Task links

- Atomic workflow unit
- Linked to a particular run
- Invocation triggers new job: stage in – exec – stage out
- Automatic job continuation
- Pattern matching conditionals
- Types – Static, Aggregate, DoNext, DoTogether, DoNTimes, Dolf, WhileDo, TaskGroup

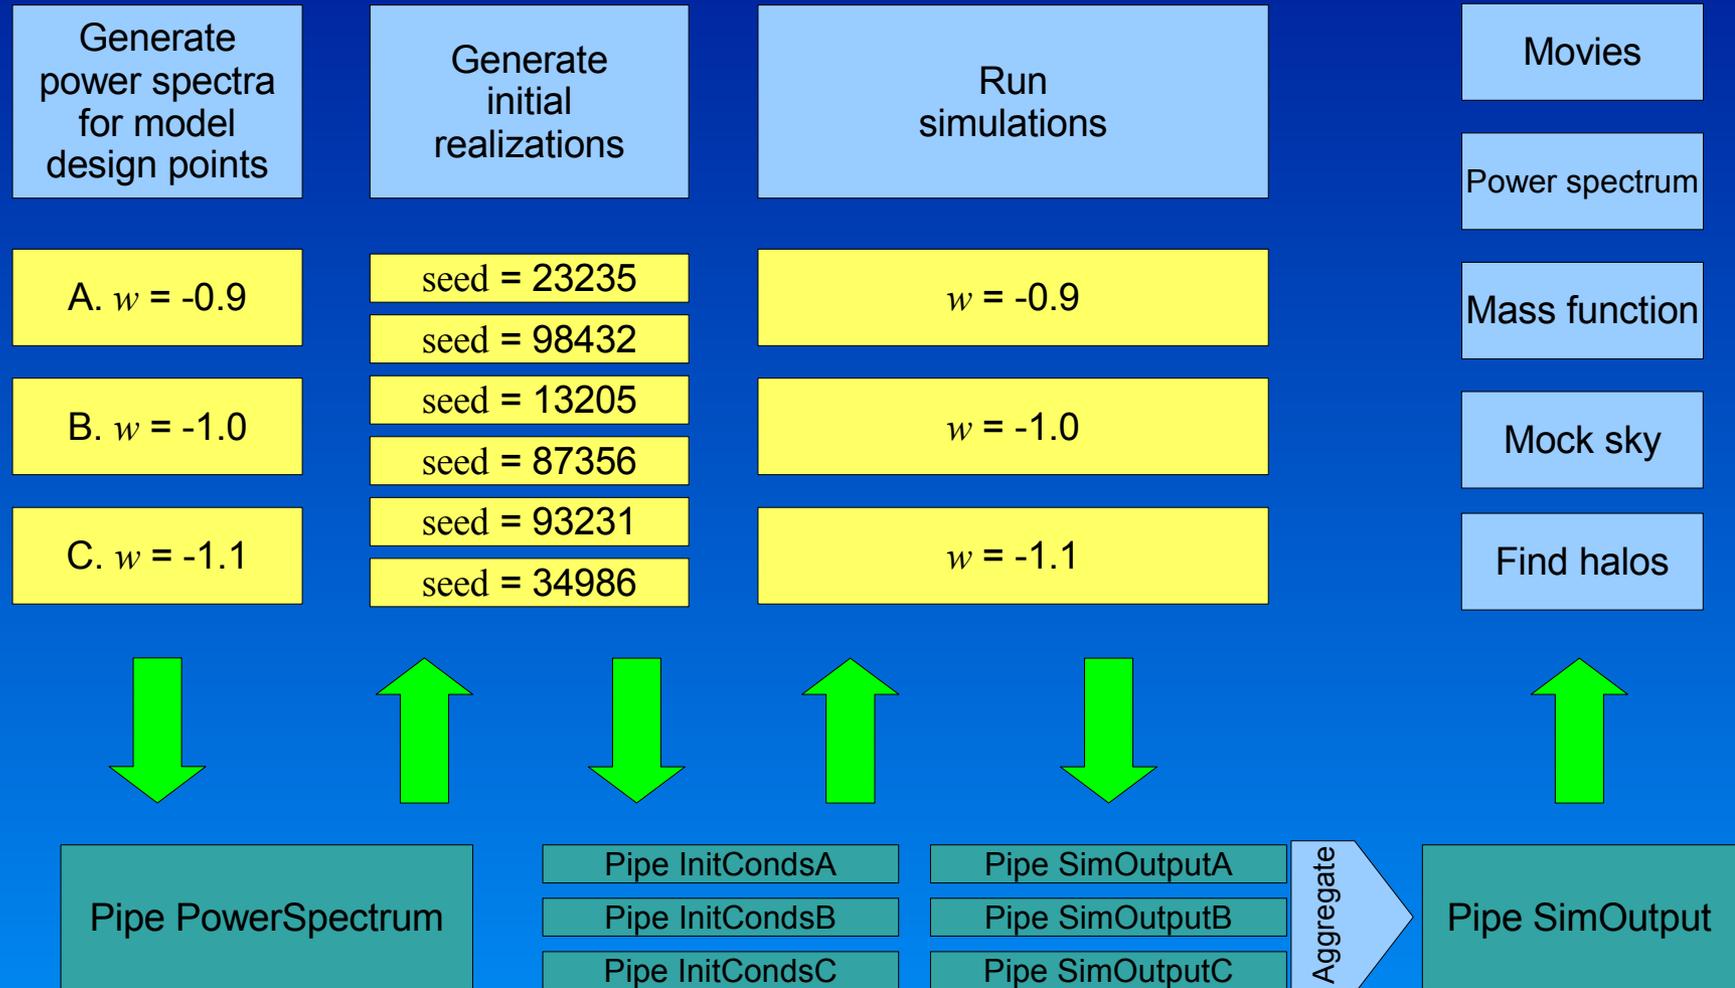
• Pipes

- Dynamically named sources / destinations of data





Example workflow





Example workflow – Teuthis implementation

The screenshot displays the Teuthis 2.0 interface. The main window is divided into two panes: 'Projects' on the left and 'Workflows' on the right. The 'Projects' pane shows a tree view under 'Cosmological simulation demo' with sub-items like 'Power spectrum generation' and 'Initial conditions generation'. The 'Workflows' pane shows a tree view under 'Cosmological simulation workflow' with sub-items like 'Task group Generate power spectra' and 'Task group Generate initial conditions'. A 'Job properties' dialog box is open in the foreground, showing details for a job with ID 'REAL000001'. The dialog is divided into several sections: 'Local job information', 'Data', 'Application', 'Actions', 'Execution', and 'Remote job information'. The 'Local job information' section includes fields for 'Local job ID', 'Comments', 'Disposition', 'Created', and 'Last modified'. The 'Data' section includes 'Src machine', 'Src files', 'Dest machine', and 'Dest path'. The 'Application' section includes 'Application', 'Executable to use', and 'Exec arguments'. The 'Actions' section includes buttons for 'View params', 'View file list', 'Initiate', 'Clone', 'View log', 'View output', 'Continue', and 'Status'. The 'Execution' section includes 'Exec machine', 'Queue', 'Account', 'Wall time', 'No. of CPUs', 'Tiling', and 'Mem/node (MB)'. The 'Remote job information' section includes 'Remote job ID', 'Submitted', and 'Run status'. The dialog also has 'Cancel', 'Apply', and 'OK' buttons at the bottom.

Projects		Workflows	
Name	Status	Name	Status
Cosmological simulation demo		Cosmological simulation workflow	
Power spectrum generation		Task group Generate power spectra	Complete
Run Model A		Do Power spectrum generation - Model A	Complete
Job MODE000001	18371 [10:28 12/12/2006] 1 CPU/00:00 (Complete) No data	Do Power spectrum generation - Model B	Complete
Run Model B		Do Power spectrum generation - Model C	Complete
Run Model C		Task group Generate initial conditions	In Group
Initial conditions generation		Do Initial conditions generation - Realization A	Complete
Run Realization A		Do Initial conditions generation - Realization B	Complete
Job REAL000001	18397 [10:28 12/12/2006] 1 CPU/00:00 (Complete) No data	Do Initial conditions generation - Realization C	Complete
Job REAL000003	18425 [10:29 12/12/2006] 1 CPU/00:00 (Complete) No data	Do Initial conditions generation - Realization D	Complete
Job REAL000003	18425 [10:29 12/12/2006] 1 CPU/00:00 (Complete) No data	Do Initial conditions generation - Realization E	Complete
Run Realization B		Do Initial conditions generation - Realization F	Complete
Run Realization C		Do Initial conditions generation - Realization G	Complete
Run Realization D		Do Initial conditions generation - Realization H	Complete
Job REAL000001	000000 [10:30 12/12/2006] 1 CPU/00:00 (Complete) No data	Do Initial conditions generation - Realization I	Staging Data Out
Job REAL000002	18562 [10:31 12/12/2006] 1 CPU/00:00 (Complete) No data	Do Initial conditions generation - Realization J	
Job REAL000003	18577 [10:31 12/12/2006] 1 CPU/00:00 (Complete) No data		
Run Realization E			
Run Realization F			
Run Realization G			
Job REAL000001			
Job REAL000002			
Job REAL000003			
Run Realization H			
Run Realization I			
Run Realization J			
Simulation			
Run Model A			
Run Model B			
Run Model C			

Local job information		Data	
Local job ID	REAL000001	Src machine	local host
Comments	Original	Src files	/home/ricker/Projects/ci/teuthis/archive/models/Model_A/ling...
Disposition	No data	Dest machine	local host
Created	Tue Dec 12 10:35:17 2006	Dest path	/home/ricker/Projects/ci/teuthis/archive/initconds/\${SM_SRC...
Last modified	Tue Dec 12 10:35:17 2006		

Application		Actions	
Application	GRAFIC1	View params	View file list
Executable to use	/home/ricker/Projects/ci/teuthis/exec/grafic1	Initiate	Clone
Exec arguments	< grafic1.in	View log	View output
		Continue	Status

Execution		Remote job information	
Exec machine	local host	Remote job ID	18787
Queue	No. of CPUs 1 Tiling 1	Submitted	Tue Dec 12 10:35:17 2006
Account	Mem/node (MB) 1000	Run status	Staging Data Out



Forthcoming work

- **More sophisticated experiment designs**
 - Latin hypercube, random samples
- **More sophisticated workflows**
 - DoTogether with multiple execution threads
 - Wait rather than fail if pipes are empty
 - GRAM job submission with Reliable File Transfer (RFT)
- **Integration with other tools**
 - External workflow engines: export workflow script and submit
 - Data management: archiving and tracking files
 - Initial conditions, analysis, visualization



Getting Teuthis

1.0 release available; 2.0 coming soon...

<http://mazama.ncsa.uiuc.edu/projects/teuthis>

TEUTHIS

About

- Documentation
- Download
- Support
- Presentations
- Related projects
- CI home
- Internal pages

Welcome to Teuthis!

Teuthis is a tool intended to improve the efficiency with which computational scientists make use of computing resources, particularly high-performance computers. It is designed especially for the needs of astrophysical simulations, but any computational task that takes a set of input parameters from a file and runs noninteractively can be managed using Teuthis.

With Teuthis you can:

- Remotely configure and build applications from local source code
- Submit and track jobs on remote computing resources
- Painlessly schedule and track multiple restart jobs
- Stage and archive data on different machines
- Create large parameter studies with a few simple operations
- Organize job metadata by purpose and disposition
- Share calculation records with collaborators

Teuthis 1.0

Name	Descr
FLASH testing	
Sedov test with file transfer Sedo	
Run A	
Job A0001	Orig
Sedov scaling test	
Run A1	
Job A1001	Orig
Job A1001 Copy	Orig
Run A2	
Job A2001	Orig
Run A4	