

# Plans and Progress in HPC

GPSC All Hands Meeting


3/29/2008

S. Klasky, S. Ethier, C. Jin



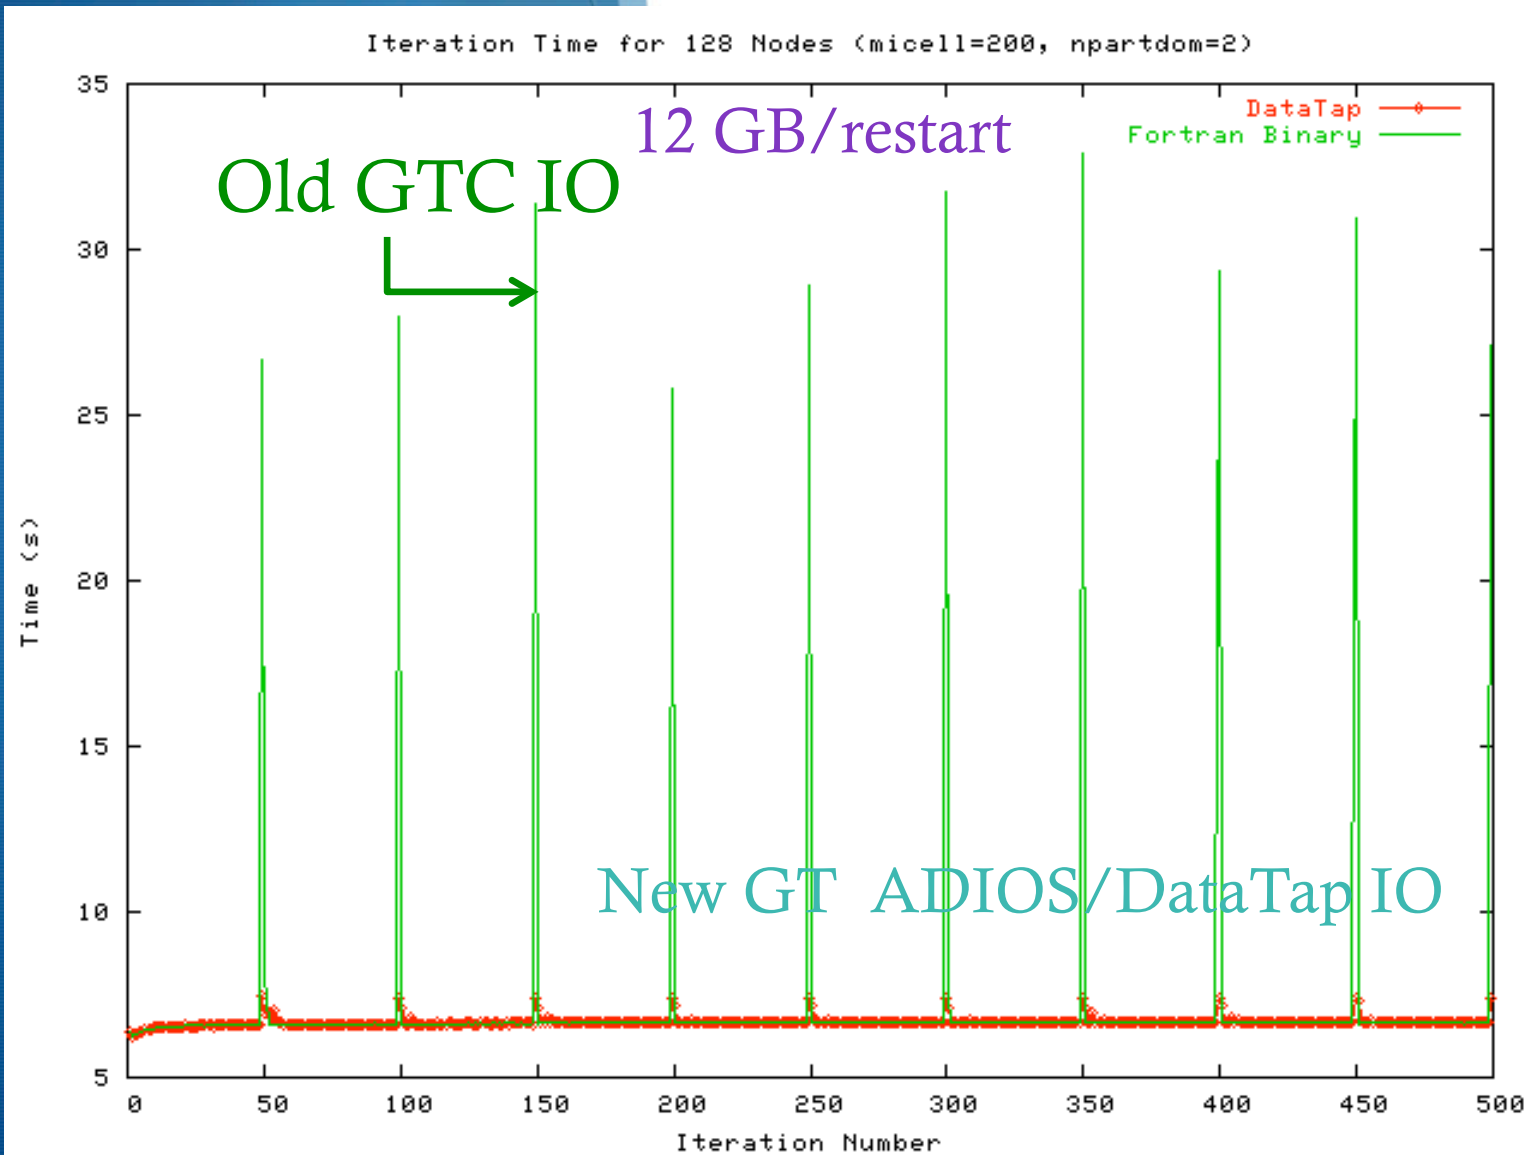
# 1 Petaflops System - Cray

- 1 Petaflops system
- 37 Gigaflops processor
- Over 27K quad-core processors. ~88K cores!
- 2 GB/core; 223 TB total
- 240 GB/s disk bandwidth
- 7.5 MW system power
- Liquid cooled



1 FY 2009: Cray XT

# IO Progress & Plans



# ADIOS

- ◆ Projected September 1 release of ADIOS at NCCS.
- ◆ Through benchmarking of ADIOS into initial test codes :  
GTC (2 versions), XGC1, Chimera, S3D with
  - ◆ MPI-IO method, Fortran IO, MPI-AIO, MPI-CIO, DataTap, Dart
  - ◆ Already see 25GB/sec with MPI-IO ADIOS method

# ORNL Work 2008

- ◆ Finish ADIOS integration into GTC framework IO modules.
- ◆ Start looking at Analysis routines.
- ◆ Dashboard analysis features.
- ◆ Dashboard collaborative features (see Klasky CPES talk).
- ◆ Place provenance collection system and PAPI information on dashboard from workflow automation.

# Dashboard movie

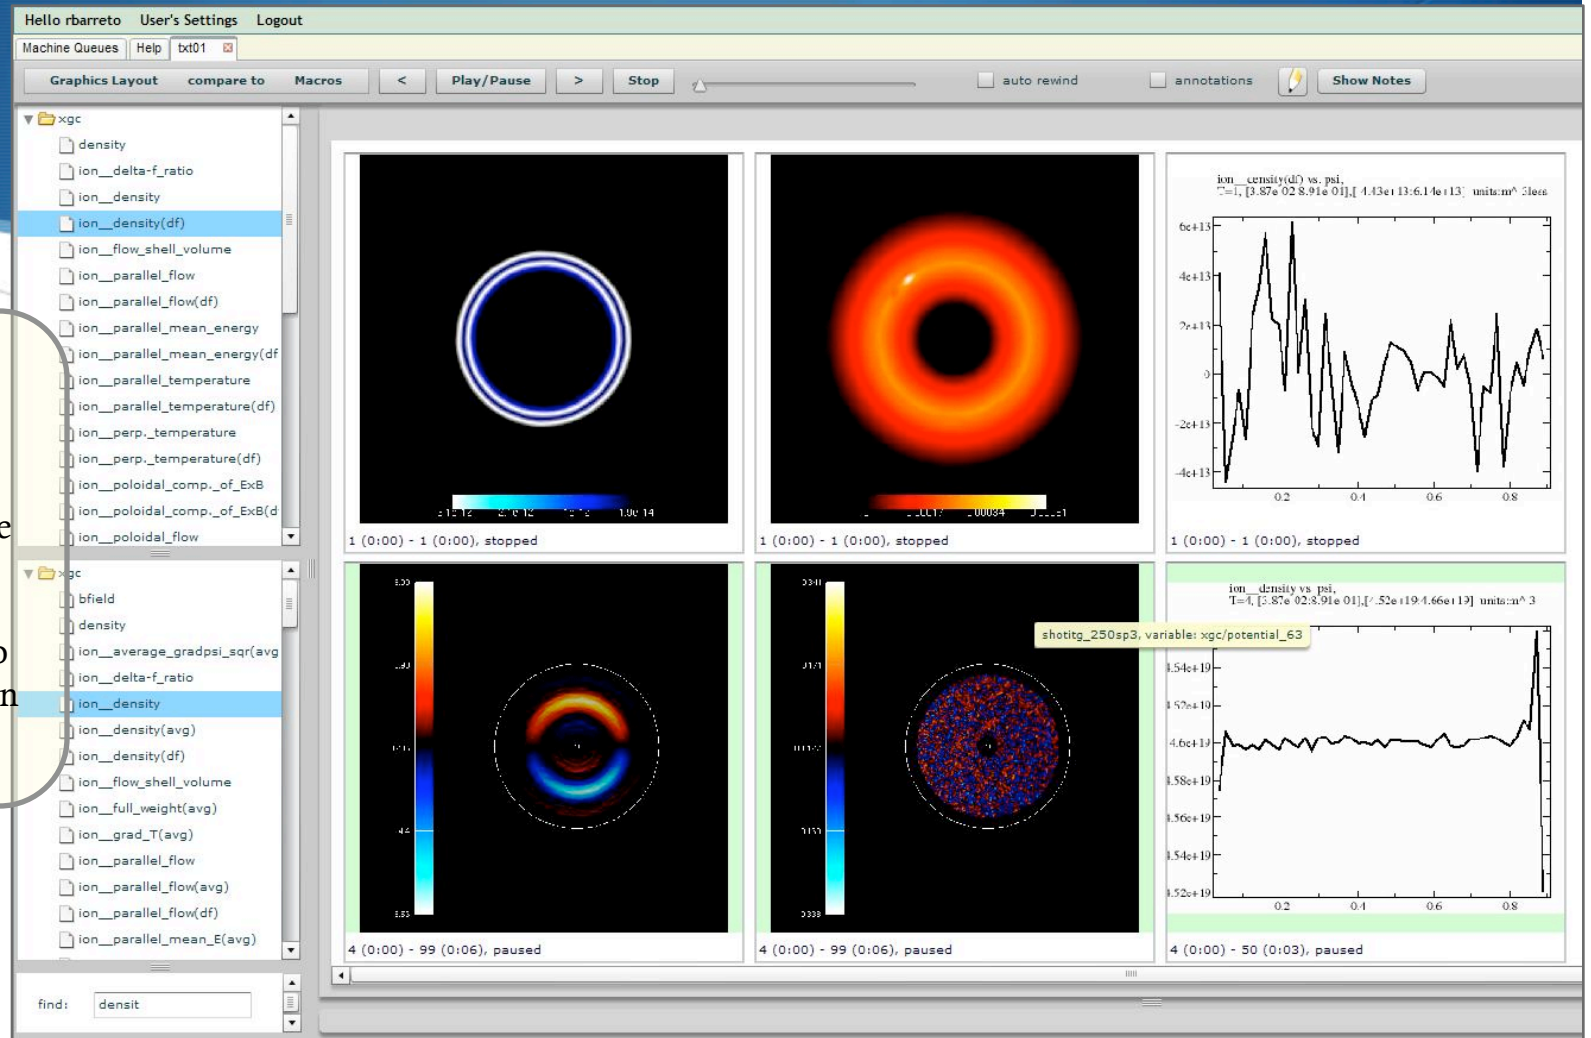


Created using **Wink**



# New Dashboard Features

- Select a shot by searching list or entering shot number
- Drop from both tree and notice tooltip and different background color to differentiate between shots



Compare shots

# New Dashboard Features

Hello rbarreto User's Settings Logout

Machine Queues Help txt01

Graphics Layout compare to Macros < Play/Pause > Stop auto rewind annotations Show Notes

ion\_parallel\_flow(df)  
ion\_parallel\_mean\_energy  
ion\_parallel\_mean\_energy(df)  
ion\_parallel\_temperature  
ion\_parallel\_temperature(df)  
ion\_perp\_temperature  
ion\_perp\_temperature(df)  
ion\_poloidal\_comp\_of\_ExB  
ion\_poloidal\_comp\_of\_ExB(df)  
ion\_poloidal\_flow  
ion\_poloidal\_flow(df)  
ion\_Radial\_current\_density  
ion\_Radial\_current\_density(df)  
ion\_Radial\_flow\_times\_grad\_f  
ion\_Radial\_flow\_times\_grad\_f(df)  
ion\_toroidal\_flow  
ion\_toroidal\_flow(df)  
ion\_v\_para\_times\_B  
ion\_v\_para\_times\_B(df)  
potential  
T\_par  
T\_perp  
v\_ExB-R  
v\_ExB-Z  
v\_par

txt  
xgc  
fort.used.eq  
input  
output.out  
xgcs

Find in text view

```
&sml_param
sml_machine=0 ! 0 circular, 1:D3D,
sml_node_file='cyclon4_coarse.1.node'
sml_ele_file='cyclon4_coarse.1.ele'
sml_bfollow_file=' !'cyclon4.bf'

sml_special=0 ! 0: normal simulation, 1: single particle simulation
sml_dt=5.00D-3 ! delta-t for one time step - unit of toroidal transit time.
sml_mstep=20000 ! totoal time step
sml_deltaf=1 ! delta-f simulation switch - incomplete
sml_turb_efield=1
sml_electron_on=0
sml_nphi_total=4
sml_canonical_maxwell=0 ! cononical maxwellian initial loading switch - incomplete
sml_bounce=2 ! Particle motion boundary condition
! 1 for edge simulation (including open field line region)
! 2 for core simulation (excluding open field line region)

sml_limiter=0 ! Limiter on/off
sml_fem_matrix=1
sml_inpsi=0.10 ! inner boundary of simulation - unit of eq_x_psi
sml_outpsi=0.90 ! outer boundary of simulation - unit of eq_x_psi
sml_dwdt_exb_only=1
sml_push_mode=1
sml_pc_order=3
sml_restart_write_period=9000
sml_restart=0
sml_initial_deltaf_noise=1D-10
sml_mode_select_on=0
sml_mode_select_n=23
sml_deltaf_f0_mode=-1
sml_zero_inner_bd=1
sml_f0_1_Ln=2.20D0
sml_f0_1_Lt=6.9D0
sml_guess_table_size=1900
sml_no_00_efield=1
sml_input_file_dir='/lustre/scratch/shku/XGC-1/inputs/'

&ptl_param
ptl_mass=2D0 ! 1 for hydrogen, 2 for deuterion
ptl_charge=1D0 ! ion charge
ptl_num=1000 ! number of particle for simulation
```

Find

Find what:  
simulation

Find

Basic find function

Drag text files from the tree

find:



# New Dashboard Features

The dashboard interface includes a top navigation bar with "Hello rbarreto", "User's Settings", and "Logout". Below this is a browser-like header with "Machine Queues", "Help", and "txt01". A control bar contains "Graphics Layout", "compare to", "Macros", and playback controls: "<", "Play/Pause", ">", "Stop", "auto rewind", "annotations", and "Show Notes".

On the left, a sidebar lists variables: "xgc/density", "xgc/ion\_\_density", "xgc/ion\_\_density(df)", "xgc/ion\_\_Radial\_current\_density", and "xgc/ion\_\_Radial\_current\_density(df)".

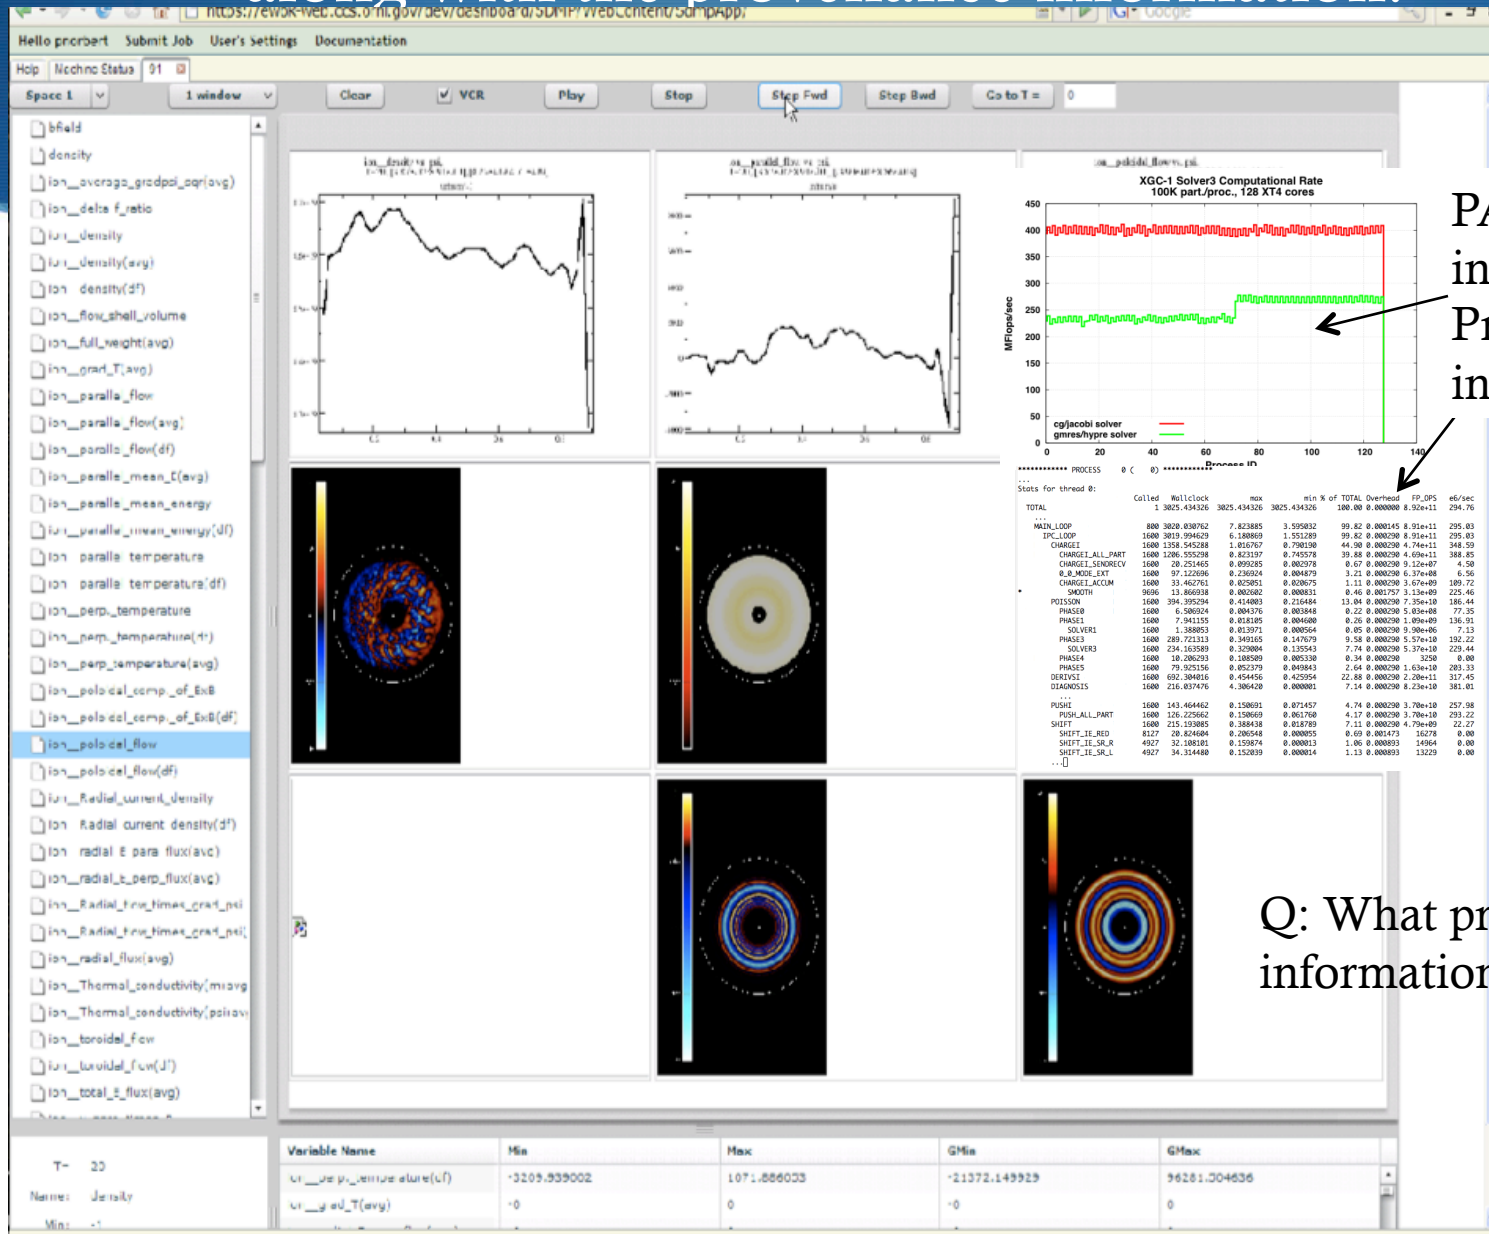
The main plot area displays "ion\_parallel\_flow vs. psi, T=1, [3.87e-02:8.91e-01],[-5.03e+02:5.04e+02] units:m/s3less". The plot shows a black oscillating line and a blue smoother line. A red circle highlights a minimum in the black line, with a green "Annotate movies" label next to it. Below the plot, it says "1 (0:00) - 1 (0:00), stopped".

An "Annotate Movie" dialog box is open on the right, featuring a toolbar with shapes and a text tool, a field for "Apply annotation from timestep 1 to timestep 1 Last timestep=1", and buttons for "Undo Add", "Erase All", "Help", "Save Annotation", "Delete Annotations", and "Close".

At the bottom, a note-taking area contains the text: "Take notes for this shots and save them in the dashboard database. I like those two, they look like the're each other upside down. (Saved on space 2)". Below this is a text editor with a font dropdown set to "Verdana", a size dropdown set to "10", and buttons for "B", "I", "U", and "Save".

Search and annotation callouts are present: "Search for a variable" points to the sidebar, and "Annotate movies" points to the dialog box. "e-scratch book" is located at the bottom right.

# Performance characterization integrated in the dashboard along with the provenance information.



PAPI  
information  
Provenance  
information

Q: What provenance information to keep?

# Machine monitoring.

WebSimMon - Mozilla Firefox 3 Beta 4

https://ewok-web.ccs.ornl.gov/

WebSimMon

Hello sklasky User's Settings Logout

Machine Queues Help demo17

View

Jaguar

showq showbf

Active Eligible Blocked

JobID	Username	Pro	rtime	stime
248028	bugget	1380	00:08:00	Thu Mar 20 14:15:21
248029	bugget	1380	00:08:00	Thu Mar 20 14:15:33
248030	bugget	1380	00:08:00	Thu Mar 20 14:15:46
248031	bugget	1864	00:08:00	Thu Mar 20 14:16:05

Phoenix

showq showbf

Active Eligible Blocked

(6 active jobs, 784 out of 1024 processors in use or 76.56%)

JobID	Username	Pro	rtime	stime
143534	joogan	576	10:07:57	Thu Mar 20 02:13:05
143576	ajohn	32	7:30:32	Thu Mar 20 11:35:40
143581	owolfe	64	5:33:42	Thu Mar 20 13:38:50
143593	fenghe	48	3:17:47	Thu Mar 20 15:37:55
143582	fenghe	48	2:02:08	Thu Mar 20 14:22:16
143572	lentz	16	1:03:46	Thu Mar 20 11:08:54

Franklin

showq showbf

Active Eligible Blocked

(100 active jobs, 19202 out of 19320 processors in use or 99.39%)

JobID	Username	Pro	rtime	stime
446807	mstewart	2	-00:01:09	Thu Mar 20 13:19:06
446782	niri	64	00:00:15	Thu Mar 20 12:45:30
446794	pkent	30	00:00:18	Thu Mar 20 13:10:33
446784	cball	26	00:02:46	Thu Mar 20 13:03:01
446806	mstewart	2	00:10:22	Thu Mar 20 13:09:37
446544	ajnonaka	16	00:10:31	Thu Mar 20 10:40:46
446559	mstewart	2	00:11:02	Thu Mar 20 11:41:17
446797	hargrove	4	00:12:37	Thu Mar 20 13:12:52
446807	vince	128	00:16:10	Thu Mar 20 13:16:25

JaguarCNL

showq showbf

Active Eligible Blocked

(17 active jobs, 7280 out of 7504 processors in use or 97.01%)

JobID	Username	Pro	rtime	stime
88827	wuof	2	-00:02:51	Thu Mar 20 16:00:08
88816	apra	412	00:14:41	Thu Mar 20 15:48:40
88803	ajnonaka	16	00:15:30	Thu Mar 20 15:18:29
88835	hagen	100	00:16:48	Thu Mar 20 16:09:47
88821	coardall	4	00:17:42	Thu Mar 20 15:50:41
88823	gshipman	16	00:20:35	Thu Mar 20 15:53:34
88804	ajnonaka	16	00:27:29	Thu Mar 20 15:30:28
88806	eendeve	24	00:27:47	Thu Mar 20 15:30:46
88774	stoltsov	12	00:28:41	Thu Mar 20 15:31:40

Ewok

showq showbf

Active Eligible Blocked

(4 active jobs, 68 out of 142 processors in use or 47.89%)

JobID	Username	Pro	rtime	stime
45930	fkelly	32	1:40:30	Thu Mar 20 14:13:30
43878	shku	2	1:50:44	Tue Mar 18 18:23:44
45944	fkelly	32	3:01:17	Thu Mar 20 15:34:17
45926	shku	2	1:18:52:31	Thu Mar 20 11:25:31

Jacquard

showq showbf

Active Eligible Blocked

(39 active jobs, 694 out of 712 processors in use or 97.47%)

JobID	Username	Pro	rtime	stime
501708	u617	8	3:04:17:00	Wed Mar 19 17:57:02
502045	alr1	18	1:23:12:30	Thu Mar 20 12:42:32
502054	schrier	2	1:23:11:54	Thu Mar 20 12:41:56
502055	schrier	2	1:23:11:54	Thu Mar 20 12:41:56
502056	schrier	2	1:23:11:54	Thu Mar 20 12:41:56
502057	schrier	2	1:23:11:54	Thu Mar 20 12:41:56
501963	pinous	16	1:21:12:46	Thu Mar 20 10:42:47
501812	dm9c	32	1:18:14:04	Thu Mar 20 07:44:06
501818	tholme	16	1:17:26:52	Thu Mar 20 06:56:54

sklasky

showstart Running Old Eligible Search Old

Machine	JobID	Shot #	Date	Notes
Jaguar	120610	120610	Thu Aug 16 08:44:42 2007	Right click to edit note or delete job.
Jaguar	98758	062701	Wed Jun 27 14:03:09 2007	Right click to edit note or delete job.
Jaguar	98305	06260707	Tue Jun 26 15:22:29 2007	Right click to edit note or delete job.
Jaguar	122365	122365	Tue Aug 21 13:43:22 2007	Right click to edit note or delete job.
Jaguar	120614	120614	Thu Aug 16 08:57:11 2007	hi scott
Jaguar	98108	001	Tue Jun 26 09:54:26 EDT 2007	bad input data
Jaguar	98131	0260701	Tue Jun 26 10:54:36 2007	98131
Jaguar	97813	001	Mon Jun 25 14:32:39 EDT 2007	excellent XGC run showing ELM!
Jaguar	98298	06260705	Tue Jun 26 15:12:15 2007	good run, high beta
Jaguar	98108	901	Tue Jun 26 09:54:40 EDT 2007	bad input data
Jaguar	98303	06260708	Tue Jun 26 15:20:27 2007	bad simulation..
Jaguar	98288	06260703	Tue Jun 26 15:01:00 2007	Right click to edit note or delete job.

Collaborators

Running Old Search Old Add/Remove

username shot number machine name

pnorbert

submit

Machine	JobID	Shot #	Date	Notes
Jaguar	118474	778	Thu Aug 9 13:16:01 2007	Right click to edit note or delete job.
Jaguar	150729	demo04	Fri Nov 9 14:43:16 2007	Right click to edit or delete job.
Jaguar	155640	demo17	Tue Dec 4 13:00:05 2007	Last succ Coupling before the tutorial

# Eventual hooks into MDS+ for

The screenshot displays the MDS+ software interface, which is used for data analysis and visualization. The interface includes a top toolbar with buttons for 'Graphics Layout', 'compare to', 'Macros', 'Play/Pause', and 'Stop'. Below the toolbar, there are several panels:

- Left Panel:** A file browser showing a directory structure with folders 'elite', 'm3d', and 'xgc'. Under 'xgc', a list of data files is shown, with 'ion\_\_density(df)' selected.
- Top-Left Plot:** A line graph titled 'ion\_\_density(2f) vs psi' showing a decreasing trend. The y-axis ranges from 0 to 4e+19. The x-axis ranges from 0.9 to 1.0. The plot is labeled with 'T=0, [9.01e 01.1.05e+00][1.52e+18.1.02e+19] unitcm^3'.
- Top-Middle Plot:** A 2D density plot titled 'density, -, l=61.51' showing a ring-like structure. A color scale on the right indicates density values from 0 to 1e+20. The plot is labeled with 'T=0, [9.01e 01.1.03e+00][5.57e 02.1.34e+20] unit:keV'.
- Top-Right Plot:** A line graph titled 'ion\_\_perp\_\_temperature vs psi' showing a decreasing trend. The y-axis ranges from 0 to 1.2. The x-axis ranges from 0.9 to 1.0. A red circle highlights a region of the plot labeled 'Strange problem'. The plot is labeled with 'T=0, [9.01e 01.1.03e+00][5.57e 02.1.34e+20] unit:keV'.
- Bottom-Left Plot:** A line graph titled 'ion\_\_density(2f) vs psi' showing a noisy signal. The y-axis ranges from 0 to 2e+19. The x-axis ranges from 0.9 to 1.0. The plot is labeled with 'T=0, [1.82e 01.1.05e+00][1.15e+18.3.96e+19] unitcm^3'.
- Bottom-Middle Plot:** A line graph titled 'ion\_\_density(2f) vs psi' showing a noisy signal. The y-axis ranges from 0 to 2e+19. The x-axis ranges from 0.9 to 1.0. The plot is labeled with 'T=0, [1.82e 01.1.05e+00][1.15e+18.3.96e+19] unitcm^3'.
- Bottom-Right Plot:** A line graph titled 'ion\_\_density(2f) vs psi' showing a decreasing trend. The y-axis ranges from 0 to 4e+19. The x-axis ranges from 0.9 to 1.0. The plot is labeled with 'T=0, [1.82e 01.1.05e+00][1.15e+18.3.96e+19] unitcm^3'.

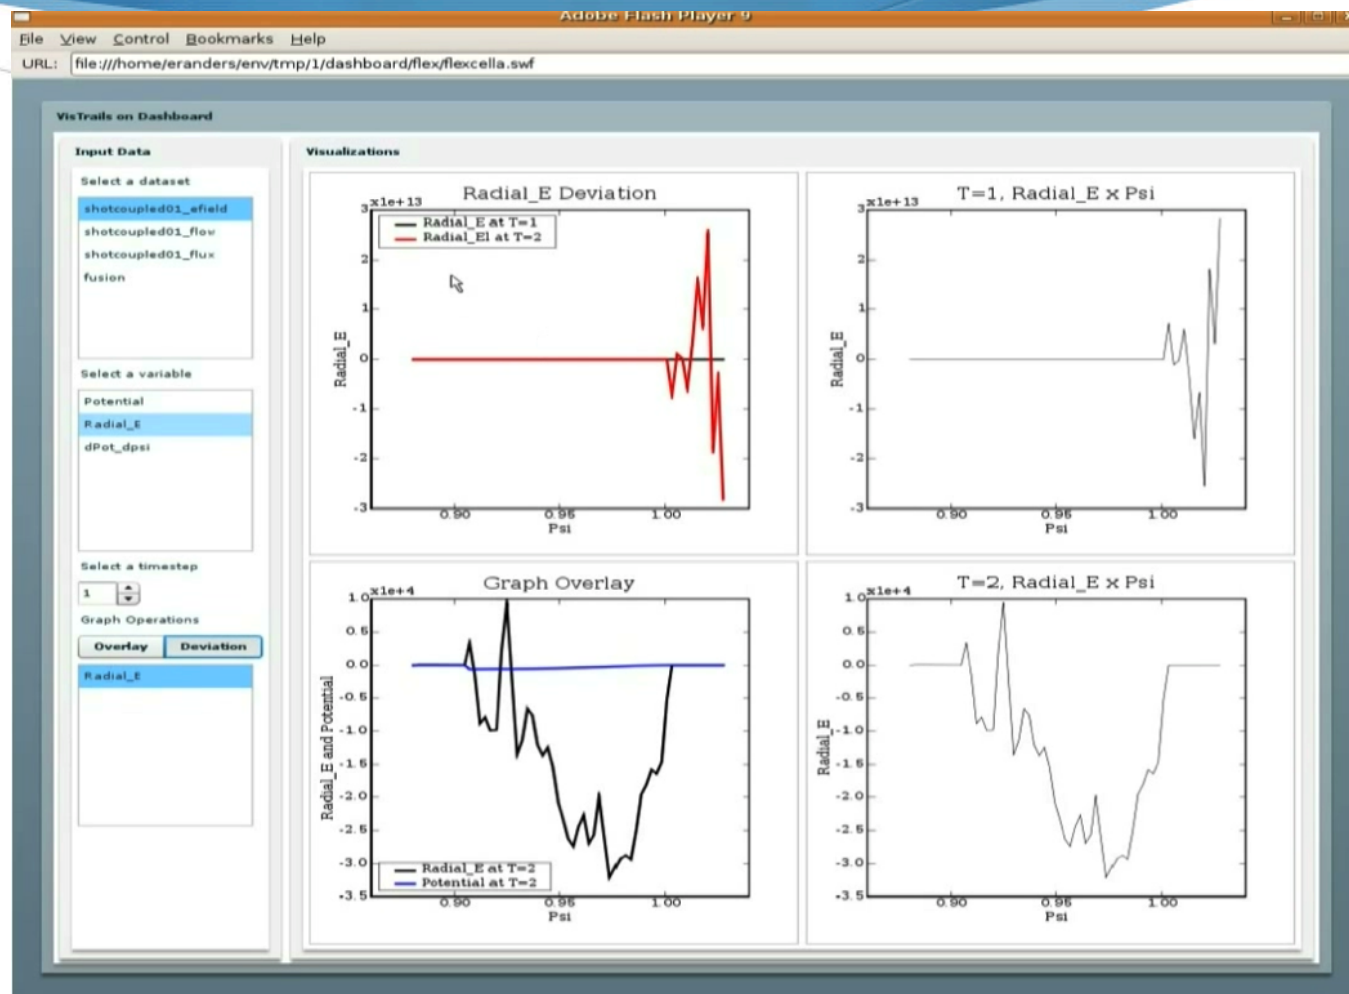
A central dialog box titled 'MDS+ data' is open, containing the following fields:

- username: sklasky
- machine: NSTX
- shot: 94754
- signal:  all

The dialog box has a 'Get' button. Below the plots, there is a text area with the text 'Compare the two runs... Lower panel on the left is a different run...'. At the bottom of the interface, there is a 'Save' button and a toolbar with various icons.

# Interactive Plotting on the dashboard

- Prototype web browser-based flash application with VisTrails ([http://vistrails.sci.utah.edu/index.php/Main\\_Page](http://vistrails.sci.utah.edu/index.php/Main_Page)) on server.





# Dashboard Post processing.

# Still need.

- ◆ Try to incorporate Ma's visualization into Dashboard environment for scaling GTC users.
- ◆ How to make GTC a more community code.

# Performance Optimization work

- ◆ Update and tune OpenMP thread-based parallelism.
- ◆ Identify regions in the code where SSE3 vectorization can be exploited.
- ◆ Optimizations for electron dynamics sub-cycling.
  - ◆ For a run with kinetic electrons most of the time is spent in the sub-cycling steps during which the electrons are moved (push phase) and then sent to the processor holding their location in space (shift phase).
  - ◆ This last step is particularly time-consuming since the electrons move fast and far.
  - ◆ We have now eliminated the shift step during the sub-cycling and do only the push steps.
  - ◆ A modified shift is then called only once at the end of sub-cycling, resulting in significant speed-up of the code (N. Wichmann).

# Performance Optimization work

- ◆ We will test the benefits of process placement at large scale on Jaguar since it has shown to lead to a 30% performance improvement on Blue Gene L, at large concurrencies (8192 cores).
- ◆ We will look into one-sided communication routines when available on Jaguar.
- ◆ Use vectorized functions from ACML.
- ◆ IO is being optimized by using ADIOS.
- ◆ Test different solvers in PETSc, including multigrid solvers, for simulations of large devices.