



Debugger Usage At NERSC

Katie Antypas
July 17, 2007



NERSC is supported by the Office of Advanced Scientific Computing Research in the Department of Energy Office of Science under contract number DE-AC02-05CH11231.





Motivation

Parallel debuggers are expensive tools for high performance computing centers.

Understanding how many users utilize them, how often, and at what concurrencies, can allow centers to better optimize resources.





Overview

- **Survey actual debugger usage at super computing center**
- **Which tools do users prefer**
- **Review of Allinea's DDT on AIX**



Support Different Types of Resources and Usage

- Large variety of applications
 - All scientific applications in DOE SC
- Heterogeneous National and International User Community
- Different types of projects
 - Single PI projects
 - Large computational science collaborations
- Range of Systems
 - Computational, storage, networking, analytics






Totalview Usage on Seaborg

- Seaborg
- SP Power3
 - ~6000 Processors
 - 380 16 processors /node

Number of times users have run Totalview on Seaborg in the past year




Number of Totalview runs	507
Number of unique users	47
Average usage/user	10.8



Totalview Usage on Bassi

- Bassi**
- Power 575
 - 888 Processors
 - 122 8 processors /node

Number of times users have run Totalview on Bassi in the past 18 months



Number of Totalview runs	1268
Number of unique users	81
Average usage/user	15.7

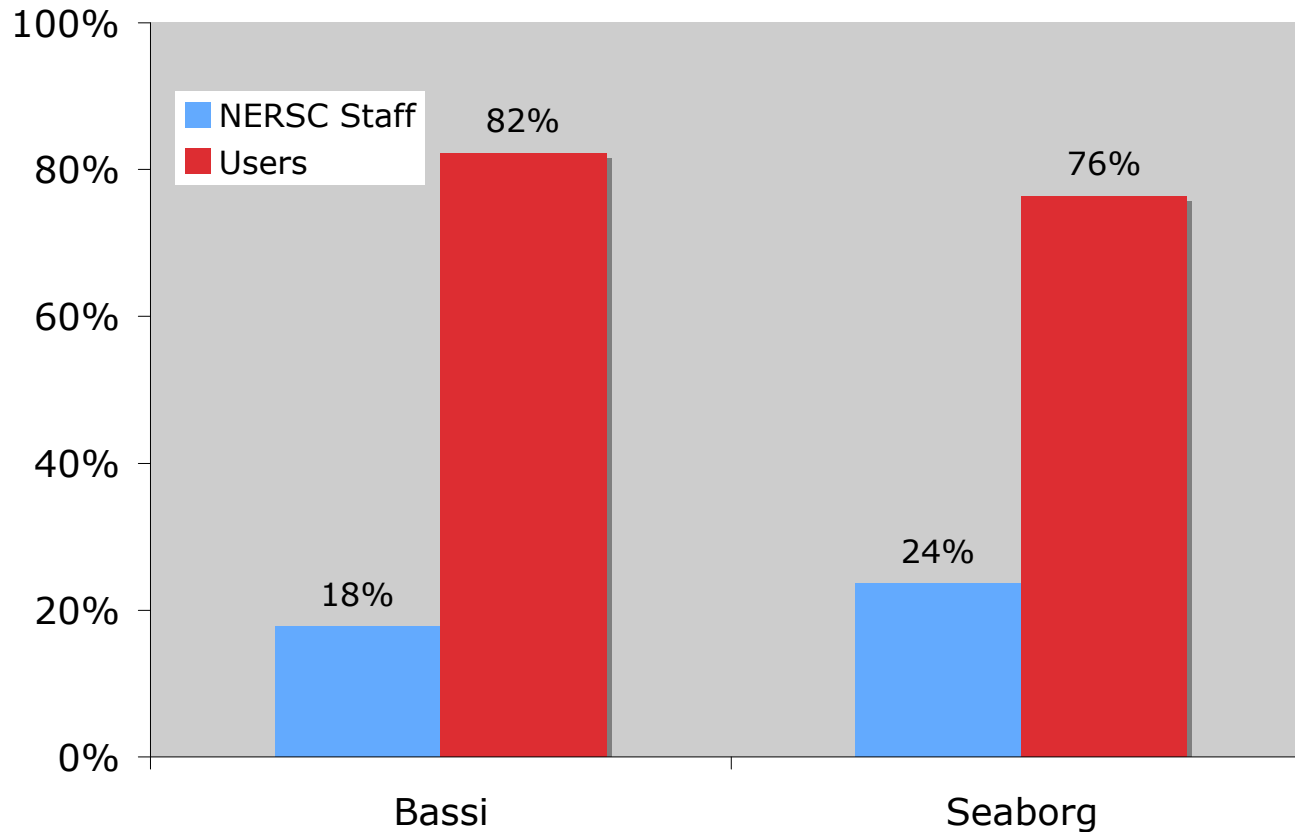


Who is using Totalview?

- **Seaborg (Power 3)**
 - More experienced users running mature production ready jobs
 - Favors larger jobs
 - More likely to use older tools
- **Bassi (Power 5)**
 - Newer users/projects
 - Small/medium sized jobs favored
- **Very roughly ~20 % of active users have run Totalview**



Internal NERSC Staff Totalview Usage



On average roughly 20 -25% of Totalview usage is from NERSC staff helping users debug applications



A word about concurrency...


- **Users are running at lower concurrencies than we expected**
- **1 - 64 processors typically - why?**
 - **In many codes an error at 512 processors can be detected at 32 processors.**
 - **Totalview runs interactively and users must wait a longer time for more nodes**
 - **Debuggers are slow at 256 + processors, can take 20 seconds to step from line to line**



Serial Debugger Usage

Number of Times Users have run GDB on Seaborg in past year

Gdb is the only serial debugger which is used at any frequency at NERSC and only by a small set up users. Totalview has become the debugger of choice at NERSC



Number of gdb uses	122
Number of unique users	15
Average usage per user	8.13

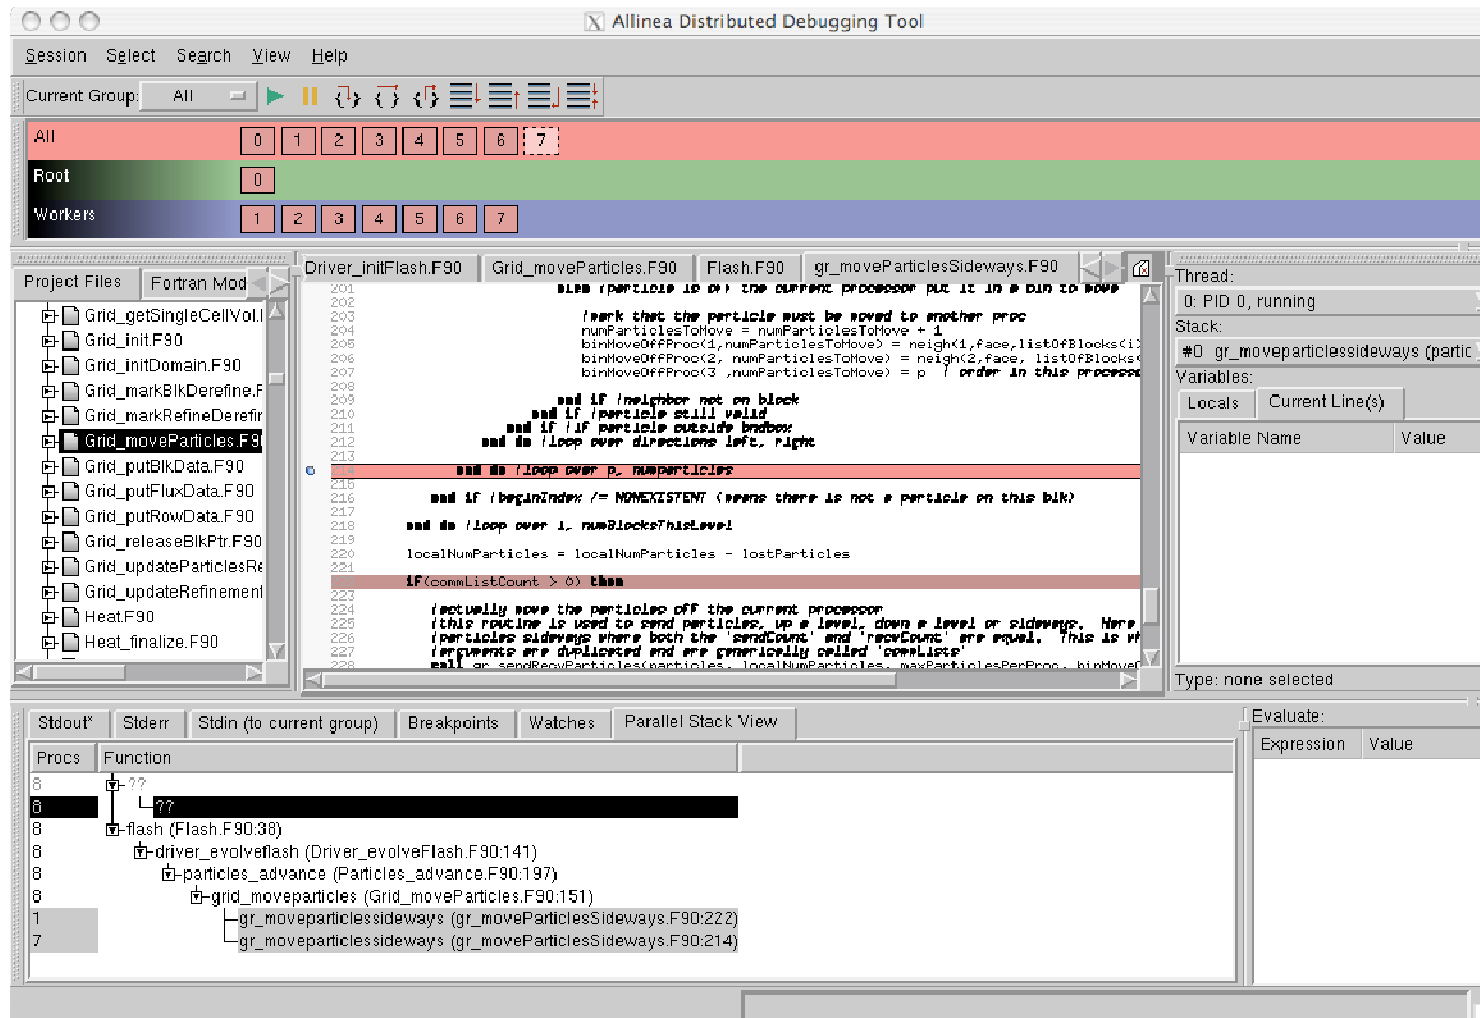


DDT Allinea

- **Parallel Debugger**
- **Typical serial debugging features**
 - set breakpoints and watches, step through program, dive into arrays, evaluate expressions, analyze core files
- **Parallel debugging features**
 - Step through processors
 - View variables across processors
 - Grouping processors Parallel Stack View



User Interface



- *Automatically pulls up the source code*
- *Aren't asked if you want to stop job now*



Parallel Stack View

- Allows user to see position of each processor in the code in the same window
- Essentially groups processors by location in code -- only reasonable strategy at high concurrencies
- Easily can find stray processor
- Can create sub-groups of processors

Stdout*	Stderr	Stdin (to current group)	Breakpoints	Watches	Parallel Stack View
Procs	Function				
7	-??				
7	-??				
7	-flash (Flash.F90:38)				
7	-driver_evolveflash (Driver_evolveFlash.F90:141)				
7	-particles_advance (Particles_advance.F90:197)				
7	-grid_moveparticles (Grid_moveParticles.F90:146)				
2	-gr_moveparticlesuptree (gr_moveParticlesUpTree.F90:218)				
1	-gr_moveparticlesuptree (gr_moveParticlesUpTree.F90:194)				
4	-gr_moveparticlesuptree (gr_moveParticlesUpTree.F90:200)				



DDT on AIX

- **AIX a recently supported platform, basic debugging functionality available, but some advanced features missing**
- **Some user defined Fortran 90 data types not supported (linked list of user defined types)**
- **Fortran 90 modules feature not supported**
- **C++ class variable not distinguished**
- **Threading support recently added in 2.0 release**



Conclusions

- While debugging tools are not used by all users, those who do use them find them invaluable
- Users at NERSC have transitioned to parallel debuggers
- Parallel debugging concurrency not as high as expected
- DDT is a promising new debugging tool, although on AIX, additional features need to be added