

Basics of Parallel Debugging


J. Melvin

`jmelvin@ices.utexas.edu`

Sustainable Horizons Institute Webinar Series

4/12/2019

Introduction



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2015 by K. Rupp

- We need debugging strategies for MPI, Threading, GPUs, etc...
- Debugging tools, GDB, Allinea DDT, etc...
- Recap:
 - Introduction to gdb/debugging (slides): <https://github.com/jamelvin/SHI-Webinar-Debugging/blob/master/Slides-DebuggingWebinar.pdf>
 - Introduction to gdb/debugging (Webinar): <https://www.youtube.com/watch?v=3p0iNcbmZFY>

- GDB (GNU Debugger) is a command line debugger (<https://www.gnu.org/software/gdb/>)
- Supports C, C++, Fortran and some others
- You may be able to use GDB with Python as well (<https://wiki.python.org/moin/DebuggingWithGdb>)
- Python has a built-in debugger called PDB which functions very similarly to GDB (<https://docs.python.org/2/library/pdb.html>)
- Other debuggers (DDT / Totalview / IDEs) typically have a more GUI based debugger but the basic commands and ideas we will discuss today should be applicable to all debuggers

Running with GDB

- ****IMPORTANT:** You need to compile with debug flags (-g or -ggdb)
- ***NOTE:** For parallel programs you may need to compile with explicit linking to mpi libraries (-l ... -L ...)
- Launch with gdb: `gdb --args* ./your_exe exe_runtime_args`
- You can also attach gdb to an already running process
- See GDB Reference card for a partial list of GDB commands
 - Execution: `run (r)`, `continue (c)`, `step (s)`, `next (n)`
 - Breakpoints: `break (b)`, `break if`, `clear`, `delete`
 - Program Stack: `backtrace (bt)`, `frame`
 - Display: `print (p)`, `display`

- Focus mainly on MPI parallelization
- Parallel debugging strategies
- Walk through examples with GDB
- A brief introduction to DDT

First Example: MPI code for Numerical Integration

$$f(x) = \begin{cases} 1 - 10x & 0 \leq x < 0.1 \\ 3x^2 - 2x + 0.17 & 0.1 \leq x < 0.6 \\ -\frac{1}{8}(x - 0.6) + 0.05 & 0.6 \leq x \leq 1.0 \end{cases}$$




Figure: The integral of this function is 0.01

Debugging Strategy: Attach to Single Process

Example file: mpiIntegrate.cpp

- Bug is occurring only on 1 processor
- Goal: Isolate the processor where the bug is occurring in gdb
 - May need to put in a hung code block for that rank

```
bool infLoop = true;

if (myRank == 7)
    while (infLoop);
```

- Attach gdb to a running process (ps ax | grep ProgramName)

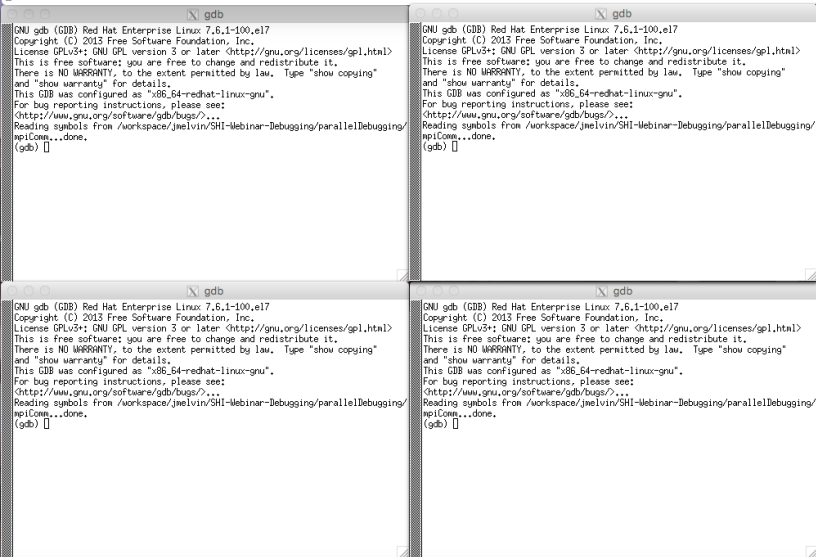
```
[jmelvin@magus parallelDebugging]$ ps ax | grep mpiIntegrate
6362 pts/3      S+   0:00 mpirun -np 8 ./mpiIntegrate
6364 ?          Rs   0:12 ./mpiIntegrate
6365 ?          Rs   0:13 ./mpiIntegrate
6366 ?          Rs   0:13 ./mpiIntegrate
6367 ?          Rs   0:13 ./mpiIntegrate
6368 ?          Rs   0:13 ./mpiIntegrate
6369 ?          Rs   0:13 ./mpiIntegrate
6370 ?          Rs   0:13 ./mpiIntegrate
6371 ?          Rs   0:13 ./mpiIntegrate
6512 pts/4      S+   0:00 grep --color=auto mpiIntegrate
```

- gdb ProgramName ProcessID

Example file: mpiComm.cpp

- Bug seems to be a result of interaction between multiple processors
- Goal: Attempt to reduce the size of your problem and use GDB to manage a small number of processors
- Run parallel program through GDB
(`mpirun -np numProcs xterm -e gdb --args ProgramName ProgramArgs`)

```
[jmelvin@magus parallelDebugging]$ mpirun -np 4 xterm -e gdb --args ./mpiComm
```



Warning: Race Condition

Example File: raceThread.c

- One issue that can arise in parallel and not serial is that of a race condition
- This can be especially difficult to debug as when you debug you alter the order of execution
- This is more likely to occur with threading and shared memory
- Some ways to spot a potential race condition
 - Deterministic code produces different answers each time you run
 - Different numbers of processors produce different answers
 - Bug goes away or changes when you run it in a debugger

```
[[jmelvin@magus parallelDebugging]$ ./raceBash.sh
The value of a is 20
The value of a is 20
The value of a is 20
The value of a is 20
The value of a is 20
The value of a is 20
The value of a is 20
The value of a is 20
The value of a is 20
The value of a is 19
```

Example File: raceThread.c

- A few important commands when using GDB with threads
- <https://sourceware.org/gdb/onlinedocs/gdb/Threads.html>
 - [info threads](#)
 - Shows you all the threads and their IDs
 - [thread idNum](#)
 - Switches debugging control to that thread

```
[(gdb) info threads
  Id  Target Id          Frame
  2   Thread 0x7ffff6b39780 (LWP 6116) "raceThread" 0x00007ffff7037311 in clone
      () from /lib64/libc.so.6
* 1   Thread 0x7ffff7fbe740 (LWP 6101) "raceThread" L_main_14__par_loop0_2_0 ()
      at raceThread.c:16
[(gdb) thread 2
[Switching to thread 2 (Thread 0x7ffff6b39780 (LWP 6116))]
```

- For debugging large parallel programs or for a more user friendly experience, commercial software like DDT
- Graphical User Interface based
- Typically available on supercomputing clusters
- <https://www.arm.com/products/development-tools/server-and-hpc/forge/ddt>
- Also can be used for debugging GPU's, OpenMP, MPI or serial codes
- Can make your life a lot easier

Reminders:

- Slides are posted in the github repository: `https://github.com/jamelvin/SHI-Webinar-Debugging/blob/master/Slides-ParallelDebuggingWebinar.pdf`
- Video of the webinar will be posted to `https://www.youtube.com/channel/UCDErMJEKVXXAdMvDXYbDsRQ/videos`

If you have questions, feel free to email me any time: jmelvin@ices.utexas.edu