

# A GDB to support debugging High Performance Computing (HPC) Applications: Upstreaming

Richard Bunt, Principal Engineer, Linaro  
Forge GDB Technical Lead



# Agenda

- Inspired by a talk at Connect '23 around the benefits of upstreaming
- Similar story for Forge's High Performance Computing (HPC) extensions to GDB
- Background
  - High Performance Computing
  - Linaro DDT - A Graphical debugger for HPC
- GDB with HPC extensions
- Downstream Pain
- Relative Upstream Bliss

# High Performance Computing (HPC)

- Parallel computing
- Examples:
  - Simulate galaxy creation, weather forecasting
  - Computational fluid dynamics or crash/impact simulations
- Commodity hardware (optimized) running Linux
- Languages {Fortran, C, C++, Python}
- Message Passing Interface (MPI)
  - Open MPI, MPICH, MVAPICH

# Linaro DDT

The screenshot shows the Linaro DDT IDE interface. The main editor displays the source code of a C program named 'hello.c'. The code includes a function 'func2()' and a 'main' function. The 'Locals' window shows the current state of local variables, including 'message' with a value of '###', 'my\_rank' with a value of 1, and 'p' with a value of 16. The 'Stacks (All)' window shows the call stack, with 'main (hello.c:141)' at the top. The 'Raw Command' window shows a command being sent to process 0: 'show version'. The output of the command is displayed in the window below the command input.

message  
my\_rank  
p

|         |     |
|---------|-----|
| message | ### |
| my_rank | 1   |
| p       | 16  |

Raw Command

Command:  Send

Command sent to process 0: show version

```
show version
&"show version\n"
~"GNU gdb (GDB) 13.1-
13dc4013e5944b096ffb4efe34e921e3a5b7f20ac.
1-review\n"
```

# GDB(s) for HPC

- Fortran (prevalent language in HPC)
- Non-GNU compiler support
- Stability
- Memory efficiency
- Third-party GPU GDBs

# Example: Limited length printing

```
(gdb) print -elements 10 bigArray  
$1 = (1, 2, 3, 4, 5, 6, 7, 8, 9, 10...)
```

```
(gdb) set max-value-size 40  
print -elements 10 bigArray  
Value requires 40000 bytes, ...
```

- HPC applications use large arrays to model physical effects multi-GB.
- GDB eagerly loads entire arrays
- Upstream status

# Example: Fortran array slicing

```
!gdb/testsuite/gdb.fortran/array-slices.f90
$1 = (1, 2, 3, 4, 5, 6, 7, 8, 9, 10)
print *,array(2:4)
$1 = (2, 3, 4)
print *,array(:3)
$1 = (1, 2, 3)
print *,array(5:3:-1)
$1 = (5, 4, 3)
print *,array4d(3:-2:-2,10:7:-2, : , -7:-10:-1)
...
```

- Array slicing syntax
- Inspect a small subset of an array
- 3 implementations
- Forge
  - 7 dimension limit
  - Limited slicing support
  - Memory efficient
- Upstream
  - Limited slicing support
- Fedora
  - Full slicing support
  - Memory inefficiency
- Upstream status

# Example: Disable source file opening

- Reading source files from 10K GDBs from a shared file system
- GDB reads source files by default
- A similar issue affects init files
- set source open [on|off]
- Forge handles file access using its scalable tree.



# Example: Max depth

```
#define N 500
struct coordinate { int a; int b; int c; };
struct coordinates { int a[N]; int b[N]; int
c[N]; };
```

```
(gdb) print -max-depth 0 -elements 2 --
coordinates_i
```

```
$5 = {...}
```

```
(gdb) print -max-depth 1 -elements 2 --
coordinates_i
```

```
$6 = {a = {...}, b = {...}, c = {...}}
```

```
(gdb) print -max-depth 2 -elements 2 --
coordinates_i
```


```
$7 = {a = {0, 1...}, b = {0, 2...}, c = {0,
3...}}
```

- Control the amount of data
- SoA, AoS
- Combined with limited length
- Forge runs with a mix of 0 and 1
- Upstream status

# The Debt

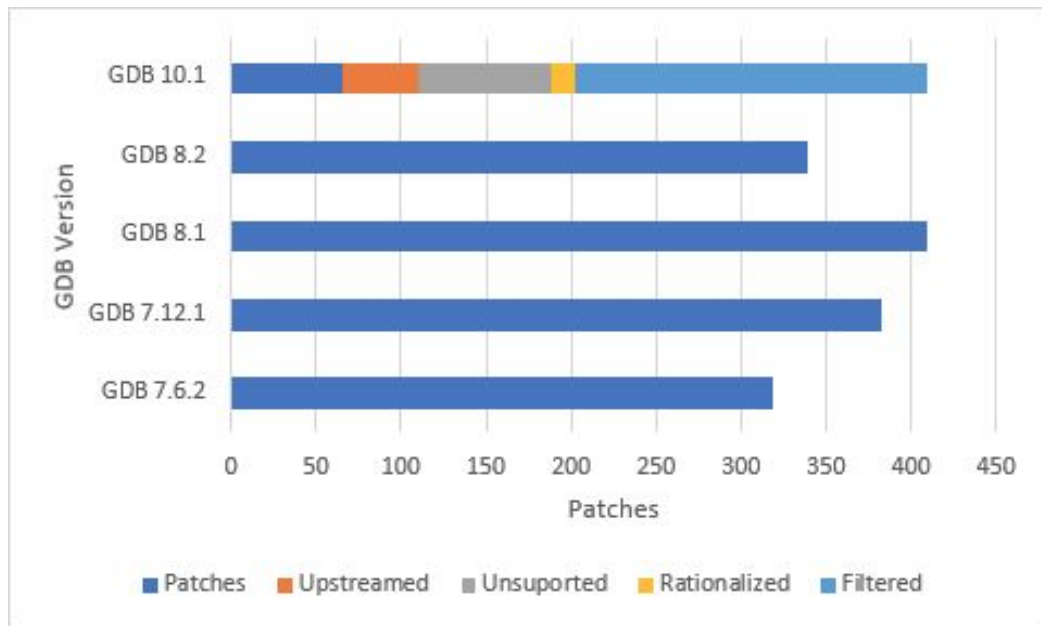
- Patches vs. Lines of Code (LOC)
- GDB 7.6.2
  - Never again!
- GDB 7.10.1
  - Abandoned
- Number of patches increasing
- GPU GDBs
- Not sustainable

Number of Patches vs GDB Version



# Patch Paydown


- Change of tack: Upstreaming
- Around GDB 9.1 that we had upstreamed enough
- Filtering



# Now

- Further upstreaming
- Python plugins

Number of Patches vs GDB Version



# Wins

- Everyone can benefit from the improvements
  - Previous examples
  - Support for isolating Python
- One Fortran array slicing implementation
- Reduced time to rebase from months to weeks
  - GDB 13
- Elided rebases altogether
  - System GPU GDBs
  - GDB 14
- Enabled customers earlier
  - Graviton 3 support case

# Future

- Projecting GDB 15 to have even fewer patches
  - Upstreams
  - GDB plugins
- Continue upstreaming

# Thank you

Any Questions?

