

# Parallel Debugging with DDT

James Willis (SciNet)

October 22, 2025

# Outline

- Software Bugs
- What is Debugging?
- Symbolic Debuggers
- What is DDT?
- Setting up DDT on Teach
  - ▶ Hello-mpi Hands-on
- Client-Server Mode
- Matrix-Matrix Multiply Hands-on Example

# Outline

- Running DDT on a GPU
- Other Useful Features of DDT
  - ▶ Attach Mode
  - ▶ Submitting Jobs to a Scheduler
  - ▶ Running Core Files

# Software Bugs

# Software Bugs

- Writing clean, efficient, error-free code is nearly impossible
- At some point you will run into situations such as:

- ▶ Compile time errors
- ▶ Segmentation faults

```
laptop:~$ gcc app.c -o app
laptop:~$ ./app
Segmentation fault
```

- ▶ Your code doesn't do what you expect
  - ▶ Incorrect results
- These are all examples of what is known as a *software bug*

# Common Symptoms

- Compile time errors
  - ▶ Code syntax errors (easy to fix)
  - ▶ Linker errors when linking against libraries
  - ▶ Cross-compilation, i.e. compiling for a different computing architecture compared to the host
  - ▶ Compilation warnings

**Always turn compiler warnings on and fix or understand them before running your code! It will save you future headaches.**

- But just because it compiles does not mean it is correct!

# Common Symptoms

- Runtime errors
  - ▶ Floating point exceptions
  - ▶ Segmentation faults
  - ▶ Aborted
  - ▶ Incorrect output (e.g. NaNs, Inf)

# Error Examples

Type	Reason
Arithmetic	Corner cases e.g. <code>sqrt(-0.0)</code> , infinities
Memory Access	Index out of range, uninitialised pointers
Logic	Infinite loop, corner cases
Misuse	Wrong input, ignored error, no initialisation
Syntax	Wrong operators/arguments
Resource Starvation	Memory leak, quota exceeded
Parallel	Race conditions, deadlock



# What is going on?

- Almost always, a condition you are sure is satisfied, is not
- But your application likely relies on many such assumptions
- First order of business is finding out what is going wrong and what assumption is not warranted
- Follow the *Fundamental Principle of Confirmation*:
  - ▶ Process of confirming, one by one, that many things you **believe** to be true about the code are **actually** true
- Debugger: program to help detect errors in other programs
- **Debugging: Methodical process of finding and fixing flaws in software**
- **You are the real debugger**

# How to Avoid Debugging

- Write better code:
  - ▶ Simple, clear, straightforward code
  - ▶ Modular (no global variables or 10,000 line functions)
  - ▶ Avoid “cute” tricks (no obfuscated C code winners)
- Improve your code/algorithm/language/API understanding
- Don't reinvent the wheel, use existing libraries
- Write (simple) tests for each part of your code
- Use version control (GIT) so you can “roll back” your code if a bug is found

# Debugging Workflow

- As soon as you are convinced there is a real problem, create the simplest test case that reproduces the bug
- This is science: model, hypothesis, experiment, conclusion
- Try a smaller problem size, turning off physical effects with options, etc. until you have a simple, fast repeatable example of the bug
- Try to narrow it down to a particular module/function/class. For Fortran, switch on bounds checking (`-fbounds-check`)
- Now you're ready to start debugging

# Ways to Debug

- Preemptive:
  - ▶ Turn on compiler warnings: fix or understand them!

```
laptop:~$ gcc/gfortran -Wall
```

- ▶ Check your assumptions (e.g. use assert)
- Inspect the exit code and read the error messages!
- Add print statements
  - ▶ **No way to debug!**

# What's wrong with using print statements?

## Strategy

- Constant cycle:
  - 1 Strategically add print statements
  - 2 Compile
  - 3 Run
  - 4 Analyse output
- Bug not found? Repeat from 1. again
- Have to remove extra code after the bug is fixed
- Rinse and repeat for each bug

## Disadvantages

- Time consuming
- Error prone
- Changes memory, timing. . .

**There's a better way!**

# Symbolic Debuggers

# Symbolic Debuggers

## Features

- ① Crash inspection
- ② Function call stack
- ③ Step through code
- ④ Automated interruption
- ⑤ Variable checking and setting

Use a graphical debugger or not?

- Local work station: graphical is convenient
- Remotely (Trillium): can be slow, but we will look at ways to improve this later

In any case, graphical and text-based debuggers use the same concepts

# Preparing Your Code for the Debugger

- Add debugging flags when compiling your code:

```
laptop:~$ gcc/g++/gfortran -g [-gstabs]
laptop:~$ icc/icpc/ifort -g [-debug parallel]
laptop:~$ nvcc -g -G --cudart shared
```

- Optional flag: switch off optimisation -O0 (sometimes symbol values are hidden to the debugger at higher optimisation levels e.g. -O2, -O3)



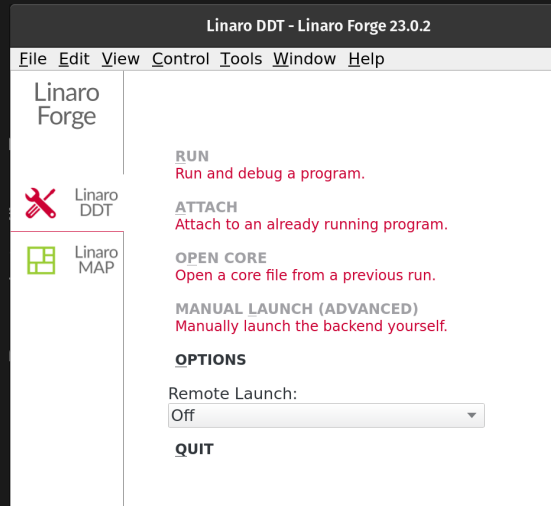
# Examples of Symbolic Debuggers

- Command-line based debuggers: GDB, LLDB
- Graphical based debuggers: **DDT**, Visual Studio, Eclipse
  - ▶ Nice, more intuitive graphical user interface
  - ▶ Front-end to command-line based tools: Same concepts
  - ▶ Need graphics support: X11 forwarding (or VNC)
- The rest of the workshop will focus on DDT

# What is DDT?

# What is DDT?

- **DDT** stands for **Distributed Debugging Tool**
- Powerful GUI-based commercial debugger by *Linaro*
- Developed for debugging parallel, multi-threaded and distributed applications
- Widely used in high-performance computing environments
- Available on Trillium and other Alliance systems (*Note: license only allows debugging up to 256 processes*)



# DDT Features

- **Key Features:**

- ▶ Parallel and distributed debugging capabilities
- ▶ Graphical user interface for intuitive navigation
- ▶ Support for multiple programming languages (e.g. C, C++, Fortran, Python)
- ▶ Supports MPI, OpenMP, threads, CUDA, ROCm and more
- ▶ Integrated performance analysis tools - *MAP*
- ▶ Memory debugging functionalities

# Launching DDT on Teach

- Load the latest software stack with a compiler and MPI module:

```
teach-login01:~$ module load StdEnv/2023
```

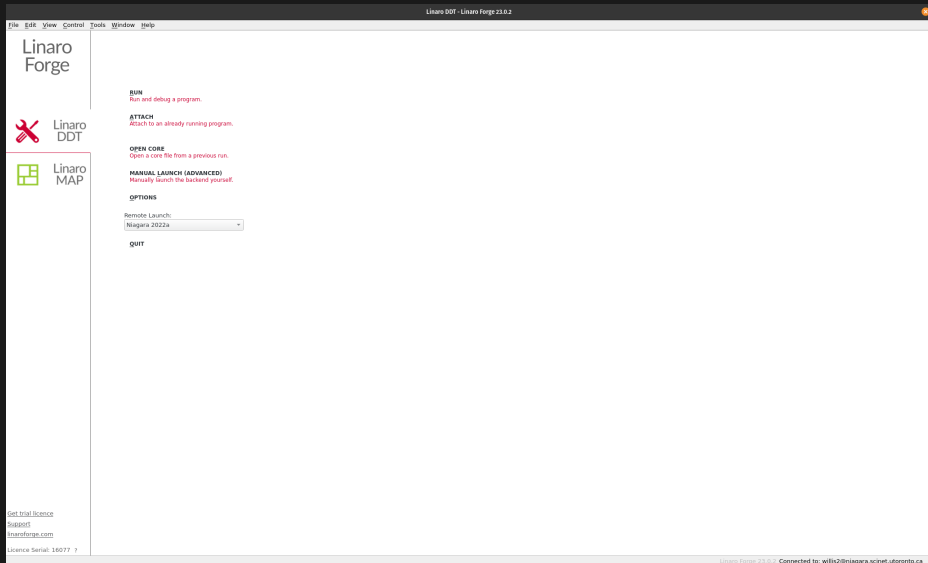
- Load DDT:

```
teach-login01:~$ module load ddt-cpu/23.1.1
```

- Start DDT with one of these commands:

```
teach-login01:~$ ddt
teach-login01:~$ ddt <exe compiled with -g flag>
teach-login01:~$ ddt <exe compiled with -g flag> <arguments>
teach-login01:~$ ddt -n <numprocs> <exe compiled with -g flag> <arguments>
```

# Launching DDT on Teach



# Creating a Debug Session

Run

Application: /gpfs/fs0/scratch/s/scinet/willis2/ddt-mmult/mmult\_c

Details

Application: /gpfs/fs0/scratch/s/scinet/willis2/ddt-mmult/mmult\_c

Arguments:

☐ stdin file:

Working Directory: /gpfs/fs0/scratch/s/scinet/willis2/ddt-mmult

☒ MPI: 4 processes, Open MPI (Compatibility)

Details

Number of Processes: 4

☐ Processes per Node 1

Implementation: Open MPI (Compatibility)

Change...

mpirun arguments

☐ OpenMP

Details

☐ CUDA

Details

☐ ROCm

Details...

☒ Memory Debugging: Fast, 1 guard page after, Backtraces, Pr

Details...

☐ Submit to Queue

Configure...

Parameters...

Environment Variables: none

Details

Plugins: none

Details

Help

Options

Run

Cancel

Memory Debugging Options

☒ Preload the memory debugging library

Language: Recommended

Note: Preloading only works for programs linked against shared libraries. If your program is statically linked, you must relink it against the dmalloc library manually.

Heap Debugging

Fast

Balanced

Thorough

Custom

Enabled Checks: -fence,free-protect,free-blank,alloc-blank

More Information

Heap Overflow/Underflow Detection

☒ Add guard pages to detect out of bounds heap access

Guard pages: 1

Add guard pages: After

Advanced

☐ Set node memory threshold at 90 percent

☐ Check heap consistency every 100 heap operations

☒ Store stack backtraces for memory allocations

☐ Only enable for these processes:

100%

Select All

x2

x0.5

1%

Help

OK

Cancel

# User Interface

Linaro DDT - Linaro Forge 23.0.2

File Edit View Control Tools Window Help

Current Group: All Focus on current: ☒ Group ☐ Process ☐ Thread ☐ Step Threads Together

All 0 1 2 3

Create Group

Project Files

Search (Ctrl+K)

Application Code

Sources

mmult.c

main(int argc, char \* argv[]): int  
mimult(int sz, double \* A): void  
mmult(int sz, int nslices, double \*,  
mwrite(int sz, double \* A, char \* fn

External Code

```
82  
83  
84 int main(int argc, char *argv[])  
85 {  
86     int mr, nproc, sz, slice;  
87     double *mat_a, *mat_b, *mat_c;  
88     char filename[32];  
89     int remainder;  
90     MPI_Status st;  
91  
92     MPI_Init (&argc, &argv);  
93     MPI_Comm_rank(MPI_COMM_WORLD, &mr); // my rank  
94     MPI_Comm_size(MPI_COMM_WORLD, &nproc); // number of processors  
95  
96     if (mr == 0)  
97     {  
98         printf(WORKED_EXAMPLE_NOTICE "\n\n");  
99     }  
100  
101     if (argc > 3 || (argc > 2 && strcmp(argv[1], "-h") != 0) )  
102     {  
103         if (mr == 0)  
104         {  
105             printf("Usage: ./mmult [-h] SIZE FILENAME\n \  
106                 \t-h: display this help message\n \  
107                 \tSIZE: size of the matrix to compute (default is %d)\n \  
108                 \tFILENAME: output matrix file name (default is %s)\n", DEFAULT_SIZE, DEFAULT_FN);  
109         }  
110  
111         return 1;  
112     }  
113     else  
114     {  
115         if (argc > 1)
```

Locals

Current Line(s)

Current Stack

Locals

Name	Value
<return value>	0
argc	1
argv	0x7fffffff008
mr	0
nproc	4199357
sz	0
slice	0
mat_a	0x7fffffff000
mat_b	0x400af0
mat_c	0x401370
filename	""
remainder	0
st	

Input/Output Breakpoints Watchpoints Stacks (All) Tracepoints Tracepoint Output Logbook Evaluate

Stacks (All)

Processes	Threads	Function
4	4	main (mmult.c:93)
4	4	progress_engine (opal_progress_threads.c:105)
4	4	progress_engine (pmix_progress_threads.c:232)
4	4	ucs_async_thread_func (thread.c:130)

Evaluate

Name	Value
slice	0
mr	0



# User Interface

Linaro DDT - Linaro Forge 23.0.2

File Edit View Control Tools Window Help

Current Group: All Focus on current: Group Process Thread Step Threads Together

All 0 1 2 3

Create Group

Project Files

Search (Ctrl+K)

Application Code

Sources

mmult.c

main(int argc, char \* argv[]): int  
munit(int sz, double \* A): void  
mmult(int sz, int nslices, double \*,  
mwrite(int sz, double \* A, char \* fn

External Code

```
82  
83  
84 int main(int argc, char *argv[])  
85 {  
86     int mr, nproc, sz, slice;  
87     double *mat_a, *mat_b, *mat_c;  
88     char filename[32];  
89     int remainder;  
90     MPI_Status st;  
91  
92     MPI_Init (&argc, &argv);  
93     MPI_Comm_rank(MPI_COMM_WORLD, &mr); // my rank  
94     MPI_Comm_size(MPI_COMM_WORLD, &nproc); // number of processors  
95  
96     if (mr == 0)  
97     {  
98         printf(WORKED_EXAMPLE_NOTICE "\n\n");  
99     }  
100  
101     if (argc > 3 || (argc > 2 && strcmp(argv[1], "-h") != 0) )  
102     {  
103         if (mr == 0)  
104         {  
105             printf("Usage: ./mmult [-h] SIZE FILENAME\n \n  
106                 \t-h: display this help message\n \n  
107                 \tSIZE: size of the matrix to compute (default is %d)\n \n  
108                 \tFILENAME: output matrix file name (default is %s)\n", DEFAULT_SIZE, DEFAULT_FN);  
109         }  
110  
111         return 1;  
112     }  
113     else  
114     {  
115         if (argc > 1)
```

Locals

Current Line(s)

Current Stack

Name Value

<return value> 0  
argc 1  
argv 0x7fffffffa008  
mr 0  
nproc 4199357  
sz 0  
slice 0  
mat\_a 0x7fffffffa000  
mat\_b 0x400af0  
mat\_c 0x401370  
filename ""  
remainder 0  
st

Input/Output Breakpoints Watchpoints Stacks (All) Tracepoints Tracepoint Output Logbook

Stacks (All)

Processes Threads Function

4 4 main (mmult.c:93)  
4 4 progress\_engine (opal\_progress\_threads.c:105)  
4 4 progress\_engine (pmix\_progress\_threads.c:232)  
4 4 ucs\_async\_thread\_func (thread.c:130)

Evaluate

Name Value

slice 0  
mr 0

DDT uses a tabbed interface

# User Interface

Linaro DDT - Linaro Forge 23.0.2

**DDT automatically finds the source code from the executable**

File Edit View Control Tools Window Help

Current Group: All Focus on current: Group Process Thread Stop Threads Together

Create Group

Project Files

Search (Ctrl+K)

Application Code

Sources

mmult.c

main(int argc, char \* argv[]): int  
munit(int sz, double \* A): void  
mmult(int sz, int nslices, double \*,  
mwrite(int sz, double \* A, char \* fn

External Code

```
82  
83  
84 int main(int argc, char *argv[])  
85 {  
86     int mr, nproc, sz, slice;  
87     double *mat_a, *mat_b, *mat_c;  
88     char filename[32];  
89     int remainder;  
90     MPI_Status st;  
91  
92     MPI_Init (&argc, &argv);  
93     MPI_Comm_rank(MPI_COMM_WORLD, &mr); // my rank  
94     MPI_Comm_size(MPI_COMM_WORLD, &nproc); // number of processors  
95  
96     if (mr == 0)  
97     {  
98         printf(WORKED_EXAMPLE_NOTICE "\n\n");  
99     }  
100  
101     if (argc > 3 || (argc > 2 && strcmp(argv[1], "-h") != 0))  
102     {  
103         if (mr == 0)  
104         {  
105             printf("Usage: ./mmult [-h] SIZE FILENAME\n \n  
106                 \t-h: display this help message\n \n  
107                 \tSIZE: size of the matrix to compute (default is %d)\n \n  
108                 \tFILENAME: output matrix file name (default is %s)\n", DEFAULT_SIZE, DEFAULT_FN);  
109         }  
110  
111         return 1;  
112     }  
113     else  
114     {  
115         if (argc > 1)
```

Locals

Current Line(s)

Current Stack

Locals

Name	Value
<return value>	0
argc	1
argv	0x7fffffffa008
mr	0
nproc	4199357
sz	0
slice	0
mat_a	0x7fffffffa000
mat_b	0x400af0
mat_c	0x401370
filename	""
remainder	0
st	

Input/Output Breakpoints Watchpoints Stacks (All) Tracepoints Tracepoint Output Logbook

Stacks (All)

Processes	Threads	Function
4	4	main (mmult.c:93)
4	4	progress_engine (opal_progress_threads.c:105)
4	4	progress_engine (pmix_progress_threads.c:232)
4	4	ucs_async_thread_func (thread.c:130)

Evaluate

Name	Value
slice	0
mr	0

# User Interface

Linaro DDT - Linaro Forge 23.0.2

File Edit View Control Tools Window Help

Current Group: All Focus on current: Group Process Thread Step Threads Together

All 0 1 2 3

Create Group

Project Files

Search (Ctrl+K)

Application Code

Sources

mmult.c

External Code

```
82
83
84
85
86 int mr, nproc, sz, slice;
87 double *mat_a, *mat_b, *mat_c;
88 char filename[32];
89 int remainder;
90 MPI_Status st;
91
92 MPI_Init (&argc, &argv);
93 MPI_Comm_rank(MPI_COMM_WORLD, &mr); // my rank
94 MPI_Comm_size(MPI_COMM_WORLD, &nproc); // number of processors
95
96 if (mr == 0)
97 {
98     printf(WORKED_EXAMPLE_NOTICE "\n\n");
99 }
100
101 if (argc > 3 || (argc > 2 && strcmp(argv[1], "-h") != 0))
102 {
103     if (mr == 0)
104     {
105         printf("Usage: ./mmult [-h] SIZE FILENAME\n \
106             \t-h: display this help message\n \
107             \tSIZE: size of the matrix to compute (default is %d)\n \
108             \tFILENAME: output matrix file name (default is %s)\n", DEFAULT_SIZE, DEFAULT_FN);
109     }
110     return 1;
111 }
112 else
113 {
114     if (argc > 1)
```

Locals

Current Line(s)

Current Stack

Locals

Name	Value
<return value>	0
argc	1
argv	0x7fffffffa008
mr	0
nproc	4199357
sz	0
slice	0
mat_a	0x7fffffffa000
mat_b	0x400af0
mat_c	0x401370
filename	""
remainder	0
st	

Input/Output Breakpoints Watchpoints Stacks (All) Tracepoints Tracepoint Output Logbook Evaluate

Stacks (All)

Processes	Threads	Function
4	4	main (mmult.c:93)
4	4	progress_engine (opal_progress_threads.c:105)
4	4	progress_engine (pmix_progress_threads.c:232)
4	4	ucs_async_thread_func (thread.c:130)

Evaluate

Name	Value
slice	0
mr	0

Can switch between process views

# User Interface

Linaro DDT - Linaro Forge 23.0.2

File Edit View Control Tools Window Help

Current Group: Group 1 Focus on current: Group Process Thread Step Threads Together

All 0 1 2 3  
Group 1 0 1  
Group 2 2 3

Create Group

Project Files

- Application Code
  - /
  - Sources
    - mmult.c
      - main: argc, char \*argv[]: int
      - main: int sz, double \*A[]: void
      - mmult: int sz, int ndices, double \*
      - main: int sz, double \*A, char \*fn
- External Code

```
1  int main(int argc, char *argv[])
2  {
3      int nr, nproc, sz, slice;
4      double *mat_a, *mat_b, *mat_c;
5      char filename[32];
6      int remainder;
7      MPI_Status st;
8
9      MPI_Init(&argc, &argv);
10     MPI_Comm_rank(MPI_COMM_WORLD, &nr); // my rank
11     MPI_Comm_size(MPI_COMM_WORLD, &nproc); // number of processors
12
13     if (nr == 0)
14     {
15         printf(WORKED_EXAMPLE_NOTICE "\n\n");
16     }
17
18     if (argc > 3 || (argc > 2 && strcmp(argv[1], "-h") != 0))
19     {
20         if (nr == 0)
21         {
22             printf("Usage: ./mmult [-h] SIZE FILENAME\n\n");
23             printf("\t-h: display this help message\n");
24             printf("\tSIZE: size of the matrix to compute (default is %d)\n",
25                   DEFAULT_SIZE);
26             printf("\tFILENAME: output matrix file name (default is %s)\n",
27                   DEFAULT_FN);
28         }
29         return 1;
30     }
31
32     if (argc == 2)
33     {
34         nr = 0;
35         nproc = 1;
36         sz = atoi(argv[1]);
37         slice = 1;
38         filename = DEFAULT_FN;
39     }
34     else if (argc == 3)
35     {
36         nr = 0;
37         nproc = atoi(argv[1]);
38         sz = atoi(argv[2]);
39         slice = 1;
40         filename = DEFAULT_FN;
41     }
42     else if (argc == 4)
43     {
44         nr = 0;
45         nproc = atoi(argv[1]);
46         sz = atoi(argv[2]);
47         slice = atoi(argv[3]);
48         filename = DEFAULT_FN;
49     }
50     else
51     {
52         nr = 0;
53         nproc = 1;
54         sz = 1024;
55         slice = 1;
56         filename = DEFAULT_FN;
57     }
58
59     if (nr != 0)
60     {
61         printf("Rank %d, nproc %d, sz %d, slice %d, filename %s\n",
62               nr, nproc, sz, slice, filename);
63     }
64
65     mat_a = malloc(sz * sz * sizeof(double));
66     mat_b = malloc(sz * sz * sizeof(double));
67     mat_c = malloc(sz * sz * sizeof(double));
68
69     if (nr == 0)
70     {
71         printf("Allocated memory for matrices\n");
72     }
73
74     MPI_File file;
75     MPI_File_open(MPI_COMM_WORLD, filename, MPI_MODE_CREATE | MPI_MODE_WRONLY,
76                  MPI_INFO_NULL, &file);
77     MPI_File_write_double(file, mat_a, sz * sz, MPI_ORDER_C, MPI_STATUS_IGNORE);
78     MPI_File_close(&file);
79
80     MPI_File file;
81     MPI_File_open(MPI_COMM_WORLD, filename, MPI_MODE_CREATE | MPI_MODE_WRONLY,
82                  MPI_INFO_NULL, &file);
83     MPI_File_write_double(file, mat_b, sz * sz, MPI_ORDER_C, MPI_STATUS_IGNORE);
84     MPI_File_close(&file);
85
86     MPI_File file;
87     MPI_File_open(MPI_COMM_WORLD, filename, MPI_MODE_CREATE | MPI_MODE_WRONLY,
88                  MPI_INFO_NULL, &file);
89     MPI_File_write_double(file, mat_c, sz * sz, MPI_ORDER_C, MPI_STATUS_IGNORE);
90     MPI_File_close(&file);
91
92     return 0;
93 }
```

Locals

Name	Value
argc	1
argv	0x7ffffff0
nr	0
nproc	4
sz	0
slice	0
mat_a	0x7ffffff0
mat_b	0x400a00
mat_c	0x401370
filename	mmult
remainder	0
st	

Input/Output Breakpoints Watchpoints Stacks (Group 1)

Processes Threads Function

1	1	main (mmult.c:98)
2	1	main (mmult.c:101)
2	2	progress_engine (opal_progress.c:101)
2	2	opal_ibevent2022_event_base_loop (event.c:1630)
2	2	epoll_dispatch (epoll.c:407)
2	2	progress_engine (pmix_progress.c:101)
2	2	opal_ibevent2022_event_base_loop (event.c:1630)
2	2	epoll_dispatch (epoll.c:407)
2	2	ucs_async_thread_func (thread.c:130)
2	2	ucs_event_set_wait (event_set.c:101)

Stacks (Group 1)

Process	Thread	Function
1	1	main (mmult.c:98)
2	1	main (mmult.c:101)
2	2	progress_engine (opal_progress.c:101)
2	2	opal_ibevent2022_event_base_loop (event.c:1630)
2	2	epoll_dispatch (epoll.c:407)
2	2	progress_engine (pmix_progress.c:101)
2	2	opal_ibevent2022_event_base_loop (event.c:1630)
2	2	epoll_dispatch (epoll.c:407)
2	2	ucs_async_thread_func (thread.c:130)
2	2	ucs_event_set_wait (event_set.c:101)

Ready Connected to: wlls3@niagara.scienet.utoronto.ca

James Willis (SciNet) Parallel Debugging with DDT October 22, 2025 28 / 58

# User Interface

Linaro DDT - Linaro Forge 23.0.2

File Edit View Control Tools Window Help

Current Group: All Focus on current: Group Process Thread Step Threads Together

0 1 2 3

Create Group

Project Files

Search (Ctrl+K)

Application Code

Sources

mmult.c

main(int argc, char \* argv[]) : int  
mini(int sz, double \* A) : void  
mmult(int sz, int nslices, double \*,  
mwrite(int sz, double \* A, char \* fn

External Code

mmult.c

```
82  
83  
84  
85  
86 int mr, nproc, sz, slice;  
87 double *mat_a, *mat_b, *mat_c;  
88  
89  
90 MPI_Status st;  
91  
92  
93  
94  
95  
96 if (mr == 0)  
97 {  
98     printf(WORKED_EXAMPLE_NOTICE "\n\n");  
99 }  
100  
101 if(argc > 3 || (argc > 2 && strcmp(argv[1], "-h") != 0) )  
102 {  
103     if(mr == 0)  
104     {  
105         printf("Usage: ./mmult [-h] SIZE FILENAME\n \  
106             \t-h: display this help message\n \  
107             \tSIZE: size of the matrix to compute (default is %d)\n \  
108             \tFILENAME: output matrix file name (default is %s)\n", DEFAULT_SIZE, DEFAULT_FN);  
109     }  
110  
111     return 1;  
112 }  
113 else  
114 {  
115     if(argc > 1)
```

Locals

Current Line(s)

Current Stack

Locals

Name	Value
<return value>	0
argc	1
argv	0x7fffffffa008
mr	0
nproc	4199357
sz	0
slice	0
mat_a	0x7fffffffa000
mat_b	0x400af0
mat_c	0x401370
filename	""
remainder	0
st	

Input/Output Breakpoints Watchpoints Stacks (All) Tracepoints Tracepoint Output Logbook

Stacks (All)

Processes	Threads	Function
4	4	main (mmult.c:93)
4	4	progress_engine (opal_progress_threads.c:105)
4	4	progress_engine (pmix_progress_threads.c:232)
4	4	ucs_async_thread_func (thread.c:130)

Evaluate

Name	Value
slice	0
mr	0

# User Interface

Linaro DDT - Linaro Forge 23.0.2

File Edit View Control Tools Window Help

Current Group: All Focus on current: Group Process Thread Stop Threads Together

All 0 1 2 3

Group 1 0 1

Group 2 2 3

Create Group

Project Files

Search (Ctrl+F)

Application Code

Sources

mmult.c

main(int argc, char \* argv): int

mmult(int sz, double \* A): void

mmult(int sz, int ndims, double \*\* A): void

mmult(int sz, double \* A, char \* fn): void

External Code

```
1  int main(int argc, char *argv[])
2  {
3      int nr, nproc, sz, slice;
4      double *mat_a, *mat_b, *mat_c;
5      char filename[32];
6      int remainder;
7      MPI_Status st;
8
9      MPI_Init(&argc, &argv);
10     MPI_Comm_rank(MPI_COMM_WORLD, &nr); // my rank
11     MPI_Comm_size(MPI_COMM_WORLD, &nproc); // number of processors
12
13     if (nr == 0)
14     {
15         printf(MORDED_EXAMPLE_NOTICE "\n\n");
16     }
17
18     if (argc > 3 || (argc > 2 && strcmp(argv[1], "-h") != 0))
19     {
20         if (nr == 0)
21         {
22             printf("Usage: ./mult [-h] SIZE FILENAME\n \
23             \t-h: display this help message\n \
24             \tSIZE: size of the matrix to compute (default is %d)\n \
25             \tFILENAME: output matrix file name (default is %s)\n", DEFAULT_SIZE, DEFAULT_FN);
26         }
27         return 1;
28     }
29 }
```

Locals Current Line(s) Current Stack

Current Line(s)

Name Value

Input/Output Breakpoints Watchpoints Stacks (All) Tracepoints Tracepoint Output Logbook

Breakpoints

Processes	Threads	File	Line	Function	Condition	Start After	Trigger Every	Stop After	Full path
<input checked="" type="checkbox"/> All	all	mmult.c	103	main		0	1	Forever	/scratch/scinet/willis2/ddt-mmult/mmult.c
<input checked="" type="checkbox"/> All	all	mmult.c	118	main		0	1	Forever	/scratch/scinet/willis2/ddt-mmult/mmult.c
<input checked="" type="checkbox"/> All	all	mmult.c	98	main		0	1	Forever	/scratch/scinet/willis2/ddt-mmult/mmult.c

Breakpoints Tab

Can suspend, jump to, delete, load, save

# User Interface

Linaro DDT - Linaro Forge 21.8.2

File Edit View Control Tools Window Help

Current Group: All Focus on current: Group Process Thread Stop Threads Together

All 0 1 2 3

Group 1 0 3

Group 2 2 3

Create Group

Project Files

Search (Ctrl+K)

mmult.c

main(int argc, char \*argv[]): int  
minkint sz, double \*A): void  
mmultint sz, int nslices, double  
minkint sz, double \*A, char \*

External Code

Headers

Sources

mmult.c  
mmult3.c  
power2.c  
address.c  
sort3.c  
alg.c  
alg2.c  
alg3.c  
alg4.c  
alg5.c  
alg6.c  
alg7.c  
alg8.c  
alg9.c  
alg10.c  
alg11.c  
alg12.c  
alg13.c  
alg14.c  
alg15.c  
alg16.c  
alg17.c  
alg18.c  
alg19.c  
alg20.c  
alg21.c  
alg22.c  
alg23.c  
alg24.c  
alg25.c  
alg26.c  
alg27.c  
alg28.c  
alg29.c  
alg30.c  
alg31.c  
alg32.c  
alg33.c  
alg34.c  
alg35.c  
alg36.c  
alg37.c  
alg38.c  
alg39.c  
alg40.c  
alg41.c  
alg42.c  
alg43.c  
alg44.c  
alg45.c  
alg46.c  
alg47.c  
alg48.c  
alg49.c  
alg50.c  
alg51.c  
alg52.c  
alg53.c  
alg54.c  
alg55.c  
alg56.c  
alg57.c  
alg58.c  
alg59.c  
alg60.c  
alg61.c  
alg62.c  
alg63.c  
alg64.c  
alg65.c  
alg66.c  
alg67.c  
alg68.c  
alg69.c  
alg70.c  
alg71.c  
alg72.c  
alg73.c  
alg74.c  
alg75.c  
alg76.c  
alg77.c  
alg78.c  
alg79.c  
alg80.c  
alg81.c  
alg82.c  
alg83.c  
alg84.c  
alg85.c  
alg86.c  
alg87.c  
alg88.c  
alg89.c  
alg90.c  
alg91.c  
alg92.c  
alg93.c  
alg94.c  
alg95.c  
alg96.c  
alg97.c  
alg98.c  
alg99.c  
alg100.c

mmult.c

```
177 {  
178     printf("r, "n");  
179 }  
180  
181 for(int i=0; i<sz/nslices; i++)  
182 {  
183     double res = 0.0;  
184     for(int j=0; j<sz; j++)  
185     {  
186         C[i*sz+j] += res;  
187     }  
188 }  
189  
190 int main(int argc, char *argv[])  
191 {
```

Locals Current Line(s) Current Stack

Current Stack

Stack Arguments

#1 main (argc=1, argv=0x7fffffa008) at /scratch/scratch/willis2/ddt-mmult/mm  
#0 mmult sz=64, nslices=4, A=0x7ffff32b6000, B=0x7ffff32a0000, C=0x7ffff32a0000

Input/Output Breakpoints Watchpoints Stacks (All) Tracepoints Tracepoint Output Logbook

Stacks (All)

Processes	Threads	Function
4	4	main (mmult.c:177)
4	4	mmult (mmult.c:181)
4	4	progress_engine (local_progress_threads.c:105)
4	4	progress_engine (pmix_progress_threads.c:232)
4	4	ucs_async_thread_func (thread.c:130)

Evaluate

Name	Value
slice	<No symbol "slice" in current context.>
nr	<No symbol "nr" in current context.>

# User Interface

Linux DDT - Linaro Forge 23.0.2

File Edit View Control Tools Window Help

Current Group: All Focus on current: Group Process Thread Stop Threads Together

All 0 1 2 3

Group 1 0 1

Group 2 2 3

Create Group

Project Files

Search (Ctrl+K)

Application Code

/

Sources

mmult.c

main(int argc, char \* argv[]) : int

main(int sz, double \* A) : void

mmult(int sz, int nlices, double \*,

mmult(int sz, double \* A, char \* fn

External Code

mmult.c

```
77 C(i*sz+) += res;
78 }
79
80 int main(int argc, char *argv[])
81 {
82     int mr, nproc, sz, slice;
83     double *mat_a, *mat_b, *mat_c;
84     char filename[32];
85     int remainder;
86     MPI_Status st;
87
88     MPI_Init(&argc, &argv);
89
90     MPI_Comm_rank(MPI_COMM_WORLD, &mr); // my rank
91     MPI_Comm_size(MPI_COMM_WORLD, &nproc); // number of processors
92
93     if (mr == 0)
94     {
95         printf(WORKED_EXAMPLE_NOTICE "\n\n");
96     }
97
98     if (argc > 3 || (argc > 2 && strcmp(argv[1], "-h") != 0))
99     {
100         if (mr == 0)
101             printf("Usage: ./mmult [-h] SIZE FILENAME\n \n");
102         printf("\t-h: display this help message\n \n");
103     }
104 }
```

Current line variables

Locals Current Line(s) Current Stack

Name	Value
mr	0

Input/Output Breakpoints Watchpoints Stacks (All) Tracepoints Tracepoint Output Logbook

Stacks (All)

Processes Threads Function

main (mmult.c:3)

progress\_engine (opal\_progress\_threads.c:105)

opal\_libevent2022\_event\_base\_loop (event.c:1630)

epoll\_dispatch (epoll.c:407)

progress\_engine (pmix\_progress\_threads.c:232)

opal\_libevent2022\_event\_base\_loop (event.c:1630)

epoll\_dispatch (epoll.c:407)

ucs\_async\_thread\_func (thread.c:130)

ucs\_event\_set\_wait (event\_set.c:198)

Evaluate

Name	Value
slice	0
mr	0

Ready. Connected to: willis2@niagara.scienet.utoronto.ca



# User Interface

Linero DDT - Linero Forge 23.0.2

File Edit View Control Tools Window Help

Current Group: All Focus on current: Group Process Thread Stop Threads Together

All 0 1 2 3

Group 1 0 1

Group 2 2 3

Create Group

Project Files

Search (Ctrl+F)

Application Code

mmult.c

main(int argc, char \*argv): int

main(int sz, double \*A): void

mmult(int sz, int nlices, double \*

mmult(int sz, double \*A, char \*fn

External Code

mmult.c

```
77 C[i*sz+j] += twa;
78 }
79 }
80
81 int main(int argc, char *argv[])
82 {
83     int mr, nproc, sz, slice;
84     double *mat_a, *mat_b, *mat_c;
85     char filename[32];
86     int remainder;
87     MPI_Status st;
88
89     MPI_Init(&argc, &argv);
90     MPI_Comm_rank(MPI_COMM_WORLD, &mr); // my rank
91     MPI_Comm_size(MPI_COMM_WORLD, &nproc); // number of processors
92
93     if (mr == 0)
94     {
95         printf(WORKED_EXAMPLE_NOTICE "\n\n");
96     }
97
98     if (argc > 3 || (argc > 2 && strcmp(argv[1], "-h") != 0))
99     {
100         if (mr == 0)
101         {
102             printf("Usage: ./mmult [-h] SIZE FILENAME\n \
103                 \t-h: display this help message\n \
104                 \n");
105         }
106     }
107 }
```

Locals

Current Line(s)

Current Stack

Name Value

<return value> 0

argc 1

argv 0x7ffffff9

mr 0

nproc 4199357

sz 0

slice 0

mat\_a 0x7ffffff0

mat\_b 0x400a0

mat\_c 0x401370

filename

remainder 0

st

Input/Output Breakpoints Watchpoints Stacks (All) Tracepoints Tracepoint Output Logbook

Stacks (All)

Processes Threads Function

4 4 main (mmult.c:1)

4 4 progress\_engine (opal\_progress\_threads.c:105)

4 4 opal\_libevent2022\_event\_base\_loop (event.c:1630)

4 4 opoll\_dispatch (epoll.c:407)

4 4 progress\_engine (pmix\_progress\_threads.c:232)

4 4 opal\_libevent2022\_event\_base\_loop (event.c:1630)

4 4 opoll\_dispatch (epoll.c:407)

4 4 ucs\_async\_thread\_func (thread.c:130)

4 4 ucs\_event\_set\_wait (event\_set.c:198)

Evaluate

Name Value

slice 0

mr 0

Ready Connected to: willis2@niagara.scienet.utoronto.ca

# User Interface

Linero DDT - Linaro Forge 23.0.2

File Edit View Control Tools Window Help

Current Group: All Focus on current: Group Process Thread Stop Threads Together

All 0 1 2 3

Group 1 0 1

Group 2 2 3

Create Group

Project Files

Search (Ctrl+F)

Application Code

mmult.c

main(int argc, char \*argv): int

main(int sz, double \*A): void

mmult(int sz, int slices, double \*

mmult(int sz, double \*A, char \*fn

External Code

mmult.c

```
77 C[i*sz+j] += twa;
78 }
79 }
80
81 int main(int argc, char *argv[])
82 {
83     int nr, nproc, sz, slice;
84     double *mat_a, *mat_b, *mat_c;
85     char filename[32];
86     int remainder;
87     MPI_Status st;
88
89     MPI_Init(&argc, &argv);
90     MPI_Comm_rank(MPI_COMM_WORLD, &nr); // my rank
91     MPI_Comm_size(MPI_COMM_WORLD, &nproc); // number of processors
92
93     if (nr == 0)
94     {
95         printf(WORKED_EXAMPLE_MESSAGE "\n");
96     }
97
98     if (argc > 3 || (argc > 2 && strcmp(argv[1], "-h") != 0))
99     {
100         if (nr == 0)
101         {
102             printf("Usage: ./mmult [-h] SIZE FILENAME\n \
103                 \t-h: display this help message\n \
104                 \n");
105         }
106     }
107 }
```

Locals

Current Line(s)

Current Stack

Name Value

<return value> 0

argc 1

argv 0x7ffffff0

nr 0

nproc 4199357

sz 0

slice 0

mat\_a 0x7ffffff0

mat\_b 0x400a00

mat\_c 0x4013f0

filename

remainder 0

st

Input/Output Breakpoints Watchpoints Stacks (All) Tracepoints Tracepoint Output Logbook

Stacks (All)

Processes Threads Function

0 4 main (mmult.c:1)

4 4 progress\_engine (opal\_progress\_threads.c:105)

4 4 opal\_libevent2022\_event\_base\_loop (event.c:1630)

4 4 opoll\_dispatch (epoll.c:407)

4 4 progress\_engine (pmix\_progress\_threads.c:232)

4 4 opal\_libevent2022\_event\_base\_loop (event.c:1630)

4 4 opoll\_dispatch (epoll.c:407)

4 4 ucs\_async\_thread\_func (thread.c:130)

4 4 ucs\_event\_set\_wait (event\_set.c:198)

Evaluate

Name Value

slice 0

nr 0

Ready Connected to: willis2@nisiagara.scienet.utoronto.ca

SciNet

ADVANCED RESEARCH COMPUTING @ U. TORONTO

# DDT Setup Demonstration

# Hands-on hello-mpi Example

- Login to Teach:

```
laptop:~$ ssh -X USERNAME@teach.scinet.utoronto.ca
```

- Load compilers, MPI library and ddt:

```
teach-login01:~$ module load StdEnv/2023 gcc openmpi ddt-cpu/23.1.1
```

- Copy examples from the course directory:

```
teach-login01:~$ cp -r /home/l/lcl_uothpc245/hpc245starter/ddt-examples .
```

- Compile MPI Hello World example, hello-mpi.c:

```
teach-login01:~/ddt-examples/ddt-hello-mpi$ mpicc -g hello-mpi.c -o hello-mpi
```

- Run ddt:

```
teach-login01:~/ddt-examples/ddt-hello-mpi$ ddt -n <numprocs> hello-mpi
```

- Experiment with different features of DDT

# Client-server Mode

- This mode can be very beneficial if you have a slow internet connection
- Keeps the bulk of the computation on Teach (*server*)
- Only sends minimal amounts of information (network traffic) to your locally running version of DDT (*client*)
- Results in a much smoother experience, avoids slow/laggy interface

# Client-server Mode Setup

## Setting up the server side

- Connect to Teach and create a startup script which will be run by the server and load the modules that your code needs:

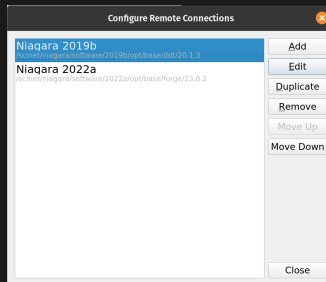
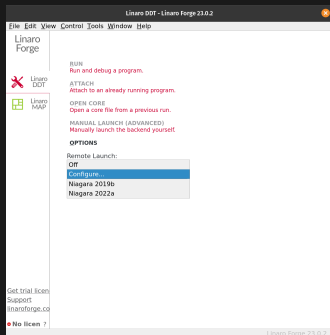
```
#!/bin/bash
module purge
module load StdEnv/2023
module load gcc openmpi
module load ddt-cpu/23.1.1
module load python/3.11.5
module load scipy-stack/2025a mpi4py/4.0.3
export OMPI_MCA_pml=ob1
export ARM_TOOLS_CONFIG_DIR=${HOME}/.arm
mkdir -p ${ARM_TOOLS_CONFIG_DIR}
```

- Name it `ddt_remote_setup.sh` and place it in `$HOME`

# Client-server Mode Setup

## Setting up the client side

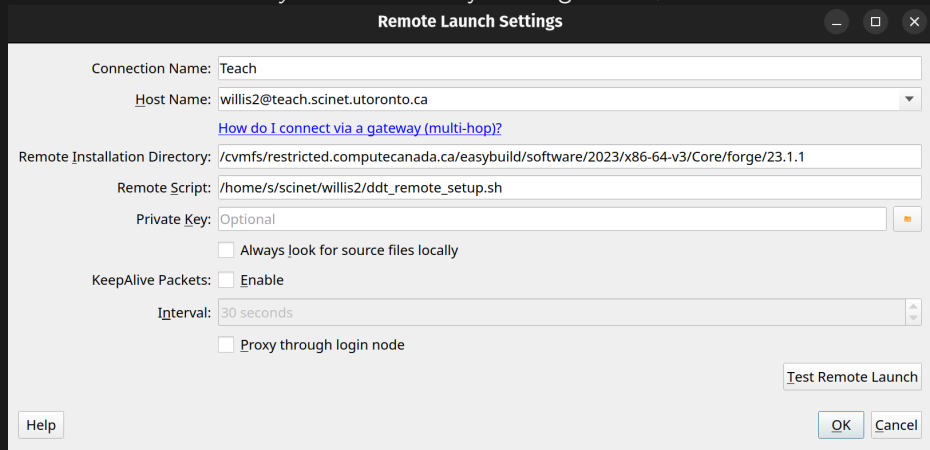
- 1 Download DDT on your local machine from Linaro and make sure the version matches the one on Teach (23.1.1): <https://www.linaroforge.com/downloadForge/>
- 2 Launch `ddt` and select *Configure* from *Remote Connections*



# Client-server Mode Setup

- 3 Click *Add* and fill out the fields as shown below

Note: *Remote Installation Directory* can be found by running `echo $EBROOTALLINEA` on Teach



The screenshot shows a window titled "Remote Launch Settings" with standard window controls (minimize, maximize, close) in the top right corner. The dialog contains several input fields and checkboxes:

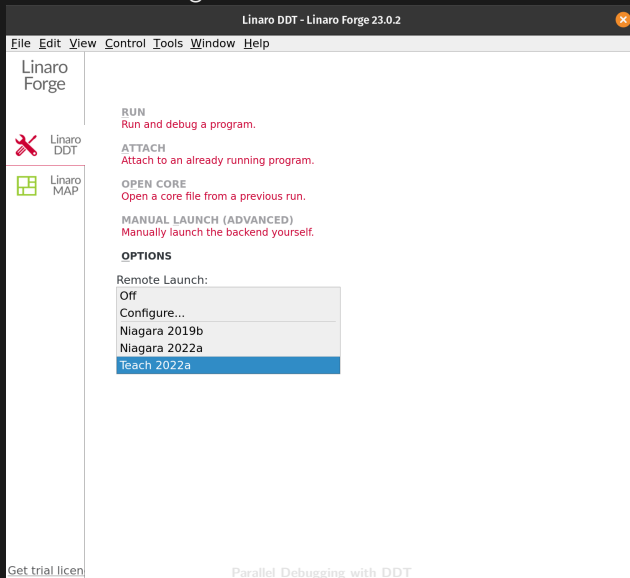
- Connection Name:** Teach
- Host Name:** willis2@teach.scinet.utoronto.ca (with a dropdown arrow)
- Remote Installation Directory:** /cvmfs/restricted.computeCanada.ca/easybuild/software/2023/x86-64-v3/Core/forge/23.1.1
- Remote Script:** /home/s/scinet/willis2/ddt\_remote\_setup.sh
- Private Key:** Optional (with a file selection icon)
- ☐ Always look for source files locally
- KeepAlive Packets:** ☐ Enable
- Interval:** 30 seconds (with a spinner control)
- ☐ Proxy through login node

At the bottom right, there is a "Test Remote Launch" button. At the bottom left, there is a "Help" button. At the bottom right, there are "OK" and "Cancel" buttons.



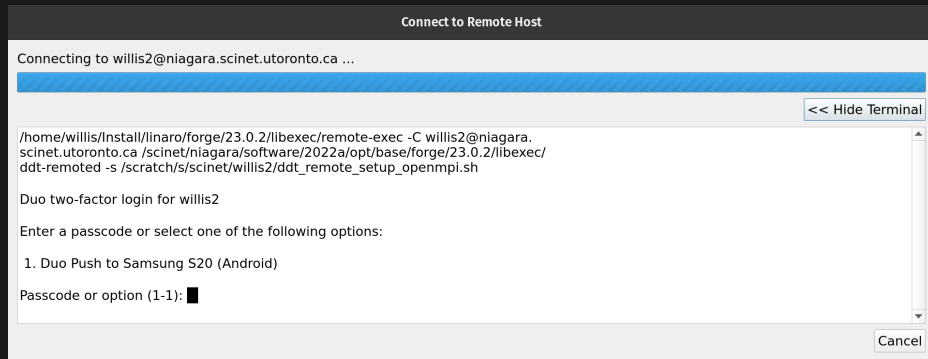
# Client-server Mode Setup

- 4 Click *OK* and now the DDT starting screen should look like this:



# Client-server Mode Setup

- 5 If you have MFA enabled follow the instructions outlined in the text box:



- More detailed instructions can be found here: [https://docs.linaroforge.com/23.0.2/html/forged/forged/connecting\\_to\\_a\\_remote\\_system/connecting\\_remotely.html](https://docs.linaroforge.com/23.0.2/html/forged/forged/connecting_to_a_remote_system/connecting_remotely.html)
- Please try setting up DDT with client-server mode

# Memory Debugging in DDT

- Memory debugging can be turned on in the *Run* window
- Causes the code to stop on an error i.e. memory corruption/leak
- Allows you to check the pointer where the memory corruption has occurred
- Can give an overall view of the memory stats/usage
- Lets look at a real example


# Matrix-Matrix Multiply Example

- Imagine we want to compute the result of this Matrix equation in parallel with MPI:

$$C = A * B + C$$

- The algorithm works as follows:
  - 1 Rank 0 initialises  $A$ ,  $B$  and  $C$
  - 2 Rank 0 sends the entire matrix  $B$ , with slices of  $A$  and  $C$  to all other ranks
  - 3 Each rank performs matrix multiplication on their domain and computes a slice of  $C$
  - 4 Rank 0 collects the slices of  $C$  from each rank and forms the final matrix  $C$
  - 5 Rank 0 writes  $C$  to a file

# Matrix-Matrix Multiply Example



# Hands-on Matrix-Matrix Multiply

- Change to the ddt-mmult directory from the course examples and compile the code:

```
teach-login01:~/ddt-mmult$ make
```

- This will build C and Fortran executables with `-g -O0` named `mmult_c` and `mmult_f`
- The example can then be run with:

```
teach-login01:~/ddt-mmult$ mpirun -np 4 ./mmult_c
```

- For python, load a python, scipy and mpi4py module before compiling the C and Fortran libraries with:


```
teach-login01:~/ddt-mmult$ ml python/3.11.5 scipy-stack/2025a mpi4py/4.0.3  
teach-login01:~/ddt-mmult$ make -f mmult_py.makefile
```

- Then run the example:

```
mpirun -np 4 python ./mmult.py
```

# Hands-on Matrix-Matrix Multiply


- Try running the code. What output do you get?
- Run the code in DDT to find out what the error is
- Note: if running with python you will need the setup shown below



or from the command line:

```
teach-login01:~/ddt-mmult$ ddt -n 4 python %allinea_python_debug% ./mmult.py
```

- Can you locate the error?
- Can you fix it?
- Hints: Try running with *memory debugging* enabled and make sure **Add guard pages** is enabled



# Matrix-Matrix Multiply Demonstration



# Running DDT on a GPU

- DDT can also debug code on GPUs
- It supports both Nvidia and AMD GPUs
- The Teach cluster has no GPUs but Trillium has Nvidia H100 GPUs available
- Let's look at how to run a simple vector addition GPU example with DDT on Trillium

# Running DDT on a GPU

- On Trillium you will need to load the ddt-gpu module instead of ddt-cpu and cuda:

```
laptop:~$ ssh USERNAME@trillium-gpu.scinet.utoronto.ca
trig-login01:~$ module load StdEnv/2023 intel/2023.2.1
trig-login01:~$ module load cuda/12.6 ddt-gpu/23.1.1
```

- Debugging on a GPU also requires extra flags “-g -G --cudart shared” during compilation:

```
trig-login01:~/ddt-vec-add$ nvcc -g -G --cudart shared vec-add.cu -o vec-add
```

- Then you can run DDT as normal, for example:

```
trig-login01:~/ddt-vec-add$ ddt ./vec-add 1000000
```


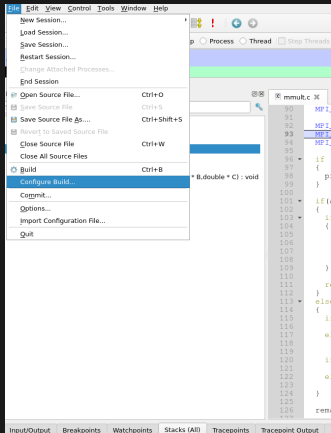
- More information on using DDT with GPUs can be found on the Linaro website here:  
[https://docs.linaroforge.com/23.0.1/html/forge/ddt/gpu\\_debugging/index.html](https://docs.linaroforge.com/23.0.1/html/forge/ddt/gpu_debugging/index.html)
- DDT GPU demonstration

# Other Useful Features of DDT

- Editing and recompiling code from within DDT GUI
- Attaching to a running job
- Submit SLURM jobs with DDT
- Running with core files
- Running DDT from a browser (Open OnDemand)

# Editing and Compiling

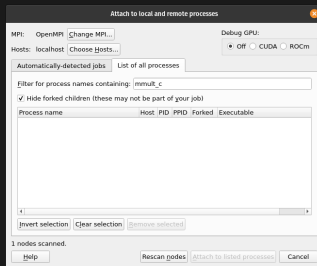
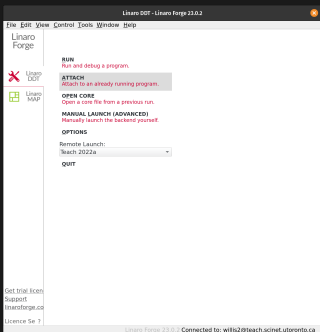
- DDT also has the ability to edit and recompile source code on-the-fly
- Making it much easier to try potential bug fixes



- Build demonstration

# Attach to a Running Job

- DDT allows you to attach to an already running job
- For example, say you have submitted a job to the scheduler on Teach and want to monitor it
- You can use the *Attach* button
- More detailed instructions can be found here: [https://docs.linaroforge.com/23.0.2/html/forge/ddt/get\\_started\\_ddt/attaching\\_to\\_running\\_programs.html#index-9](https://docs.linaroforge.com/23.0.2/html/forge/ddt/get_started_ddt/attaching_to_running_programs.html#index-9)



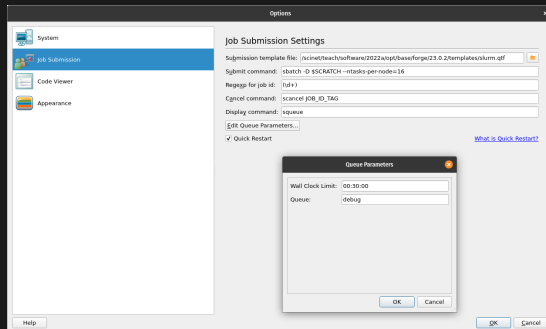

• Attach demonstration

# Submit SLURM Jobs

- DDT allows you to submit jobs directly to the SLURM scheduler on Teach
- The job will be monitored in the queue and as soon as it runs DDT will attach to the job

## Setup

- Click *Run -> Submit to Queue -> Configure*



# Running .core Files

- When your code terminates unexpectedly it will generate what is known as a *core dump*
- A core dump is a set of files ending in `.core` per process running
- Each `.core` file contains the process's address space (memory) at the time of the crash
- DDT allows you to run with the core files showing the state of the code at the time of the crash
- Can be useful if your job fails after running for a long time
- If no core dump is generated check that `ulimit -c` is set to `unlimited`, this sets the maximum size a `.core` file can be
- Demonstration

# Running DDT from a browser (Open OnDemand)

- DDT can also be run inside the new Open OnDemand portal via the Remote Desktop application
- This allows you to run DDT from your browser
- Once you are in the remote desktop, you can launch DDT as you would on Trillium or Teach
- This is a great way to run DDT if you are having issues with X11 forwarding or can't install DDT on your local machine
- More information on Open OnDemand can be found here:  
[https://docs.scinet.utoronto.ca/index.php/Open\\_OnDemand\\_Quickstart](https://docs.scinet.utoronto.ca/index.php/Open_OnDemand_Quickstart)
- Demonstration



# Summary

- DDT is a powerful graphical debugger
- Supports parallel debugging in multiple languages (e.g. C, C++, Fortran, Python)
- Supports MPI, OpenMP, threads, CUDA and more
- DDT documentation: <https://docs.linaroforge.com/23.0.2/html/forgeddt/index.html>

## Support

Questions? Need help?

Don't be afraid to contact us! We are here to help.

- Email to [support@scinet.utoronto.ca](mailto:support@scinet.utoronto.ca) or to [trillium@computecanada.ca](mailto:trillium@computecanada.ca)

# References

- Slide 19: Linaro DDT
- Slide 38: Matrix-Matrix Multiply Worked Example
- Slide 44: Client-Server Mode
- Slide 51: Attach Mode
- Slide 52: Submitting to a Queue