

Debugging

Marcelo Ponce
SciNet HPC Consortium
University of Toronto

July 15, 2016



Outline

- ▶ Debugging Basics
- ▶ Debugging with the command line: GDB
- ▶ Memory debugging with the command line: valgrind
- ▶ (Parallel) Debugging with DDT

Debugging basics

Debugging basics

Debugging basics

Help, my program doesn't work!

Debugging basics

Help, my program doesn't work!

```
$ gcc -O3 answer.c  
$ ./a.out  
Segmentation fault
```

Debugging basics

Help, my program doesn't work!

```
$ gcc -O3 answer.c  
$ ./a.out  
Segmentation fault
```

a miracle occurs

Debugging basics

Help, my program doesn't work!



a miracle occurs



My program works brilliantly!

```
$ gcc -O3 answer.c  
$ ./a.out  
Segmentation fault
```


Debugging basics

Help, my program doesn't work!

```
$ gcc -O3 answer.c  
$ ./a.out  
Segmentation fault
```

a miracle occurs

My program works brilliantly!

```
$ gcc -O3 answer.c  
$ ./a.out  
42
```

Debugging basics

Help, my program doesn't work!

```
$ gcc -O3 answer.c  
$ ./a.out  
Segmentation fault
```

a miracle occurs

My program works brilliantly!

```
$ gcc -O3 answer.c  
$ ./a.out  
42
```

- ▶ Unfortunately, “miracles” are not yet supported by SciNet.

Debugging basics

Help, my program doesn't work!

```
$ gcc -O3 answer.c  
$ ./a.out  
Segmentation fault
```

a miracle occurs

My program works brilliantly!

```
$ gcc -O3 answer.c  
$ ./a.out  
42
```

- ▶ Unfortunately, “miracles” are not yet supported by SciNet.

Debugging:

Methodical process of finding and fixing flaws in software

Common symptoms

Errors at compile time

Common symptoms

Errors at compile time

- ▶ Syntax errors: easy to fix
- ▶ Library issues
- ▶ Cross-compiling
- ▶ Compiler warnings

Common symptoms

Errors at compile time

- ▶ Syntax errors: easy to fix
- ▶ Library issues
- ▶ Cross-compiling
- ▶ Compiler warnings

Always switch this on, and fix or understand them!

Common symptoms

Errors at compile time

- ▶ Syntax errors: easy to fix
- ▶ Library issues
- ▶ Cross-compiling
- ▶ Compiler warnings

Always switch this on, and fix or understand them!

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

Common symptoms

Errors at compile time

- ▶ Syntax errors: easy to fix
- ▶ Library issues
- ▶ Cross-compiling
- ▶ Compiler warnings

Always switch this on, and fix or understand them!

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

Runtime errors

Common symptoms

Errors at compile time

- ▶ Syntax errors: easy to fix
- ▶ Library issues
- ▶ Cross-compiling
- ▶ Compiler warnings

Always switch this on, and fix or understand them!

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

Runtime errors

- ▶ Floating point exceptions
- ▶ Segmentation fault
- ▶ Aborted
- ▶ Incorrect output (nans)

Common issues

Arithmetic	corner cases (<code>sqrt(-0.0)</code>), infinities
Memory access	Index out of range, uninitialized pointers.
Logic	Infinite loop, corner cases
Misuse	wrong input, ignored error, no initialization
Syntax	wrong operators/arguments
Resource starvation	memory leak, quota overflow
Parallel	race conditions, deadlock

What is going on?

- ▶ Almost always, a condition you are sure is satisfied, is not.
- ▶ But your programs likely relies on many such assumptions.
- ▶ First order of business is finding out what goes wrong, and what assumption is not warranted.
- ▶ *Debugger*: program to help detect errors in other programs.
- ▶ **You are the real debugger.**

How to avoid debugging:

- ▶ Write better code.
 - ▶ Simpler, clear, straightforward code.
 - ▶ Modularity (no global variables or 10,000-line functions)
 - ▶ Avoid 'cute' tricks (no obfuscated C code winners)
- ▶ Don't write code, use existing libraries
- ▶ Write (simple) tests for each part of your code
- ▶ Use version control so you can 'roll back'.

Debugging Workflow

First things first:

- ▶ As soon as you are convinced there is a real problem, create the simplest situation in which it reproducibly occurs.
- ▶ 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

Ways to debug

- ▶ Preemptive:
 - ▶ Turn on compiler warnings: fix or understand them!
`$ gcc/gfortran -Wall`
 - ▶ Check your assumptions (e.g. use **assert**).

Ways to debug

- ▶ Preemptive:
 - ▶ Turn on compiler warnings: fix or understand them!
`$ gcc/gfortran -Wall`
 - ▶ Check your assumptions (e.g. use **assert**).
- ▶ Inspect the exit code and read the error messages!

Ways to debug

- ▶ Preemptive:
 - ▶ Turn on compiler warnings: fix or understand them!
`$ gcc/gfortran -Wall`
 - ▶ Check your assumptions (e.g. use **assert**).
- ▶ Inspect the exit code and read the error messages!
- ▶ Use a debugger

Ways to debug

- ▶ Preemptive:
 - ▶ Turn on compiler warnings: fix or understand them!
`$ gcc/gfortran -Wall`
 - ▶ Check your assumptions (e.g. use **assert**).
- ▶ Inspect the exit code and read the error messages!
- ▶ Use a debugger
- ▶ Add print statements

Ways to debug

- ▶ Preemptive:
 - ▶ Turn on compiler warnings: fix or understand them!
`$ gcc/gfortran -Wall`
 - ▶ Check your assumptions (e.g. use `assert`).
- ▶ Inspect the exit code and read the error messages!
- ▶ Use a debugger
- ▶ Add print statements ← **No way to debug!**

What's wrong with using print statements?

Strategy

What's wrong with using print statements?

Strategy

- ▶ Constant cycle:

What's wrong with using print statements?

Strategy

- ▶ Constant cycle:
 1. strategically add print statements

What's wrong with using print statements?

Strategy

- ▶ Constant cycle:
 1. strategically add print statements
 2. compile

What's wrong with using print statements?

Strategy

- ▶ Constant cycle:
 1. strategically add print statements
 2. compile
 3. run

What's wrong with using print statements?

Strategy

- ▶ Constant cycle:
 1. strategically add print statements
 2. compile
 3. run
 4. analyze output

What's wrong with using print statements?

Strategy

- ▶ Constant cycle:
 1. strategically add print statements
 2. compile
 3. run
 4. analyze output *bug not found?*

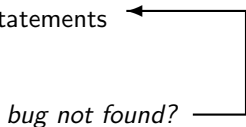
What's wrong with using print statements?

Strategy

► Constant cycle:

1. strategically add print statements
2. compile
3. run
4. analyze output

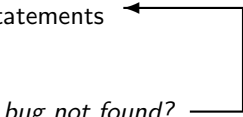
bug not found?



What's wrong with using print statements?

Strategy

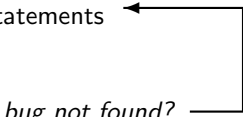
- ▶ Constant cycle:
 1. strategically add print statements
 2. compile
 3. run
 4. analyze output

bug not found? 
- ▶ Removing the extra code after the bug is fixed

What's wrong with using print statements?

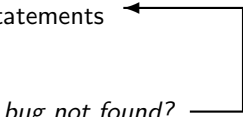
Strategy

- ▶ Constant cycle:
 1. strategically add print statements
 2. compile
 3. run
 4. analyze output

bug not found?
 - ▶ Removing the extra code after the bug is fixed
 - ▶ Repeat for each bug
- 

What's wrong with using print statements?

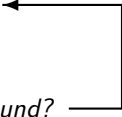
Strategy

- ▶ Constant cycle:
 1. strategically add print statements
 2. compile
 3. run
 4. analyze output
 - ▶ Removing the extra code after the bug is fixed
 - ▶ Repeat for each bug
- bug not found?* 

Problems with this approach

What's wrong with using print statements?

Strategy

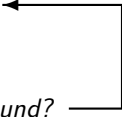
- ▶ Constant cycle:
 1. strategically add print statements
 2. compile
 3. run
 4. analyze output
 - ▶ Removing the extra code after the bug is fixed
 - ▶ Repeat for each bug
- 
- bug not found?*

Problems with this approach

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

What's wrong with using print statements?


Strategy

- ▶ Constant cycle:
 1. strategically add print statements
 2. compile
 3. run
 4. analyze output
 - ▶ Removing the extra code after the bug is fixed
 - ▶ Repeat for each bug
- 
- bug not found?*

Problems with this approach

- ▶ Time consuming
- ▶ Error prone
- ▶ Changes memory, timing... **There's a better way!**

Symbolic debuggers



Symbolic debuggers

Features

Symbolic debuggers

Features

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

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?

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 (SciNet): can be slow

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

Symbolic debuggers

Preparing the executable

- ▶ Add required compilation flags:
 - \$ gcc/g++/gfortran -g [-gstabs]
 - \$ icc/icpc/ifort -g [-debug parallel]
 - \$ nvcc -g -G
- ▶ Optional: switch off optimization -O0


Symbolic debuggers

Preparing the executable

- ▶ Add required compilation flags:
 - \$ gcc/g++/gfortran -g [-gstabs]
 - \$ icc/icpc/ifort -g [-debug parallel]
 - \$ nvcc -g -G
- ▶ Optional: switch off optimization -O0

Command-line based symbolic debuggers: gdb

GDB



What is GDB?

- ▶ Free, GNU license, symbolic debugger.
- ▶ Available on many systems.
- ▶ Been around for a while, but still developed and up-to-date
- ▶ Text based, but has a '-tui' option.

```
$ module load gcc/4.7.2
$ gcc -Wall -g -O0 example.c -o example
$ module load gdb/7.6
$ gdb -tui example
...
(gdb)_
```

GDB basic building blocks



Demonstration of GDB features

- ▶ We will look at the features of gdb using a running example.
- ▶ Example reads integers from command line and sums them.
- ▶ There's a C and a Fortran version.

```
$ ssh USER@login.scinet.utoronto.ca -X
$ ssh gpc01 -X
$ qsub -l nodes=1:ppn=8,walltime=8:00:00 -I -X -qteach
$ cp -r /scinet/course/ss2016/debug $SCRATCH
$ source $SCRATCH/debug/code/setup
$ cd $SCRATCH/debug/code/bugexample
$ make bugexample #(or make bugexample_f)
```

GDB building block #1: Inspect crashes

Inspecting core files

GDB building block #1: Inspect crashes

Inspecting core files

Core = file containing state of program after a crash

GDB building block #1: Inspect crashes

Inspecting core files

Core = file containing state of program after a crash

- ▶ needs max core size set (`ulimit -c <number>`)
- ▶ gdb reads with `gdb <executable> <corefile>`
- ▶ it will show you where the program crashed

GDB building block #1: Inspect crashes

Inspecting core files

Core = file containing state of program after a crash

- ▶ needs max core size set (`ulimit -c <number>`)
- ▶ gdb reads with `gdb <executable> <corefile>`
- ▶ it will show you where the program crashed

No core file?

GDB building block #1: Inspect crashes

Inspecting core files

Core = file containing state of program after a crash

- ▶ needs max core size set (`ulimit -c <number>`)
- ▶ gdb reads with `gdb <executable> <corefile>`
- ▶ it will show you where the program crashed

No core file?

- ▶ can start gdb as `gdb <executable>`
- ▶ type **run** to start program
- ▶ gdb will show you where the program crashed if it does.

GDB building block #1: Inspect crashes

Inspecting core files

Core = file containing state of program after a crash

- ▶ needs max core size set (`ulimit -c <number>`)
- ▶ gdb reads with `gdb <executable> <corefile>`
- ▶ it will show you where the program crashed

No core file?

- ▶ can start gdb as `gdb <executable>`
- ▶ type **run** to start program
- ▶ gdb will show you where the program crashed if it does.

Related gdb commands

run	run the executable from the start
list	list code lines (where current execution is, or range)

GDB building block #1: Inspect crashes

```
$ ulimit -c 1024
```

```
$ ./bugexample #(or ./bugexample_f)
```

Give some integers as command-line arguments

```
$ ./bugexample 1 3 5
```

```
Segmentation fault (core dumped)
```

GDB building block #1: Inspect crashes

```
$ ulimit -c 1024
```

```
$ ./bugexample #(or ./bugexample_f)
```

Give some integers as command-line arguments

```
$ ./bugexample 1 3 5
```

```
Segmentation fault (core dumped)
```

```
$ gdb ./bugexample core.2387 # core number varies
```

GDB building block #1: Inspect crashes

```
$ ulimit -c 1024
$ ./bugexample #(or ./bugexample_f)
Give some integers as command-line arguments
$ ./bugexample 1 3 5
Segmentation fault (core dumped)

$ gdb ./bugexample core.2387 # core number varies
```

```
GNU gdb (GDB) 7.6
```

```
Copyright (C) 2013 Free Software Foundation, Inc.
```

```
...
```

```
Reading symbols from debug/code/bugexample/bugexample...done.
```

```
[New LWP 3817]
```

```
warning: Can't read pathname for load map: Input/output error.
```

```
Core was generated by './bugexample 1 3 5'.
```

```
Program terminated with signal 11, Segmentation fault.
```

```
#0  0x4007d5 in sum_integers (n=3, a=0x4) at intlittestools.c:30
```

```
30      s += a[i];
```

```
(gdb)
```

GDB building block #1: Inspect crashes

...

Program terminated with signal 11, Segmentation fault.

```
#0 0x4007d5 in sum_integers (n=3, a=0x4) at intlisttools.c:30
30          s += a[i];
```

This points at the line where the error is detected.

GDB building block #1: Inspect crashes

...

Program terminated with signal 11, Segmentation fault.

```
#0 0x4007d5 in sum_integers (n=3, a=0x4) at intlisttools.c:30
30          s += a[i];
```

This points at the line where the error is detected. More context:

GDB building block #1: Inspect crashes

...

Program terminated with signal 11, Segmentation fault.

```
#0 0x4007d5 in sum_integers (n=3, a=0x4) at intlisttools.c:30
30          s += a[i];
```

This points at the line where the error is detected. More context:

```
(gdb) list
```

GDB building block #1: Inspect crashes

...

Program terminated with signal 11, Segmentation fault.

```
#0 0x4007d5 in sum_integers (n=3, a=0x4) at intlisttools.c:30
30      s += a[i];
```

This points at the line where the error is detected. More context:

```
(gdb) list
```

```
25 /* Compute the sum of the array of integers */
26 int sum_integers(int n, int* a)
27 {
28     int i, s;
29     for (i=0; i<n; i++)
30         s += a[i];
31     return s;
32 }
```

```
(gdb)
```


GDB building block #2: Function call stack

Interrupting program

- ▶ Press Ctrl-C while program is running in gdb
- ▶ gdb will show you where the program was.

GDB building block #2: Function call stack

Interrupting program

- ▶ Press Ctrl-C while program is running in gdb
- ▶ gdb will show you where the program was.

Stack trace

- ▶ From what functions was this line reached?
- ▶ What were the arguments of those function calls?

GDB building block #2: Function call stack

Interrupting program

- ▶ Press Ctrl-C while program is running in gdb
- ▶ gdb will show you where the program was.

Stack trace

- ▶ From what functions was this line reached?
- ▶ What were the arguments of those function calls?

`gdb` commands

<code>backtrace</code>	function call stack
<code>continue</code>	continue
<code>down</code>	go to called function
<code>up</code>	go to caller

GDB building block #2: Function call stack

...

```
(gdb) list
```

```
25 /* Compute the sum of the array of integers */
```

```
26 int sum_integers(int n, int* a)
```

```
27 {
```

```
28     int i, s;
```

```
29     for (i=0; i<n; i++)
```

```
30         s += a[i];
```

```
31     return s;
```

```
32 }
```

```
(gdb) backtrace
```

```
#0 0x4007d5 in sum_integers (n=3,a=0x4) at intlittestools.c:30
```

```
#1 0x40082a in process (argc=4,argv=0x7fff0b89ce58) at process.c
```

```
#2 0x4006d3 in main (argc=4,argv=0x7fff0b89ce58) at bugexample.c
```

```
(gdb)
```

GDB building block #3: Variables

Checking a variable

- ▶ Can print the value of a variable
- ▶ Can keep track of variable (print at prompt)
- ▶ Can stop the program when variable changes
- ▶ Can change a variable (“what if . . .”)

GDB building block #3: Variables

Checking a variable

- ▶ Can print the value of a variable
- ▶ Can keep track of variable (print at prompt)
- ▶ Can stop the program when variable changes
- ▶ Can change a variable (“what if ...”)

`gdb` commands

<code>print</code>	print variable
<code>display</code>	print at every prompt
<code>set variable</code>	change variable
<code>watch</code>	stop if variable changes

GDB building block #3: Variables

Remember: We were looking at a seg fault in `s += a[i]`.

```
(gdb) print i
```

```
0
```

```
(gdb) print a[0]
```

```
Cannot access memory at address 0x4
```

```
(gdb) print a
```

```
0x4
```

```
(gdb) up
```

```
#1 0x00000000040082a in process (argc=4, argv=0x7fff0b89ce58)
```

```
11 int s = sum_integers(n, arg);
```

```
(gdb) print arg
```

```
$1 = (int *) 0x4
```

```
(gdb) list
```

```
7 void process(int argc, char** argv)
```

```
8 {
```

```
9     int* arg = read_integer_arguments(argc, argv);
```

```
10     int n = argc-1;
```

```
11     int s = sum_integers(n, arg);
```

```
12     print_integers(n, arg);
```

```
13     printf("Sum of integers is: %d\n", s);
```

GDB building block #4: Automatic interruption

Breakpoints

- ▶ **break** [**file:**]**<line>**|**<function>**
- ▶ each breakpoint gets a number
- ▶ when run, automatically stops there
- ▶ can add conditions, temporarily remote breaks, etc.

GDB building block #4: Automatic interruption

Breakpoints

- ▶ **break** [**file:**]**<line>**|**<function>**
- ▶ each breakpoint gets a number
- ▶ when run, automatically stops there
- ▶ can add conditions, temporarily remove breaks, etc.

Related gdb commands

delete	unset breakpoint
condition	break if condition met
disable	disable breakpoint
enable	enable breakpoint
info breakpoints	list breakpoints
tbreak	temporary breakpoint

GDB building block #4: Automatic interruption

```
...
(gdb) list
7 void process(int argc, char** argv)
8 {
9     int* arg = read_integer_arguments(argc, argv);
10    int n = argc-1;
11    int s = sum_integers(n, arg);
12    print_integers(n, arg);
13    printf("Sum of integers is: %d\n", s);
14    free(arg);
15 }
(gdb) break read_integer_arguments
Breakpoint 1 at 0x4006ec: file intlisttools.c, line 8.
```

GDB building block #4: Automatic interruption

...

```
(gdb) list
```

```
7 void process(int argc, char** argv)
8 {
9     int* arg = read_integer_arguments(argc, argv);
10    int n = argc-1;
11    int s = sum_integers(n, arg);
12    print_integers(n, arg);
13    printf("Sum of integers is: %d\n", s);
14    free(arg);
15 }
```

```
(gdb) break read_integer_arguments
```

```
Breakpoint 1 at 0x4006ec: file intlisttools.c, line 8.
```

```
(gdb) run 1 3 5
```

GDB building block #4: Automatic interruption

...

```
(gdb) list
```

```
7 void process(int argc, char** argv)
8 {
9     int* arg = read_integer_arguments(argc, argv);
10    int n = argc-1;
11    int s = sum_integers(n, arg);
12    print_integers(n, arg);
13    printf("Sum of integers is: %d\n", s);
14    free(arg);
15 }
```

```
(gdb) break read_integer_arguments
```

```
Breakpoint 1 at 0x4006ec: file intlisttools.c, line 8.
```

```
(gdb) run 1 3 5
```

```
Starting program: debug/code/bugexample/bugexample 1 3 5
```

```
Breakpoint 1, read_integer_arguments (n=4, a=0x7fffffff9b8)
    at intlisttools.c:8
```

```
8     int* result = malloc(sizeof(int)*(n-1));
```

```
(gdb)
```

GDB building block #5: Step through code

Stepping through code

- ▶ Line-by-line
- ▶ Choose to step into or over functions
- ▶ Can show surrounding lines or use **-tui**

GDB building block #5: Step through code

Stepping through code

- ▶ Line-by-line
- ▶ Choose to step into or over functions
- ▶ Can show surrounding lines or use **-tui**

`gdb` commands

list	list part of code
next	continue until next line
step	step into function
finish	continue until function end
until	continue until line/function

GDB building block #5: Step through code

...

```
(gdb) list 6,14
```

```
6 int* read_integer_arguments(int n, char** a)
```

```
7 {
```

```
8     int* result = malloc(sizeof(int)*(n-1));
```

```
9     int i;
```

```
10     /* convert every argument, but skip '0', because it is ju
```

```
11         executable name */
```

```
12     for (i=1;i<n;i++)
```

```
13         result[i] = atoi(a[i]);
```

```
14 }
```

```
(gdb) display result
```

```
1: result = (int *) 0x0
```

```
(gdb) next
```

```
12     for (i=1;i<n;i++)
```

```
1: result = (int *) 0x601010
```

```
(gdb) until 14
```

GDB building block #5: Step through code

```
(gdb) until 14
read_integer_arguments (n=4,a=0x7fffffff9b8) at intlisttools.c:
14 }
1: result = (int *) 0x601010
(gdb) finish
Run till exit from #0  read_integer_arguments (n=4,
      a=0x7fffffff9b8) at intlisttools.c:14
0x00000000040080c in process (argc=4, argv=0x7fffffff9b8)
      at process.c:9
9      int* arg = read_integer_arguments(argc, argv);
Value returned is $3 = (int *) 0x4
(gdb)
```


GDB building block #5: Step through code

```
(gdb) until 14
read_integer_arguments (n=4,a=0x7fffffff9b8) at intlisttools.c:
14 }
1: result = (int *) 0x601010
(gdb) finish
Run till exit from #0  read_integer_arguments (n=4,
      a=0x7fffffff9b8) at intlisttools.c:14
0x00000000040080c in process (argc=4, argv=0x7fffffff9b8)
      at process.c:9
9      int* arg = read_integer_arguments(argc, argv);
Value returned is $3 = (int *) 0x4
(gdb)
```

He, why is the result variable equal to **0x601010** while the value returned is **0x4**?

GDB building block #5: Step through code

```
(gdb) until 14
read_integer_arguments (n=4,a=0x7fffffff9b8) at intlisttools.c:
14 }
1: result = (int *) 0x601010
(gdb) finish
Run till exit from #0  read_integer_arguments (n=4,
      a=0x7fffffff9b8) at intlisttools.c:14
0x00000000040080c in process (argc=4, argv=0x7fffffff9b8)
      at process.c:9
9      int* arg = read_integer_arguments(argc, argv);
Value returned is $3 = (int *) 0x4
(gdb)
```

He, why is the result variable equal to **0x601010** while the value returned is **0x4**?

Contradicts your assumption of what the program does.

The program is always right, you are wrong.

GDB building block #5: Step through code

Why is the result variable equal to `0x601010` while the value returned is `0x4`?

GDB building block #5: Step through code

Why is the result variable equal to **0x601010** while the value returned is **0x4**?

```
(gdb) list read_integer_arguments,+7
7 {
8     int* result = malloc(sizeof(int)*(n-1));
9     int i;
10    /* convert every argument, but skip '0', because it is ju
11       executable name */
12    for (i=1;i<n;i++)
13        result[i] = atoi(a[i]);
14 }
```

GDB building block #5: Step through code

Why is the result variable equal to **0x601010** while the value returned is **0x4**?

```
(gdb) list read_integer_arguments,+7
7 {
8     int* result = malloc(sizeof(int)*(n-1));
9     int i;
10    /* convert every argument, but skip '0', because it is ju
11       executable name */
12    for (i=1;i<n;i++)
13        result[i] = atoi(a[i]);
14 }
```

Aargh! Forgot the return statement!

GDB building block #5: Step through code

Why is the result variable equal to **0x601010** while the value returned is **0x4**?

```
(gdb) list read_integer_arguments,+7
7 {
8     int* result = malloc(sizeof(int)*(n-1));
9     int i;
10    /* convert every argument, but skip '0', because it is ju
11       executable name */
12    for (i=1;i<n;i++)
13        result[i] = atoi(a[i]);
14 }
```

Aargh! Forgot the return statement!

Feeling like an idiot is a common side-effect of debugging.

GDB command summary

help	h	print description of
run	r	run from the start (+args)
backtrace/where	ba	function call stack
list	l	list code lines
break	b	set breakpoint
delete	d	delete breakpoint
continue	c	continue
step	s	step into function
next	n	continue until next line
print	p	print variable
finish	fin	continue until function end
set variable	set var	change variable
down	do	go to called function
tbreak	tb	set temporary breakpoint
until	unt	continue until line/function
up	up	go to caller
watch	wa	stop if variable changes
quit	q	quit gdb

Memory Debugging



Memory Checking: Valgrind

- ▶ Memory errors do not always give segfaults
- ▶ Commonly have to go *way* out of bounds to get a segfault.
- ▶ Write into other variable - hard to find problem.
- ▶ **Valgrind** - intercepts each memory call and checks them.
- ▶ Finds illegal accesses, uninitialized values, memory leaks.
- ▶ Warning: Quite verbose, typically, and, if you use external libraries, sometimes false positives. debugging too.



Valgrind example

```
$ valgrind ./bugexample 1 3 5
==909== Memcheck, a memory error detector
==909== Copyright (C) 2002-2013, and GNU GPL'd, by Julian Seward
==909== Using Valgrind-3.9.0 and LibVEX; rerun with -h for copyr
==909== Command: ./bugexample 1 3 5
==909==
==909== Invalid write of size 4
==909==    at 0x400741: read_integer_arguments (intltools.c:1
==909==    by 0x40080B: process (process.c:9)
==909==    by 0x4006D2: main (bugexample.c:12)
==909== Address 0x51c304c is 0 bytes after a block of size 12 a
==909==    at 0x4C2636D: malloc (vg_replace_malloc.c:291)
==909==    by 0x4006FF: read_integer_arguments (intltools.c:8
==909==    by 0x40080B: process (process.c:9)
==909==    by 0x4006D2: main (bugexample.c:12)
==909==
==909== Invalid read of size 4
...
```

Valgrind example (continued)

```
==909== HEAP SUMMARY:
==909==      in use at exit: 12 bytes in 1 blocks
==909==    total heap usage: 1 allocs, 0 frees, 12 bytes allocated
==909==
==909== LEAK SUMMARY:
==909==    definitely lost: 12 bytes in 1 blocks
==909==    indirectly lost: 0 bytes in 0 blocks
==909==    possibly lost: 0 bytes in 0 blocks
==909==    still reachable: 0 bytes in 0 blocks
==909==    suppressed: 0 bytes in 0 blocks
==909== Rerun with --leak-check=full to see details of leaked memory
==909==
==909== For counts of detected and suppressed errors, rerun with
==909== ERROR SUMMARY: 2 errors from 2 contexts (suppressed: 6 from 6)
Segmentation fault
$ _
```

Valgrind recommendations

- ▶ Using valgrind on mature codes often shows lots of errors. Now, some may not be an issue (e.g. dead code or false positives from libraries), but hard to know.
- ▶ So: start using valgrind early in development.
- ▶ Program modularly, and create small unit tests, on which you can comfortably use valgrind.
- ▶ Apart from this basic valgrind usage, there are other tools available with valgrind to deal cache performance, to get more detailed memory leak information, to detect race conditions, etc. (some of which we'll discuss later).

Graphical symbolic debuggers



Graphical symbolic debuggers

Features

- ▶ Nice, more intuitive graphical user interface
- ▶ Front to command-line based tools: Same concepts
- ▶ Need graphics support: X forwarding (or VNC)

Graphical symbolic debuggers

Features

- ▶ Nice, more intuitive graphical user interface
- ▶ Front to command-line based tools: Same concepts
- ▶ Need graphics support: X forwarding (or VNC)

Available on SciNet: ddd and ddt

- ▶ ddd

```
$ module load gcc ddd
```

```
$ ddd <executable compiled with -g flag>
```

- ▶ ddt

```
$ module load ddt
```

```
$ ddt <executable compiled with -g flag>
```

(more later)

Graphical symbolic debuggers - ddd

The screenshot displays the DDD graphical debugger interface. The main window shows a C program with OpenMP parallelism:

```
float f=0.0;
int i, th;
#pragma omp parallel for default(none) private(i,th) shared(f)
for (i = 0; i<100; i++) {
    double g;
    th = omp_get_thread_num();
    printf("%d\n",th);
    g = sqrt(0.25*i+th);
    f += g;
}
printf("result = %f\n", f);
```

A red stop sign icon is visible on the left side of the code editor. The top status bar shows the current thread: `3: th` at address `0x41401940`.

The **Threads** window shows the following threads:

Thread ID	Address
4 Thread 0x41e02940	() at add.c:17
3 Thread 0x41401940	() at add.c:17
2 Thread 0x40a00940	() at add.c:17
1 Thread 0x2aaaab8d3d20	() at add.c:17

The **Control Panel** on the right includes buttons for **Run**, **Interrupt**, **Step**, **Next**, **Until**, **Cont**, **Up**, **Undo**, **Edit**, **Stepi**, **Nexti**, **Finish**, **Kill**, **Down**, **Redo**, and **Make**.

The bottom status bar shows: `Display 3: th (enabled, scope main.omp_fn.0, address 0x41401074)`.

Graphical symbolic debuggers - ddt

Alinea DDT v3.1 (on gpc-f102n084)

Session Control Search View Help

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

All 0 1 2 3

Create Group

Project Files

Search (Ctrl+F)

```
95 p.runtime = ini.get_double("runtime", 1.0e5);
96 p.dt = ini.get_double("dt", 0.2);
97 p.dc = ini.get_double("dc", 2.0);
98 p.l[0] = ini.get_double("lx", 10);
99 p.l[1] = ini.get_double("ly", 10);
100 p.l[2] = ini.get_double("lz", 10);
101 p.n[0] = ini.get_long("nx", 10);
102 p.n[1] = ini.get_long("ny", 10);
103 p.n[2] = ini.get_long("nz", 10);
104
105 cout << "l = "
106 << p.l[0] << " ";
107 << p.l[1] << " ";
108 << p.l[2] << " \n";
109 << "n = "
110 << p.n[0] << " ";
111 << p.n[1] << " ";
112 << p.n[2] << " \n";
113
114 // points per processor
115 double ppp = (p.n[0]*p.n[1]*p.n[2])/size;
116 n.dim[0] = n.dim[1] = n.dim[2] = 1;
```

Locals

Variable Name	Value
DdtOverDx2	
argc	2
argv	0x7fffffcf
comm	
coords	
dfield	0x17
dims	
field	0x7ffff6e2
fullInn	
ini	
lastt	14073729
negProc	
negSlabin	
negSlabOut	
npoints	14073735
nthrds	2
oldprogress	-1342464
origin	
p	
periods	

Input/Output* Breakpoints Watchpoints Stacks Tracepoints Tracepoint Output

Stacks

Processes	Threads	Function
1	1	+ kmp_launch_monitor
1	1	+ kmp_launch_worker
1	1	+bb_openib_async_thread
1	1	main (diff3d.cc:105)
1	1	+service_thread_start

Evaluate

Expression	Value
	<No symbol '*' in current context.>

DDT



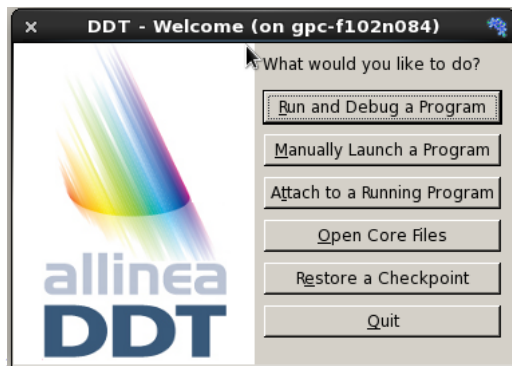
- ▶ “Distributed Debugging Tool”
- ▶ Powerful GUI-based commercial debugger by *Allinea*.
- ▶ Supports C, C++ and Fortran
- ▶ Supports MPI, OpenMP, threads, CUDA and more
- ▶ Available on all SciNet clusters (GPC, TCS, ARC, P7)
- ▶ Available on SHARCNET’s kraken, requin, orca and monk.
- ▶ Part of the “Allinea Forge” suite, which also includes a ‘profiler’ called *MAP*.

Launching ddt

- ▶ Load your compiler and MPI modules.
- ▶ Load the ddt module: `$ module load ddt`
- ▶ Start ddt with one of these:
 - `$ ddt`
 - `$ ddt <executable compiled with -g flag>`
 - `$ ddt <executable compiled with -g flag> <arguments>`
- ▶ First time: create config file: OpenMPI (skip other steps)
- ▶ Then gui for setting up debug session.

Launching ddt

- ▶ Load your compiler and MPI modules.
- ▶ Load the ddt module: `$ module load ddt`
- ▶ Start ddt with one of these:
 - `$ ddt`
 - `$ ddt <executable compiled with -g flag>`
 - `$ ddt <executable compiled with -g flag> <arguments>`
- ▶ First time: create config file: OpenMPI (skip other steps)
- ▶ Then gui for setting up debug session.



Run and Debug a Program (session setup)

The image shows two overlapping windows from the DDT (Data Display Tool) interface. The background window is titled "DDT - Run (on gpc-f102n084)" and contains configuration options for running a program. The foreground window is titled "Memory Debugging Options (on gpc-f102n084)" and provides settings for memory debugging.

DDT - Run (on gpc-f102n084) Configuration:

- Application:** /home/s/scinet/rzon/Code/diff3d/diff3d
- Application:** /home/s/scinet/rzon/Code/diff3d/diff3d
- Arguments:** (empty)
- Input File:** (empty)
- Working Directory:** (empty)
- MPI:** 2 processes, OpenMPI
- Number of processes:** 2
- Implementation:** OpenMPI, no queue
- mpirun arguments:** (empty)
- OpenMP:** 4 threads
- Number of OpenMP threads:** 4
- CUDA**
- Memory Debugging:** Minimal, No guard pages, Backtraces, Preload
- Environment Variables:** none
- Plugins:** none

Memory Debugging Options (on gpc-f102n084) Configuration:

- Preload the memory debugging library:** Language: C++, threads
- 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:**
 - Minimal** (fewest tests, picks up invalid pointers passed to memory functions)
 - Runtime** (fast, basic tests including fence-post checking, null handling)
 - Low** (adds minimal heap checking, overwriting of allocated/freed space)
 - Medium** (adds full heap checking, always relocates block on realloc)
 - High** (adds checking for arguments to common functions)
 - Custom:** (empty)
- Heap Overflow/Underflow Detection:**
 - Add guard pages to detect out of bounds heap access**
 - Guard pages:** 1
 - Add guard pages:** After
- Advanced:**
 - Specify heap-check interval:** 100
 - Store stack backtraces for memory allocations**
 - Only enable for these processes:**
 - 0-1 (100%)
 - Select All
 - x2
 - x0.5
 - 1%

Run and Debug a Program (session setup)

DDT - Run (on gpc-f102n084)

Application: /home/s/scinet/rzon/Code/diff3d/diff3d

Details ^

Application: /home/s/scinet/rzon/Code/diff3d/diff3d

Arguments:

Input File:

Working Directory:

MPI: 2 processes, OpenMPI

Details ^

Number of processes: 2

Implementation: OpenMPI, no queue

Change...

mpirun arguments

OpenMP: 4 threads

Details ^

Number of OpenMP threads: 4

CUDA

Details v

Memory Debugging: Minimal, No guard pages, Backtraces, Preload

Details...

Environment Variables: none

Details v

Plugins: none

Details v

Run

Cancel

Memory Debugging Options (o

Preload the memory debugging library: Lang

Note: Preloading only works for programs linked a program is statically linked, you must relink it again

Heap Debugging

- Minimal** (fewest tests, picks up invalid pointers)
- Runtime** (fast, basic tests including fence-post)
- Low** (adds minimal heap checking, overwriting)
- Medium** (adds full heap checking, always relo
- High** (adds checking for arguments to commo
- Custom:**

Heap Overflow/Underflow Detection

Add guard pages to detect out of bounds hea

Guard pages: 1 Add guard pages: Aft

Advanced

- Specify heap-check interval:** 100
- Store stack backtraces for memory allocations**
- Only enable for these processes:**

0-1

100%

Select

Application: /home/s/scinet/rzon/Code/diff3d/diff3d

Details ▲

Application: /home/s/scinet/rzon/Code/diff3d/diff3d

Arguments:

Input File:

Working Directory:

 MPI: 2 processes, OpenMPI

Details ▲

Number of processes: 2

Implementation: OpenMPI, no queue

Change...

mpirun arguments

 OpenMP: 4 threads

Details ▲

Number of OpenMP threads: 4

 CUDA

Details ▼

 Memory Debugging: Minimal, No guard pages, Backtraces, Preload

Details...

Environment Variables: none

Details ▼

Plugins: none

Details ▼

Memory De Preload the memory de**Note:** Preloading only works if the program is statically linked

Heap Debugging

 Minimal (fewest tests, Runtime (fast, basic te Low (adds minimal he Medium (adds full hea High (adds checking f Custom:

Heap Overflow/Underflow

 Add guard pages to c

Guard pages: 1

Advanced

 Specify heap-check in Store stack backtrace Only enable for these

0-1

Application: /home/s/scinet/rzon/Code/diff3d/diff3d

Details ▲

Application: /home/s/scinet/rzon/Code/diff3d/diff3d

Arguments:

Input File:

Working Directory:

 MPI: 2 processes, OpenMPI

Details ▲

Number of processes: 2

Implementation: OpenMPI, no queue

Change...

mpirun arguments

 OpenMP: 4 threads

Details ▲

Number of OpenMP threads: 4

 CUDA

Details ▼

 Memory Debugging: Minimal, No guard pages, Backtraces, Preload

Details...

Environment Variables: none

Details ▼

Plugins: none

Details ▼

 Preload**Note:** Pr
program

Heap D

 Mini Run Low Med High Cust

Heap O

 Add

Guard p

Advanc

 Spe Sto

Application: /home/s/scinet/rzon/Code/diff3d/diff3d

Details ▲

Application: /home/s/scinet/rzon/Code/diff3d/diff3d

Arguments:

Input File:

Working Directory:

MPI: 2 processes, OpenMPI

Details ▲

Number of processes: 2

Implementation: OpenMPI, no queue

Change...

mpirun arguments

OpenMP: 4 threads

Details ▲

Number of OpenMP threads: 4

CUDA

Details ▼

Memory Debugging: Minimal, No guard pages, Backtraces, Preload

Details...

ff3d Details ▲

f3d ▼ 

▼

▼ 

▼ 

Details ▲

▼

Details ▲

Details ▼

es, Backtraces, Preload Details...

Memory Debugging Options

Preload the memory debugging library: []

Note: Preloading only works for programs link program is statically linked, you must relink it a

Heap Debugging

- Minimal (fewest tests, picks up invalid poin
- Runtime (fast, basic tests including fence-p
- Low (adds minimal heap checking, overwr
- Medium (adds full heap checking, always r
- High (adds checking for arguments to com
- Custom: []

Heap Overflow/Underflow Detection

Add guard pages to detect out of bounds

Guard pages: [1] Add guard pages:

Advanced



Memory Debugging Options (on gpc-f102n084)



Preload the memory debugging library: Language: C++, threads

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

- Minimal (fewest tests, picks up invalid pointers passed to memory functions)
- Runtime (fast, basic tests including fence-post checking, null handling)
- Low (adds minimal heap checking, overwriting of allocated/freed space)
- Medium (adds full heap checking, always relocates block on realloc)
- High (adds checking for arguments to common functions)
- Custom:

Heap Overflow/Underflow Detection

Add guard pages to detect out of bounds heap access

Guard pages: Add guard pages: After

Advanced

Run and Debug a Program (session setup)

The image shows the DDT (Data Display Tool) interface for running a program. The main window is titled "DDT - Run (on gpc-f102n084)". It contains several configuration sections:

- Application:** /home/s/scinet/rzon/Code/diff3d/diff3d
- Arguments:** (empty)
- Input File:** (empty)
- Working Directory:** (empty)
- MPI:** 2 processes, OpenMPI. Number of processes: 2. Implementation: OpenMPI, no queue. mpirun arguments: (empty).
- OpenMP:** 4 threads. Number of OpenMP threads: 4.
- CUDA:** (unchecked).
- Memory Debugging:** Minimal, No guard pages, Backtraces, Preload.
- Environment Variables:** none.
- Plugins:** none.

The "Memory Debugging Options" sub-dialog is open, showing the following settings:

- Preload the memory debugging library:** Language: C++, threads
- 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:**
 - Minimal (fewest tests, picks up invalid pointers passed to memory functions)
 - Runtime (fast, basic tests including fence-post checking, null handling)
 - Low (adds minimal heap checking, overwriting of allocated/freed space)
 - Medium (adds full heap checking, always relocates block on realloc)
 - High (adds checking for arguments to common functions)
 - Custom: (empty)
- Heap Overflow/Underflow Detection:**
 - Add guard pages to detect out of bounds heap access
 - Guard pages: 1 Add guard pages: After
- Advanced:**
 - Specify heap-check interval: 100
 - Store stack backtraces for memory allocations
 - Only enable for these processes: 0-1 100% Select All x2 x0.5 1%

Buttons for "Run", "Cancel", "OK", and "Cancel" are visible at the bottom of the dialog boxes.

User interface (1)

The screenshot displays the Alinea DDT v3.1 (on gpc-f102n084) user interface. The top bar shows session control options like 'Session Control', 'Search', 'View', and 'Help'. Below this is a toolbar with various icons. The 'Current Group' is set to 'All', and 'Focus on current' is set to 'Group'. A 'Step Threads Together' checkbox is present. The main area shows a hierarchical view of the application's execution state:

- All** (blue bar): 0, 1, 2, 3
- Root** (green bar): 0
- Workers** (yellow bar): 1, 2, 3

The 'Create Group' button is visible below the group view. The 'Project Files' pane on the left shows a tree view of files, with 'diff3d.cc' selected. The code editor displays the following code:

```
74     }
75     // MPI::COMM_WORLD.Abort(1);
76 }
77
78 const int nthreads = get_num_threads();
79 const int root = 0;
80 const int size = MPI::COMM_WORLD.Get_size();
81 int rank = MPI::COMM_WORLD.Get_rank();
82
83 cerr << "nthreads=" << nthreads << endl;
84
85 //#include "mpidebug.ch"
86
87 mpiCommit<Parameters>();
88
```

The 'Locals' pane shows the current line(s) and the current stack. The 'Current Line(s)' pane shows the variable 'rank' with a value of 32767. The 'Stacks' pane shows the current stack with the following entries:

Processes	Threads	Function
4	4	gomp_thread_start (team.c:120)
4	4	main (diff3d.cc:81)
4	4	mxm_event_cleanup

The 'Input/Output*', 'Breakpoints', 'Watchpoints', 'Stacks', 'Tracepoints', and 'Tracepoint Output' panes are visible at the bottom. The 'Evaluate' pane is also present. The status bar at the bottom right shows 'Ready'.

User interface (2)

Alinea DDT v3.1 (on gpc-f102n084)

Session Control Search View Help

Current Group: All

DDT uses a tabbed-document interface.

All: 0 1 2 3

Root: 0

Workers: 1 2 3

Create Group

Project Files: diff3d.cc

```
74 }
75 // MPI::COMM_WORLD.Abort(1);
76 }
77
78 const int nthreads = get_num_threads();
79 const int root = 0;
80 const int size = MPI::COMM_WORLD.Get_size();
81 int rank = MPI::COMM_WORLD.Get_rank();
82
83 cerr << "nthreads=" << nthreads << endl;
84
85 // #include "mpidebug.ch"
86
87 mpiCommit<Parameters>();
88
```

Locals: Current Line(s) Current Stack

Variable Name	Value
MPI::COMM_...	
rank	32767

Type: none selected

Input/Output* Breakpoints Watchpoints Stacks Tracepoints Tracepoint Output Evaluate

Stacks

Processes	Threads	Function
4	4	gomp_thread_start (team.c:120)
4	4	main (diff3d.cc:81)
4	4	mxm_event_cleanup

Ready

User interface (3)

The screenshot displays the Alinea DDT v3.1 (on gpc-f102n084) interface. A callout box with a blue border and white background contains the text: "When the session begins, DDT automatically finds source code from information compiled in the executable." An arrow points from this box to line 81 of the code in the diff3d.cc editor. The code editor shows the following snippet:

```
74     }  
75     // MPI::COMM_WORLD.Abort(1);  
76     }  
77  
78     const int nthrds = get_num_threads();  
79     const int root = 0;  
80     const int size = MPI::COMM_WORLD.Get_size();  
81     int rank = MPI::COMM_WORLD.Get_rank();  
82  
83     cerr << "nthrds=" << nthrds << endl;  
84  
85     // #include "mpidebug.ch"  
86  
87     mpiCommit<Parameters>();  
88
```

The stack window at the bottom shows the following information:

Processes	Threads	Function
4	4	gomp_thread_start (team.c:120)
4	4	main (diff3d.cc:81)
4	4	mxm_event_cleanup

The Locals window shows the current line(s) and variable values:

Variable Name	Value
MPI::COMM_...	
rank	32767

User interface (4)

The screenshot shows the Allinea DDT v3.1 (on gpc-f102n084) interface. At the top, there is a menu bar with 'Session', 'Control', 'Search', 'View', and 'Help'. Below the menu bar is a toolbar with various icons. A status bar at the top indicates 'Current Group: All' and 'Focus on current: Group Process Thread Step Threads Together'. The main area displays a tree view of process groups: 'All' (blue bar with sub-groups 0, 1, 2, 3), 'Root' (green bar with sub-group 0), and 'Workers' (yellow bar with sub-groups 1, 2, 3). An arrow points from the 'Workers' group to a callout box. The callout box contains the following text:

Process Control and Process Groups:

- ▶ Can group process together.
- ▶ Predefined groups All, Root, Workers. (Session→options, automatically create)
- ▶ Can create, delete modify groups (drag drop, right click stacks, ...)

The background interface also shows a 'Project Files' list with 'diff3d.cc' selected, a 'Search (Ctrl+K)' field, and a 'Stacks' panel at the bottom showing process and thread information.

User interface (5)

The screenshot displays the Alinea DDT v3.1 (on gpc-f102n084) interface. A callout box with a blue border and white background contains the text: "Different colour coding for each group's current source line." An arrow points from this box to line 81 in the code editor, which is highlighted in blue. Other lines in the editor are highlighted in different colors: line 74 is light blue, line 76 is light green, line 78 is light yellow, line 79 is light orange, line 80 is light red, and line 82 is light purple. The interface also shows a "Current Line(s)" panel with a table of variable names and values, and a "Stacks" panel showing the current stack of processes and threads.

Current Line(s)

Variable Name	Value
MPI::COMM_...	
rank	32767

Type: none selected

Stacks

Processes	Threads	Function
4	4	gomp_thread_start (team.c:120)
4	4	main (diff3d.cc:81)
4	4	mxm_event_cleanup

User interface (6)

The screenshot displays the Alinea DDT v3.1 (on gpc-f102n084) interface. At the top, there is a menu bar with 'Session Control', 'Search', 'View', and 'Help'. Below the menu is a toolbar with various icons. A dropdown menu shows 'Current Group: All' and 'Focus on current: Group Process Thread Step Threads Together'. Below this are four buttons labeled '0', '1', '2', and '3'. A blue callout box with the text 'Session Control Dialog: Control program execution, e.g., play/continue, pause, step into, step over, step out' is overlaid on the interface. The main area is divided into several panes: a left sidebar with a file tree showing 'del_opv.cc', 'del_opvnt.cc', 'delete.c', 'diff3d.cc', 'distances.c', and 'divtf3.c'; a central code editor showing C code with line numbers 77-88; a right pane showing 'Variable Name' and 'Value' with 'rank' set to '32767'; and a bottom pane with tabs for 'Input/Output*', 'Breakpoints', 'Watchpoints', 'Stacks', 'Tracepoints', and 'Tracepoint Output'. The 'Stacks' tab is active, showing a table with columns for 'Processes', 'Threads', and 'Function'. The stack entries are: 'gomp_thread_start (team.c:120)', 'main (diff3d.cc:81)', and 'mxm_event_cleanup'. The status bar at the bottom right shows 'Ready'.

Session Control Dialog:
Control program execution, e.g., play/continue,
pause, step into, step over, step out

Processes	Threads	Function
4	4	gomp_thread_start (team.c:120)
4	4	main (diff3d.cc:81)
4	4	mxm_event_cleanup

User interface (7)

The screenshot displays the Allinea DDT v3.1 (on gpc-f102n084) interface. At the top, there is a menu bar with 'Session', 'Control', 'Search', 'View', and 'Help'. Below the menu is a toolbar with various icons. A status bar shows 'Current Group: All' and 'Focus on current: Group Process Thread Step Threads Together'. The main area is divided into several sections: 'All' (0, 1, 2, 3), 'Root' (0), and 'Workers' (1, 2, 3). Below these is a 'Create Group' button. The 'Project Files' section shows 'diff3d.cc' with tabs for 'Locals', 'Current Line(s)', and 'Current Stack'. A search bar is present. The 'Breakpoints Tab' is highlighted with a blue box and contains the text: 'Breakpoints Tab Can suspend, jump to, delete, load, save'. Below this is a code editor showing lines 85-87 of 'diff3d.cc'. The bottom section is divided into 'Input/Output*', 'Breakpoints', 'Watchpoints', 'Stacks', 'Tracepoints', and 'Tracepoint Output'. The 'Stacks' section is expanded, showing a table with columns 'Processes', 'Threads', and 'Function'. The table contains three rows: 'gomp_thread_start (team.c:120)', 'main (diff3d.cc:81)', and 'mxm_event_cleanup'. The 'main' row is highlighted in blue. The bottom right corner shows 'Ready' and the Allinea logo.

Allinea DDT v3.1 (on gpc-f102n084)

Session Control Search View Help

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

All 0 1 2 3

Root 0

Workers 1 2 3

Create Group

Project Files diff3d.cc Locals Current Line(s) Current Stack

Search (Ctrl+K)

Breakpoints Tab
Can suspend, jump to, delete, load, save

del_opv.cc
del_opvnt.cc
delete.c
diff3d.cc
distances.c
divtf3.c

85 // #include "mpidebug.ch"
86
87 mpiCommit<Parameters>();

Type: none selected

Input/Output* Breakpoints Watchpoints Stacks Tracepoints Tracepoint Output Evaluate

Stacks Expression Value

Processes	Threads	Function
4	4	gomp_thread_start (team.c:120)
4	4	main (diff3d.cc:81)
4	4	mxm_event_cleanup

Ready

User interface (8)

The screenshot shows the Alinea DDT v3.1 (on gpc-f102n084) interface. At the top, the 'Focus on current' dropdown menu is open, showing options for 'Group', 'Process', and 'Thread'. A callout box with a blue border and white background contains the text: *Focus:* Choose between Group, process or thread. Below the callout, the 'Current Line(s)' window shows the source code for `diff3d.cc`, with line 81 highlighted: `int rank = MPI::COMM_WORLD.Get_rank();`. The 'Stacks' window at the bottom shows the current stack of frames, with `main (diff3d.cc:81)` selected. The 'Current Stack' window shows the variable `rank` with a value of `32767`.

Session Control Search View Help

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

All 0 1 2 3

Root 0

Workers

Create Group

Project Files

Search (Ctrl+K)

del_opv.cc

del_opvnt.cc

delete.c

diff3d.cc

distances.c

divtf3.c

```
75 // MPI::COMM_WORLD.Abort(1);
76 }
77
78 const int nthreads = get_num_threads();
79 const int root = 0;
80 const int size = MPI::COMM_WORLD.Get_size();
81 int rank = MPI::COMM_WORLD.Get_rank();
82
83 cerr << "nthreads=" << nthreads << endl;
84
85 //#include "mpidebug.ch"
86
87 mpiCommit<Parameters>();
88
```

Current Line(s)

Variable Name	Value
MPI::COMM_...	
rank	32767

Type: none selected

Input/Output* Breakpoints Watchpoints Stacks Tracepoints Tracepoint Output Evaluate

Stacks

Processes	Threads	Function
4	4	gomp_thread_start (team.c:120)
4	4	main (diff3d.cc:81)
4	4	mxm_event_cleanup

Expression Value

Ready

User interface (9)

The screenshot shows the Alinea DDT v3.1 (on gpc-f102n084) interface. A callout box with a blue border and white background is overlaid on the left side, containing the following text:

Stacks: Current and Parallel

- ▶ Tree of functions (merged)
- ▶ Click on branch to see source
- ▶ Hover to see process ranks
- ▶ Use to gather processes in new groups

The interface includes a menu bar (Session Control, Search, View, Help), a toolbar, and a main workspace. The workspace is divided into several panels:

- Current Line(s):** Shows the current line of code being executed, with a variable name and value.
- Current Stack:** Shows the current stack of functions, with the current function highlighted.
- Stacks:** A table showing the stack of processes and threads, with the current process and thread highlighted.

The **Current Line(s)** panel shows the following code:

```
82  
83 cerr << "nthreads=" << nthreads << endl;  
84  
85 // #include "mpidebug.ch"  
86  
87 mpiCommit<Parameters>();  
88
```

The **Current Stack** panel shows the following stack:

Variable Name	Value
MPI::COMM_...	
rank	32767

The **Stacks** panel shows the following stack:

Processes	Threads	Function
4	4	gomp_thread_start (team.c:120)
4	4	main (diff3d.cc:81)
4	4	mxm_event_cleanup

The **Input/Output*** panel shows the following output:

```
Type: none selected
```

User interface (9)

The screenshot shows the Alinea DDT v3.1 (on gpc-f102n084) interface. A callout box with a blue border and white background contains the following text:

Stacks: Current and Parallel

- ▶ Tree of functions (merged)
- ▶ Click on branch to see source
- ▶ Hover to see process ranks
- ▶ Use to gather processes in new groups

The interface includes a menu bar (Session Control, Search, View, Help), a toolbar, and a main window with several panels:

- Current Line(s):** Shows the current line of code being executed.
- Current Stack:** A table showing the current stack of function calls.
- Stacks:** A table showing the stack of processes and threads.
- Locals:** A table showing the current line of code and the current stack.
- Input/Output*, Breakpoints, Watchpoints, Stacks, Tracepoints, Tracepoint Output, Evaluate:** A row of tabs for different debugging features.

The **Current Stack** panel shows the following data:

Variable Name	Value
MPI::COMM_...	
rank	32767

The **Stacks** panel shows the following data:

Processes	Threads	Function
4	4	gomp_thread_start (team.c:120)
4	4	main (diff3d.cc:81)
4	4	mxm_event_cleanup

The **Locals** panel shows the following data:

Current Line(s)	Current Stack
Current Line(s)	
Variable Name	Value
MPI::COMM_...	
rank	32767

The **Input/Output*** panel shows the following data:

Expression	Value

User interface (10)

The screenshot displays the Alinea DDT v3.1 (on gpc-f102n084) interface. At the top, there is a menu bar with 'Session', 'Control', 'Search', 'View', and 'Help'. Below the menu bar is a toolbar with various icons. A blue box highlights the 'Current line variables' section, which shows a tree view of variables for 'All', 'Root', and 'Workers' groups. The 'All' group contains variables 0, 1, 2, and 3. The 'Root' group contains variable 0. The 'Workers' group contains variables 1, 2, and 3. Below this is a 'Create Group' button. The main area is divided into three panes: 'Project Files' on the left, a code editor in the center, and 'Locals' on the right. The 'Project Files' pane shows a tree view of files including 'del_opv.cc', 'del_opvnt.cc', 'delete.c', 'diff3d.cc', 'distances.c', and 'divtf3.c'. The code editor shows the source code for 'diff3d.cc', with line 81 highlighted: `int rank = MPI::COMM_WORLD.Get_rank();`. The 'Locals' pane shows the current line(s) and their values, with 'rank' having a value of 32767. Below the code editor is a 'Type: none selected' label. At the bottom, there is a 'Stacks' pane with columns for 'Processes', 'Threads', and 'Function'. The stack shows three entries: 'gomp_thread_start (team.c:120)', 'main (diff3d.cc:81)', and 'mxm_event_cleanup'. The 'main' entry is highlighted in blue. To the right of the 'Stacks' pane is an 'Evaluate' pane with columns for 'Expression' and 'Value'.

Current line variables

Current Line(s)

Variable Name	Value
MPI::COMM_...	
rank	32767

Type: none selected

Processes	Threads	Function
4	4	gomp_thread_start (team.c:120)
4	4	main (diff3d.cc:81)
4	4	mxm_event_cleanup

User interface (11)

The screenshot displays the Alinea DDT v3.1 (on gpc-f102n084) interface. At the top, there is a menu bar with 'Session', 'Control', 'Search', 'View', and 'Help'. Below the menu bar is a toolbar with various icons. The main area is divided into several panels:

- Current Group:** A dropdown menu showing 'All'.
- Process Grouping:** A section with colored bars representing different process groups: 'All' (blue), 'Root' (green), and 'Workers' (yellow). Below these are buttons for '0', '1', '2', and '3'.
- Create Group:** A button labeled 'Create Group'.
- Project Files:** A tree view on the left showing a project structure with files like 'del_opv.cc', 'del_opvnt.cc', 'delete.c', 'diff3d.cc', 'distances.c', and 'divtf3.c'. 'diff3d.cc' is selected.
- Code Editor:** A central window showing the source code of 'diff3d.cc'. The current line is 81, which is highlighted in blue. The code includes MPI-related functions and a loop for threads.
- Locals Window:** A window on the right showing the current line(s) and the current stack. It displays a table with 'Variable Name' and 'Value'. The variable 'rank' is shown with a value of 32767.
- Stacks Window:** A window at the bottom showing the current stack. It has columns for 'Processes', 'Threads', and 'Function'. The stack shows three frames: 'gomp_thread_start (team.c:120)', 'main (diff3d.cc:81)', and 'mxm_event_cleanup'.

A blue box with the text 'Local variables for process' is overlaid on the interface, with an arrow pointing to the 'rank' variable in the locals window.

User interface (12)

The screenshot displays the Alinea DDT v3.1 (on gpc-f102n084) user interface. The top menu bar includes Session, Control, Search, View, and Help. Below the menu is a toolbar with various icons. The main window is divided into several panes:

- Session Control:** Shows the current group as 'All' and a list of groups: All (0, 1, 2, 3), Root (0), and Workers (1, 2, 3).
- Project Files:** Lists files like del_opv.cc, del_opvnt.cc, delete.c, diff3d.cc (selected), distances.c, and divtf3.c.
- Code Editor:** Displays the source code for diff3d.cc, with line 81 highlighted: `int rank = MPI::COMM_WORLD.Get_rank();`
- Locals:** Shows the current line(s) and the value of the variable 'rank' as 32767.
- Stacks:** Shows the current stack of frames, including gomp_thread_start (team.c:120), main (diff3d.cc:81) (selected), and mxm_event_cleanup.
- Evaluate Window:** A blue box labeled 'Evaluate window' is overlaid on the interface, with an arrow pointing to the 'Evaluate' button in the bottom right pane.

The 'Evaluate' window is currently empty, showing the 'Expression' and 'Value' columns.

First Demonstration DDT

```
$ cd $SCRATCH/debug/code  
$ source setup  
$ cd bugexample  
$ make  
$ ddt bugexample
```

Other features of DDT (1)

- ▶ Some of the user-modified parameters and windows are saved by right-clicking and selecting a save option in the corresponding window (Groups; Evaluations)
- ▶ DDT can load and save sessions.
- ▶ *Find* and *Find in Files* in the Search menu.
- ▶ *Goto line* in Search menu (or Ctrl-G)
- ▶ Synchronize processes in group: Right-click, “Run to here” .
- ▶ View multiple source codes simultaneously: Right-click, “Split”
- ▶ Right-click power!

Other features of DDT (2)

- ▶ Signal handling: SEGV, FPE, PIPE,ILL
- ▶ Support for Fortran modules
- ▶ Change data values in evaluate window
- ▶ Examine pointers (vector, reference, dereference)
- ▶ Multi-dimensional arrays
- ▶ Viewer

Other features of DDT (3)


Memory debugging

- ▶ Select “memory debug” in Run window
- ▶ Stops on error (before crash or corruption)
- ▶ Check pointer (right click in evaluate)
- ▶ View, overall memory stats

Demonstration Memory Debugging with DDT

```
$ cd $SCRATCH/debug/code  
$ source setup  
$ cd ex4  
$ make  
$ ddt ex4
```

Parallel debugging



Parallel debugging - 1 Shared memory

Use gdb for

- ▶ Tracking each thread's execution and variables
- ▶ OpenMP serialization: `p omp_set_num_threads(1)`
- ▶ Stepping into OpenMP block: `break` at first line!
- ▶ Thread-specific breakpoint: `b <line> thread <n>`

Parallel debugging - 1 Shared memory

Use gdb for

- ▶ Tracking each thread's execution and variables
- ▶ OpenMP serialization: `p omp_set_num_threads(1)`
- ▶ Stepping into OpenMP block: `break` at first line!
- ▶ Thread-specific breakpoint: `b <line> thread <n>`

Use helgrind for

- ▶ Finding race conditions:

```
$ module load valgrind
$ valgrind --tool=helgrind <exe> &> out
$ grep <source> out
```

where `<source>` is the name of the source file where you suspect race conditions (valgrind reports a lot more)

Shared memory debugging with DDT

Or use DDT:

Thread debugging example

```
$ cd $SCRATCH/debug/code  
$ source setup  
$ cd ex5  
$ make  
$ ddt ex5
```

Parallel debugging - 2 Distributed memory

Multiple MPI processes

- ▶ Your code is running on different cores!
- ▶ Where to run debugger?
- ▶ Where to send debugger output?
- ▶ Much going on at same time.
- ▶ No universal free solution.

Parallel debugging - 2 Distributed memory

Multiple MPI processes

- ▶ Your code is running on different cores!
- ▶ Where to run debugger?
- ▶ Where to send debugger output?
- ▶ Much going on at same time.
- ▶ No universal free solution.

Good approach:

1. Write your code so it can run in serial: perfect that first.
2. Deal with communication, synchronization and deadlock on *smaller* number of MPI processes/threads.
3. Only then try full size.

Parallel debugging demands specialized tools: ddt

Demonstration MPI Debugging with DDT

```
$ cd $SCRATCH/debug/code  
$ source setup  
$ cd ex2  
$ make  
$ ddt ex2
```

Detecting deadlock with DDT

Message Queue

- ▶ View → show message queue
- ▶ produces both a graphical view and table for active communications
- ▶ Helps to find e.g. deadlocks

DDT - Message Queues

Select queues to show

- Send
- Receive
- Unexpected

Ranks

- Show local ranks
- Show global ranks

Select communicator

- MPI_COMM_WORLD
- MPI_COMM_SELF
- MPI_COMM_NULL

Demonstration MPI Message Queue in DDT

```
$ cd $SCRATCH/debug/code
```

```
$ source setup
```

```
$ cd ex3
```

```
$ make
```

```
$ ddt ex3
```


Useful references

- ▶ N Matloff and PJ Salzman
The Art of Debugging with GDB, DDD and Eclipse
- ▶ *GDB*: sources.redhat.com/gdb
- ▶ *DDT*: www.allinea.com/knowledge-center/tutorials
- ▶ *SciNet Wiki*: wiki.scinethpc.ca: Tutorials & Manuals