

allinea



Leaders in parallel software development tools

Accelerated Debugging

Allinea DDT and OpenACC on the XK6

David Lecomber

CTO, Allinea Software

david@allinea.com

www.allinea.com

Some history

- A long time ago (2007) in a galaxy far, far away....
 - The CUDA programming model is introduced
 - Powerful, efficient and C based
 - Understood and adopted by new groups of experts
 - Existing codes modified to extract SIMD parallelism and introduce CUDA kernels
 - Performance of codes is optimized by overlapping device and host, or rearranging memory usage inside device
 - The first CUDA bug is created

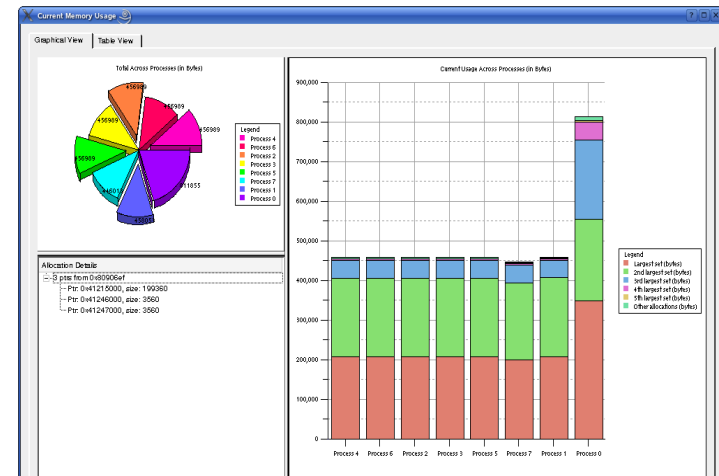

Allinea Software



- HPC development tools company
 - Flagship product Allinea DDT
 - The leading debugger in parallel computing
 - The scalable debugger
 - Record holder for debugging software on largest machines
 - Production use at extreme scale – and desktop
 - Wide customer base
 - Blue-chip engineering, government and academic research
 - Strong collaborative relationships with customers and partners
 - Leaders in performance and usability

Alinea DDT in a nutshell

- Graphical source level debugger for
 - Parallel, multi-threaded, scalar or hybrid code
 - C, C++, F90, Co-Array Fortran, UPC
- Strong feature set
 - Memory debugging
 - Data analysis
- Managing concurrency
 - Emphasizing differences
 - Collective control



May 2011 - CUG 2011

- Petascale debugging becomes real
- Allinea DDT 3.0 – lightning speed – 100,000 cores and beyond
 - Record holder for largest machines
 - Debugging at scale becomes fast, simple and **routine**
 - Production use at extreme scale on Cray XE and XT systems

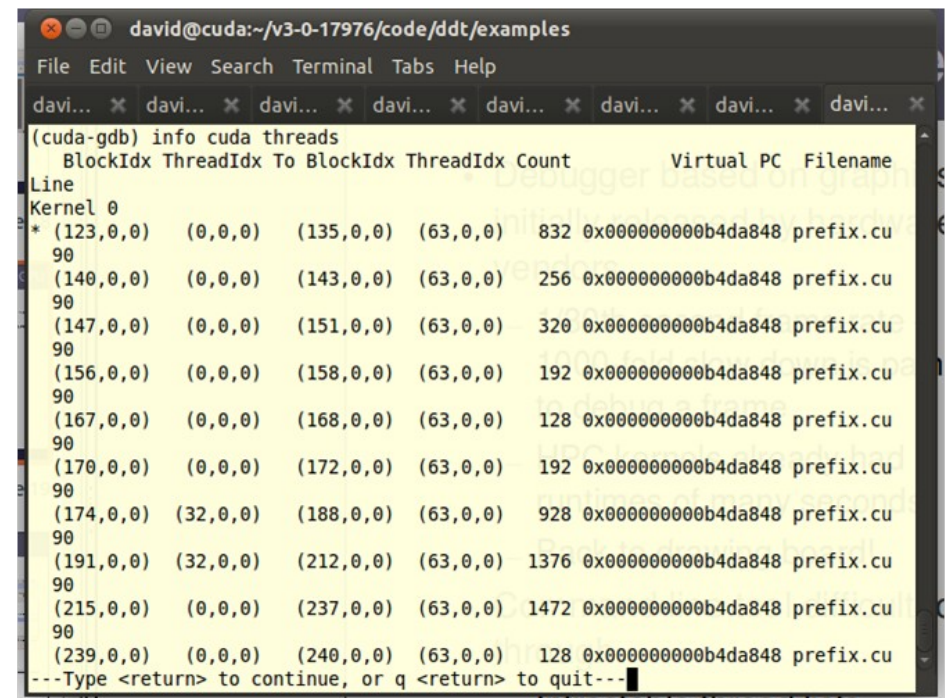
Processes	Function
150120	_start
150120	__libc_start_main
150120	main
150120	pop (POP.f90:81)
150120	initialize_pop (initial.f90:119)
150120	init_communicate (communicate.f90:87)
150119	create_ocn_communicator (communicate.f90:300)
1	create_ocn_communicator (communicate.f90:303)

Lift-off – beyond petascale

- Fairbanks, May 2011 – Cray XK6 announced
 - Large GPU systems firmly on the agenda
 - Allinea and ORNL collaborate to ensure GPU applications debuggable at scale
 - Petascale debugging, but with GPUs
 - Core needs identified and key features and performance specified
 - How would the XK6 be programmed?
 - Candidate pragma languages to remove CUDA burden

How do we fix GPU bugs?

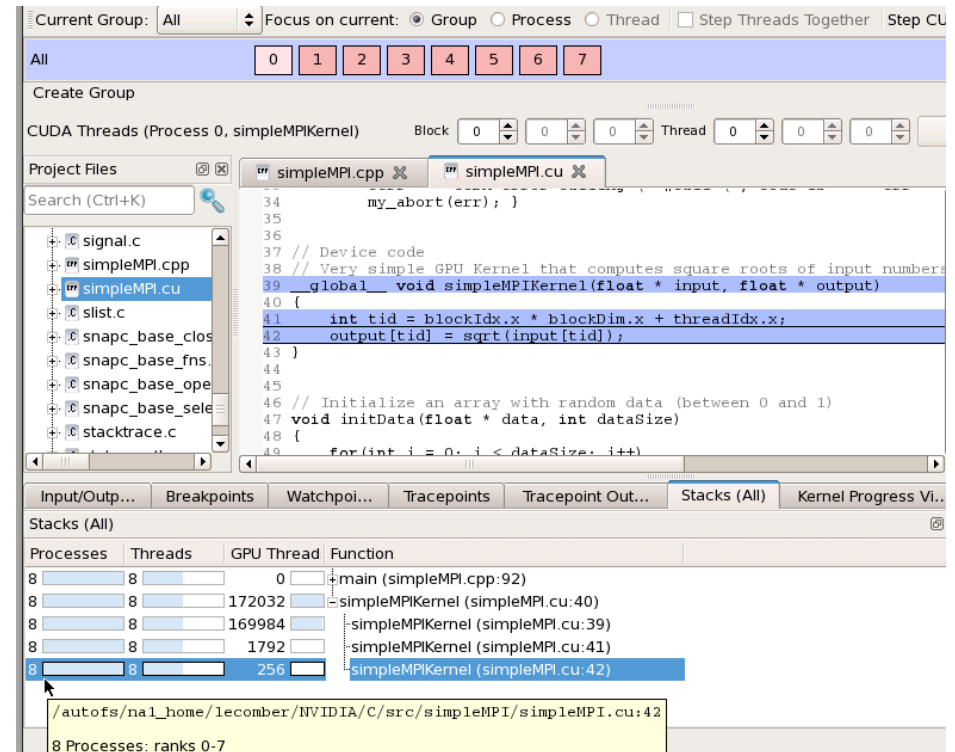
- Print statements
 - Too intrusive
- Command line debugger?
 - A good start:
 - Variables, source code
 - Large thread counts overwhelming
 - Too complex
- A graphical debugger...



```
 david@cuda:~/v3-0-17976/code/ddt/examples
File Edit View Search Terminal Tabs Help
davi... x davi... x davi... x davi... x davi... x davi... x davi... x davi... x
(cuda-gdb) info cuda threads
BlockIdx ThreadIdx To BlockIdx ThreadIdx Count Virtual PC Filename
Line
Kernel 0
* (123,0,0) (0,0,0) (135,0,0) (63,0,0) 832 0x00000000b4da848 prefix.cu
90
(140,0,0) (0,0,0) (143,0,0) (63,0,0) 256 0x00000000b4da848 prefix.cu
90
(147,0,0) (0,0,0) (151,0,0) (63,0,0) 320 0x00000000b4da848 prefix.cu
90
(156,0,0) (0,0,0) (158,0,0) (63,0,0) 192 0x00000000b4da848 prefix.cu
90
(167,0,0) (0,0,0) (168,0,0) (63,0,0) 128 0x00000000b4da848 prefix.cu
90
(170,0,0) (0,0,0) (172,0,0) (63,0,0) 192 0x00000000b4da848 prefix.cu
90
(174,0,0) (32,0,0) (188,0,0) (63,0,0) 928 0x00000000b4da848 prefix.cu
90
(191,0,0) (32,0,0) (212,0,0) (63,0,0) 1376 0x00000000b4da848 prefix.cu
90
(215,0,0) (0,0,0) (237,0,0) (63,0,0) 1472 0x00000000b4da848 prefix.cu
90
(239,0,0) (0,0,0) (240,0,0) (63,0,0) 128 0x00000000b4da848 prefix.cu
---Type <return> to continue, or q <return> to quit---
```


GPU debugging with Alinea DDT

- Almost like debugging a CPU – we can still:
 - Run through to a crash
 - Step through and observe
- CPU-like debugging features
 - Double click to set breakpoints
 - Hover the mouse for more information
 - Step a warp, block or kernel
 - Follow threads through the kernel
- Simultaneously debugs CPU code
- CUDA Memcheck feature detects read/write errors



Examining GPU data

- Debugger reads host and device memory
 - Shows all memory classes: shared, constant, local, global, register..
 - Able to examine variables
 - ... or plot larger arrays directly from device memory



Overviews of GPUs

Attribute Name	Value
Ranks 0,21,35,98	
gf100	2 Devices
IDs	0-1
Compute Capability	sm_20
Number of SMs	14
Warps per SM	48
Lanes per Warp	32
Registers per Lane	64
Ranks 1-20,22-34,36-55,57-97,99-119	No Device

- Device overview shows system properties
 - Helps optimize grid sizes
 - Handy for bug fixing – and detecting hardware failure!
- Kernel progress view
 - Shows progress through kernels
 - Click to select a thread

Kernel	Progress
simpleM...	<div style="width: 10%;"><div style="width: 5%; background-color: green;"></div><div style="width: 5%; background-color: red;"></div></div>

Kernels: 7
CUDA thread: <<<(1080,0,0),(0,0,0)>>>
Dimensions: <<<(10000,1,1),(256,1,1)>>>

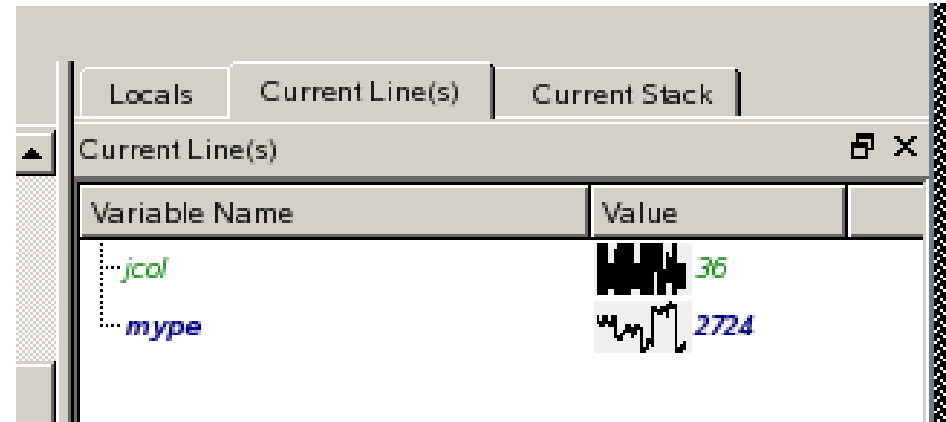
A New Hope

- Seattle, November 2011
 - CAPS, Cray, NVIDIA and PGI announce new standard for accelerator programming
 - Access CUDA compute power easily
 - A common standard
 - Allinea supports debugging Cray OpenACC compiler



November 2011

- Allinea DDT 3.1 – innovation for all scales
 - Sparklines – automatic data comparison
 - Offline debugging – scalable hands-free
 - Static analysis – automatic detection of coding errors
 - Cray XK6 support
 - Cray UPC, CoArray and OpenACC support



```
29
30  threads = calloc(sizeof(pthread_t), nthreads);
31  ids = calloc(sizeof(int), nthreads);
32
33  init_mutex();
34
35  pthread_mutex_lock(mutley);
36  for (i = 0; i < nthreads; ++i) {
37      ids[i] = i;
38      pthread_create (threads + i, NULL, &thread,
39                      );
40  pthread_mutex_unlock(mutley);
41  for (i = 0; i < nthreads; ++i)
42      pthread_join (threads[i], NULL);
43
44  return 0;
45
46  error Memory leak: threads
47  error Memory leak: ids
48  void *q)
49
50  volatile int busy = 0;
51  volatile int locker = 0; /* to be amended by
52  int i, j;
53  double k = 1;
54  int tid = *(int*) q;
55  usleep(rand() % 31);
56
```

OpenACC debugging

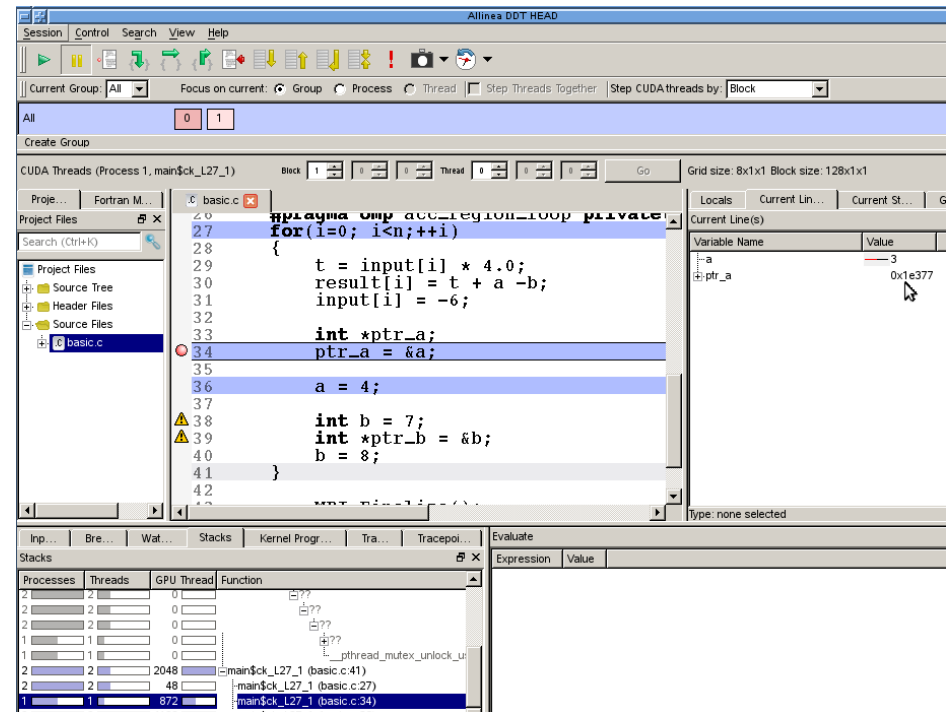
```
41 plot.height = 25;
42 plot.xoffset = 0;
43 plot.yoffset = 0;
44
45 int j;
46 float dist;
47 float total;
48 coords_3d temp;
49
50 #pragma omp acc region loop reduction(+:total)
51 for (j=0; j<N; ++j) {
52     temp = nodes[j];
53     dist = distance(temp,plot.origin);
54
55     coords_3d* ptr = &nodes[j];
56
57     total = dist;
58 }
59
60 printf("total = %f\n",total);
61
62 return (0);
63 }
```

Threads	GPU Thread	Function
1	128	main\$ck_L51_1 (reduction.c:58)
1	96	main\$ck_L51_1 (reduction.c:51)
1	10	main\$ck_L51_1 (reduction.c:52)
1	22	main\$ck_L51_1 (reduction.c:58)
1	0	main (reduction.c:58)

- On device debugging with Allinea DDT
 - Variables – arrays, pointers, full F90 and C support
 - Set breakpoints and step warps and blocks
- Requires Cray compiler for on-device debugging
 - Other compilers to follow
- Identical to CUDA
 - Full warp/block/kernel controls

XK6 Status

- Interlagos update
 - No problem – tested at scale!
- Last known major GPU issues fixed
 - Driver fixes deployed April 2012
 - 3-way debug sessions to diagnose and fix a device-hang
- Full OpenACC and CUDA support



What's next?

- Large Cray XK6 systems in (or almost in) place
 - ORNL Titan – 300,000 CPU cores
 - NCSA Blue Waters – 380,000 CPU cores
- Allinea DDT chosen for both systems – at scale
 - Watch out for improvements over next 6 months!
 - Kepler support for system upgrades

