Debugging Tools

Charlene Yang

July 18, 2018

2018 CSGF Annual Review

Outline

- General on debugging
  - print statements, debuggers, HPC/parallel debuggers
  - GDB, Eclipse, CodeBlocks, **DDT**, TotalView, **CUDA-GDB**, CUDA-MEMCHECK
  - Common bugs

- Debugging on Titan
  - **DDT**: Reverse connect via remote client
  - Demo/Hands-on: an MPI code, ring.c
  - **CUDA-GDB**: Command line
  - Demo/Hands-on: a CUDA code, bitreverse.cu

General on debugging

- Your first debugger: printf
  - print a value/message at a certain location of the code
  - bisection to narrow down the buggy area
  - **slow and long process, not scalable**

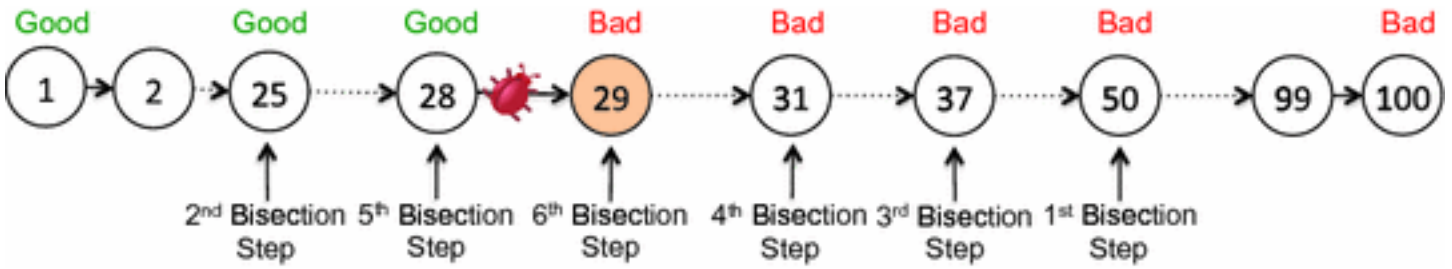

General on debugging

- Debuggers can work for us!
  - GDB, Eclipse, CodeBlocks -- command line vs GUI
  - view source, navigate between functions
  - set breakpoints, step through code
  - examine variables/arrays, dereference pointers, set watches
  - **save manpower on tracking code, focus on finding bugs**
  - **make random bugs (more) deterministic**

A bug that occurs 1% of the time is many times harder to fix than one that occurs 100% of the time

General on debugging

- Talk to a colleague
  - The very act of explaining your thinking can help diagnose the problem

- Rubber Duck Debugging
  - https://rubberduckdebugging.com

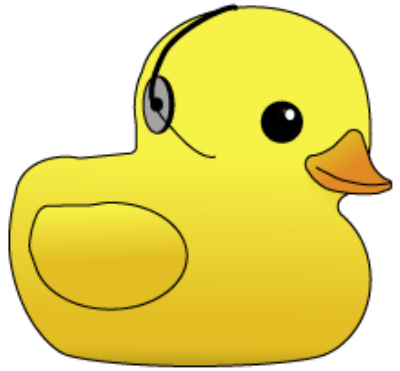

Ultimate goal is to resolve the discrepancy between what you think your code is doing and what the machine is actually doing

General on debugging

- Parallel debugging:
  - massive parallelism, large, distributed data, graphics lagging, batch system
  - DDT, TotalView, Valgrind, CUDA-GDB, CUDA-MEMCHECK
  - **Arm Forge - DDT**:
    - parallel stack view, step through processes/threads/warps, simultaneously or separately
    - sparkline data comparison, array viewer for multi-dimension arrays, pattern searching
    - submit to queue, remote client, offline debugging
  - **CUDA-GDB**:
    - https://darkdust.net/files/GDB%20Cheat%20Sheet.pdf
    - https://docs.nvidia.com/cuda/cuda-gdb/index.html
- Tools have comparable features, and debugging skills are transferrable

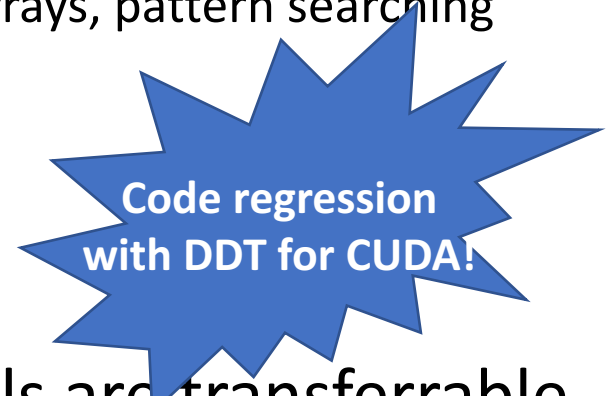

General on debugging

- Common bugs
  - memory leaks, free resources twice
  - uninitialized variable, out of bound read/write
  - NAN’s, datatype overflow
  - deadlocks, race conditions

- wrong results, wrong results every time, wrong results some time
- ‘randomly’ wrong

DDT on Titan

- Compiling
  - C/C++/Fortran code: -g (-O0)
  - CUDA code: -g -G
- X11 forwarding
  - Direct X11: slo……w
  - NoMachine: https://www.olcf.ornl.gov/software_package/freenx/
- **Reverse connect**
  - Remote Client: https://developer.arm.com/products/software-development-tools/hpc/downloads/download-arm-forge
- Demo/Hands-on
  - MPI code: ring.c
  - Copy from $PROJWORK/csc261/cjyang/csgf

Reverse Connect

- Download the Remote Client 18.2
- https://developer.arm.com/products/software-development-tools/hpc/downloads/download-arm-forge

Platform	Operating System/Distribution Version
Mac OS/X	Mountain Lion+ 64-bit (AMD/Intel)
Windows	XP+ 64-bit (AMD/Intel) **Note:** For more information, see Installing Arm Forge Remote Client on Windows.
Linux	Use the Full Install section above. All Linux installs also function as remote clients.

- Older versions are also available
- Make sure remote version matches **EXACTLY** the version you're using on Titan!!

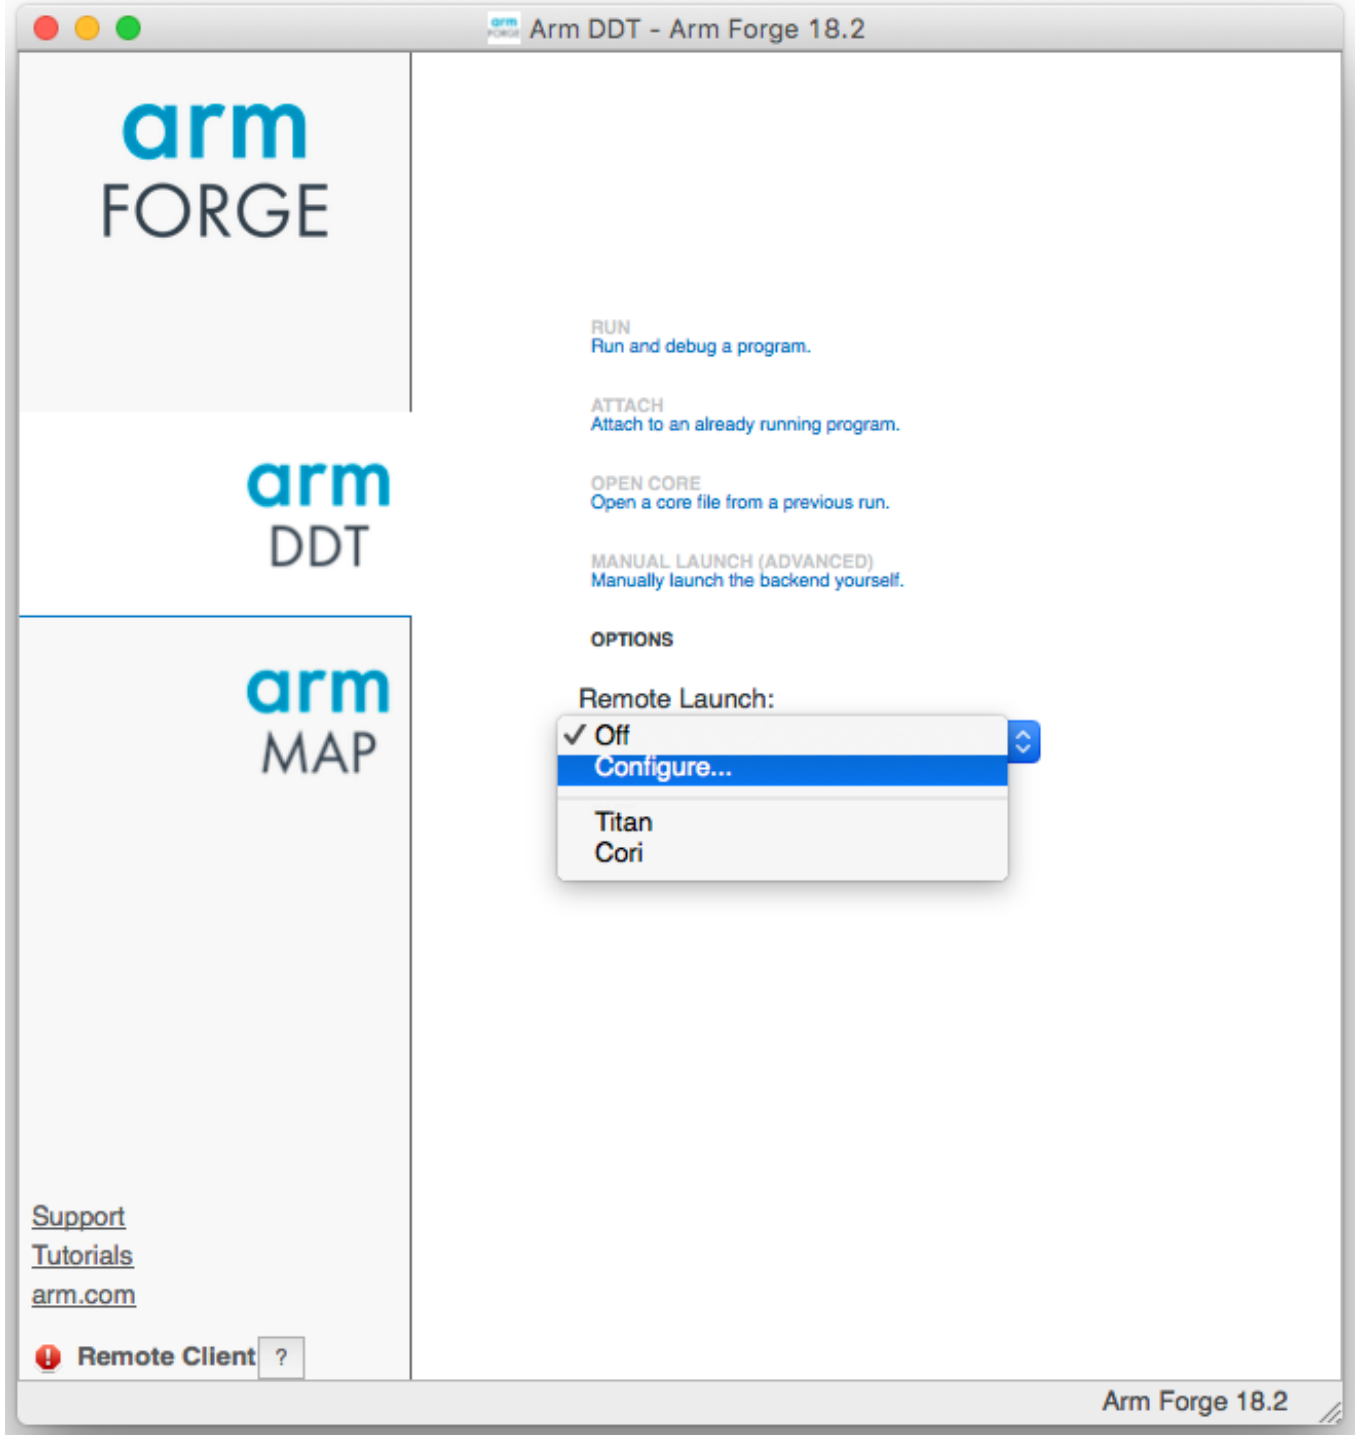
Arm DDT - Arm Forge 18.2
arm FORGE
arm DDT
arm MAP
RUN
Run and debug a program.
ATTACH
Attach to an already running program.
OPEN CORE
Open a core file from a previous run.
MANUAL LAUNCH (ADVANCED)
Manually launch the backend yourself.
OPTIONS
Remote Launch:
Off
Configure...
Titan
Cori
Support
Tutorials
arm.com
Remote Client
?
Arm Forge 18.2

- Connection Name: Titan
- Host Name: <*username*>@titan.ccs.ornl.gov
- Remote Install Dir: /sw/xk6/forge/18.2/sles11_binary

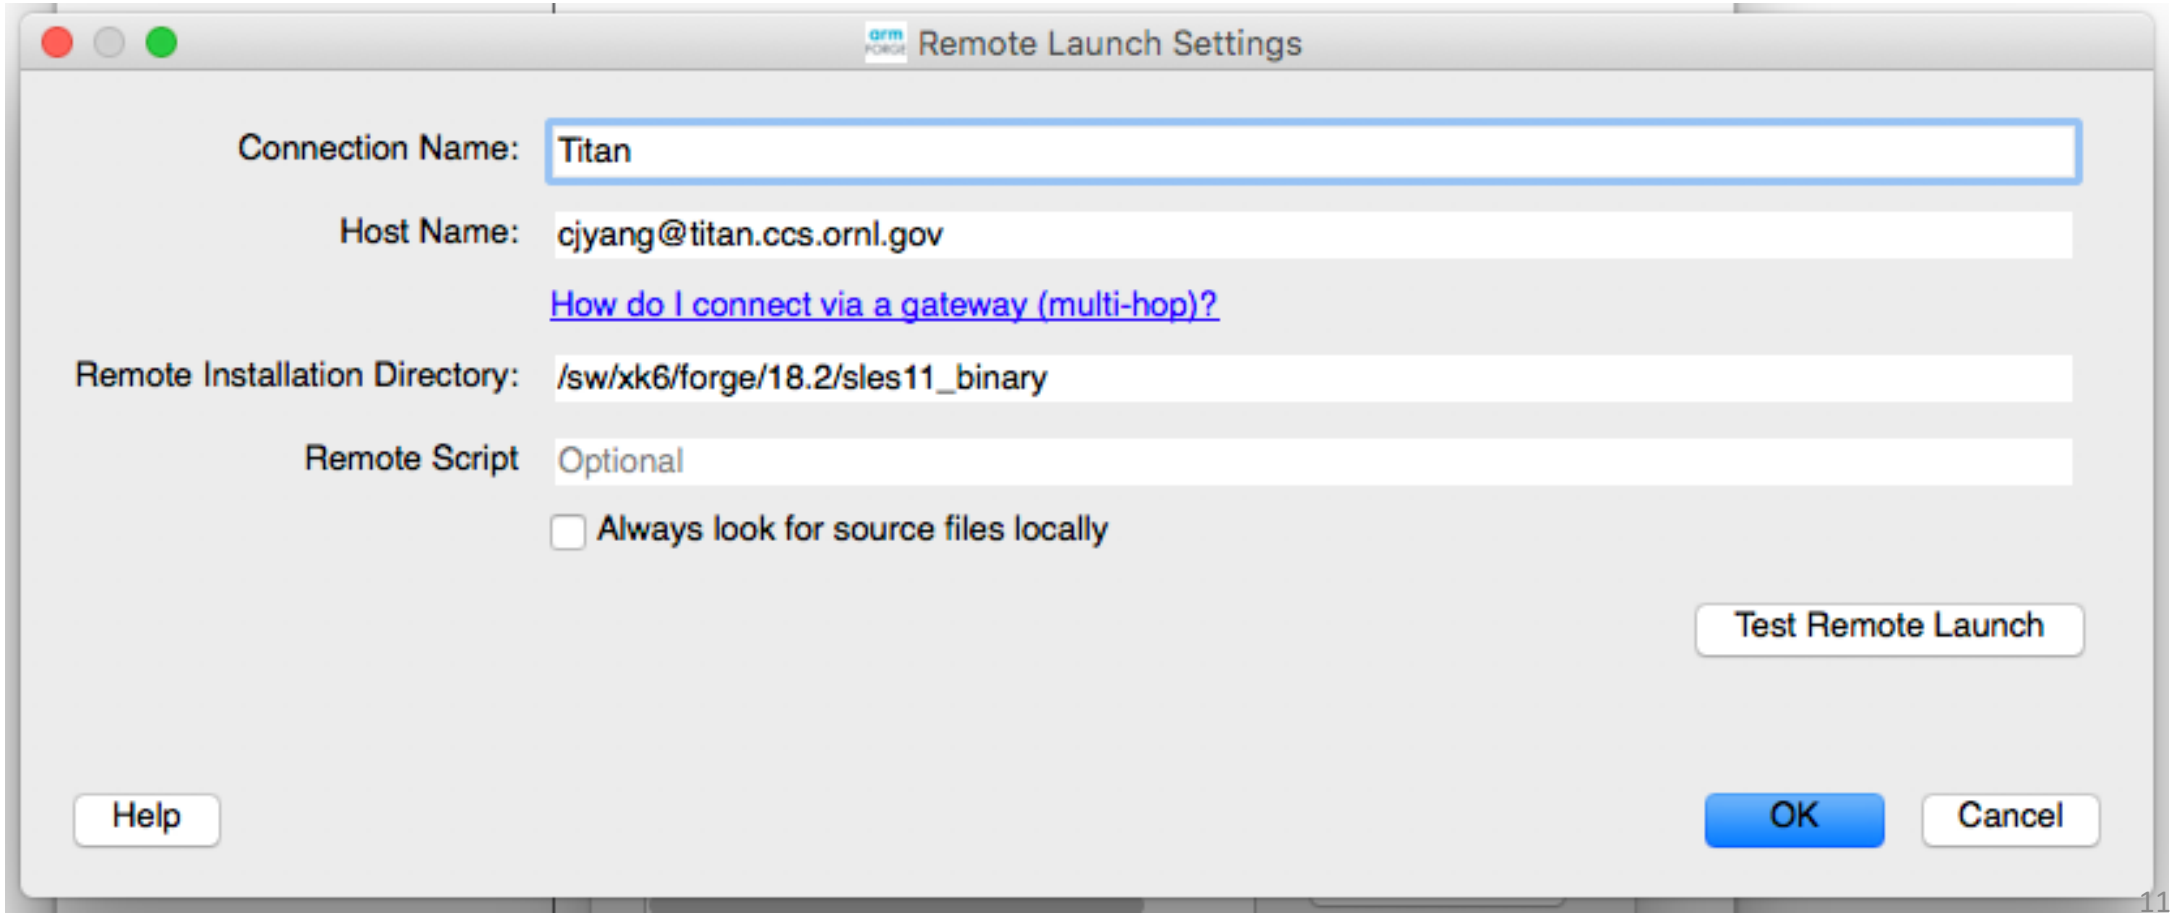

- Enter your password + RSA
- Connected to Titan!

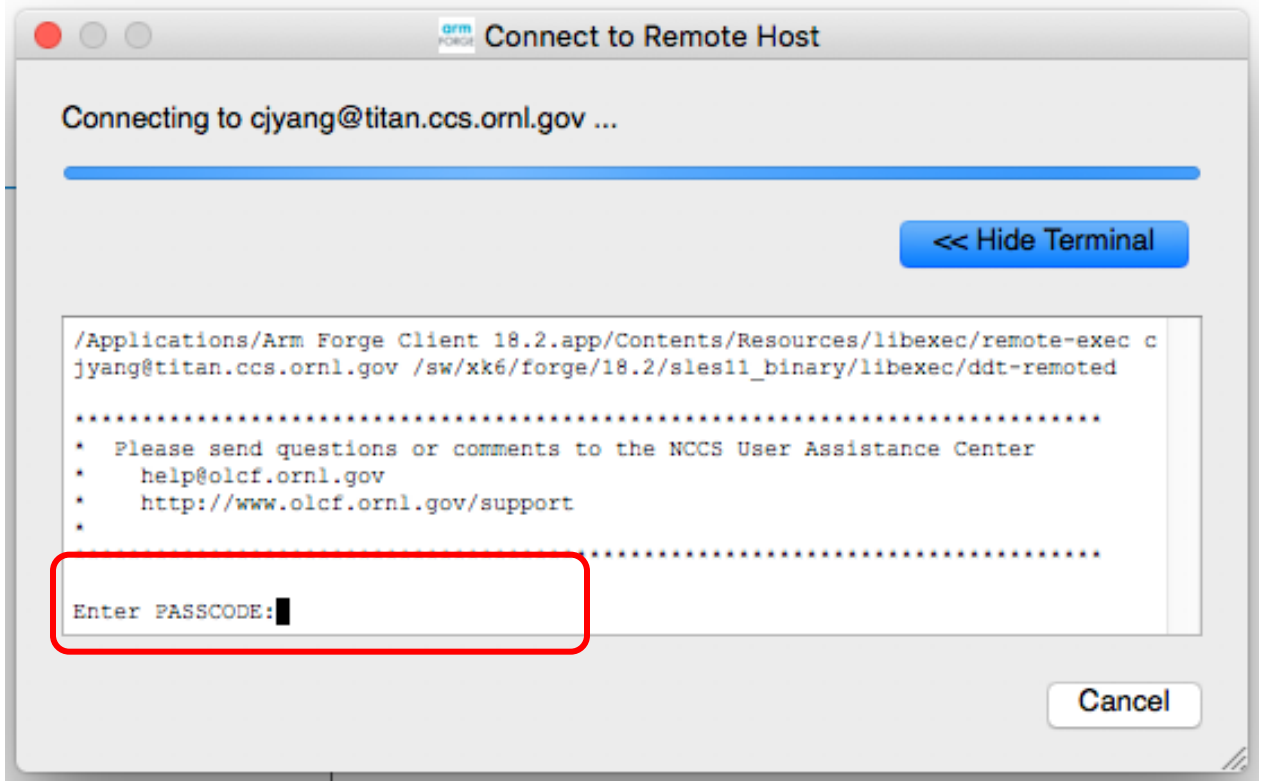

DDT on Titan

- In a Terminal, login to Titan (credentials + RSA)
- Start an interactive job
  - qsub -I -A csc261 -X -l walltime=30:00 -l nodes=1
- cd $PROJWORK/csc261/$USER/
- cp -r $PROJWORK/csc261/cjyang/csgf .
- cd csgf/
- cc **-g** ring.c -o ring_c
- module load forge/18.2
- ddt --connect aprun -n 8 ./ring_c

aprun -n 8 ./ring_c

pe=0: send

one iteration of

```
while 1
   recv
   …
   send
```

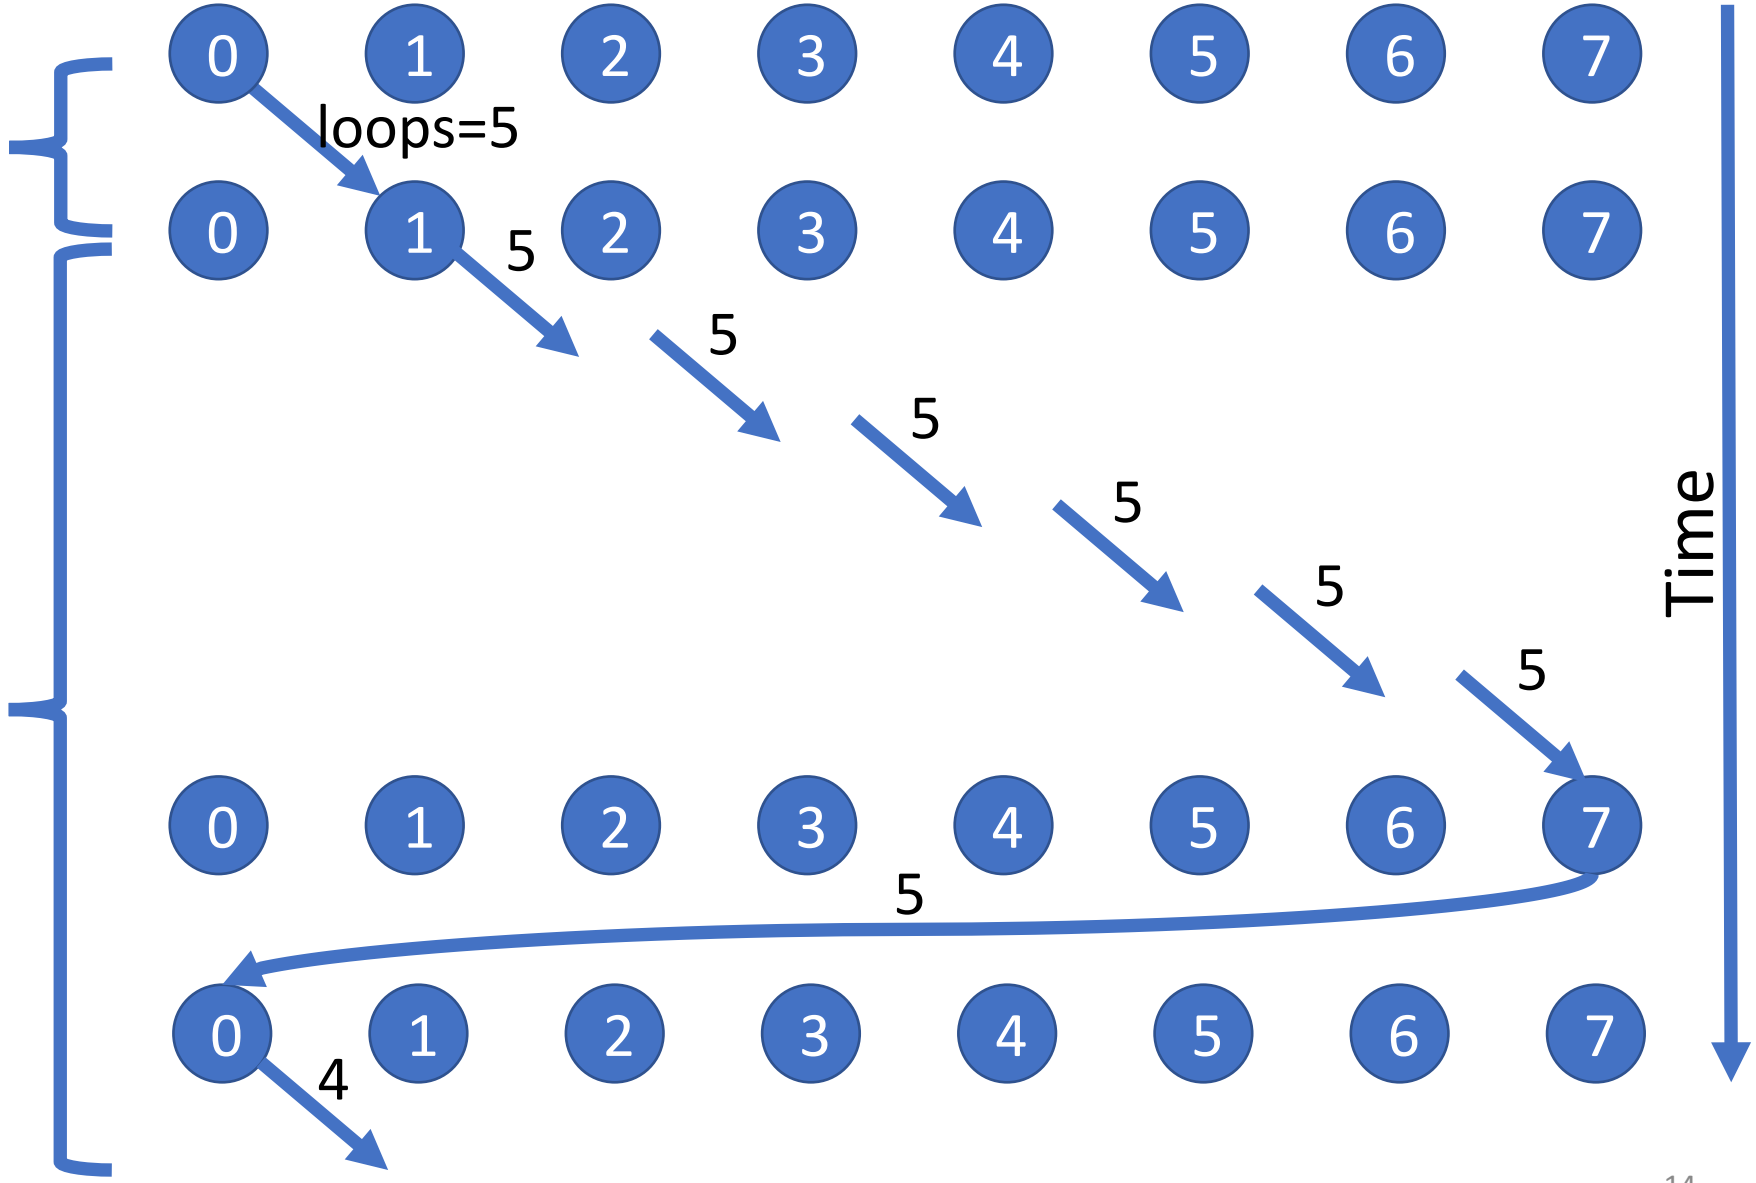

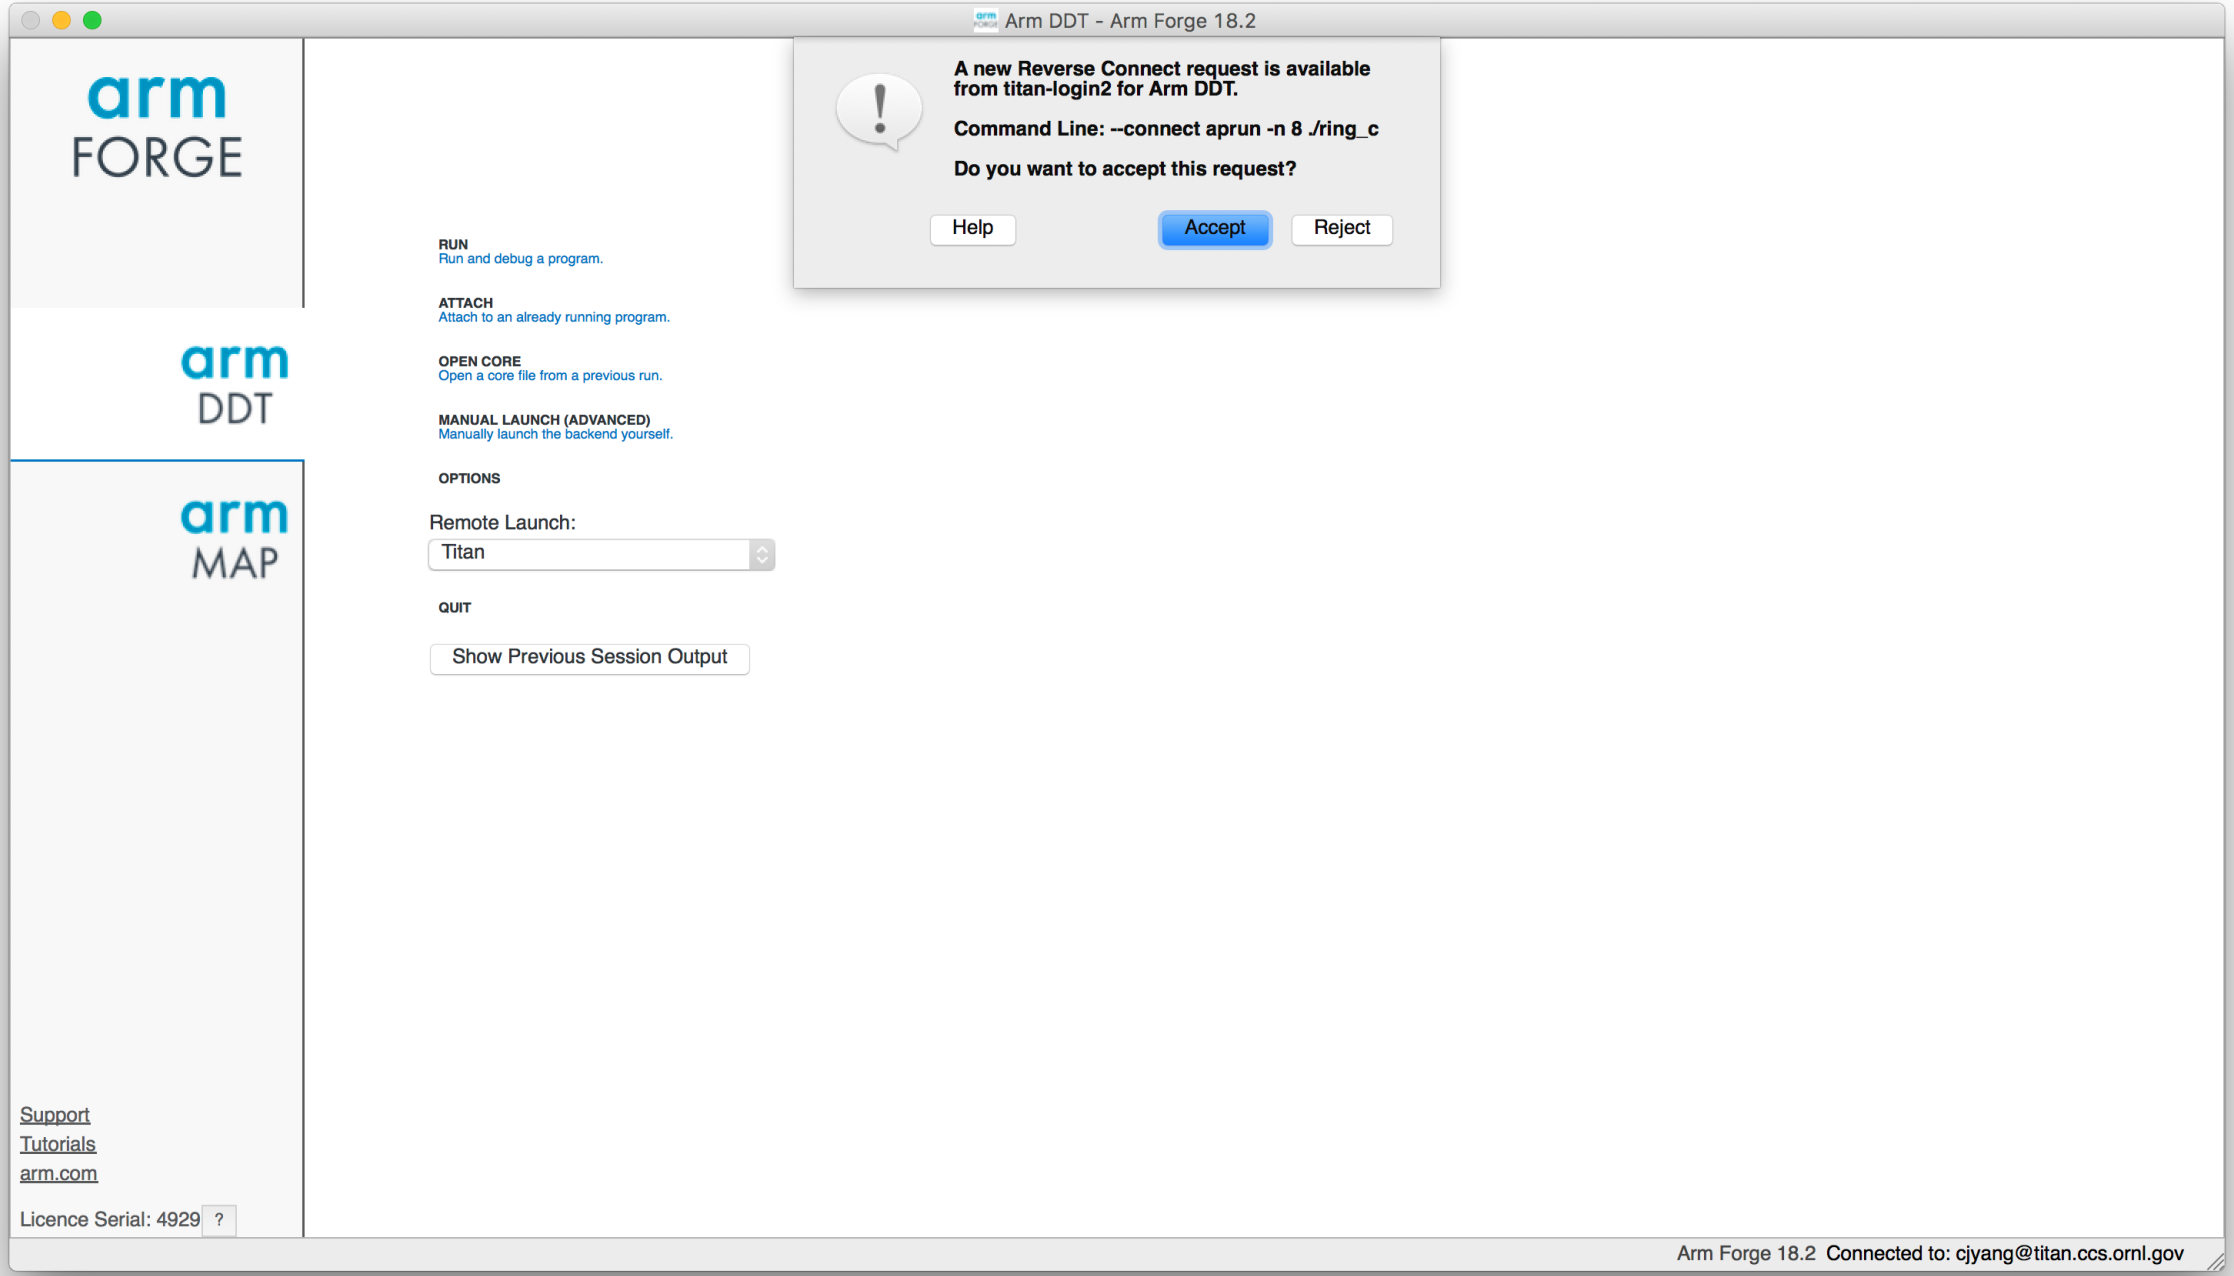
Arm DDT - Arm Forge 18.2
arm FORGE
arm DDT
arm MAP
A new Reverse Connect request is available from titan-login2 for Arm DDT.
Command Line: --connect aprun -n 8 ./ring_c
Do you want to accept this request?
Help
Accept
Reject
RUN
Run and debug a program.
ATTACH
Attach to an already running program.
OPEN CORE
Open a core file from a previous run.
MANUAL LAUNCH (ADVANCED)
Manually launch the backend yourself.
OPTIONS
Remote Launch:
Titan
QUIT
Show Previous Session Output
Support
Tutorials
arm.com
Licence Serial: 4929
?
Arm Forge 18.2 Connected to: cjyang@titan.ccs.ornl.gov

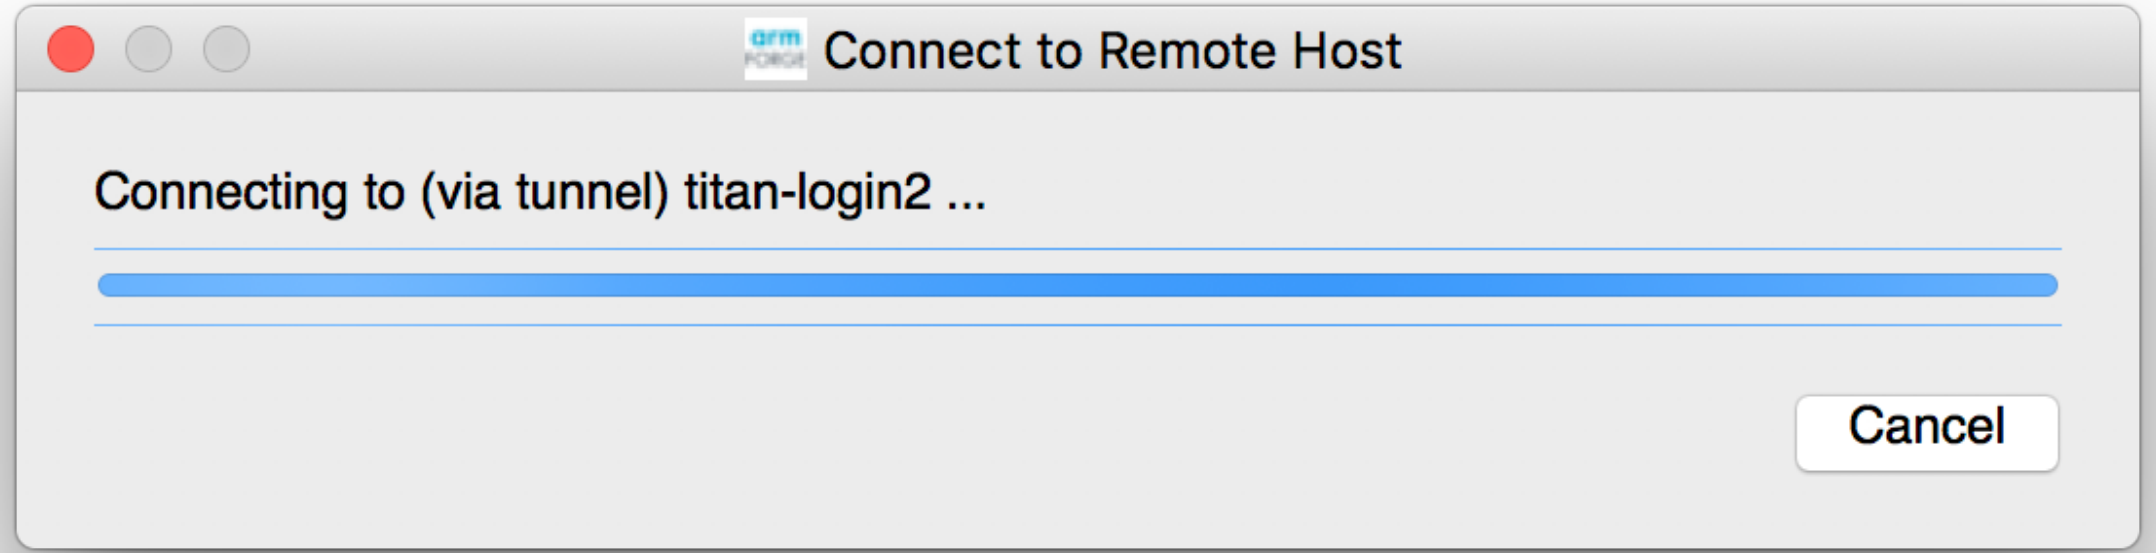
Connect to Remote Host
Connecting to (via tunnel) titan-login2 ...
Cancel

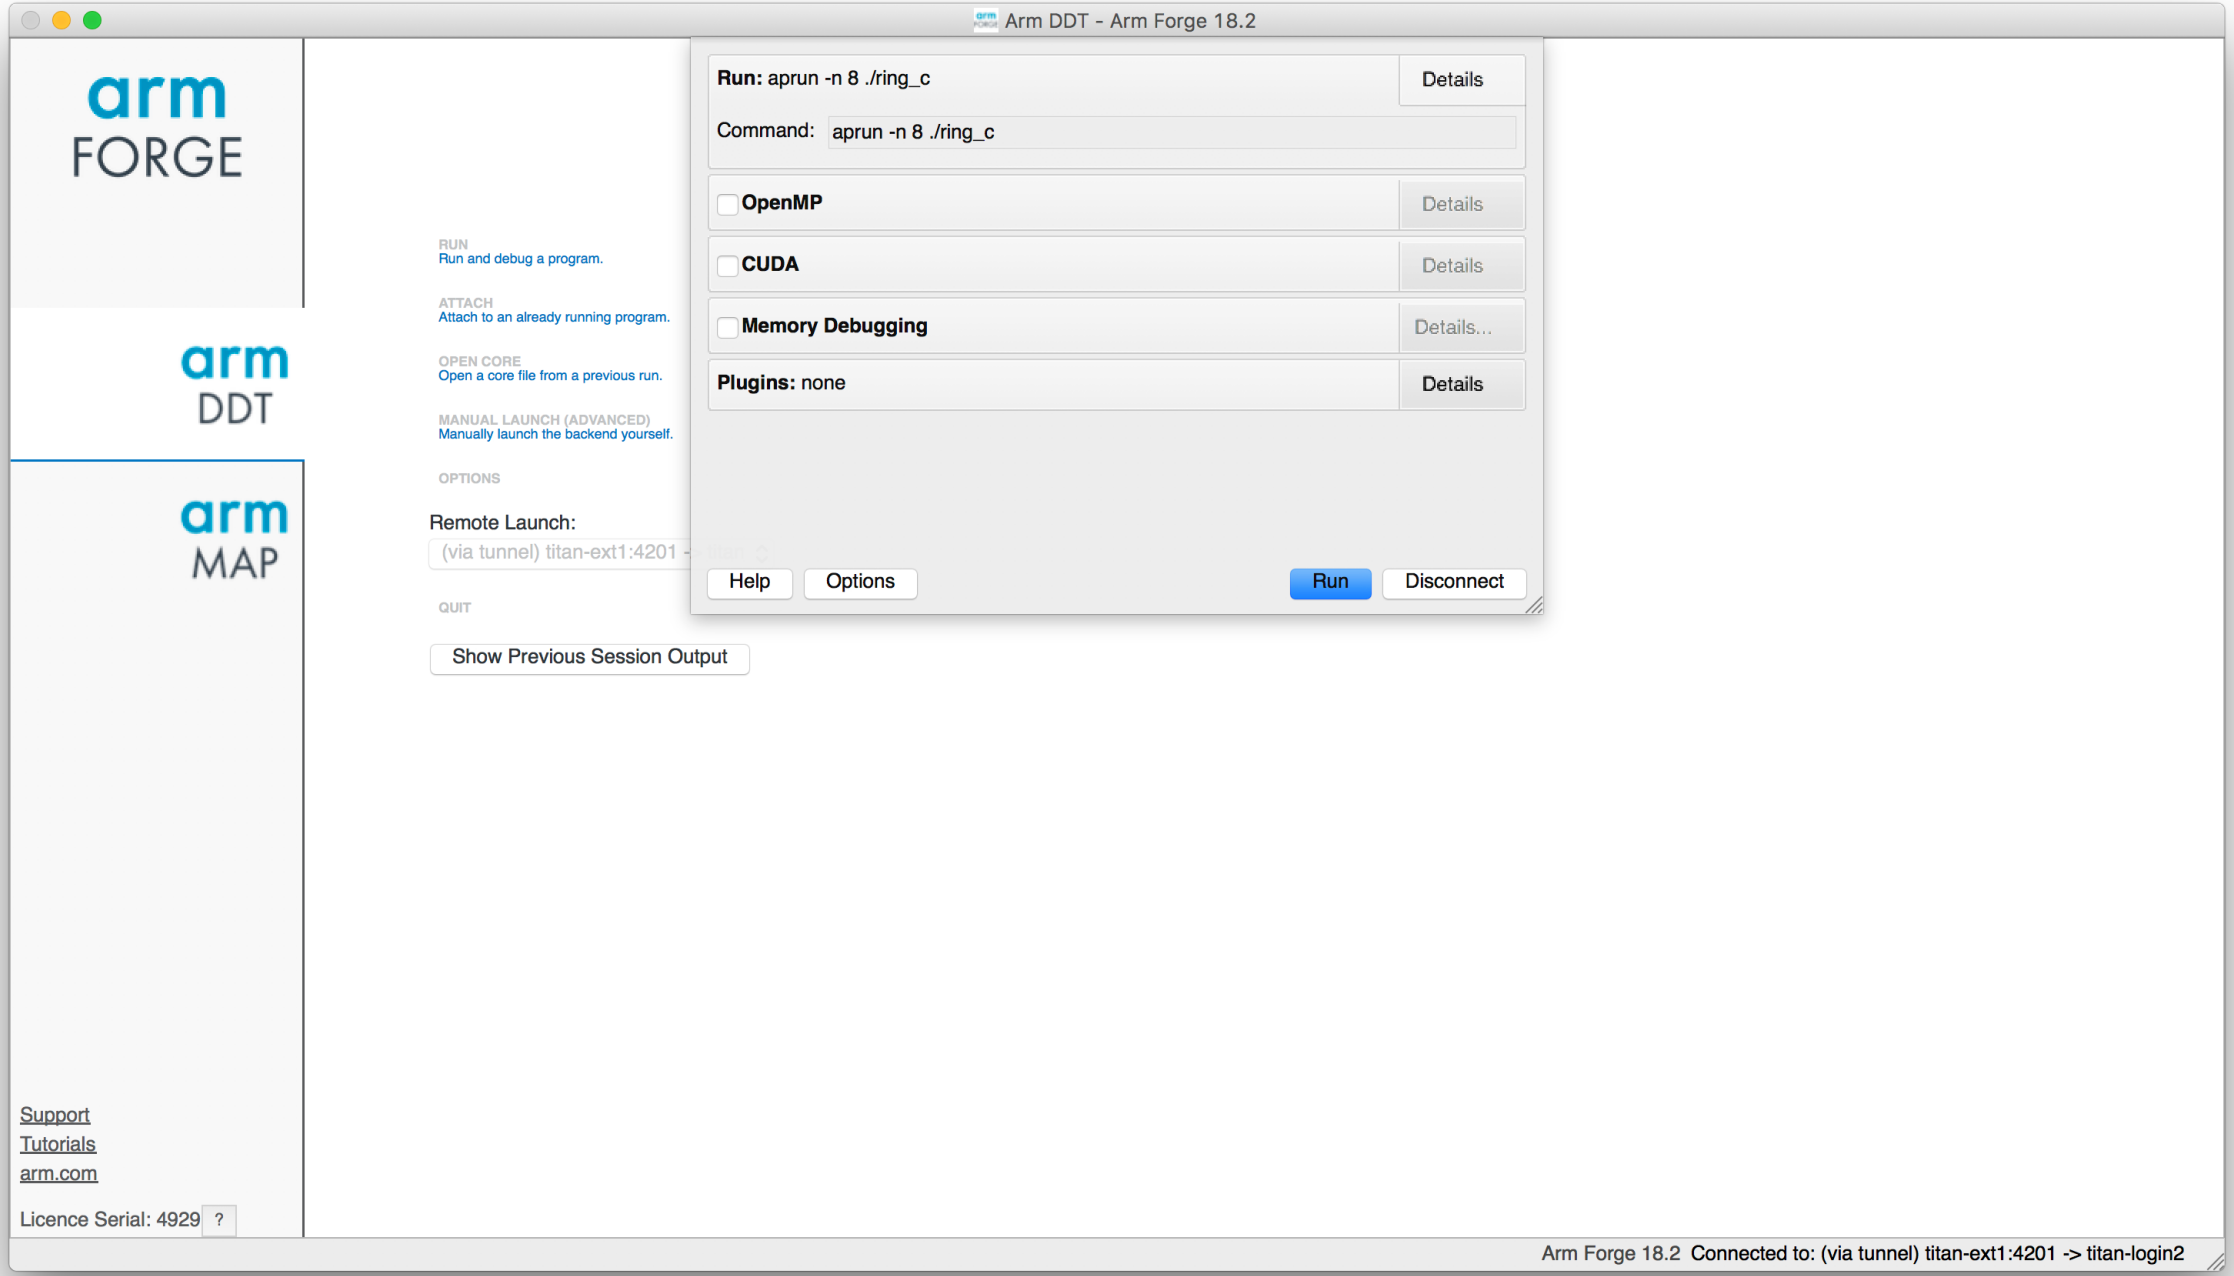
Arm DDT - Arm Forge 18.2
arm FORGE
arm DDT
arm MAP
RUN
Run and debug a program.
ATTACH
Attach to an already running program.
OPEN CORE
Open a core file from a previous run.
MANUAL LAUNCH (ADVANCED)
Manually launch the backend yourself.
OPTIONS
Remote Launch:
(via tunnel) titan-ext1:4201
QUIT
Show Previous Session Output
Support
Tutorials
arm.com
Licence Serial: 4929
?
Run: aprun -n 8 ./ring_c
Details
Command: aprun -n 8 ./ring_c
OpenMP
Details
CUDA
Details
Memory Debugging
Details...
Plugins: none
Details
Help
Options
Run
Disconnect
Arm Forge 18.2 Connected to: (via tunnel) titan-ext1:4201 -> titan-login2

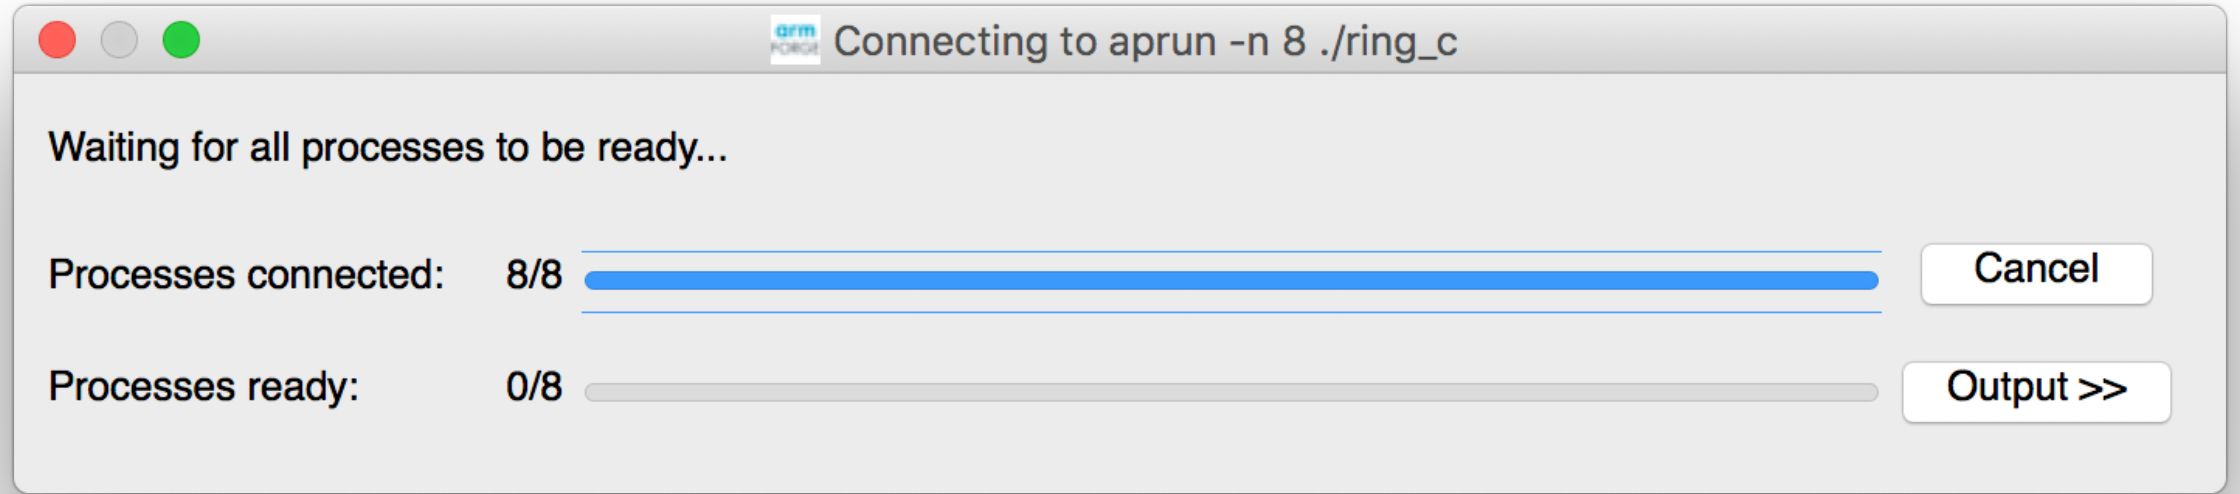
Connecting to aprun -n 8 ./ring_c
Waiting for all processes to be ready...
Processes connected: 8/8
Cancel
Processes ready: 0/8
Output >>

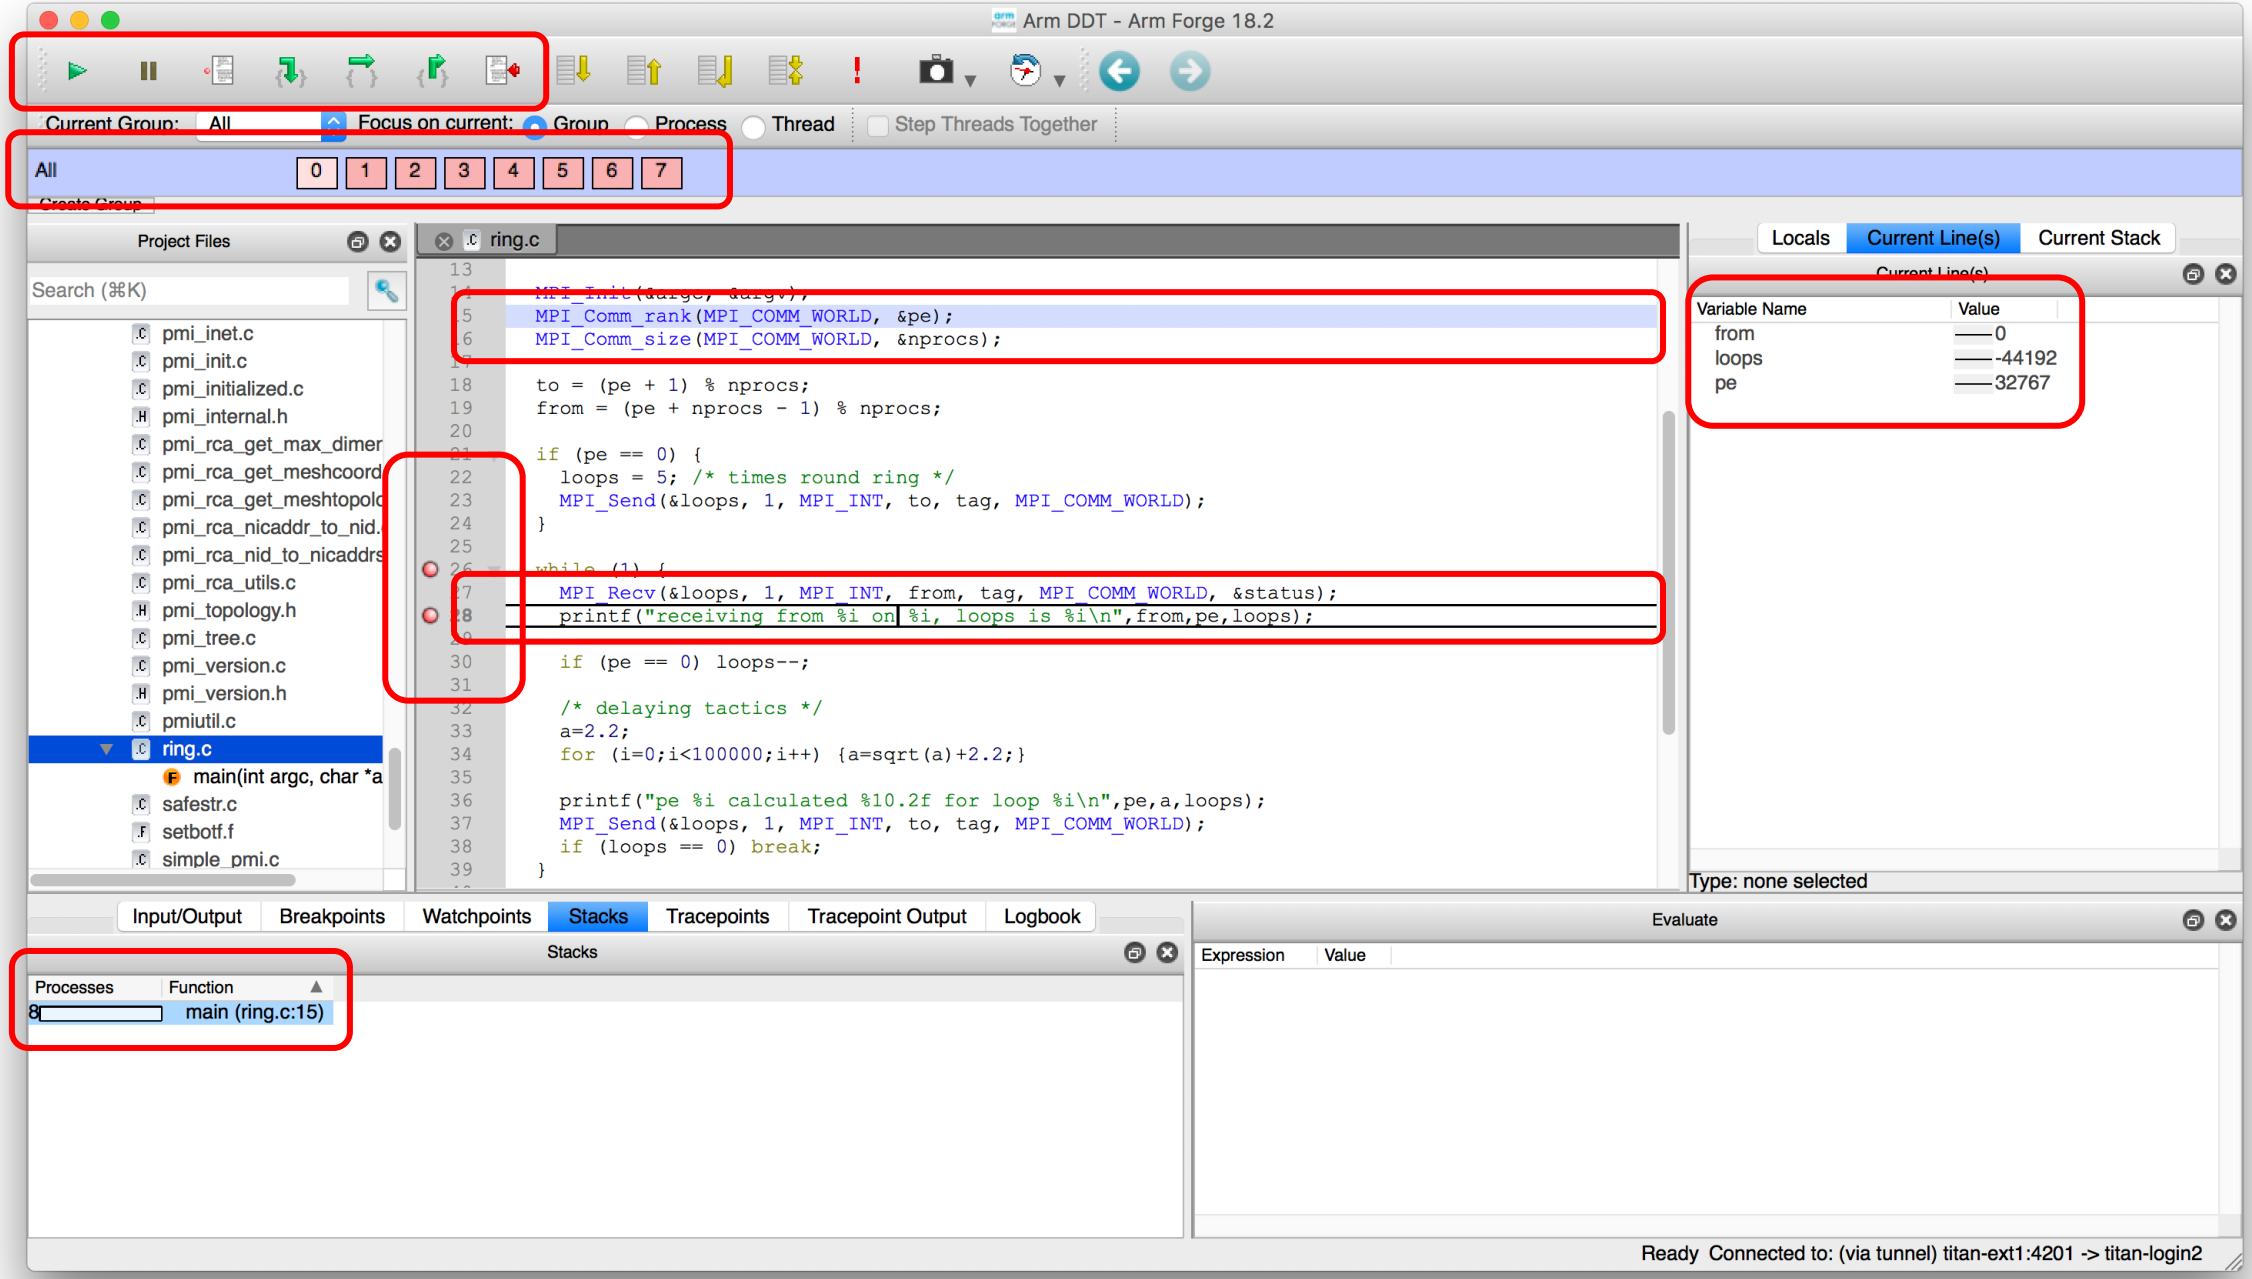
Arm DDT - Arm Forge 18.2
Current Group: All
Focus on current: Group Process Thread
Step Threads Together
All 0 1 2 3 4 5 6 7
Create Group
Project Files
Search (⌘K)
pmi_inet.c
pmi_init.c
pmi_initialized.c
pmi_internal.h
pmi_rca_get_max_dimer
pmi_rca_get_meshcoord
pmi_rca_get_meshtopolo
pmi_rca_nicaddr_to_nid.
pmi_rca_nid_to_nicaddrs
pmi_rca_utils.c
pmi_topology.h
pmi_tree.c
pmi_version.c
pmi_version.h
pmiutil.c
ring.c
main(int argc, char *a
safestr.c
setbotf.f
simple_pmi.c
ring.c
13
14 MPI_Init(&argc, &argv);
15 MPI_Comm_rank(MPI_COMM_WORLD, &pe);
16 MPI_Comm_size(MPI_COMM_WORLD, &nprocs);
17
18 to = (pe + 1) % nprocs;
19 from = (pe + nprocs - 1) % nprocs;
20
21 if (pe == 0) {
22 loops = 5; /* times round ring */
23 MPI_Send(&loops, 1, MPI_INT, to, tag, MPI_COMM_WORLD);
24 }
25
26 while (1) {
27 MPI_Recv(&loops, 1, MPI_INT, from, tag, MPI_COMM_WORLD, &status);
28 printf("receiving from %i on %i, loops is %i\n",from,pe,loops);
29
30 if (pe == 0) loops--;
31
32 /* delaying tactics */
33 a=2.2;
34 for (i=0;i<100000;i++) {a=sqrt(a)+2.2;}
35
36 printf("pe %i calculated %10.2f for loop %i\n",pe,a,loops);
37 MPI_Send(&loops, 1, MPI_INT, to, tag, MPI_COMM_WORLD);
38 if (loops == 0) break;
39 }
Locals
Current Line(s)
Current Stack
Current Line(s)
Variable Name
Value
from
0
loops
-44192
pe
32767
Type: none selected
Input/Output
Breakpoints
Watchpoints
Stacks
Tracepoints
Tracepoint Output
Logbook
Stacks
Processes
Function
8
main (ring.c:15)
Evaluate
Expression
Value
Ready Connected to: (via tunnel) titan-ext1:4201 -> titan-login2

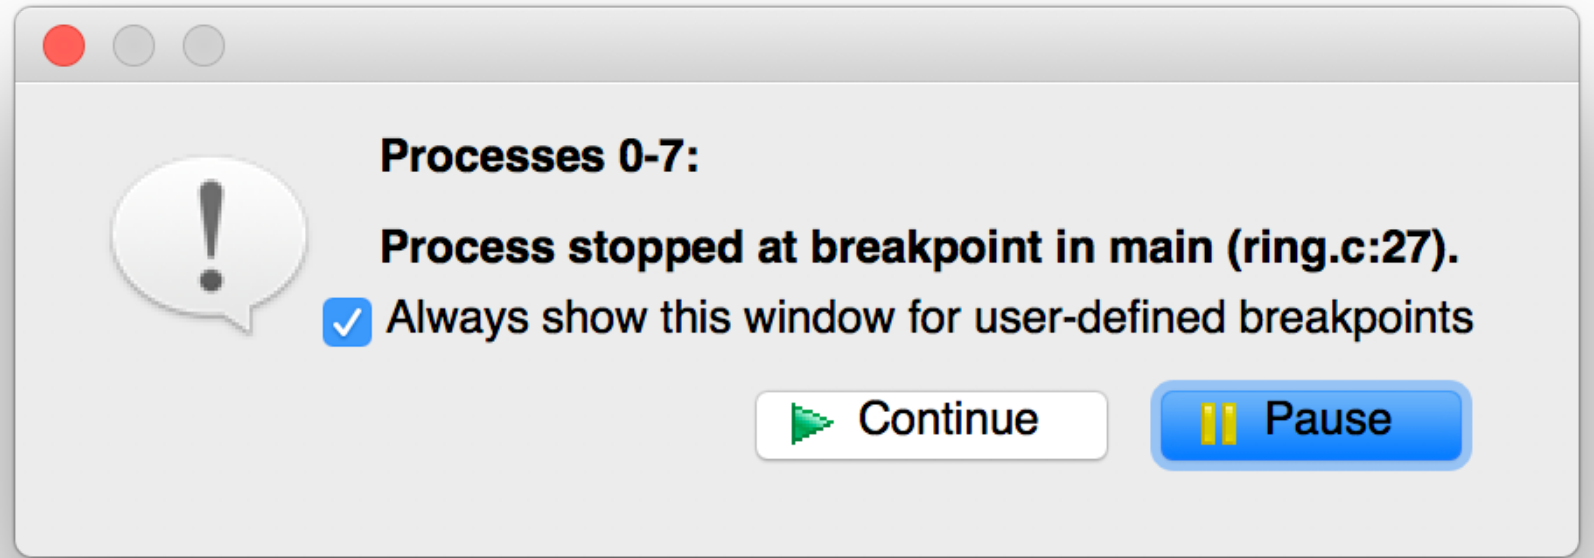
Processes 0-7:
Process stopped at breakpoint in main (ring.c:27).
Always show this window for user-defined breakpoints
Continue
Pause

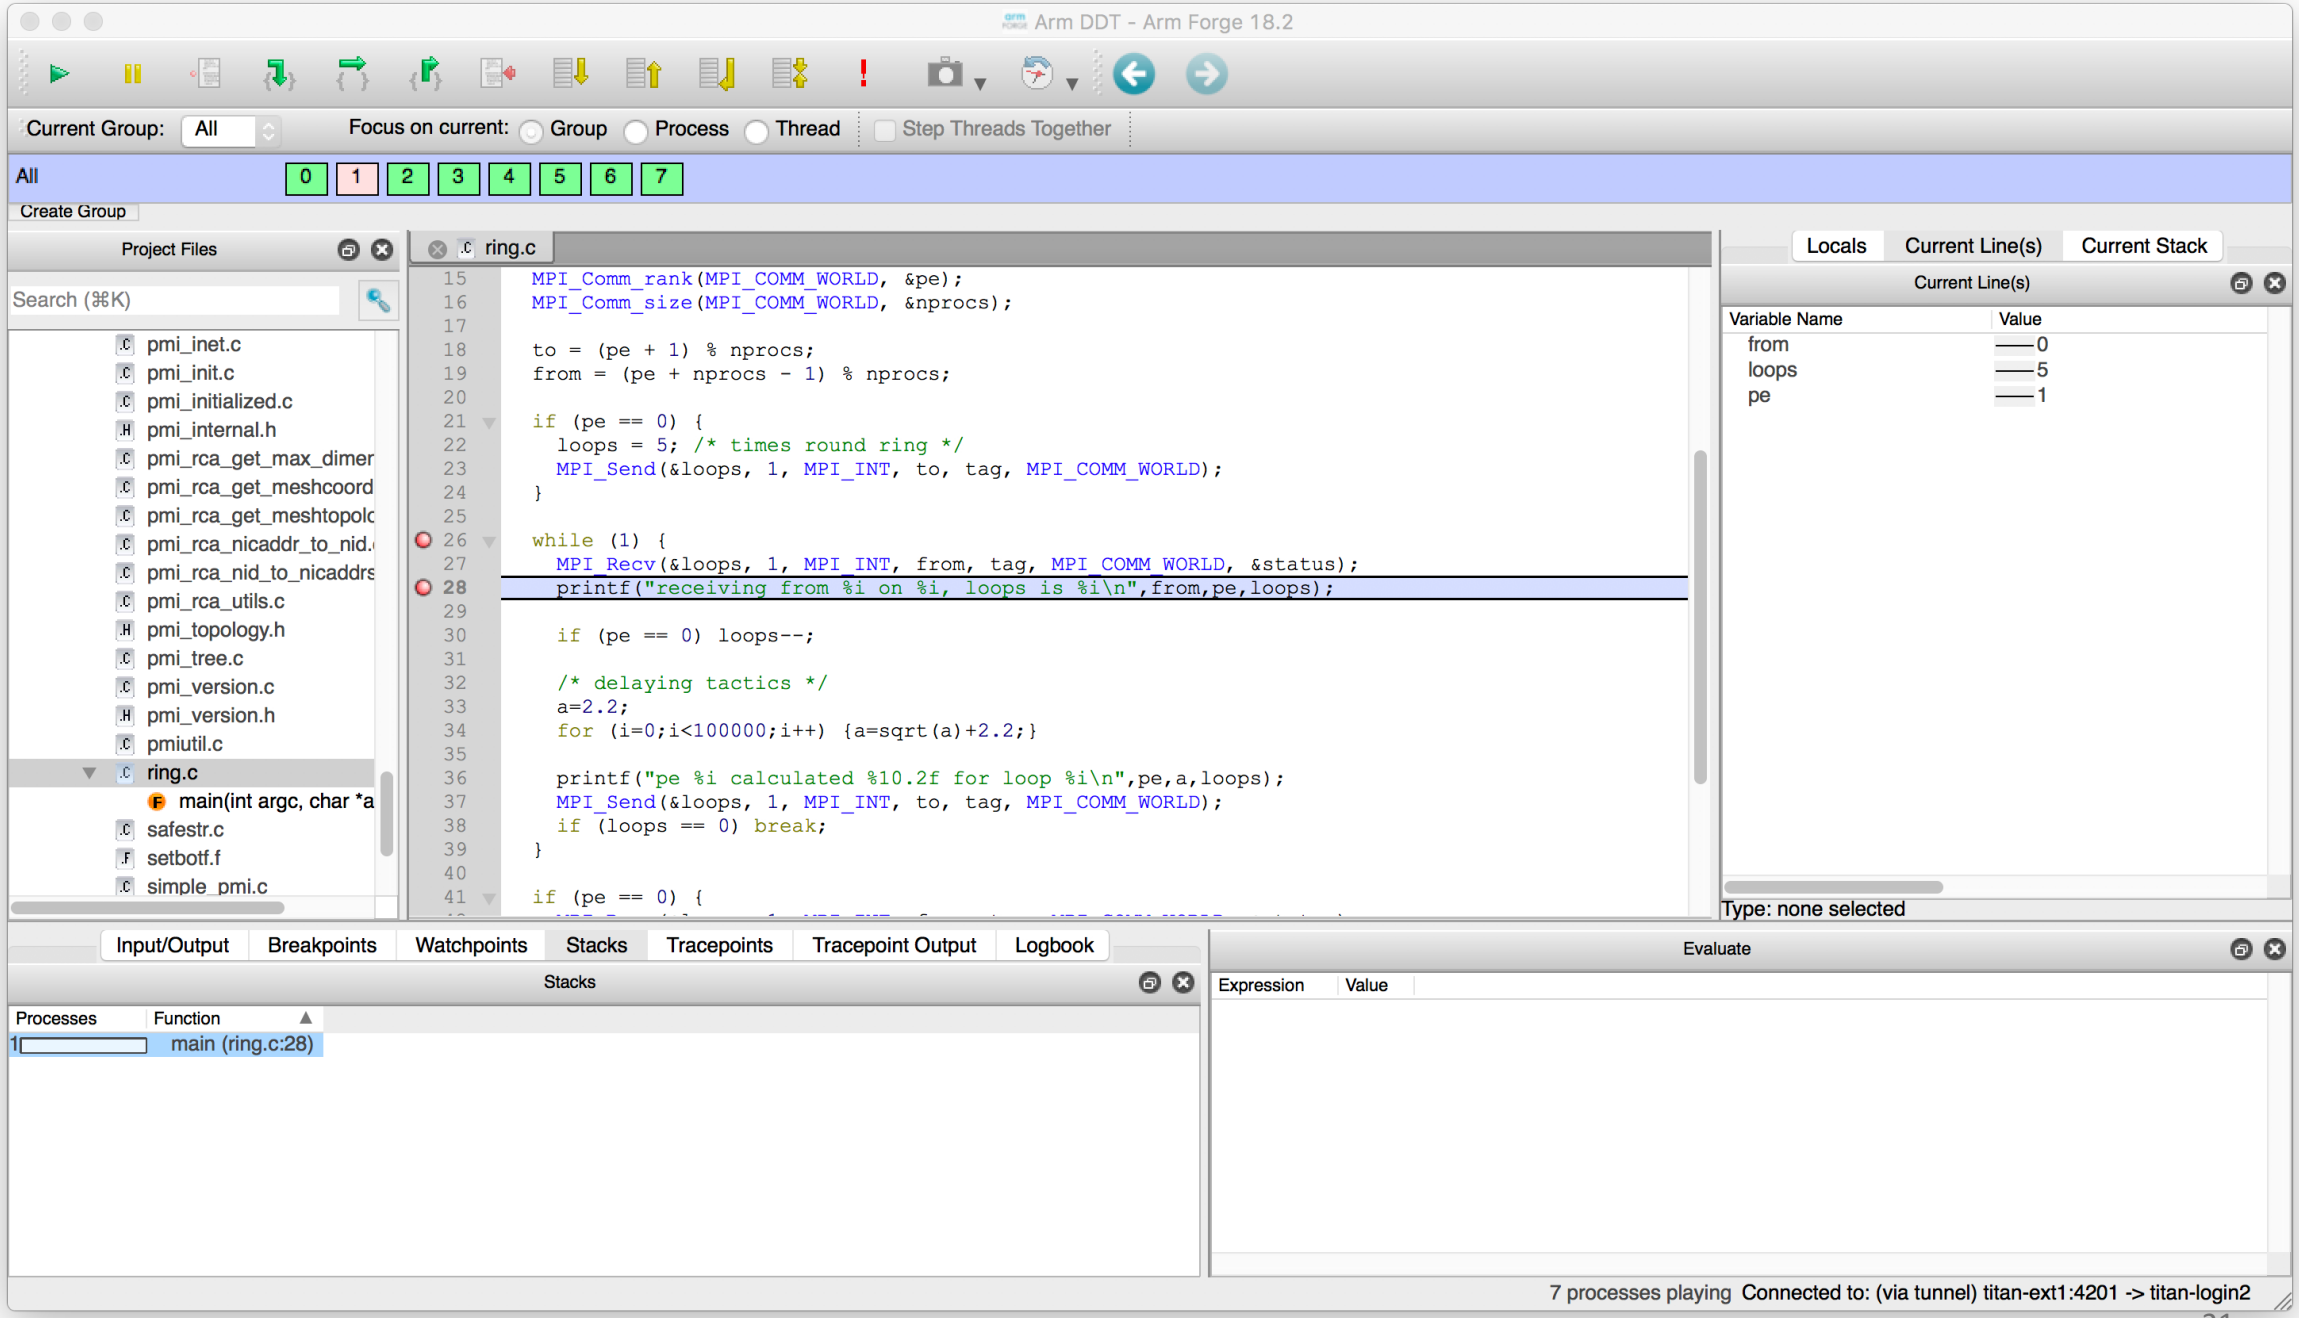
Arm DDT - Arm Forge 18.2
Current Group: All
Focus on current: Group Process Thread
Step Threads Together
All 0 1 2 3 4 5 6 7
Create Group
Project Files
Search (⌘K)
pmi_inet.c
pmi_init.c
pmi_initialized.c
pmi_internal.h
pmi_rca_get_max_dimer
pmi_rca_get_meshcoord
pmi_rca_get_meshtopolc
pmi_rca_nicaddr_to_nid.
pmi_rca_nid_to_nicaddrs
pmi_rca_utils.c
pmi_topology.h
pmi_tree.c
pmi_version.c
pmi_version.h
pmiutil.c
ring.c
main(int argc, char *a
safestr.c
setbotf.f
simple_pmi.c
ring.c
15 MPI_Comm_rank(MPI_COMM_WORLD, &pe);
16 MPI_Comm_size(MPI_COMM_WORLD, &nprocs);
17
18 to = (pe + 1) % nprocs;
19 from = (pe + nprocs - 1) % nprocs;
20
21 if (pe == 0) {
22 loops = 5; /* times round ring */
23 MPI_Send(&loops, 1, MPI_INT, to, tag, MPI_COMM_WORLD);
24 }
25
26 while (1) {
27 MPI_Recv(&loops, 1, MPI_INT, from, tag, MPI_COMM_WORLD, &status);
28 printf("receiving from %i on %i, loops is %i\n",from,pe,loops);
29
30 if (pe == 0) loops--;
31
32 /* delaying tactics */
33 a=2.2;
34 for (i=0;i<100000;i++) {a=sqrt(a)+2.2;}
35
36 printf("pe %i calculated %10.2f for loop %i\n",pe,a,loops);
37 MPI_Send(&loops, 1, MPI_INT, to, tag, MPI_COMM_WORLD);
38 if (loops == 0) break;
39 }
40
41 if (pe == 0) {
Locals Current Line(s) Current Stack
Current Line(s)
Variable Name Value
from 0
loops 5
pe 1
Type: none selected
Input/Output Breakpoints Watchpoints Stacks Tracepoints Tracepoint Output Logbook
Stacks
Processes Function
1 main (ring.c:28)
Evaluate
Expression Value
7 processes playing Connected to: (via tunnel) titan-ext1:4201 -> titan-login2

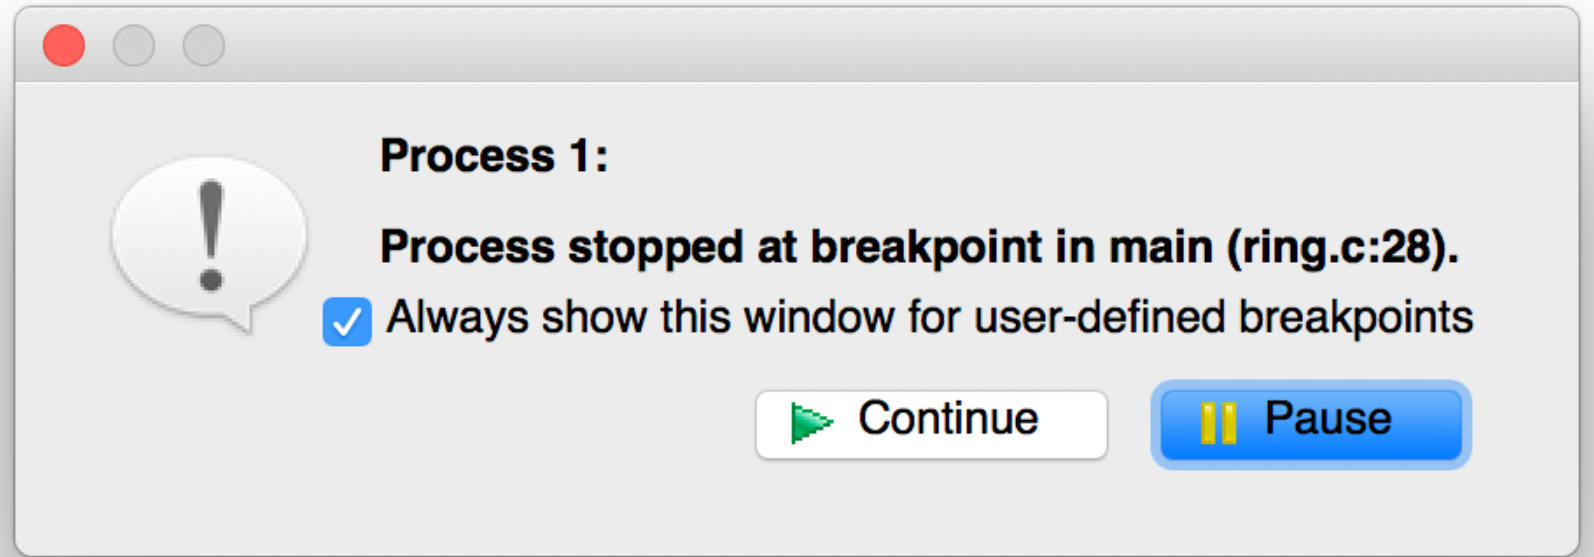
Process 1:
Process stopped at breakpoint in main (ring.c:28).
Always show this window for user-defined breakpoints
Continue
Pause

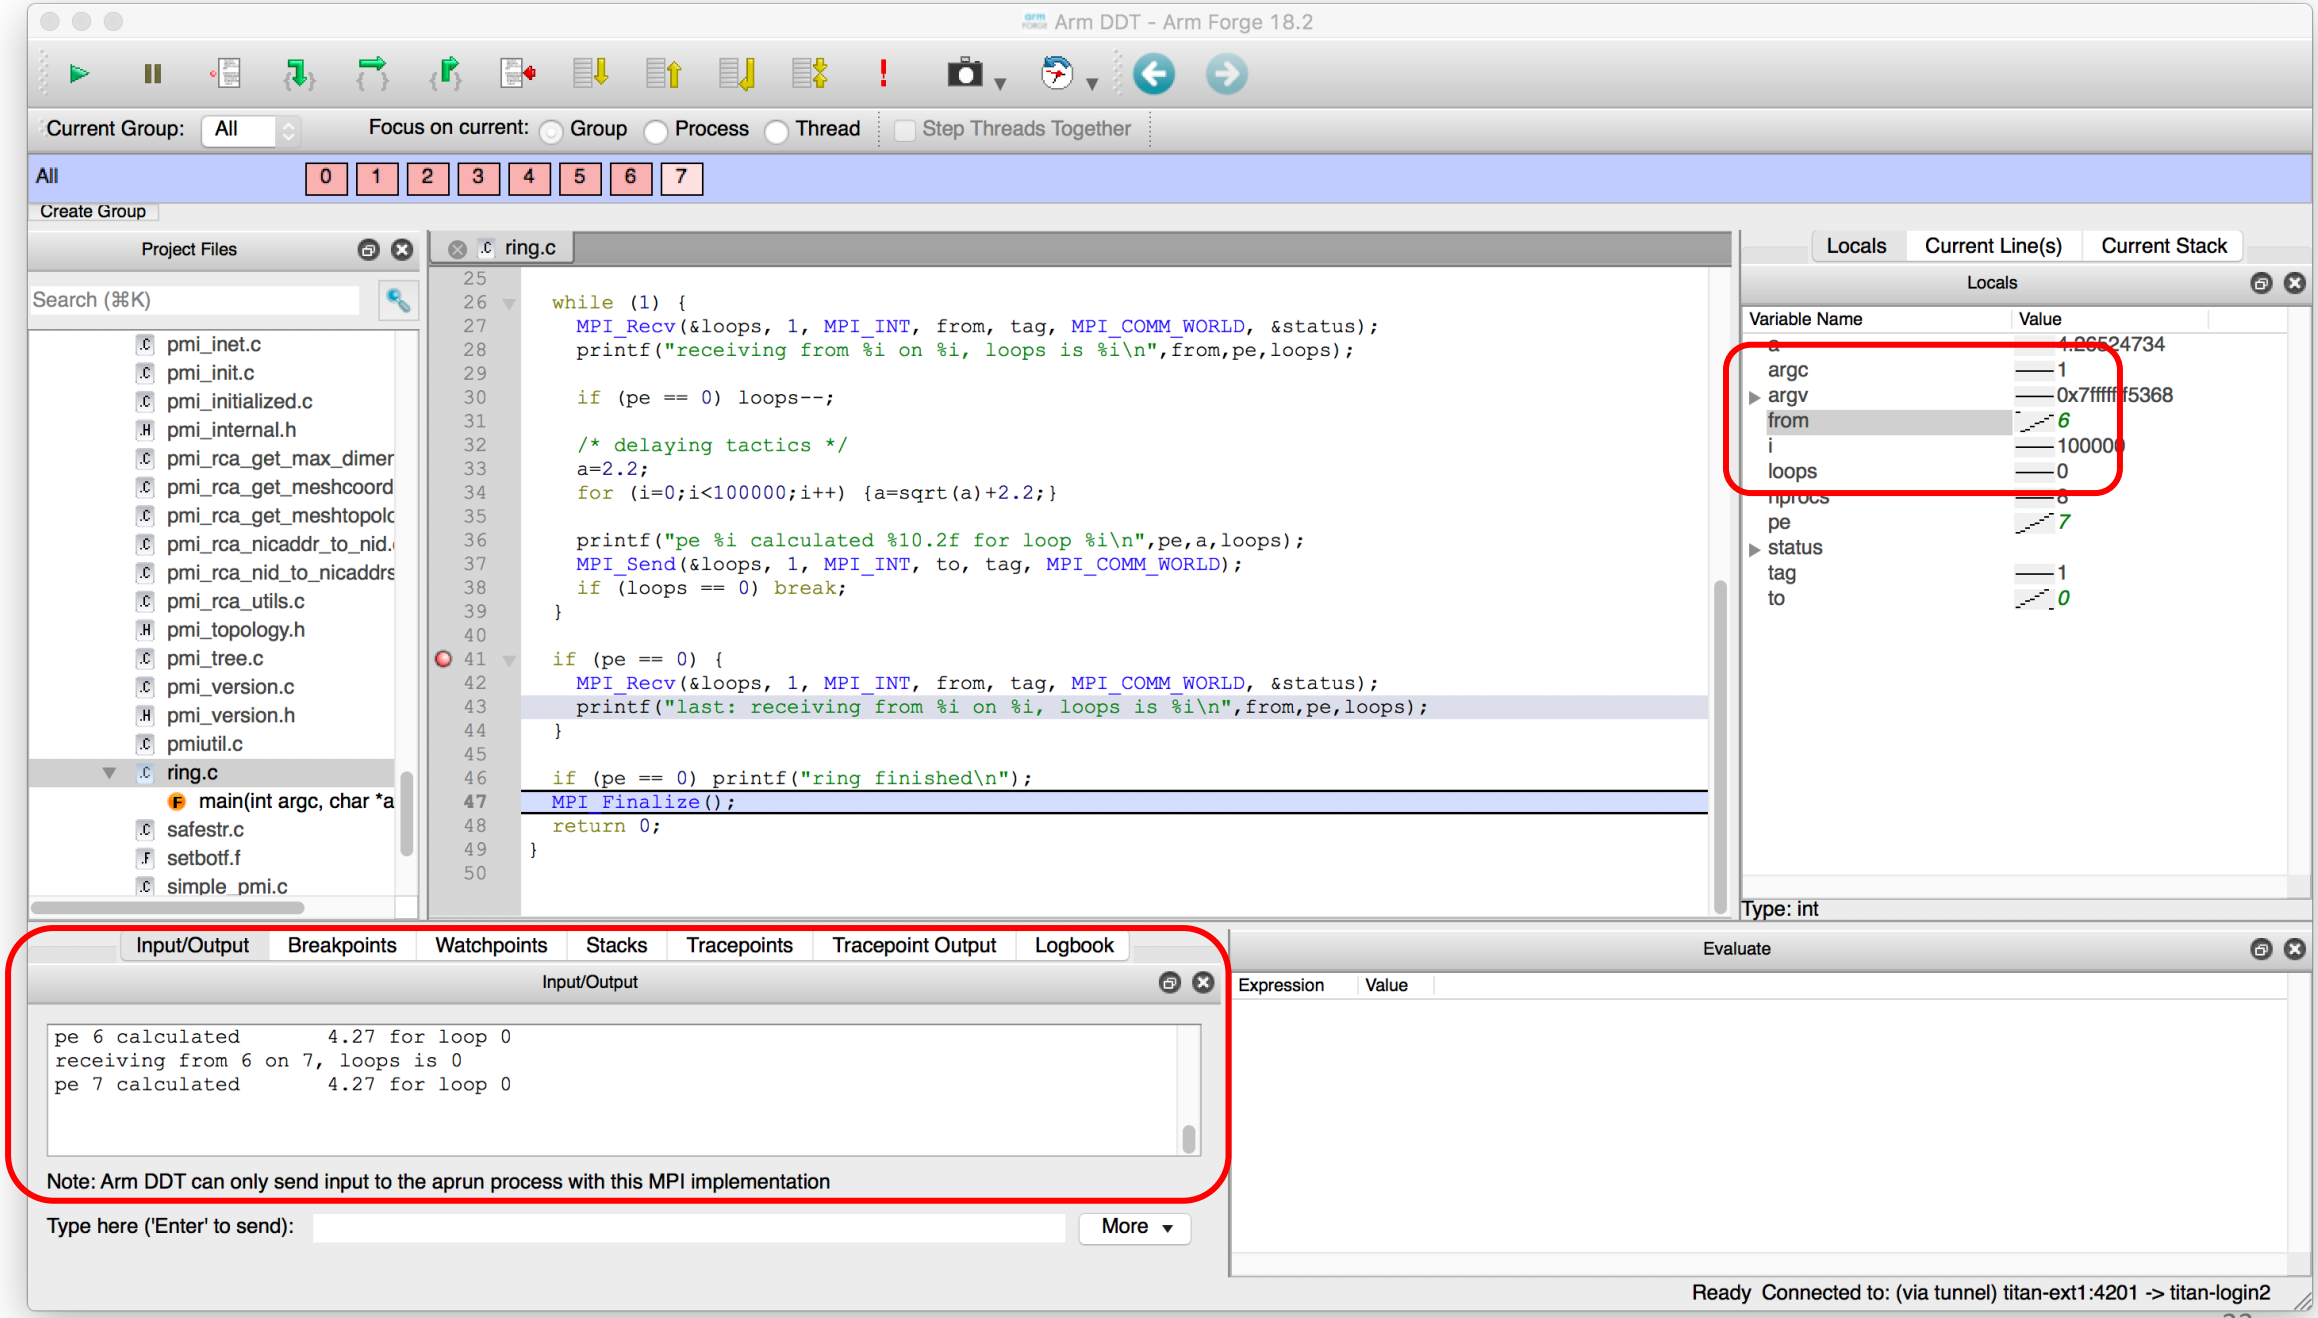
Arm DDT - Arm Forge 18.2
Current Group: All
Focus on current: Group Process Thread
Step Threads Together
All 0 1 2 3 4 5 6 7
Create Group
Project Files
Search (⌘K)
pmi_inet.c
pmi_init.c
pmi_initialized.c
pmi_internal.h
pmi_rca_get_max_dimer
pmi_rca_get_meshcoord
pmi_rca_get_meshtopolc
pmi_rca_nicaddr_to_nid.
pmi_rca_nid_to_nicaddrs
pmi_rca_utils.c
pmi_topology.h
pmi_tree.c
pmi_version.c
pmi_version.h
pmiutil.c
ring.c
main(int argc, char *a
safestr.c
setbotf.f
simple_pmi.c
ring.c
while (1) {
MPI_Recv(&loops, 1, MPI_INT, from, tag, MPI_COMM_WORLD, &status);
printf("receiving from %i on %i, loops is %i\n",from,pe,loops);
if (pe == 0) loops--;
/* delaying tactics */
a=2.2;
for (i=0;i<100000;i++) {a=sqrt(a)+2.2;}
printf("pe %i calculated %10.2f for loop %i\n",pe,a,loops);
MPI_Send(&loops, 1, MPI_INT, to, tag, MPI_COMM_WORLD);
if (loops == 0) break;
}
if (pe == 0) {
MPI_Recv(&loops, 1, MPI_INT, from, tag, MPI_COMM_WORLD, &status);
printf("last: receiving from %i on %i, loops is %i\n",from,pe,loops);
}
if (pe == 0) printf("ring finished\n");
MPI_Finalize();
return 0;
}
Locals
Current Line(s)
Current Stack
Locals
Variable Name
Value
argc 1
argv 0x7fffff5368
from 6
i 100000
loops 0
pe 7
status
tag 1
to 0
Type: int
Input/Output
Breakpoints
Watchpoints
Stacks
Tracepoints
Tracepoint Output
Logbook
Input/Output
pe 6 calculated 4.27 for loop 0
receiving from 6 on 7, loops is 0
pe 7 calculated 4.27 for loop 0
Note: Arm DDT can only send input to the aprun process with this MPI implementation
Type here ('Enter' to send):
More
Evaluate
Expression
Value
Ready Connected to: (via tunnel) titan-ext1:4201 -> titan-login2

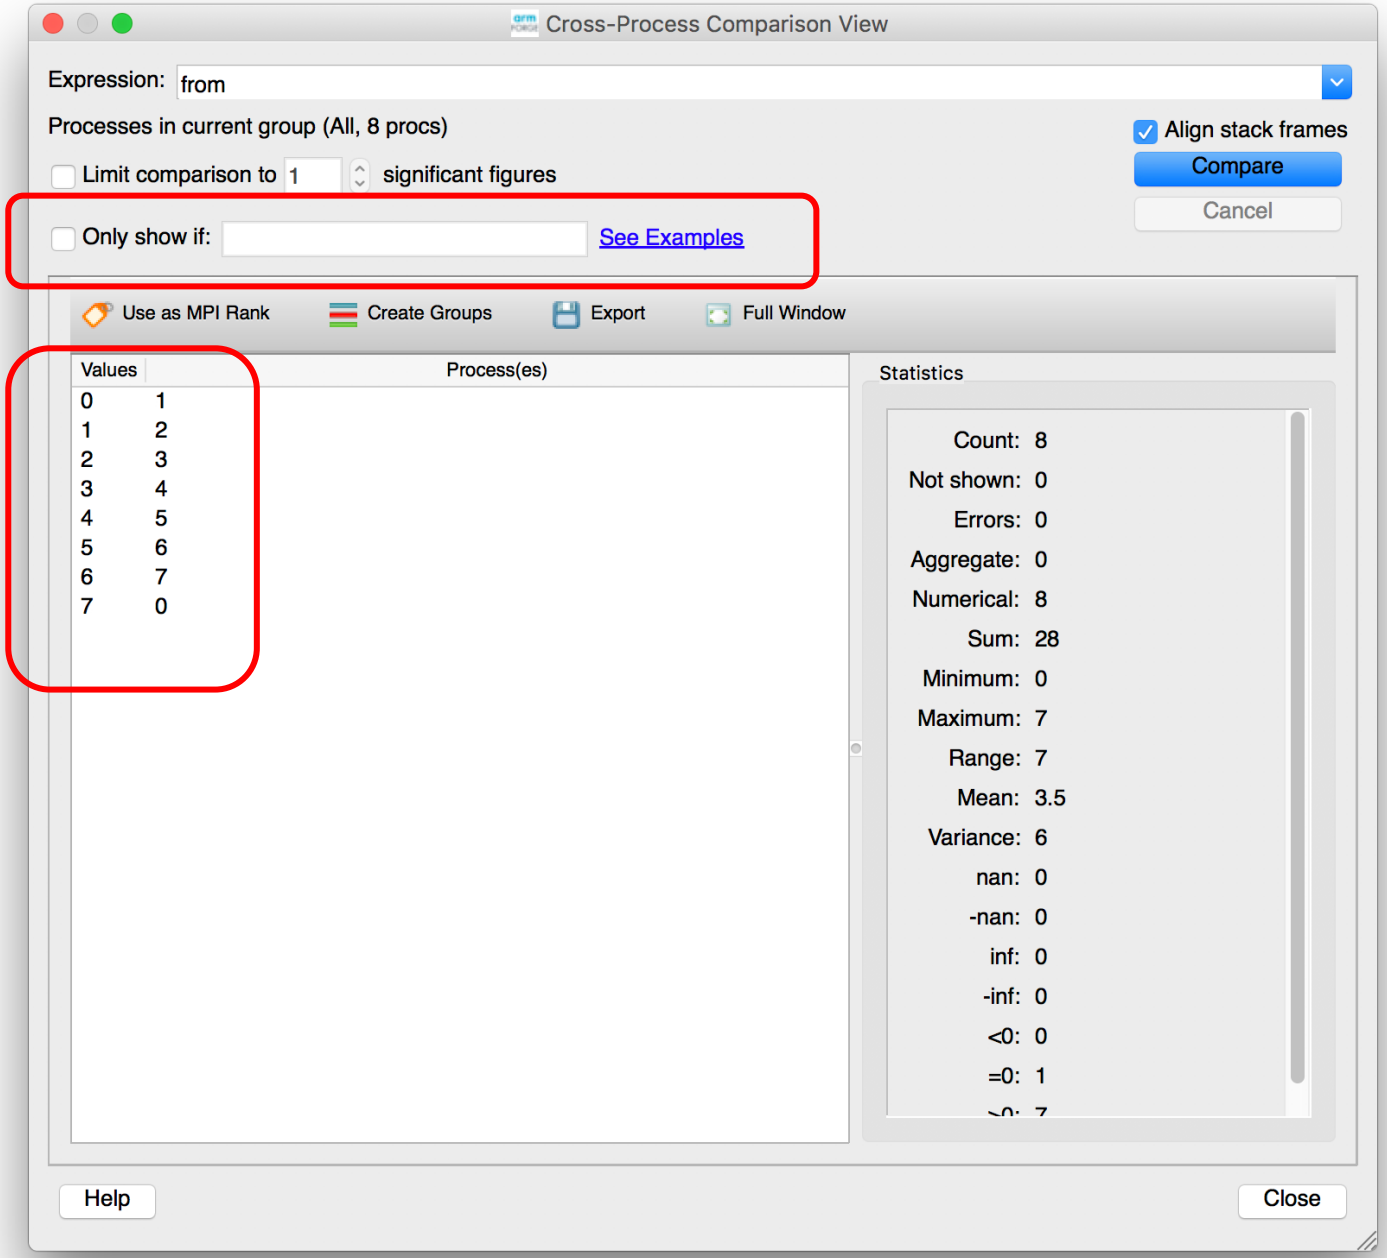
Cross-Process Comparison View
Expression: from
Processes in current group (All, 8 procs)
Limit comparison to 1 significant figures
Only show if:
See Examples
Align stack frames
Compare
Cancel
Use as MPI Rank
Create Groups
Export
Full Window
Values
Process(es)
0 1
1 2
2 3
3 4
4 5
5 6
6 7
7 0
Statistics
Count: 8
Not shown: 0
Errors: 0
Aggregate: 0
Numerical: 8
Sum: 28
Minimum: 0
Maximum: 7
Range: 7
Mean: 3.5
Variance: 6
nan: 0
-nan: 0
inf: 0
-inf: 0
<0: 0
=0: 1
Help
Close

CUDA-GDB on Titan

- module load cudatoolkit
- cd $PROJWORK/csc261/$USER/csgf
- nvcc **-g -G** bitreserve.cu -o bitreverse
- aprun -n 1 cuda-gdb ./bitreverse

Common commands

- run/continue
- break
- next
- info cuda threads/warps/kernels
- cuda warp n
- quit

https://darkdust.net/files/GDB%20Cheat%20Sheet.pdf
https://docs.nvidia.com/cuda/cuda-gdb/index.html

```
 1 #include <stdio.h>
 2 #include <stdlib.h>
 3
 4 // Simple 8-bit bit reversal Compute test
 5
 6 #define N 256
 7
 8 __global__ void bitreverse(void *data) {
 9    unsigned int *idata = (unsigned int*)data;
10    extern __shared__ int array[];
11
12    array[threadIdx.x] = idata[threadIdx.x];
13
14    array[threadIdx.x] = ((0xf0f0f0f0 & array[threadIdx.x]) >> 4) |
15                         ((0x0f0f0f0f & array[threadIdx.x]) << 4);
16    array[threadIdx.x] = ((0xcccccccc & array[threadIdx.x]) >> 2) |
17                         ((0x33333333 & array[threadIdx.x]) << 2);
18    array[threadIdx.x] = ((0xaaaaaaaa & array[threadIdx.x]) >> 1) |
19                          ((0x55555555 & array[threadIdx.x]) << 1);
20
21    idata[threadIdx.x] = array[threadIdx.x];
22 }
23
```

```
24 int main(void) {
25     void *d = NULL; int i;
26     unsigned int idata[N], odata[N];
27
28     for (i = 0; i < N; i++)
29         idata[i] = (unsigned int)i;
30
31     cudaMalloc((void**)&d, sizeof(int)*N);
32     cudaMemcpy(d, idata, sizeof(int)*N,
33                cudaMemcpyHostToDevice);
34
35     bitreverse<<<1, N, N*sizeof(int)>>>(d);
36
37     cudaMemcpy(odata, d, sizeof(int)*N,
38                cudaMemcpyDeviceToHost);
39
40     for (i = 0; i < N; i++)
41        printf("%u -> %u\n", idata[i], odata[i]);
42
43     cudaFree((void*)d);
44     return 0;
45 }
```

```
(cuda-gdb) b 12

Breakpoint 1 at 0x40365c: file bitreverse.cu, line
12.

(cuda-gdb) b 21

Note: breakpoint 1 also set at pc 0x40365c.

Breakpoint 2 at 0x40365c: file bitreverse.cu, line
21.

(cuda-gdb) r

Starting program: /lustre/atlas2/csc261/proj-
shared/cjyang/csgf/bitreverse

[Thread debugging using libthread_db enabled]

Using host libthread_db library
"/lib64/libthread_db.so.1".

[New Thread 0x2aaaacdfa700 (LWP 1515)]

[New Thread 0x2aaaacffb700 (LWP 1516)]

[Switching focus to CUDA kernel 0, grid 1, block
(0,0,0), thread (0,0,0), device 0, sm 13, warp 0,
lane 0]

Thread 1 "bitreverse" hit Breakpoint 1,
bitreverse<<<(1,1,1),(256,1,1)>>> (

    data=0xb05580000) at bitreverse.cu:12

12   array[threadIdx.x] = idata[threadIdx.x];
```

```
(cuda-gdb) info cuda threads

  BlockIdx ThreadIdx To BlockIdx ThreadIdx
Count         Virtual PC      Filename  Line

Kernel 0

*  (0,0,0)   (0,0,0)     (0,0,0) (255,0,0)   256
0x0000000000b07a98 bitreverse.cu    12

(cuda-gdb) n

14   array[threadIdx.x] = ((0xf0f0f0f0 &
array[threadIdx.x]) >> 4) |

(cuda-gdb) n

16   array[threadIdx.x] = ((0xcccccccc &
array[threadIdx.x]) >> 2) |

(cuda-gdb) info cuda threads

  BlockIdx ThreadIdx To BlockIdx ThreadIdx
Count         Virtual PC      Filename  Line

Kernel 0

*  (0,0,0)   (0,0,0)     (0,0,0)  (31,0,0)    32
0x0000000000b07f60 bitreverse.cu    16

   (0,0,0)  (32,0,0)     (0,0,0) (255,0,0)   224
0x0000000000b07a98 bitreverse.cu    12

(cuda-gdb) cuda warp

warp 0
```

```
(cuda-gdb) cuda warp 2

[Switching focus to CUDA kernel 0, grid 1, block
(0,0,0), thread (64,0,0), device 0, sm 13, warp 2,
lane 0]

12   array[threadIdx.x] = idata[threadIdx.x];

(cuda-gdb) n

14   array[threadIdx.x] = ((0xf0f0f0f0 &
array[threadIdx.x]) >> 4) |

(cuda-gdb) n

16   array[threadIdx.x] = ((0xcccccccc &
array[threadIdx.x]) >> 2) |

(cuda-gdb) info cuda threads

  BlockIdx ThreadIdx To BlockIdx ThreadIdx
Count         Virtual PC      Filename  Line

Kernel 0

   (0,0,0)   (0,0,0)     (0,0,0)  (31,0,0)    32
0x0000000000b07f60 bitreverse.cu    16

   (0,0,0)  (32,0,0)     (0,0,0)  (63,0,0)    32
0x0000000000b07a98 bitreverse.cu    12

*  (0,0,0)  (64,0,0)     (0,0,0)  (95,0,0)    32
0x0000000000b07f60 bitreverse.cu    16

   (0,0,0)  (96,0,0)     (0,0,0) (255,0,0)   160
0x0000000000b07a98 bitreverse.cu    12
```

```
(cuda-gdb) c

Continuing.

[Switching focus to CUDA kernel 0, grid 1, block
(0,0,0), thread (0,0,0), device 0, sm 13, warp 0,
lane 0]


Thread 1 "bitreverse" hit Breakpoint 2,
bitreverse<<<(1,1,1),(256,1,1)>>> (

    data=0xb05580000) at bitreverse.cu:21

21   idata[threadIdx.x] = array[threadIdx.x];

(cuda-gdb) info cuda threads

  BlockIdx ThreadIdx To BlockIdx ThreadIdx
Count         Virtual PC      Filename  Line

Kernel 0

*  (0,0,0)   (0,0,0)     (0,0,0) (255,0,0)   256
0x0000000000b084e0 bitreverse.cu    21

(cuda-gdb) print array[0]@12

$1 = {0, 128, 64, 192, 32, 160, 96, 224, 16, 144,
80, 208}

(cuda-gdb) c

Continuing.

0 -> 0

1 -> 128 ……
```

Questions?