allinea

High performance tools to debug, profile, and analyze your applications

Illuminating OpenMP + MPI Performance

Beau Paisley, bpaisley@allinea.com

allinea FORGE allinea DDT allinea MAP allinea PERFORMANCE REPORTS

HPC means being able to work productively on remote machines

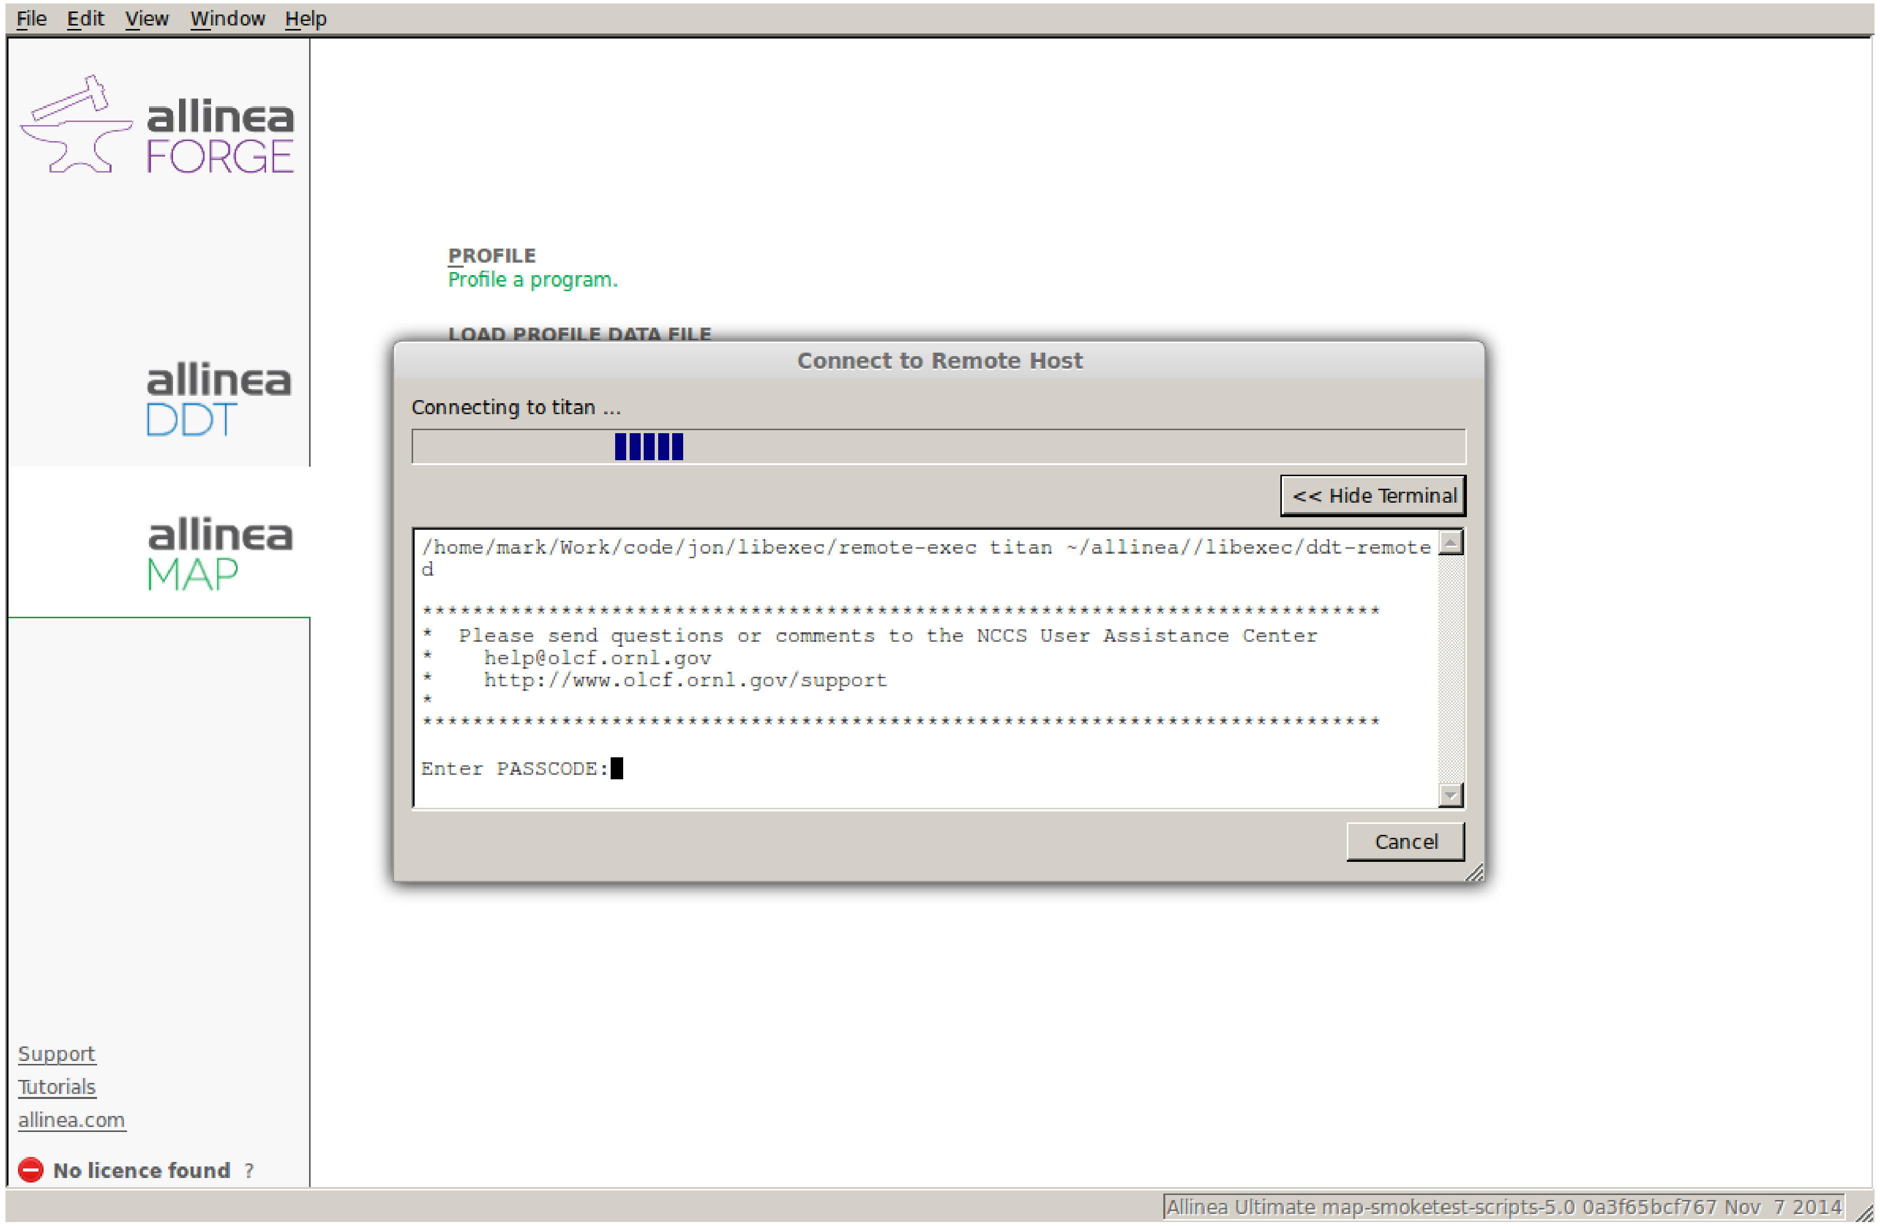

Connecting to titan ...

```
/home/mark/Work/code/jon/libexec/remote-exec titan ~/allinea//libexec/ddt-remote
d

*****************************************************************************
*  Please send questions or comments to the NCCS User Assistance Center
*    help@olcf.ornl.gov
*    http://www.olcf.ornl.gov/support
*
*****************************************************************************

Enter PASSCODE:
```

Calculating pi with 1 thread, 1 process

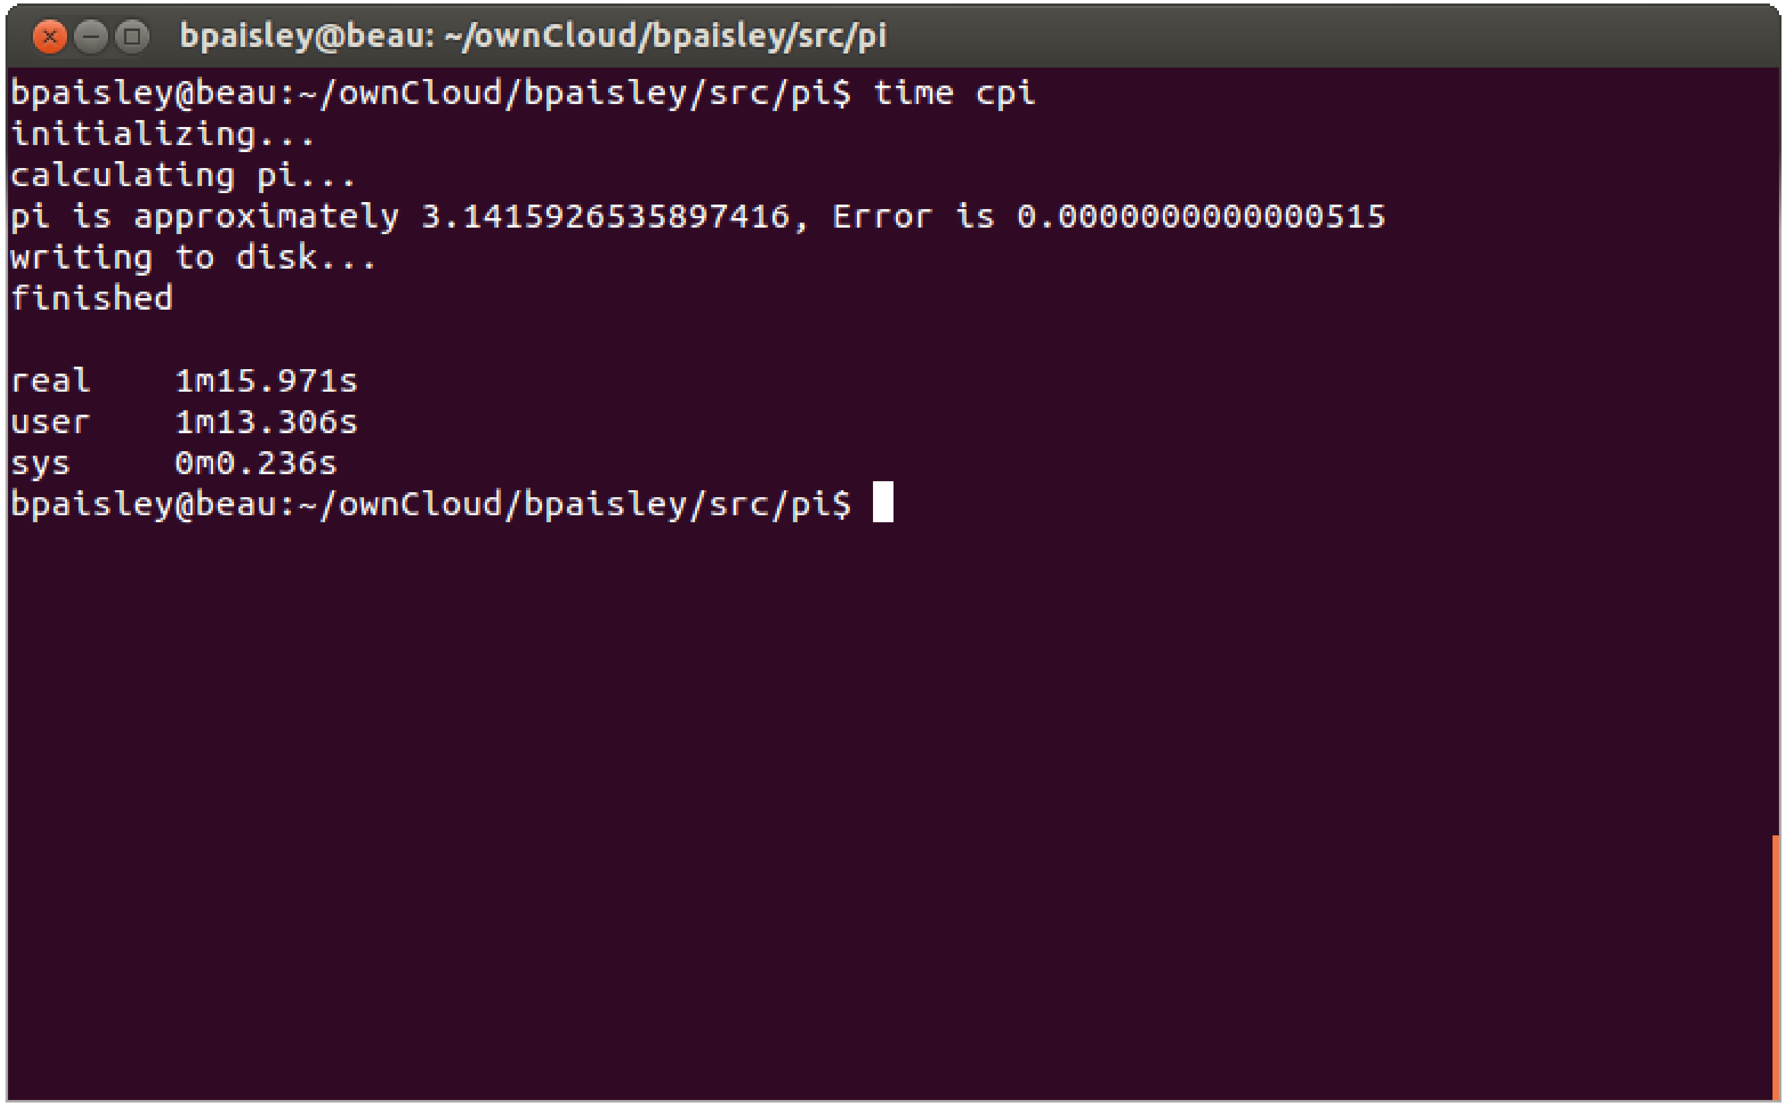
bpaisley@beau: ~/ownCloud/bpaisley/src/pi

```
bpaisley@beau:~/ownCloud/bpaisley/src/pi$ time cpi
initializing...
calculating pi...
pi is approximately 3.1415926535897416, Error is 0.0000000000000515
writing to disk...
finished

real	1m15.971s
user	1m13.306s
sys	0m0.236s
bpaisley@beau:~/ownCloud/bpaisley/src/pi$ 
```

Profile with 1 process, 1 thread

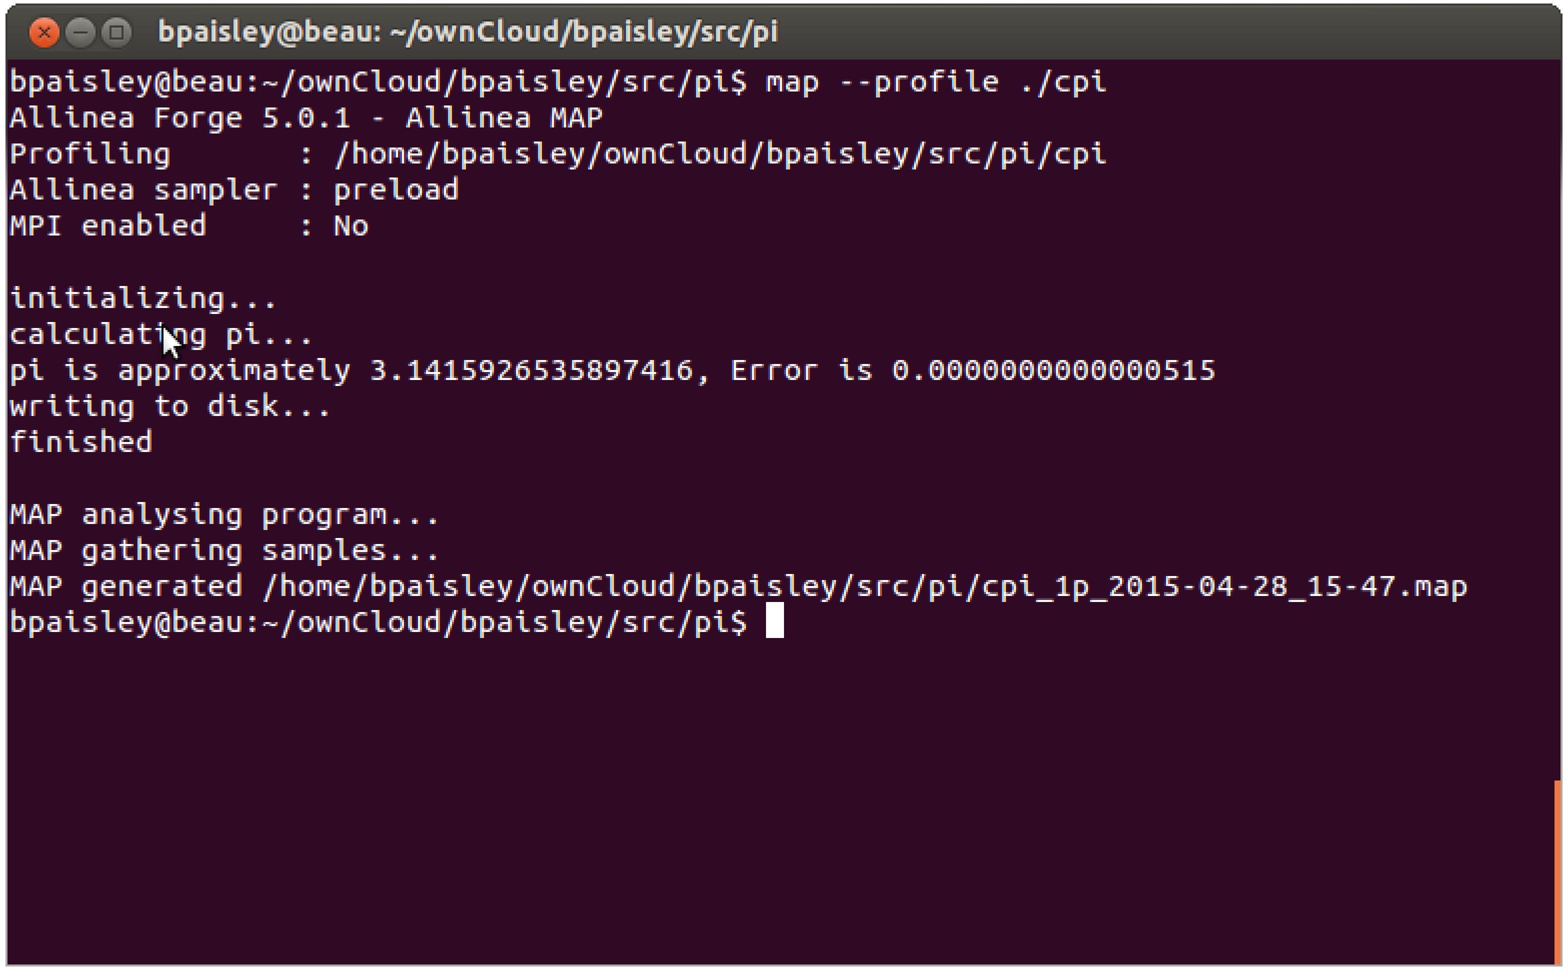

bpaisley@beau: ~/ownCloud/bpaisley/src/pi

```
bpaisley@beau:~/ownCloud/bpaisley/src/pi$ map --profile ./cpi
Allinea Forge 5.0.1 - Allinea MAP
Profiling        : /home/bpaisley/ownCloud/bpaisley/src/pi/cpi
Allinea sampler : preload
MPI enabled      : No

initializing...
calculating pi...
pi is approximately 3.1415926535897416, Error is 0.0000000000000515
writing to disk...
finished

MAP analysing program...
MAP gathering samples...
MAP generated /home/bpaisley/ownCloud/bpaisley/src/pi/cpi_1p_2015-04-28_15-47.map
bpaisley@beau:~/ownCloud/bpaisley/src/pi$
```

Using MAP to analyze our 1 process, 1 thread run

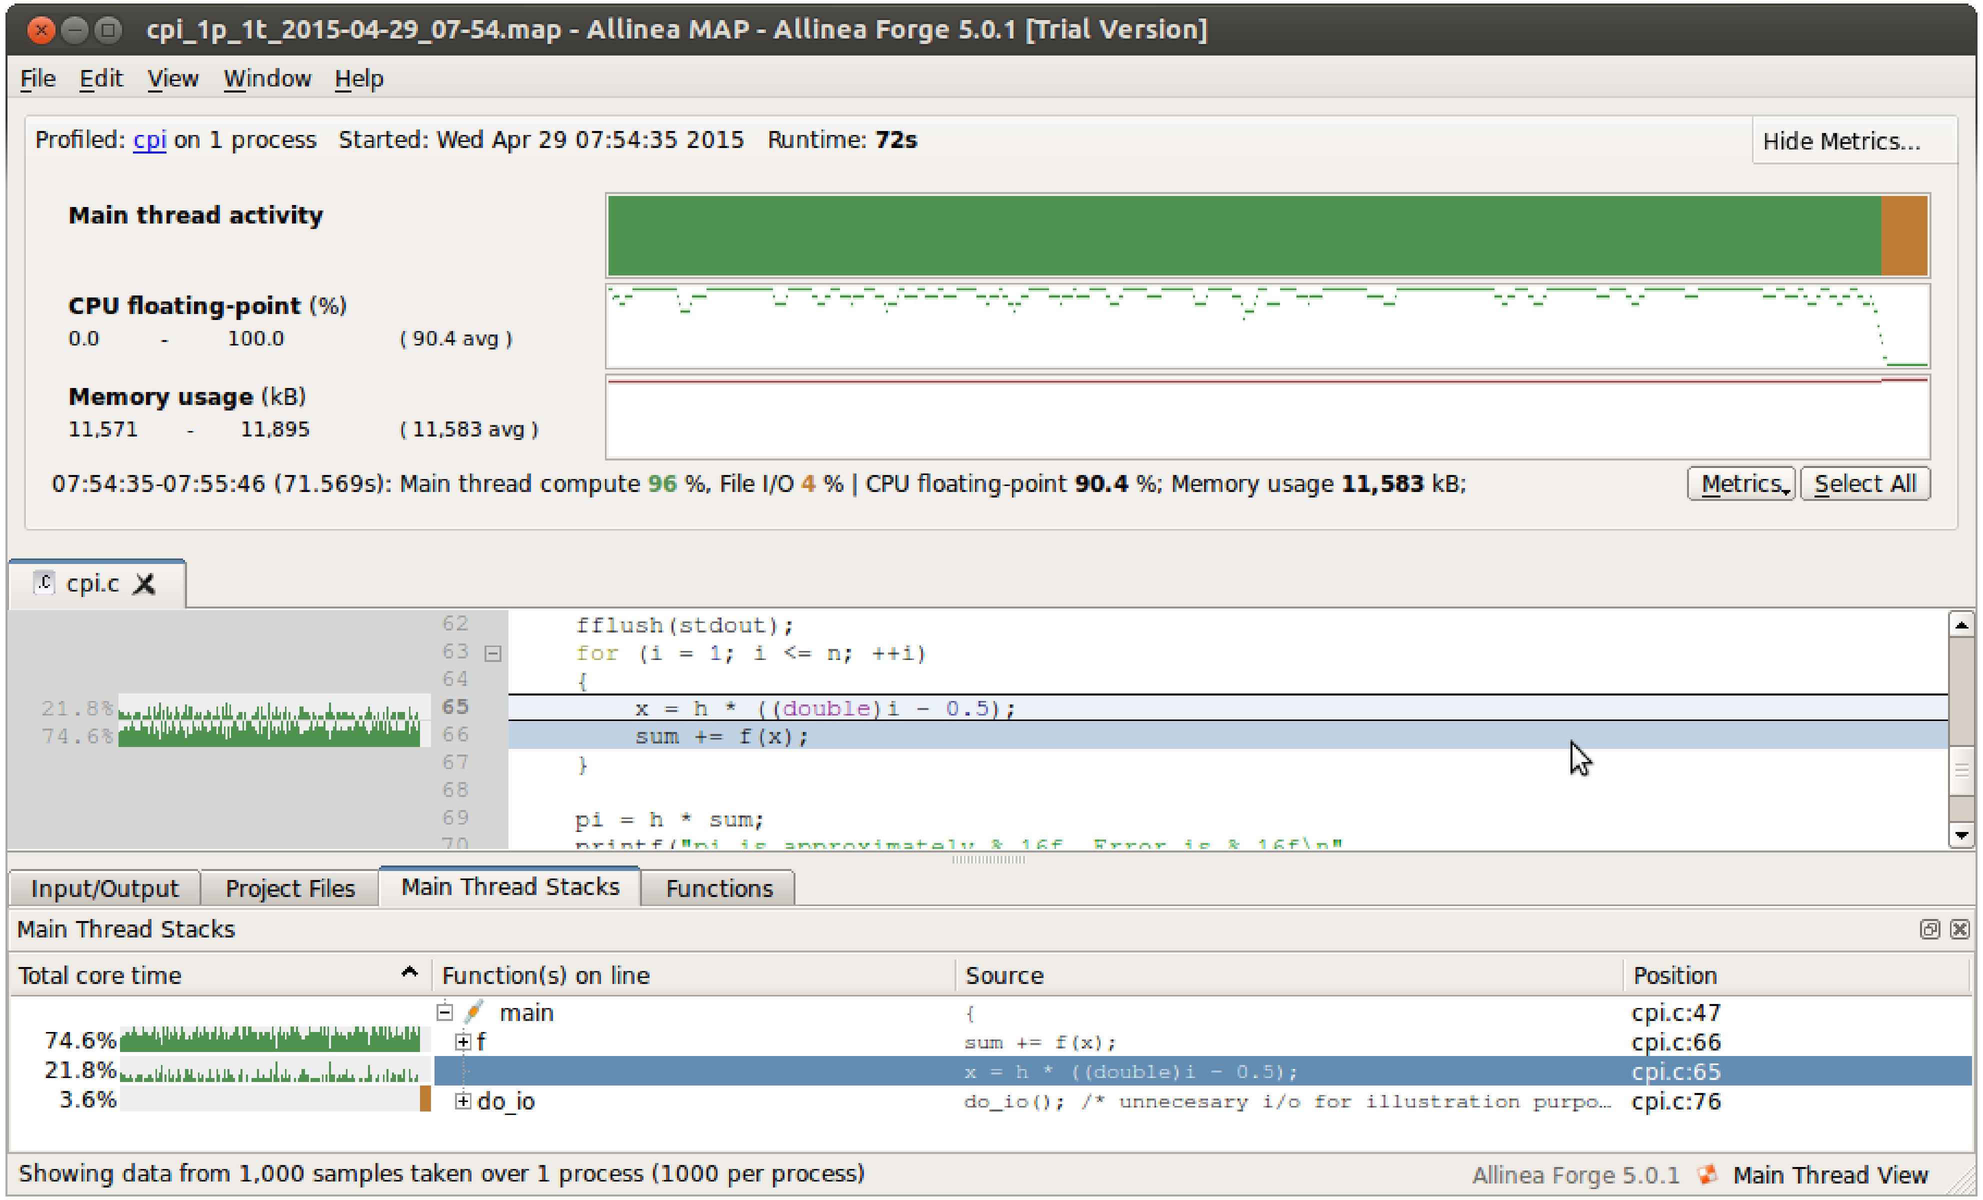
cpi_1p_1t_2015-04-29_07-54.map - Allinea MAP - Allinea Forge 5.0.1 [Trial Version]
File Edit View Window Help
Hide Metrics...
Metrics
Select All
cpi.c
Input/Output
Project Files
Main Thread Stacks
Functions

Profiled: cpi on 1 process Started: Wed Apr 29 07:54:35 2015 Runtime: **72s**

Main thread activity

CPU floating-point (%)
0.0 - 100.0 (90.4 avg)

Memory usage (kB)
11,571 - 11,895 (11,583 avg)

07:54:35-07:55:46 (71.569s): Main thread compute **96** %, File I/O **4** % | CPU floating-point **90.4** %; Memory usage **11,583** kB;

```
62          fflush(stdout);
63          for (i = 1; i <= n; ++i)
64          {
65  21.8%       x = h * ((double)i - 0.5);
66  74.6%       sum += f(x);
67          }
68
69          pi = h * sum;
70          printf("pi is approximately %.16f, Error is %.16f\n"
```

Main Thread Stacks

Total core time	Function(s) on line	Source	Position
	main	{	cpi.c:47
74.6%	f	sum += f(x);	cpi.c:66
21.8%		x = h * ((double)i - 0.5);	cpi.c:65
3.6%	do_io	do_io(); /* unnecesary i/o for illustration purpo…	cpi.c:76

Showing data from 1,000 samples taken over 1 process (1000 per process) Allinea Forge 5.0.1 Main Thread View

allinea FORGE

Zooming in on the IO portion of our run

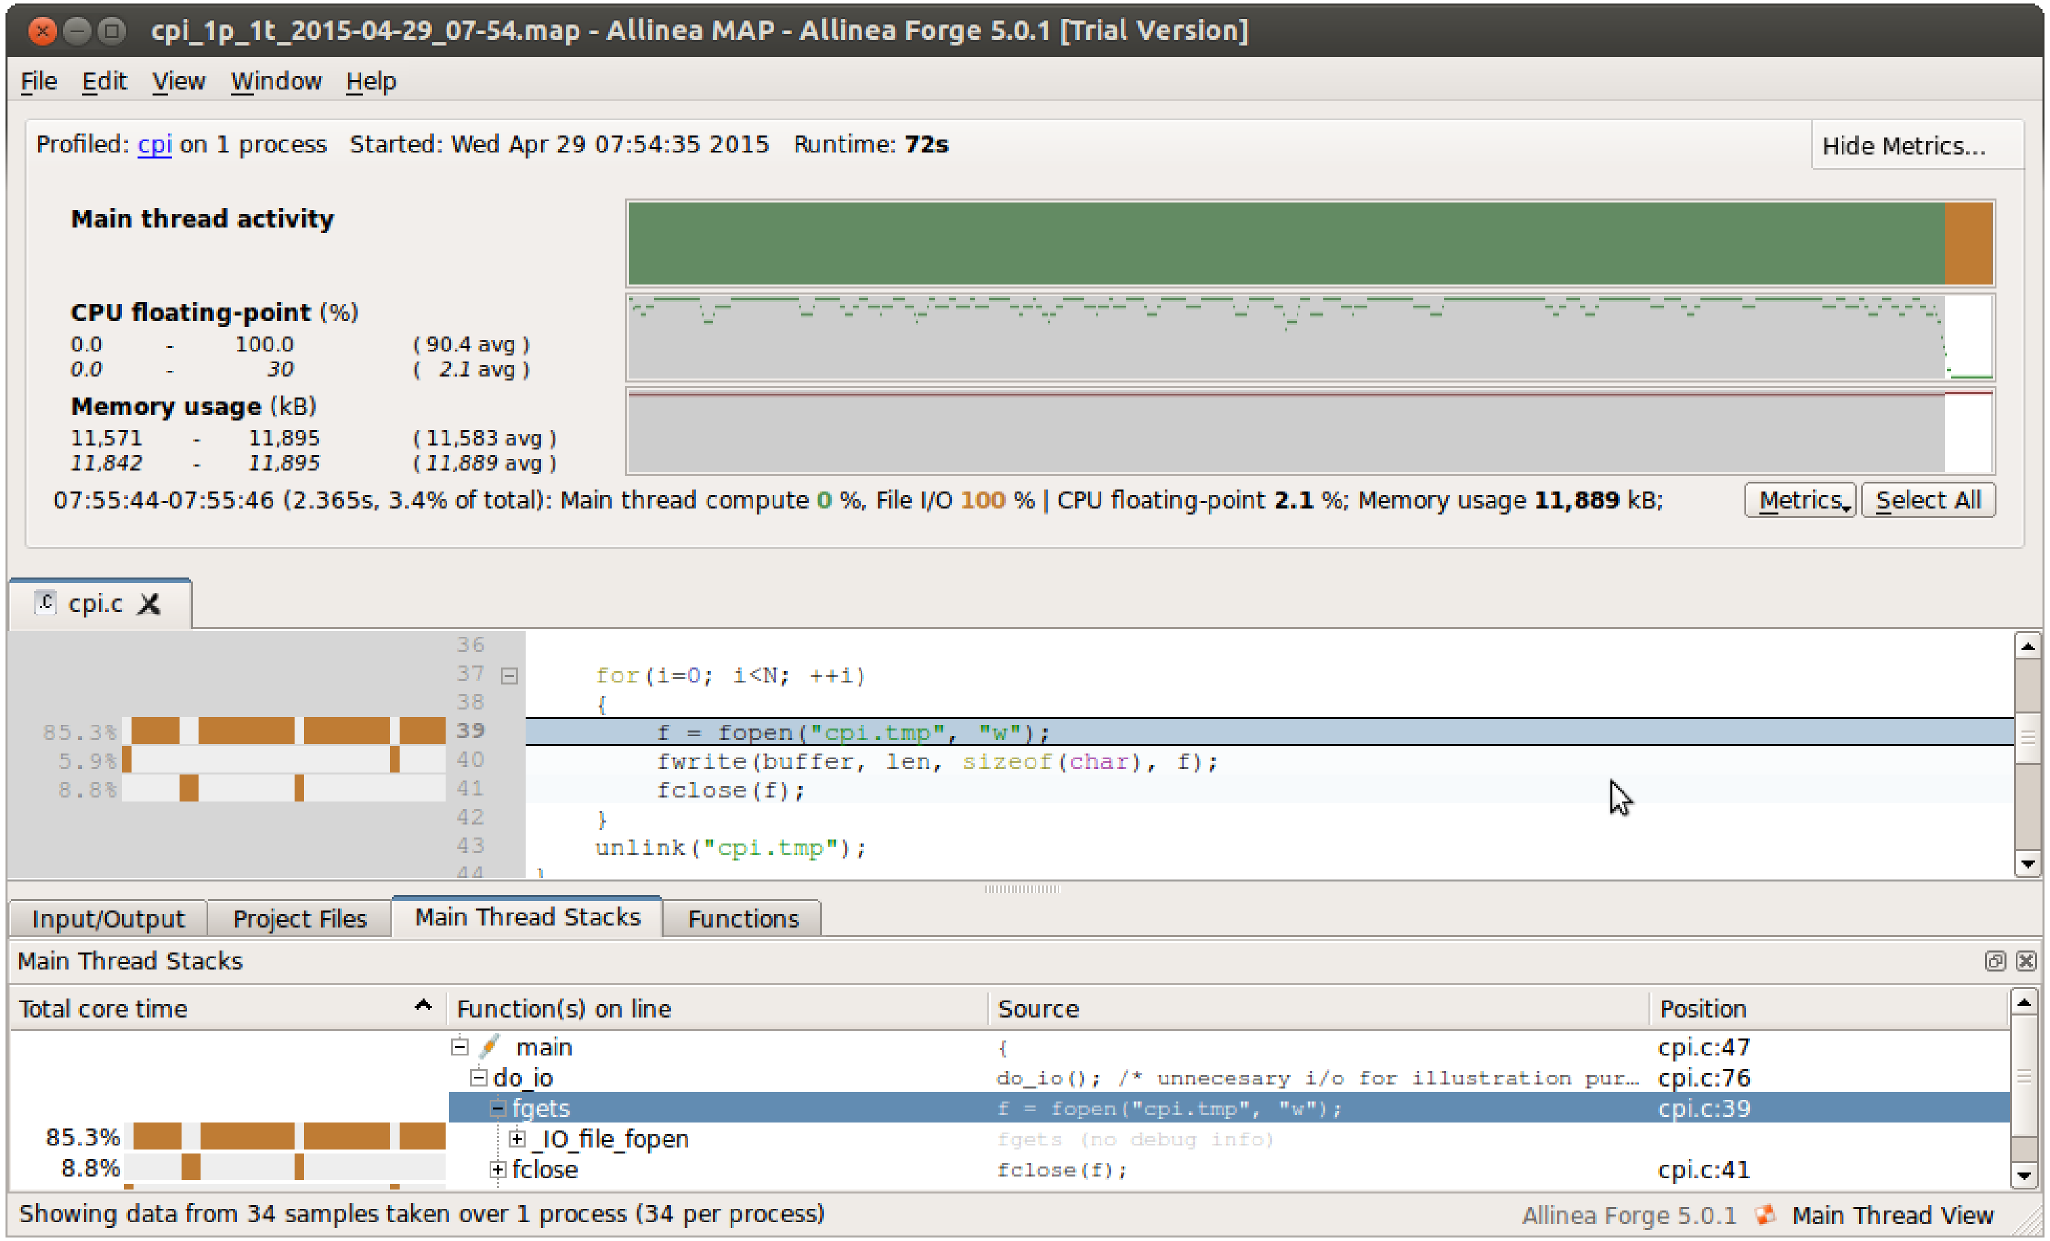

cpi_1p_1t_2015-04-29_07-54.map - Allinea MAP - Allinea Forge 5.0.1 [Trial Version]

File Edit View Window Help

Profiled: cpi on 1 process Started: Wed Apr 29 07:54:35 2015 Runtime: **72s**

Hide Metrics...

Main thread activity

CPU floating-point (%)
0.0 - 100.0 (90.4 avg)
0.0 - 30 (2.1 avg)

Memory usage (kB)
11,571 - 11,895 (11,583 avg)
11,842 - 11,895 (11,889 avg)

07:55:44-07:55:46 (2.365s, 3.4% of total): Main thread compute **0** %, File I/O **100** % | CPU floating-point **2.1** %; Memory usage **11,889** kB;

Metrics Select All

cpi.c

```
        for(i=0; i<N; ++i)
        {
85.3%       f = fopen("cpi.tmp", "w");
 5.9%       fwrite(buffer, len, sizeof(char), f);
 8.8%       fclose(f);
        }
        unlink("cpi.tmp");
```

Input/Output | Project Files | Main Thread Stacks | Functions

Main Thread Stacks

Total core time	Function(s) on line	Source	Position
	main	{	cpi.c:47
	do_io	do_io(); /* unnecesary i/o for illustration pur…	cpi.c:76
	fgets	f = fopen("cpi.tmp", "w");	cpi.c:39
85.3%	_IO_file_fopen	fgets (no debug info)	
8.8%	fclose	fclose(f);	cpi.c:41

Showing data from 34 samples taken over 1 process (34 per process)

Allinea Forge 5.0.1 Main Thread View

allinea FORGE

Always be thinking about vectorization

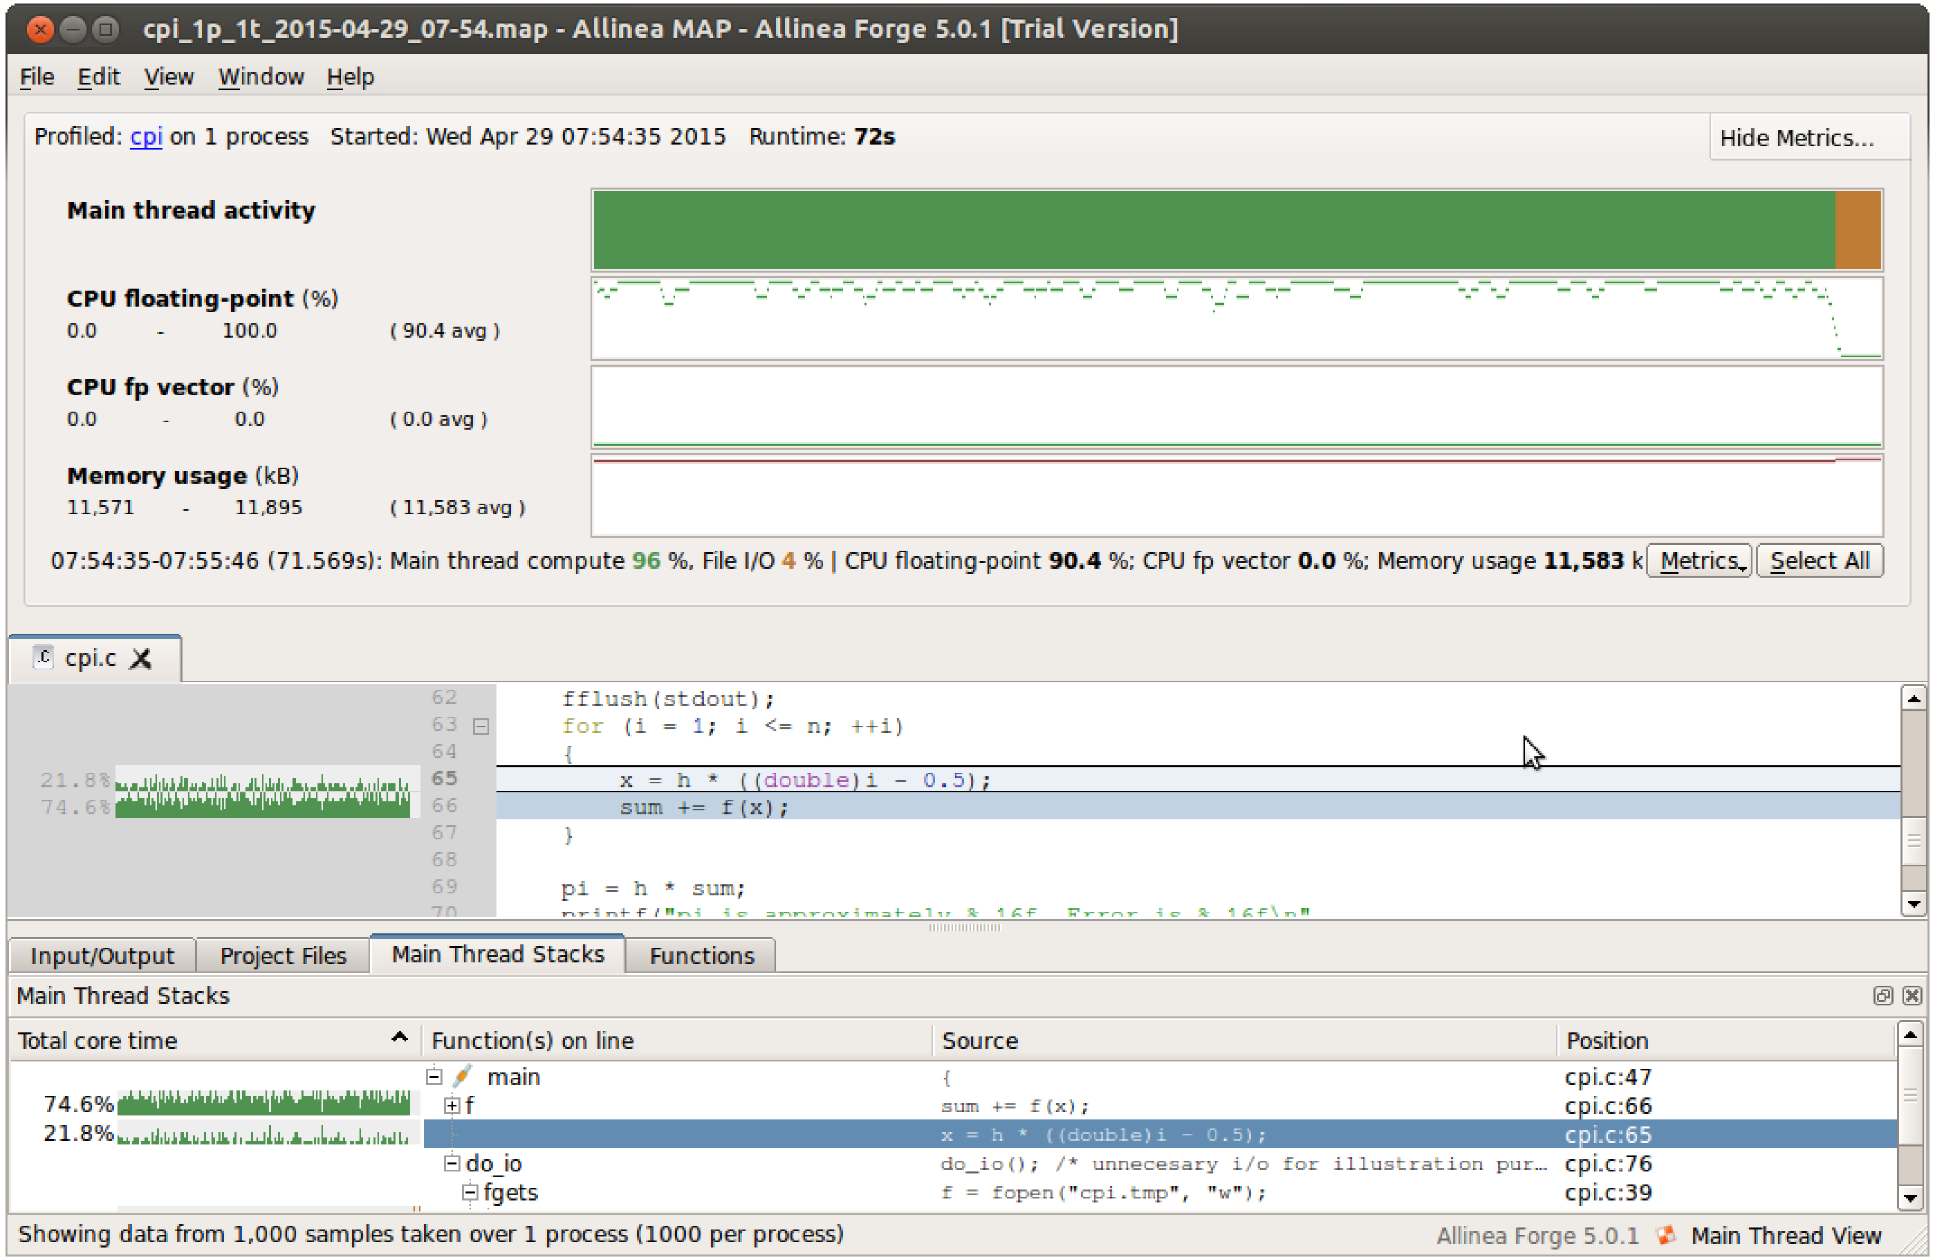

Profiled: cpi on 1 process Started: Wed Apr 29 07:54:35 2015 Runtime: **72s**

Main thread activity				
CPU floating-point (%)	0.0	-	100.0	(90.4 avg)
CPU fp vector (%)	0.0	-	0.0	(0.0 avg)
Memory usage (kB)	11,571	-	11,895	(11,583 avg)

07:54:35-07:55:46 (71.569s): Main thread compute **96** %, File I/O **4** % | CPU floating-point **90.4** %; CPU fp vector **0.0** %; Memory usage **11,583** k

```
                        fflush(stdout);
                        for (i = 1; i <= n; ++i)
                        {
21.8%                       x = h * ((double)i - 0.5);
74.6%                       sum += f(x);
                        }

                        pi = h * sum;
```

Main Thread Stacks

Total core time	Function(s) on line	Source	Position
	main	{	cpi.c:47
74.6%	f	sum += f(x);	cpi.c:66
21.8%		x = h * ((double)i - 0.5);	cpi.c:65
	do_io	do_io(); /* unnecesary i/o for illustration pur…	cpi.c:76
	fgets	f = fopen("cpi.tmp", "w");	cpi.c:39

Showing data from 1,000 samples taken over 1 process (1000 per process)

allinea FORGE

While still connected to the server we can switch to the debugger

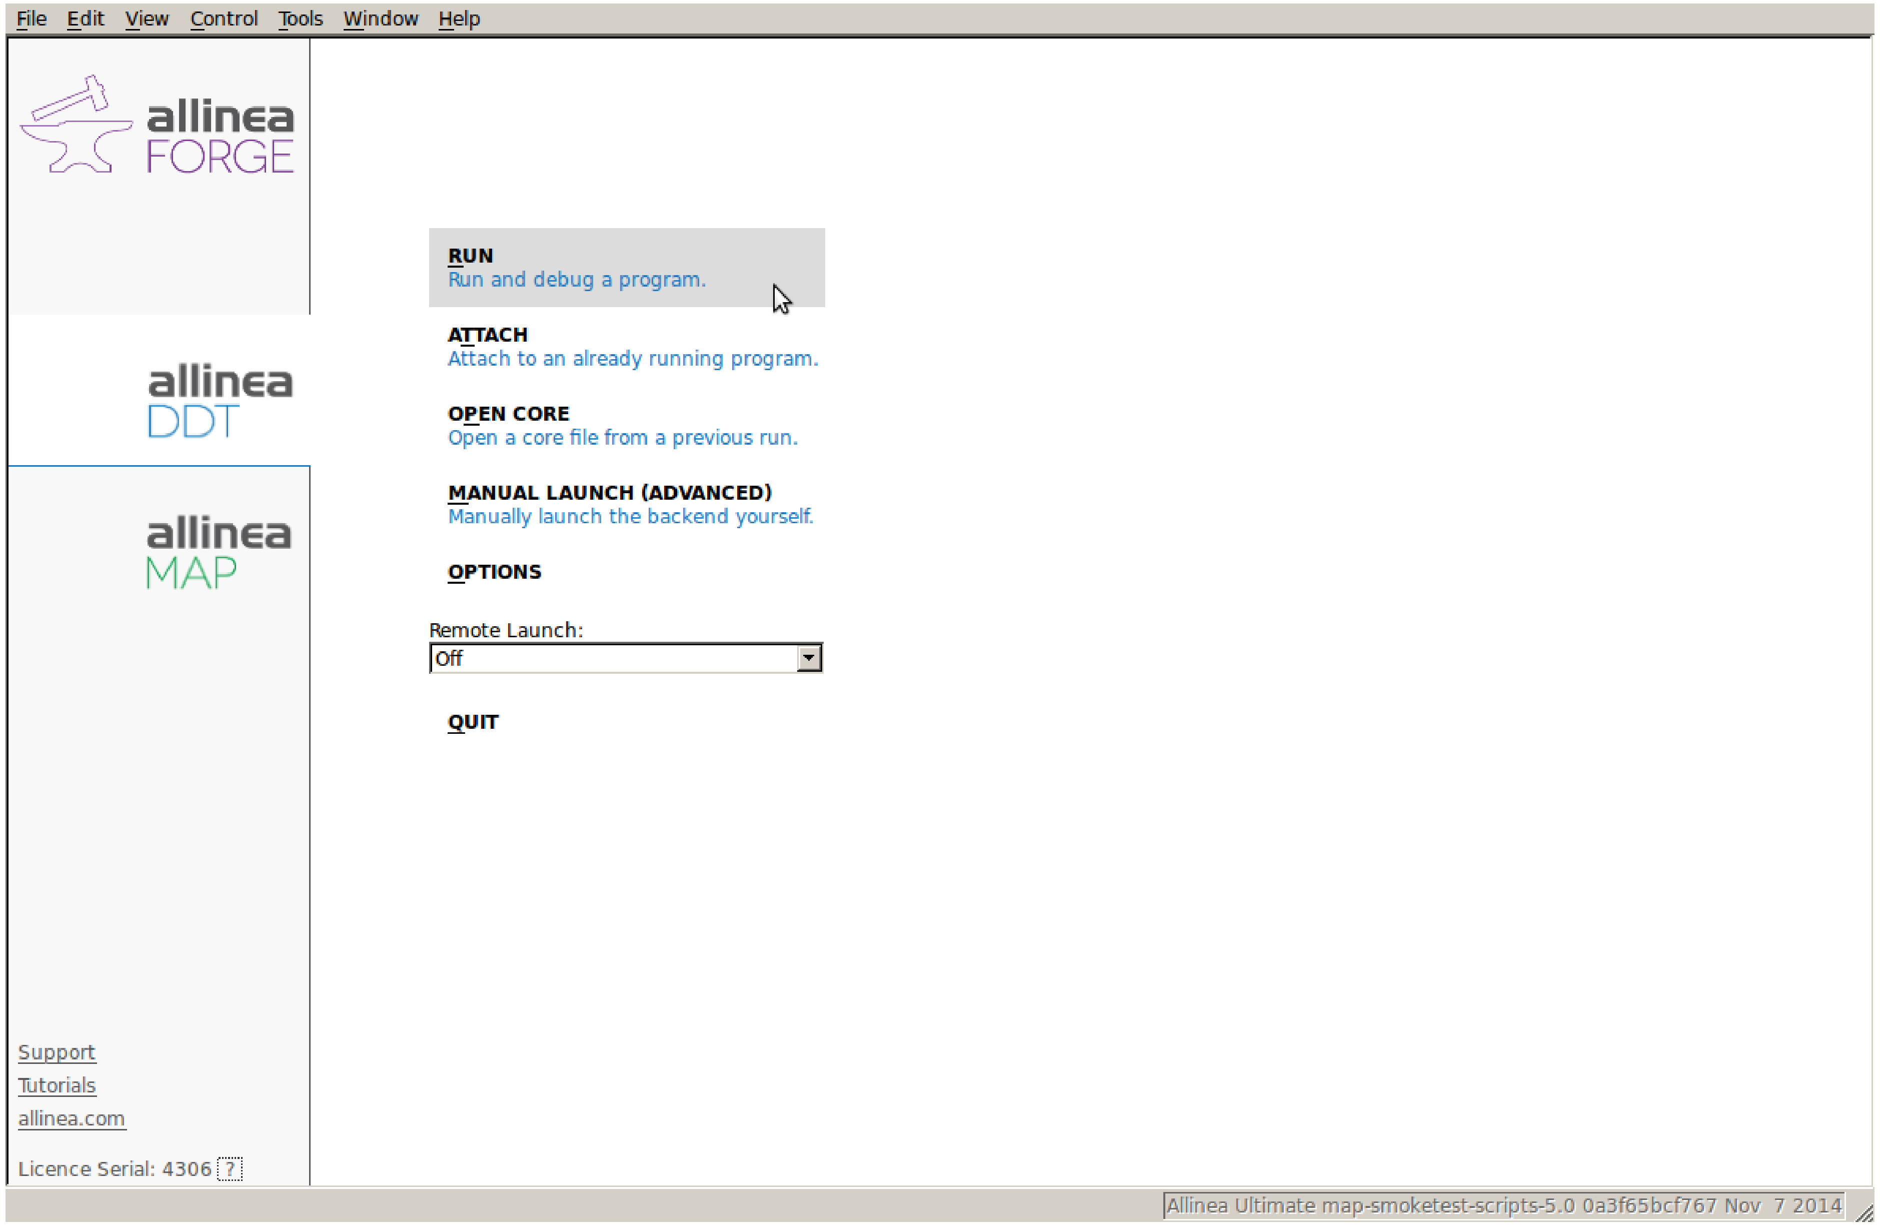

File Edit View Control Tools Window Help

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

QUIT

Support
Tutorials
allinea.com

Licence Serial: 4306

Allinea Ultimate map-smoketest-scripts-5.0 0a3f65bcf767 Nov 7 2014

Adding some OpenMP pragmas to multithread

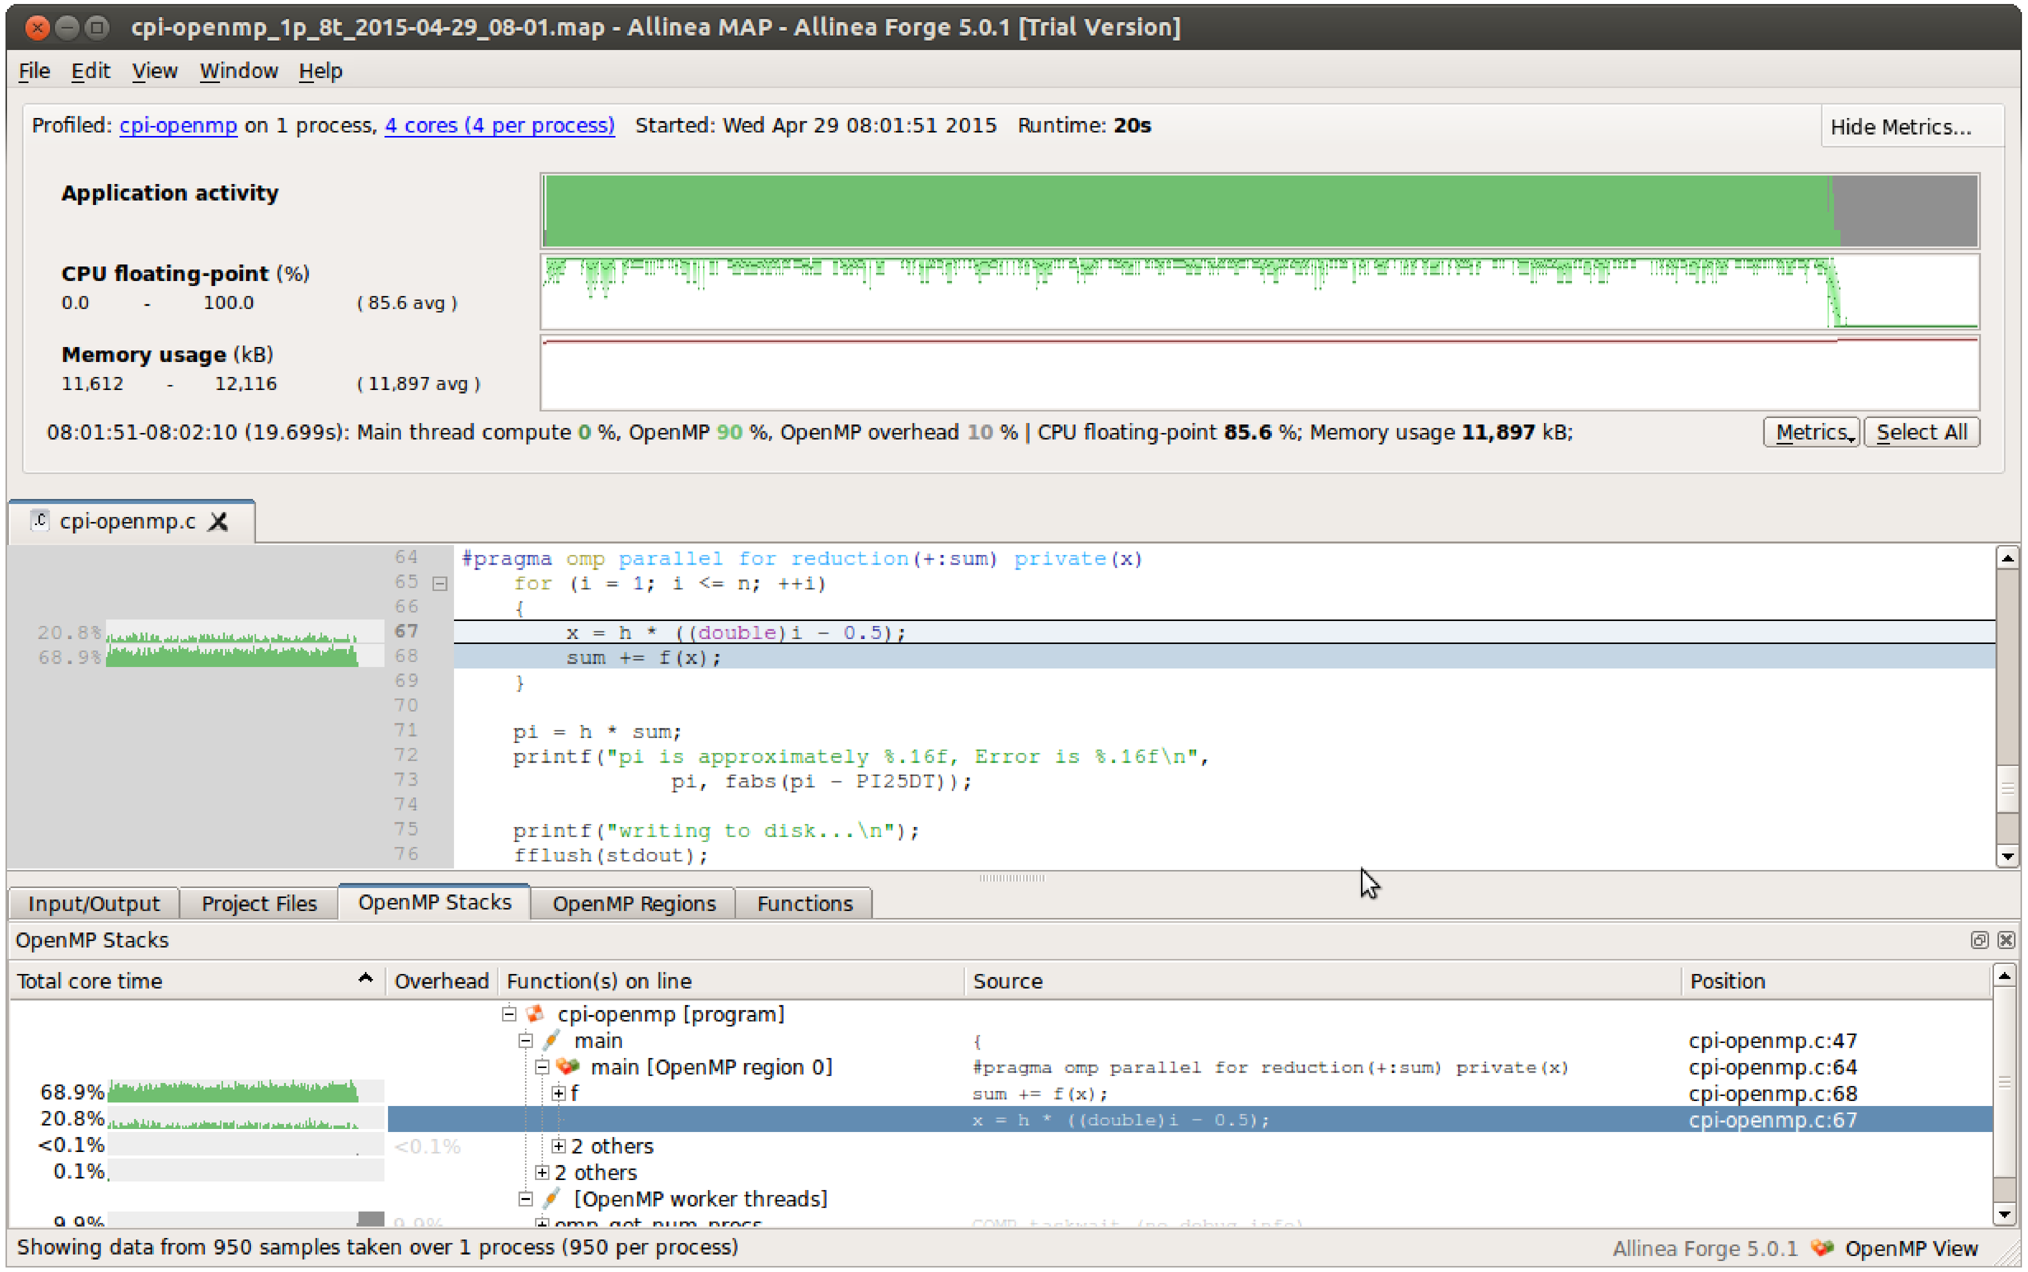
cpi-openmp_1p_8t_2015-04-29_08-01.map - Allinea MAP - Allinea Forge 5.0.1 [Trial Version]
File Edit View Window Help
Hide Metrics...
Metrics
Select All
cpi-openmp.c
Input/Output
Project Files
OpenMP Stacks
OpenMP Regions
Functions

Profiled: cpi-openmp on 1 process, 4 cores (4 per process) Started: Wed Apr 29 08:01:51 2015 Runtime: **20s**

Application activity

CPU floating-point (%)
0.0 - 100.0 (85.6 avg)

Memory usage (kB)
11,612 - 12,116 (11,897 avg)

08:01:51-08:02:10 (19.699s): Main thread compute 0 %, OpenMP 90 %, OpenMP overhead 10 % | CPU floating-point **85.6** %; Memory usage **11,897** kB;

```
#pragma omp parallel for reduction(+:sum) private(x)
    for (i = 1; i <= n; ++i)
    {
20.8%       x = h * ((double)i - 0.5);
68.9%       sum += f(x);
    }

    pi = h * sum;
    printf("pi is approximately %.16f, Error is %.16f\n",
                 pi, fabs(pi - PI25DT));

    printf("writing to disk...\n");
    fflush(stdout);
```

OpenMP Stacks

Total core time	Overhead	Function(s) on line	Source	Position
		cpi-openmp [program]		
		main	{	cpi-openmp.c:47
		main [OpenMP region 0]	#pragma omp parallel for reduction(+:sum) private(x)	cpi-openmp.c:64
68.9%		f	sum += f(x);	cpi-openmp.c:68
20.8%			x = h * ((double)i - 0.5);	cpi-openmp.c:67
<0.1%	<0.1%	2 others		
0.1%		2 others		
		[OpenMP worker threads]		
9.9%	9.9%	omp_get_num_procs	GOMP_taskwait (no debug info)	

Showing data from 950 samples taken over 1 process (950 per process)

Allinea Forge 5.0.1 OpenMP View

allinea FORGE

Adding MPI support and submitting to and HPC batch system

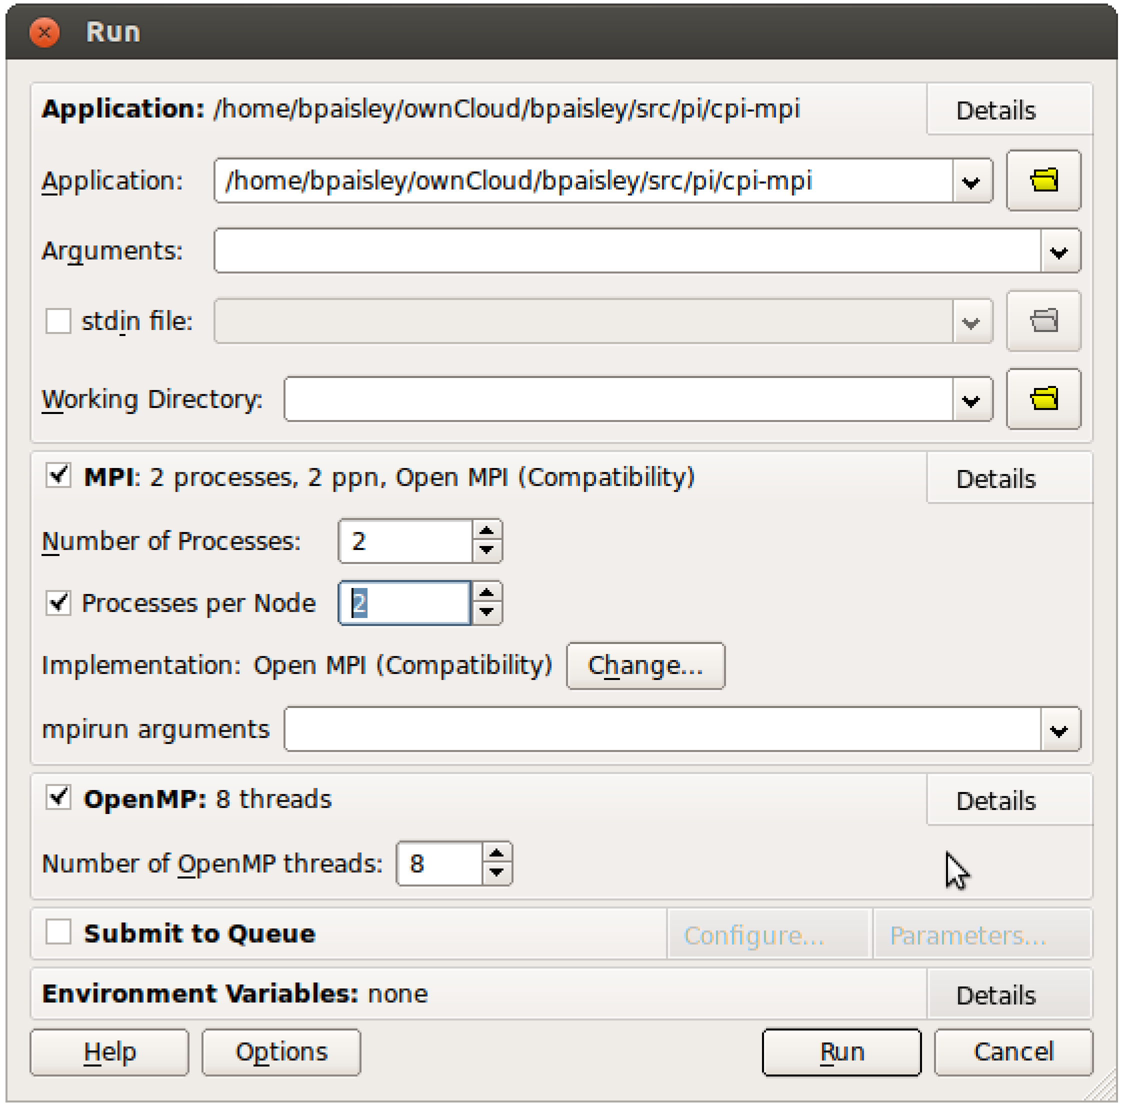

Run

Application: /home/bpaisley/ownCloud/bpaisley/src/pi/cpi-mpi — Details

Application: /home/bpaisley/ownCloud/bpaisley/src/pi/cpi-mpi

Arguments:

☐ stdin file:

Working Directory:

☑ **MPI:** 2 processes, 2 ppn, Open MPI (Compatibility) — Details

Number of Processes: 2

☑ Processes per Node 2

Implementation: Open MPI (Compatibility) Change...

mpirun arguments

☑ **OpenMP:** 8 threads — Details

Number of OpenMP threads: 8

☐ **Submit to Queue** — Configure... Parameters...

Environment Variables: none — Details

Help Options Run Cancel

Analyzing our OpenMP + MPI results

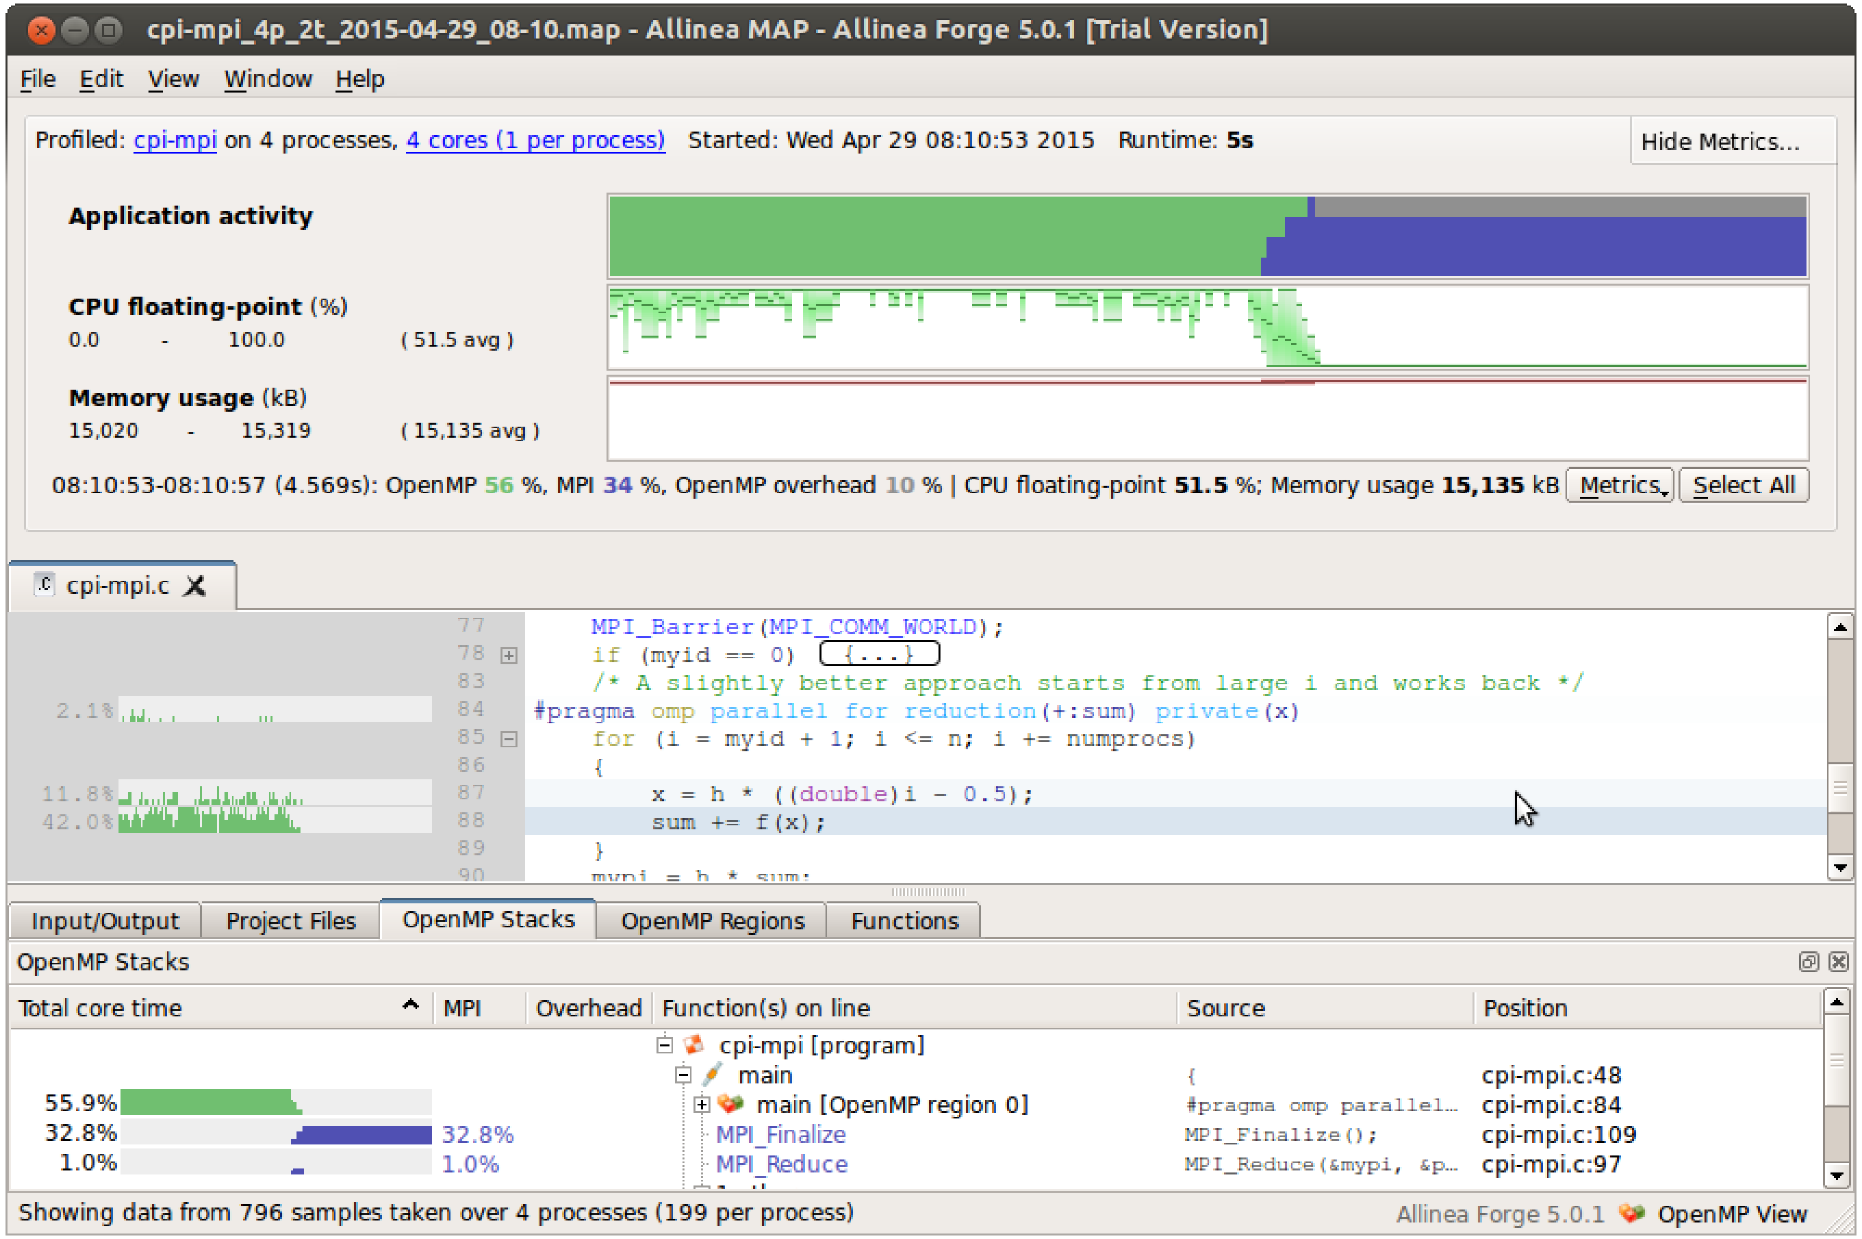

Profiled: cpi-mpi on 4 processes, 4 cores (1 per process) Started: Wed Apr 29 08:10:53 2015 Runtime: **5s**

Application activity

CPU floating-point (%)
0.0 - 100.0 (51.5 avg)

Memory usage (kB)
15,020 - 15,319 (15,135 avg)

08:10:53-08:10:57 (4.569s): OpenMP 56 %, MPI 34 %, OpenMP overhead 10 % | CPU floating-point **51.5** %; Memory usage **15,135** kB

```
        MPI_Barrier(MPI_COMM_WORLD);
        if (myid == 0) {...}
        /* A slightly better approach starts from large i and works back */
2.1%  #pragma omp parallel for reduction(+:sum) private(x)
        for (i = myid + 1; i <= n; i += numprocs)
        {
11.8%       x = h * ((double)i - 0.5);
42.0%       sum += f(x);
        }
        mypi = h * sum;
```

OpenMP Stacks

Total core time	MPI	Overhead	Function(s) on line	Source	Position
			cpi-mpi [program]		
			main	{	cpi-mpi.c:48
55.9%			main [OpenMP region 0]	#pragma omp parallel…	cpi-mpi.c:84
32.8%	32.8%		MPI_Finalize	MPI_Finalize();	cpi-mpi.c:109
1.0%	1.0%		MPI_Reduce	MPI_Reduce(&mypi, &p…	cpi-mpi.c:97

Showing data from 796 samples taken over 4 processes (199 per process) Allinea Forge 5.0.1 OpenMP View

allinea FORGE

Why all this synchronization?

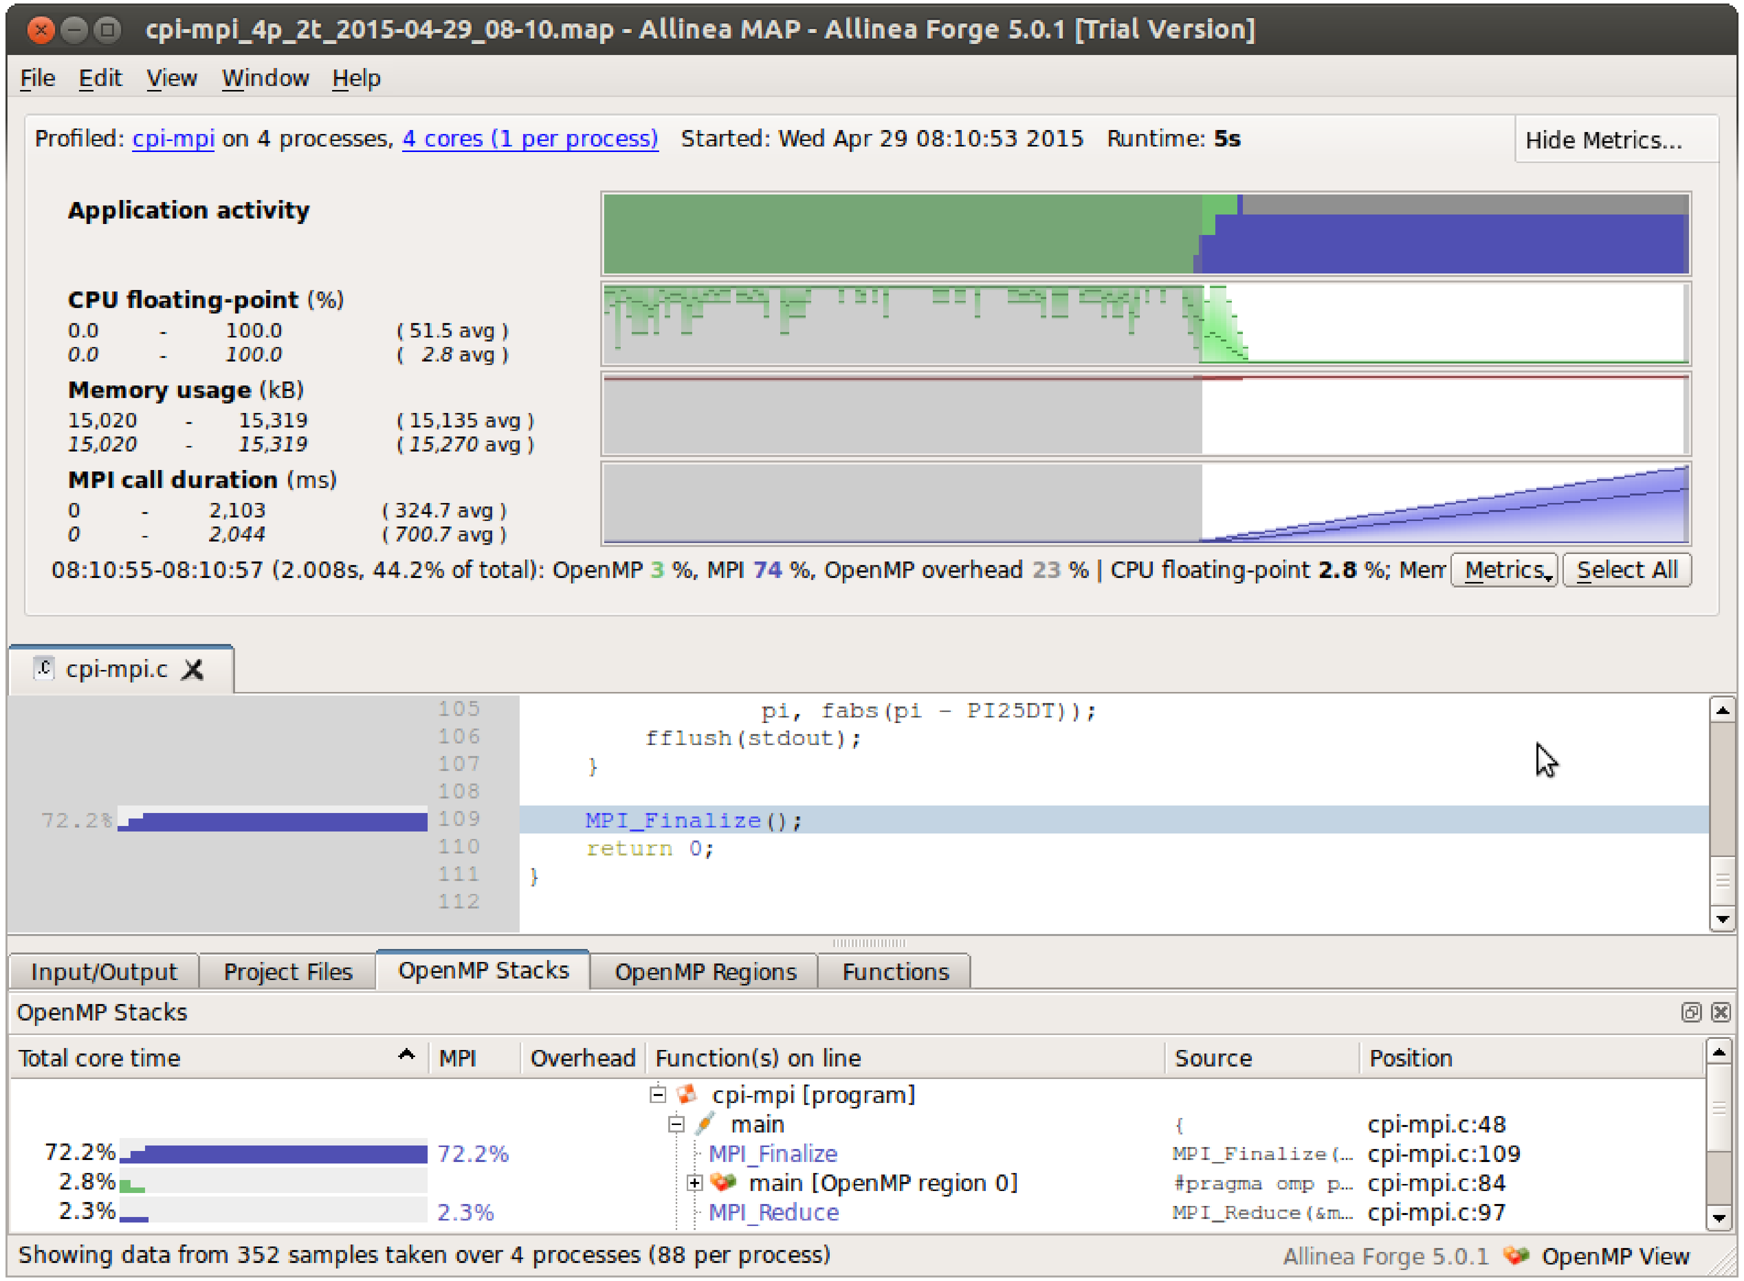

cpi-mpi_4p_2t_2015-04-29_08-10.map - Allinea MAP - Allinea Forge 5.0.1 [Trial Version]

Profiled: cpi-mpi on 4 processes, 4 cores (1 per process) Started: Wed Apr 29 08:10:53 2015 Runtime: **5s**

Application activity

CPU floating-point (%)			
0.0	-	100.0	(51.5 avg)
0.0	-	*100.0*	(*2.8* avg)
Memory usage (kB)			
15,020	-	15,319	(15,135 avg)
15,020	-	*15,319*	(*15,270* avg)
MPI call duration (ms)			
0	-	2,103	(324.7 avg)
0	-	*2,044*	(*700.7* avg)

08:10:55-08:10:57 (2.008s, 44.2% of total): OpenMP 3 %, MPI 74 %, OpenMP overhead 23 % | CPU floating-point **2.8** %; Mem

cpi-mpi.c

```
72.2%  105                pi, fabs(pi - PI25DT));
       106          fflush(stdout);
       107      }
       108
       109      MPI_Finalize();
       110      return 0;
       111  }
       112
```

OpenMP Stacks

Total core time	MPI	Overhead	Function(s) on line	Source	Position
			cpi-mpi [program]		
			main	{	cpi-mpi.c:48
72.2%	72.2%		MPI_Finalize	MPI_Finalize(...	cpi-mpi.c:109
2.8%			main [OpenMP region 0]	#pragma omp p...	cpi-mpi.c:84
2.3%	2.3%		MPI_Reduce	MPI_Reduce(&m...	cpi-mpi.c:97

Showing data from 352 samples taken over 4 processes (88 per process) Allinea Forge 5.0.1 OpenMP View

allinea FORGE

The workflow of pi from serial to HPC

allinea FORGE

allinea DDT

allinea MAP

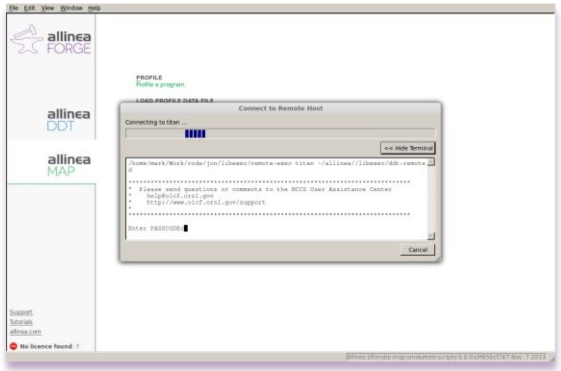

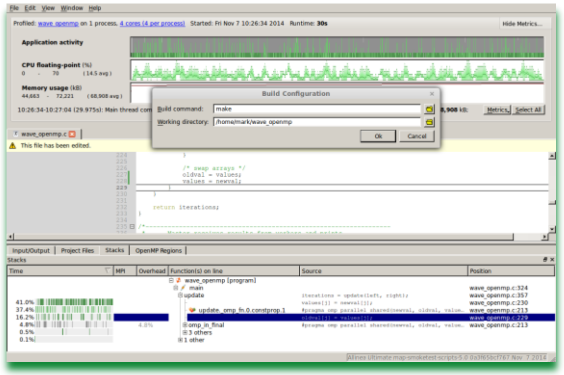

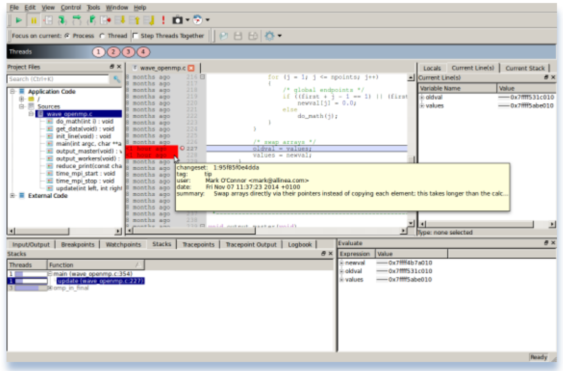

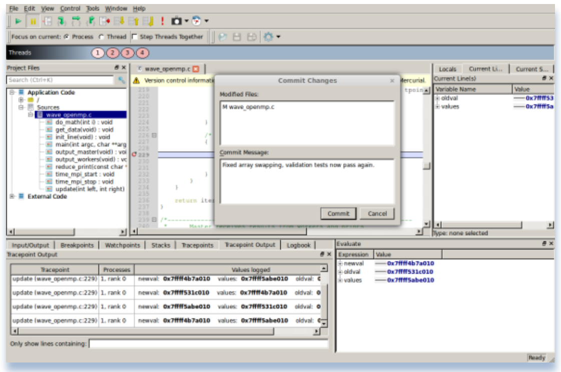

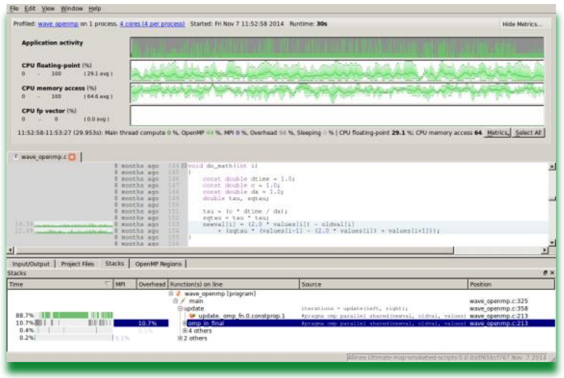

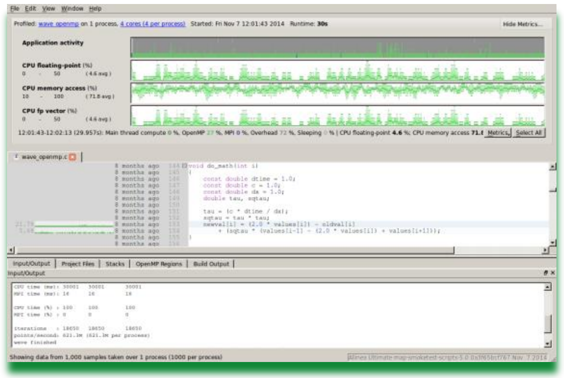

allinea

https://www.allinea.com/products/downloads/free-trial

allinea FORGE

allinea DDT

allinea MAP

allinea PERFORMANCE REPORTS