

Advanced Topics:


Additional Resources

Performance tools and exotic architectures



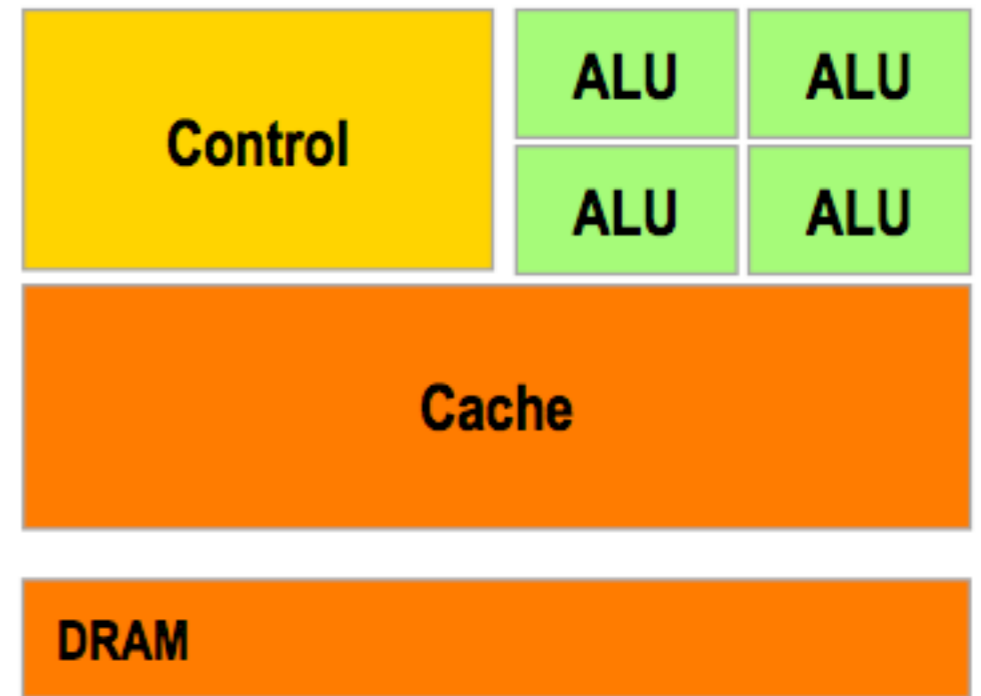
GPUs

- Don't look now, but your graphics card is possibly more powerful than your CPU.




GPU

- The triumphant return of the massively data-parallel type machine (Connection Machine, MasPar..)

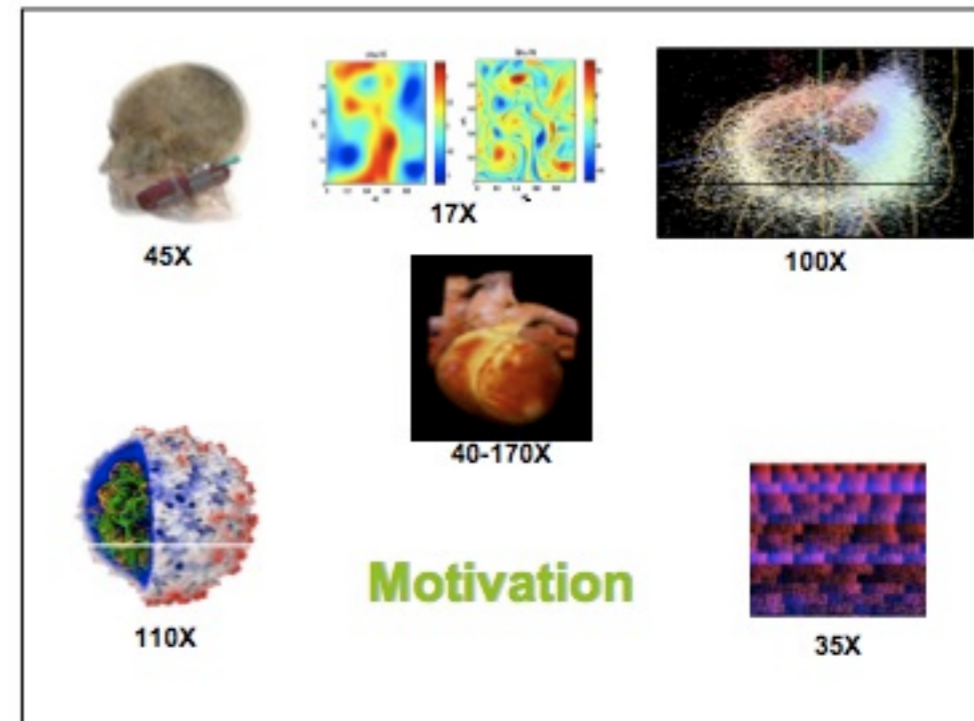


CPU



Programming GPUs

- Used to be pretty bad; put array in as 'textures', have each point in your grid be a vertex that maps the texture...
- Much better now: CUDA (NVidia), OpenCL (NVidia, Apple, AMD; coming soon). BLAS, FFT libraries...



NVidia SC2007
tutorial slides



GPU Computing



- Processors execute computing threads
- Thread Execution Manager issues threads
- 128 Thread Processors
- Parallel Data Cache accelerates processing



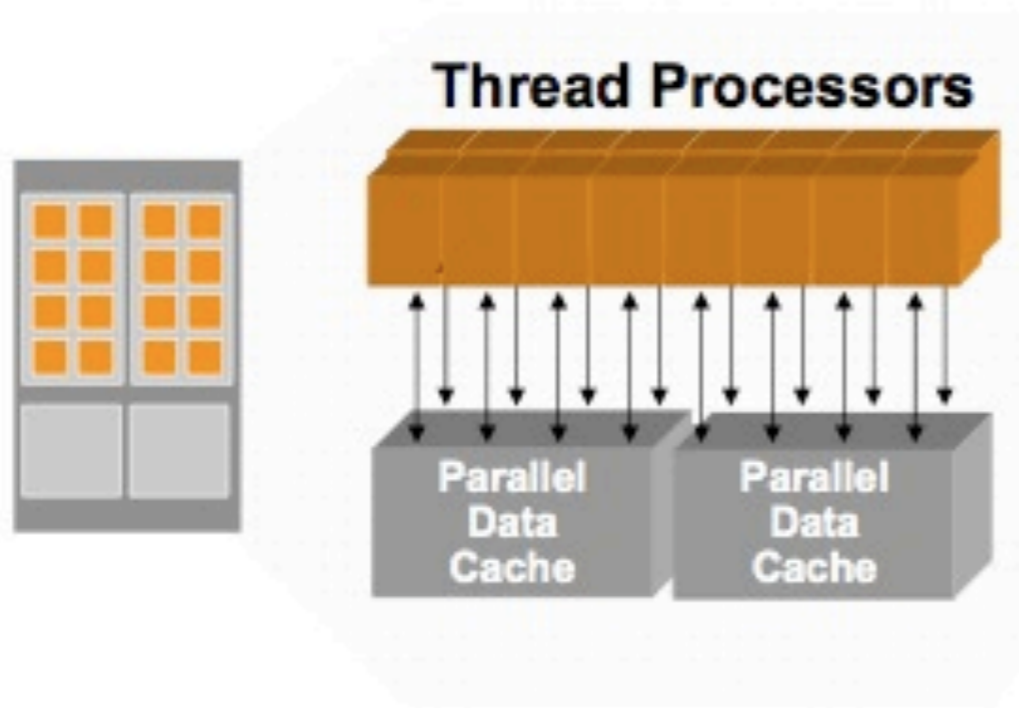
S05: High Performance Computing with CUDA



Thread Processor Group




- 128, 1.35 GHz processors
- 16KB Parallel Data Cache per group
- Scalar architecture
- IEEE 754 Precision



S05: High Performance Computing with CUDA



Execution Model



Multiple levels of parallelism

- **Thread block**
 - Up to 512 threads per block
 - Communicate through shared memory
 - Threads guaranteed to be resident
 - threadIdx, blockIdx
 - __syncthreads()
- **Grid of thread blocks**
 - f<<<nblocks, nthreads>>>(a,b,c)

C-Code Example to Add Arrays



CPU C program

```
void add_matrix_cpu
    (float *a, float *b, float *c, int N)
{
    int i, j, index;
    for (i=0;i<N;i++) {
        for (j=0;j<N;j++) {
            index =i+j*N;
            c[index]=a[index]+b[index];
        }
    }
}

void main()
{
    .....
    add_matrix(a,b,c,N);
}
```

CUDA C program

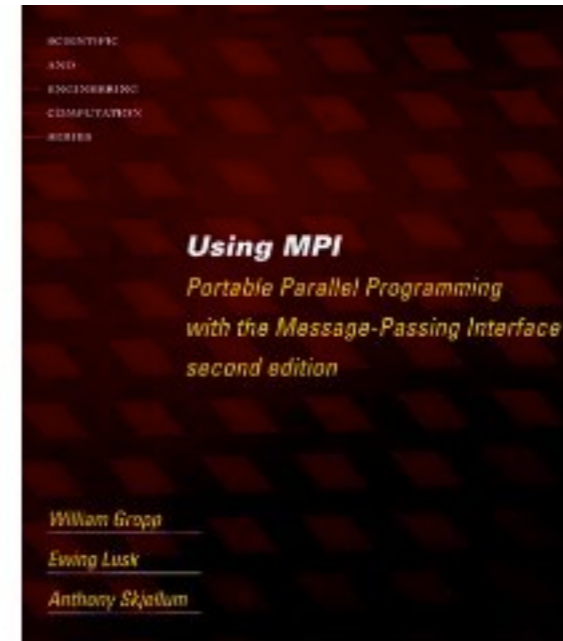
```
__global__ void add_matrix_gpu
    (float *a, float *b, float *c, int N)
{
    int i=blockIdx.x*blockDim.x+threadIdx.x;
    int j=blockIdx.y*blockDim.y+threadIdx.y;
    int index =i+j*N;
    if( i <N && j <N) c[index]=a[index]+b[index];
}

void main()
{
    dim3 dimBlock (blocksize,blocksize);
    dim3 dimGrid (N/dimBlock.x,N/dimBlock.y);
    add_matrix_gpu<<<dimGrid,dimBlock>>>(a,b,c,N);
}
```



Books We Like

- Using MPI, 2e; Gropp, Lusk, Skjellum
- Using OpenMP; Kuck, Chapman, Jost, and van der Pas



Parallel Frameworks

- Parallel infrastructure for codes for scientific computing
- Meshing, AMR, IO, etc.



Chombo

- C++ environment for AMR solution of finite difference equations
- Many bits built in (elliptic solvers, ...)




[http://seesar.lbl.gov/
ANAG/chombo/](http://seesar.lbl.gov/ANAG/chombo/)



Samrai

- C++ AMR environment
- Used for many different applications



[https://
computation.llnl.gov/
casc/SAMRAI/](https://computation.llnl.gov/casc/SAMRAI/)



Cactus

- C-based environment for finite-difference equations
- Originally developed for Numerical GR, but somewhat more general purpose



NATIONAL SCIENCE FOUNDATION

<http://www.cactuscode.org>

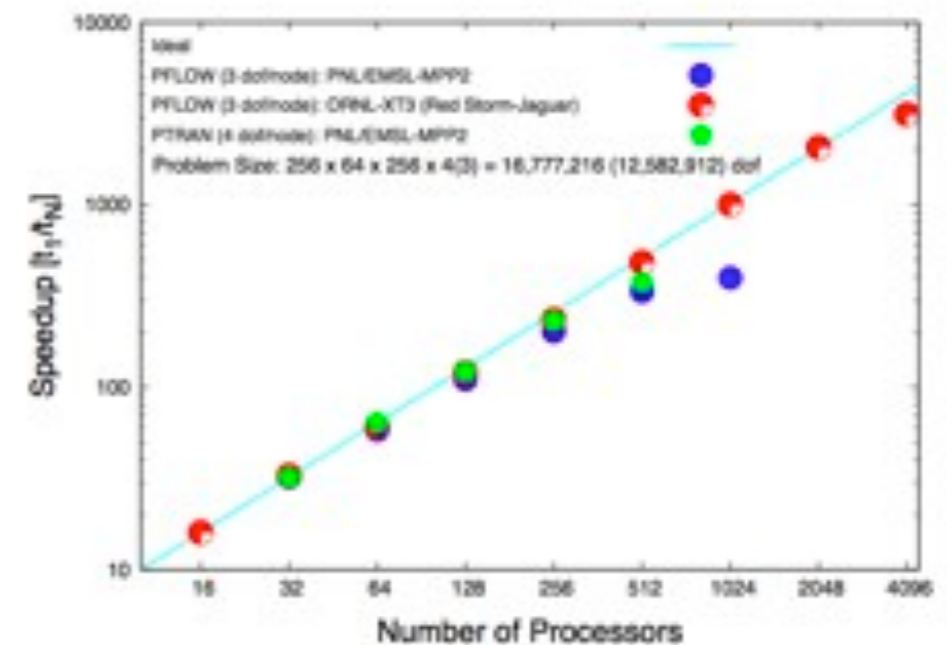

Parallel Libraries

- Libraries for doing common computations in parallel




PETSc

- Parallel vectors, matrices on top of MPI
- Parallel preconditioners, Krylov subspace methods, Newton-type nonlinear solvers, ODE solvers
- Easy to try lots of different solvers with a given code
- Developed/used by US DOE labs; won't go away anytime soon, actively developed
- 'Takes over' your application -- have to make your matrices PETSc matrices



<http://www-unix.mcs.anl.gov/petsc/petsc-2/>



Other Linear-Algebraish things

- Trilinos/Epetra; subset of blas/lapack <http://trilinos.sandia.gov/packages/>
- Aztec; iterative solver for sparse systems <http://www.cs.sandia.gov/CRF/aztec1.html>



FFTW

- Fastest Fourier Transform in the West
- MPI **or** OpenMP
- <http://www.fftw.org>

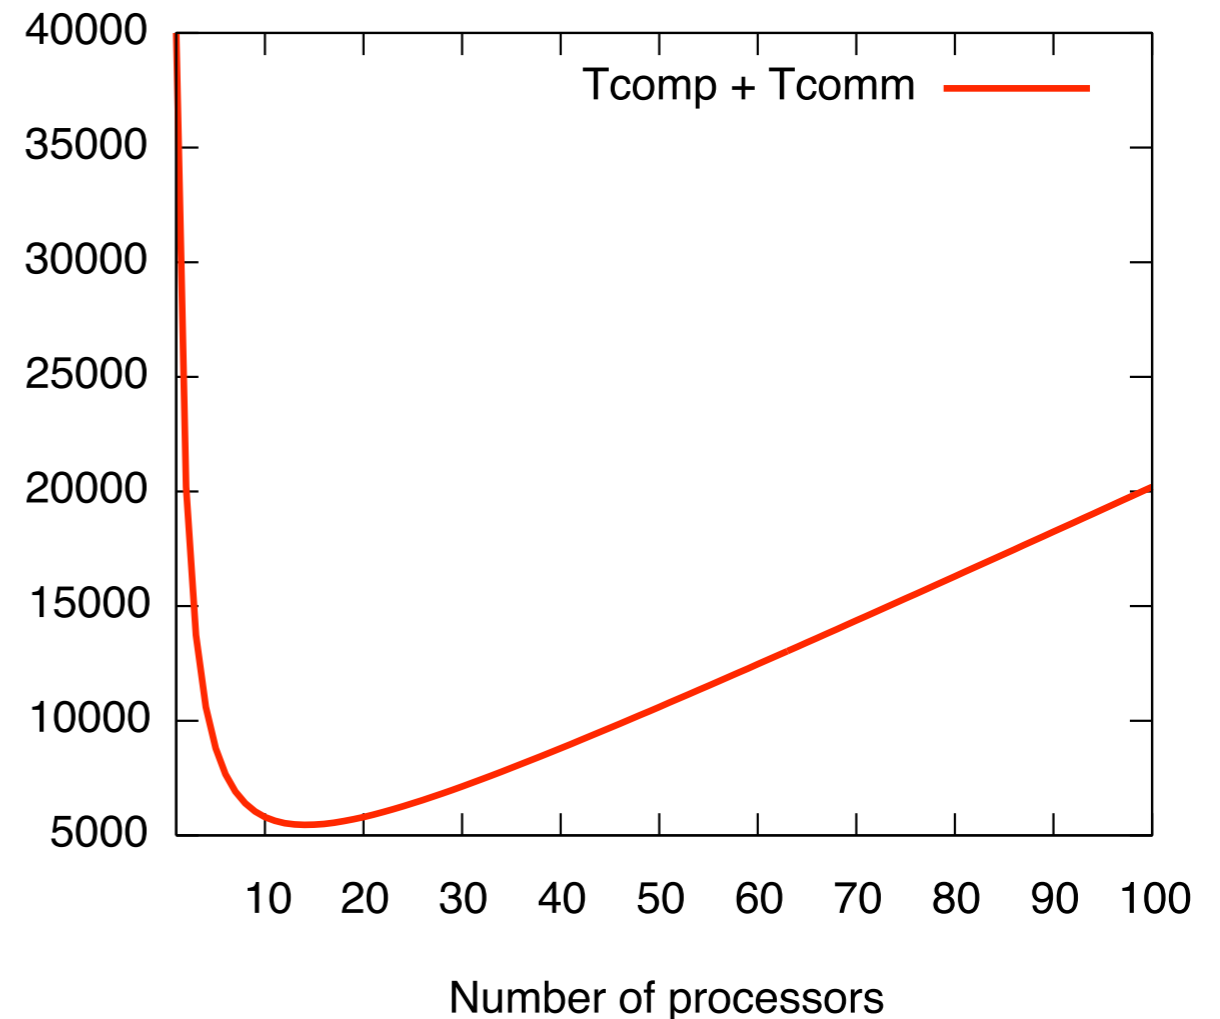
FFTW



Performance Tools

Time (some units)

- Parallel Programs are *complicated*
- When something gives worse than expected performance, how do you know what to fix?



Measuring Performance

- You can't improve what you can't measure
- You can't even tell there's a problem if you're not measuring.



http://commons.wikimedia.org/wiki/File:Plastic_tape_measure.jpg

First line of defence:

- `/bin/time` (or `/usr/bin/time`)
- by-section timing (eg, like `tick()`, `tock()`)

```
$ /usr/bin/time -p ls
Makefile          pca_mpi_utils.c
pca_mpi_utils.h   pca_utils.c
pca_utils.h
real              0.00
user              0.00
sys               0.00
```

```
$ ./nbody --nsteps=10 --nooutput
mydata has 1500 particles in a 3-
dimensional space.
Simulation is 1
Tock registers    6.5011e-02s
Tock registers    7.2196e-01s
```



gprof

- Talked about by J. Sievers
- Works almost anywhere
- Sampling (as vs. tracing)
- With long enough run, can give meaningful stats line-by-line
- installed on TCS, GPC

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <time.h>
4
5 double calc_pi(long n) {
6     long in = 0;
7     long out = 0;
8     long i;
9     double x,y;
10
11     for (i=0; i<n; i++) {
12         x = drand48();
13         y = drand48();
14         if (x*x+y*y < 1) {
15             in++;
16         } else {
17             out++;
18         }
19     }
20
21     return 4.*(double)in/(double)(in+out);
22 }
23
24 int main(int argc, char **argv) {
25     long n, defaultn=100000;
26     double pi;
27     time_t t;
28
29     /* seed random number generator */
30     srand48(time(&t));
31
32     /* get number of tries */
33     if (argc < 2 || (n=atoi(argv[1]))<1) {
34         n = defaultn;
35         printf("Using default n = %ld\n", n);
36     }
37
38     pi = $ gcc -g -pg -o pi pi.c
39     prin $ ./pi
40     retu [...]
41     $ gprof --line pi gmon.out
42 }
```

Each sample counts as 0.01 seconds.

% cumulative	self	self	self	total	
time	seconds	seconds	calls	Ts/call	Ts/call name
70.31	0.70	0.70			calc_pi (pi.c:14 @ 40078b)
14.27	0.84	0.14			calc_pi (pi.c:17 @ 4007bc)
5.10	0.89	0.05			calc_pi (pi.c:11 @ 4007c1)
4.08	0.93	0.04			calc_pi (pi.c:15 @ 4007b5)
3.06	0.96	0.03			calc_pi (pi.c:13 @ 400781)
2.55	0.98	0.03			calc_pi (pi.c:12 @ 400777)
1.53	1.00	0.02			calc_pi (pi.c:11 @ 40076d)
0.00	1.00	0.00	1	0.00	0.00 calc_pi (pi.c:5 @ 40074c)



OpenSpeedShop

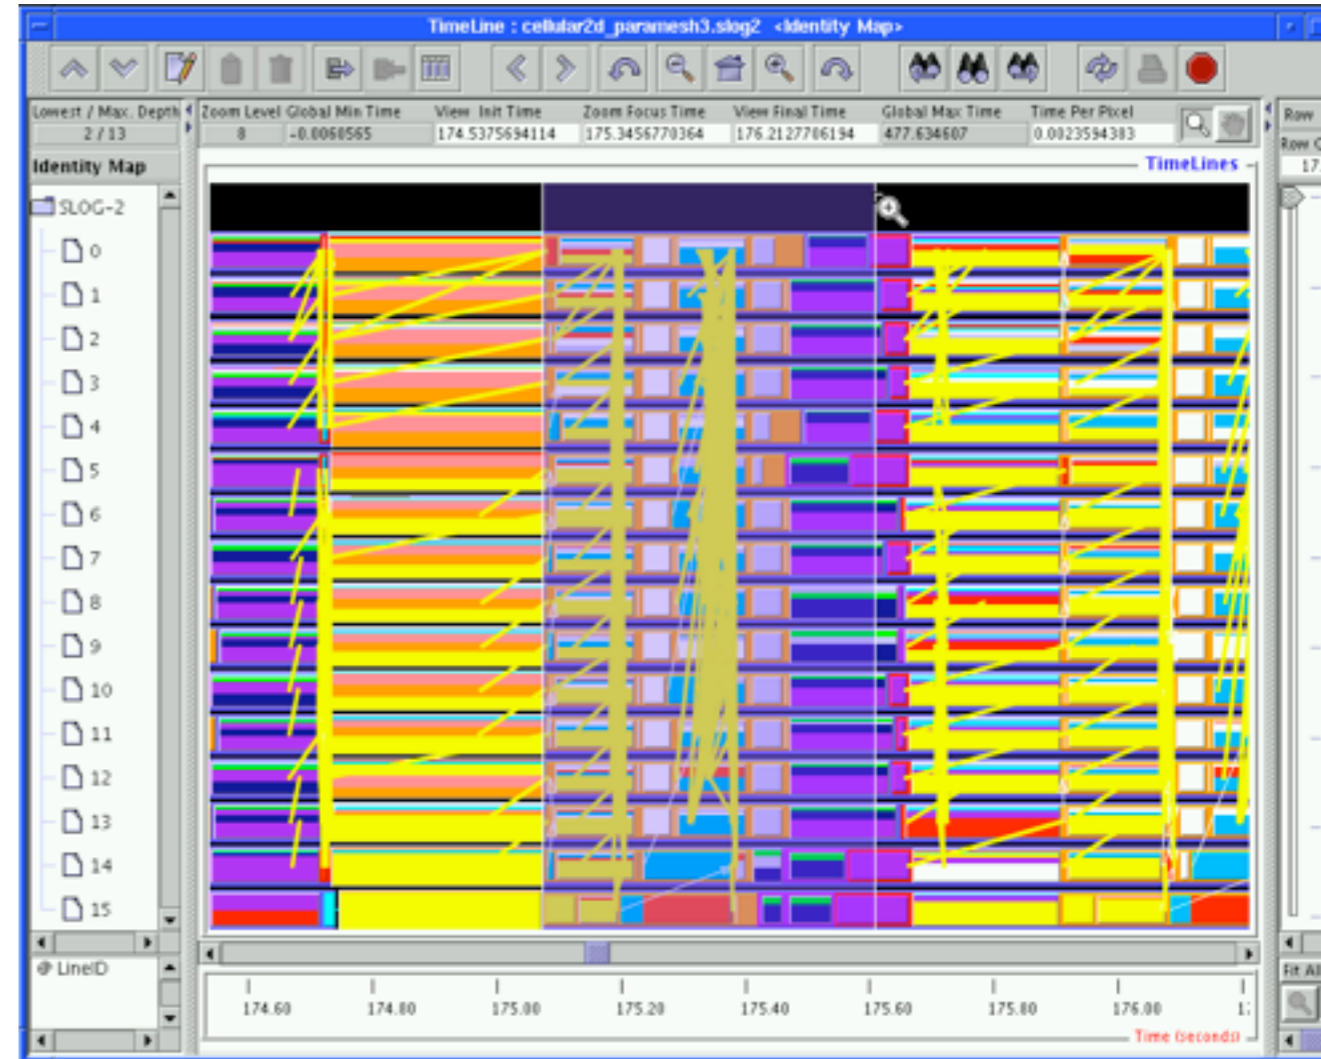
- <http://www.openspeedshop.org>
- Linux
- ‘experiments’
- Saves data, makes it very easy to compare results as you’re optimizing.
- Installed on GPC (module load openspeedshop)
- `$ openss -f ./icos pcsamp`



```
Exclusive CPL /home/ldursi/Codes/athena3.1/src/fr_states_plm.c
1.640000    dWl[n] = pWl[n] - pWl[n-1];
0.600000    dWr[n] = pWr[n+1] - pWr[n];
3.740000    if (dWl[n]*dWr[n] > 0.0) {
4.070000    dWg[n] = 2.0*dWl[n]*dWr[n]/(dWl[n]+dWr[n]);
} else {
1.360000    dWg[n] = 0.0;
}
}
}
/*--- Step 3. -----
* Project differences in primitive variables along characteristics */
0.010000    for (n=0; n<NWAVE; n++) {
1.930000    dac[n] = lem[n][0]*dWc[0];
0.520000    dal[n] = lem[n][0]*dWl[0];
0.860000    dar[n] = lem[n][0]*dWr[0];
2.030000    dag[n] = lem[n][0]*dWg[0];
16.560000   for (m=1; m<NWAVE; m++) {
4.940000    dac[n] += lem[n][m]*dWc[m];
5.790000    dal[n] += lem[n][m]*dWl[m];
3.080000    dar[n] += lem[n][m]*dWr[m];
}
}
}
```

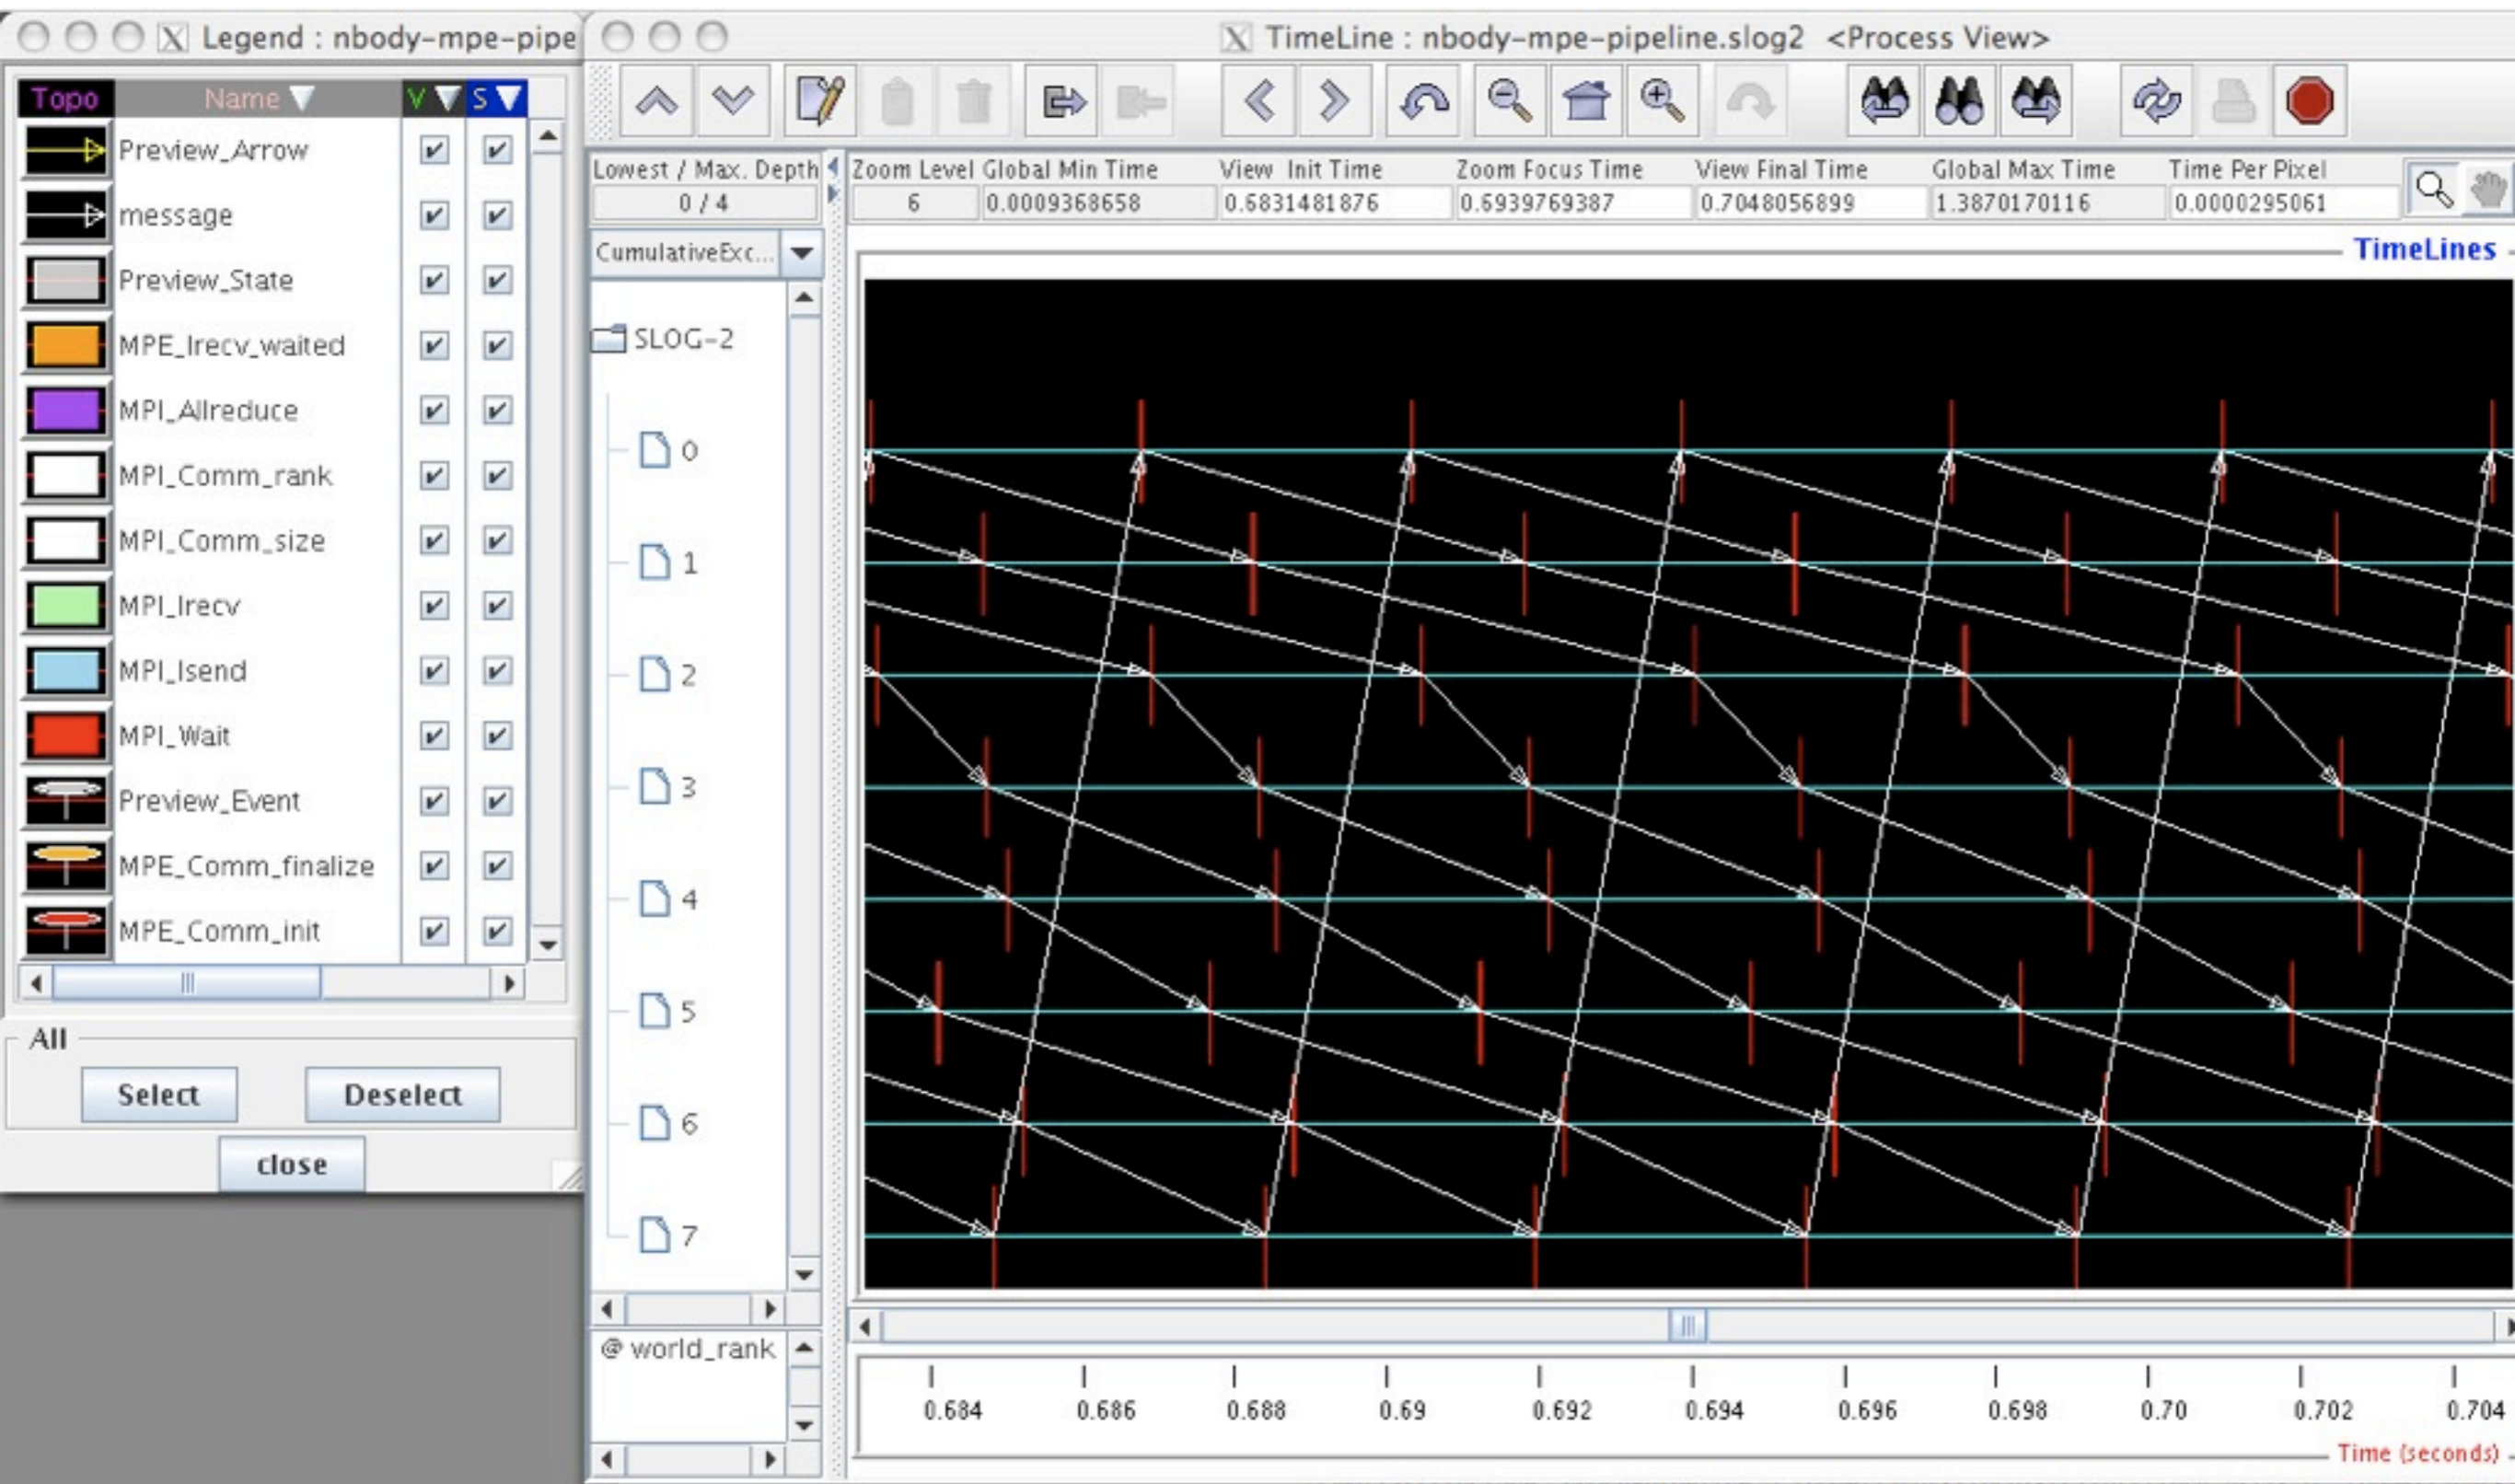
Jumpshot/MPE

- MPI has built in support for logging, so there are lots of MPI tools
- Logging big programs not easy!
- Comes with some installs of mpich.
- Installed on TCS, soon to be on GPC



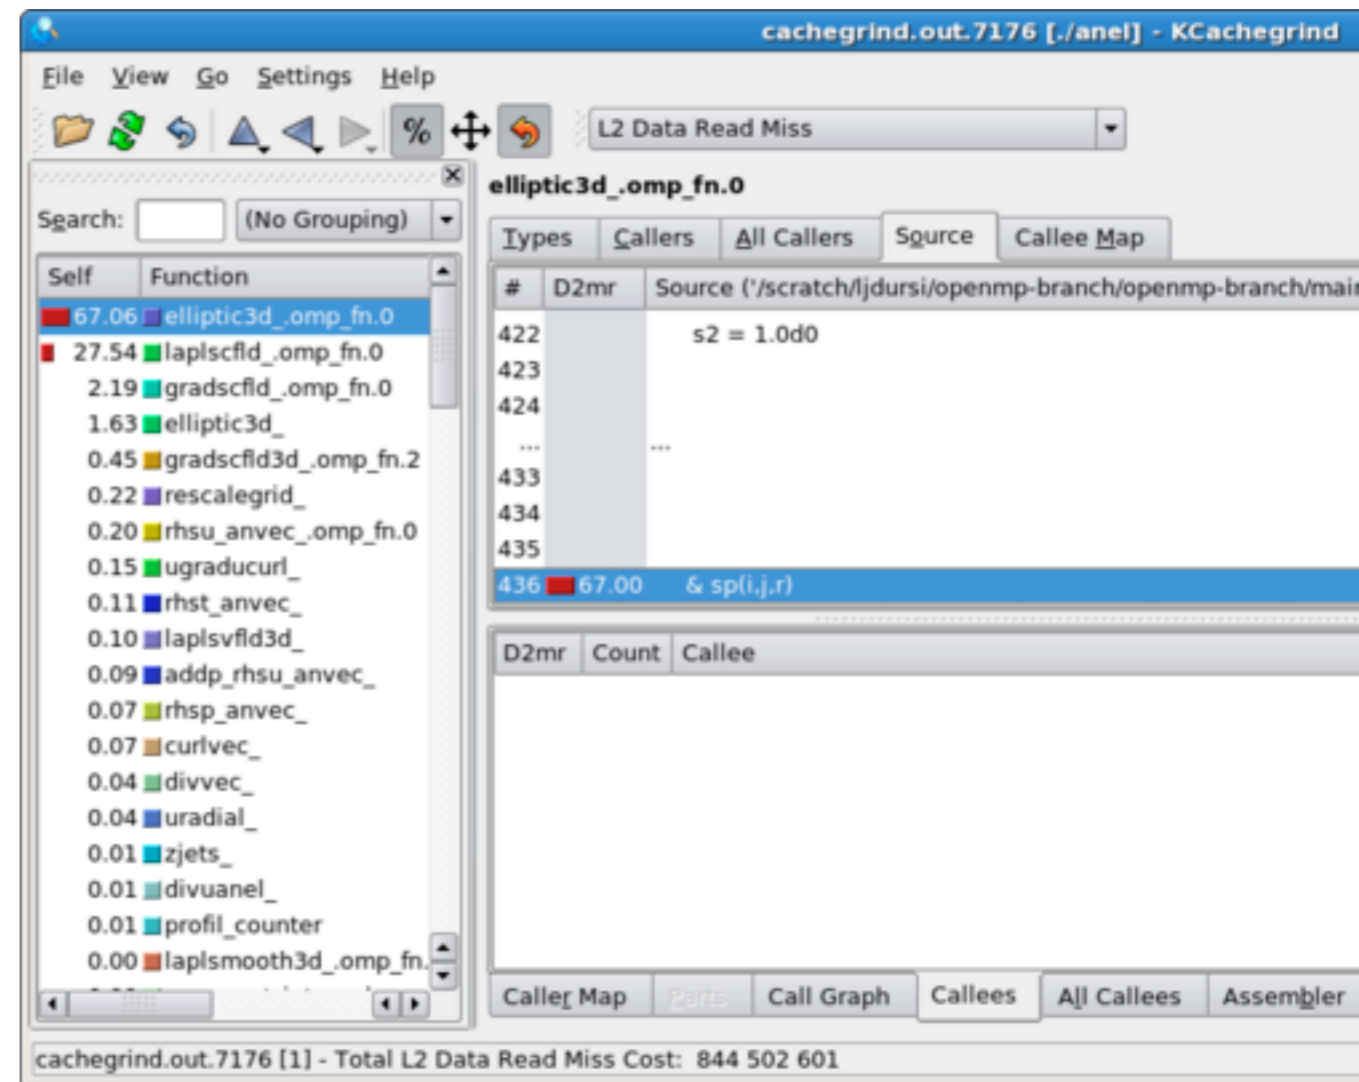
<http://www-unix.mcs.anl.gov/perfvis/>





cachegrind (valgrind)

- Part of valgrind package
- Emulator
- Exquisitely detailed description of cache performance
- Linux tool; installed on GPC



The screenshot shows the KCachegrind application window. The title bar reads "cachegrind.out.7176 [./anel] - KCachegrind". The interface includes a menu bar (File, View, Go, Settings, Help), a toolbar with navigation icons, and a search box. A list of functions is displayed on the left, with "elliptic3d_omp_fn.0" selected, showing a cost of 67.06. The right pane shows a detailed view of this function, including a table of L2 Data Read Misses (D2mr) with columns for line number, count, and source code. The source code snippet shows "s2 = 1.0d0" and "& sp(i,j,r)". The bottom status bar indicates "Total L2 Data Read Miss Cost: 844 502 601".

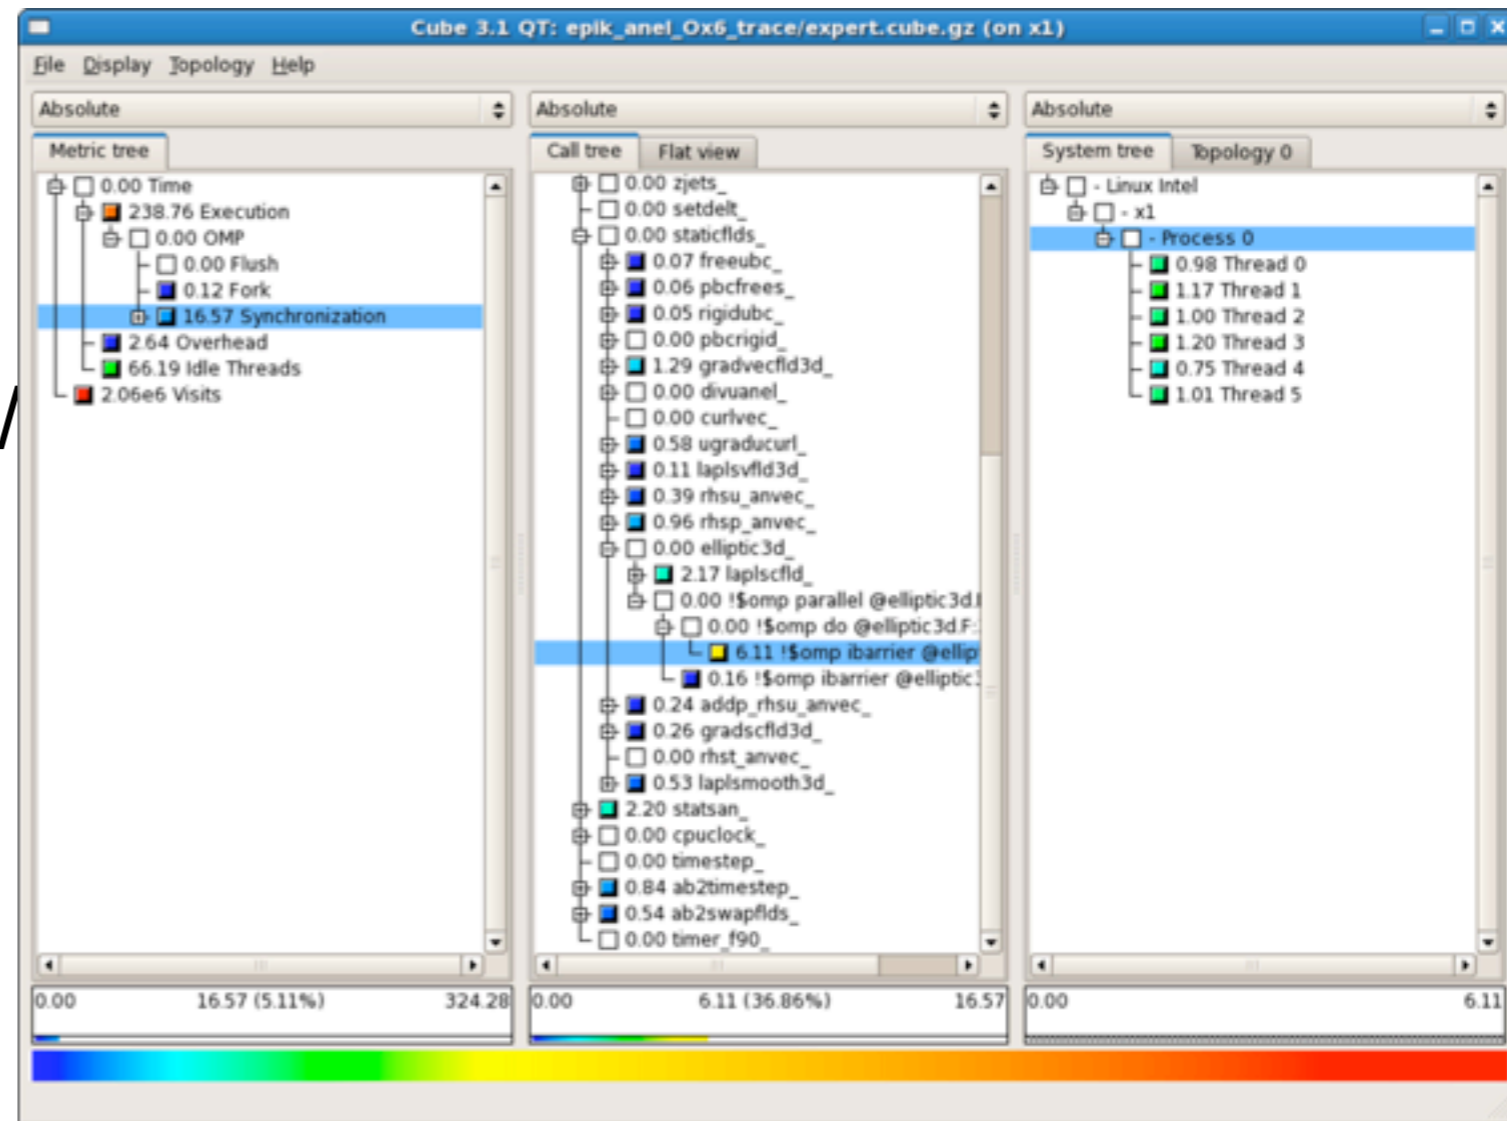
```
$ valgrind --tool=cachegrind  
myprogram myprogram_arguments
```

```
$ kcachegrind ./cachegrind.out
```



Scalasca

- <http://www.fz-juelich.de/jsc/scalasca/>
- Auto-instrument code
- Then run
- Sampling + tracing
- OpenMP + MPI
- Knows parallel issues
- Installed on TCS, GPC (module load scalasca)





Us!

- With specific questions, email support@scinet
- For longer discussions, make an appointment and come talk with us

