

CISC 372: Parallel Computing

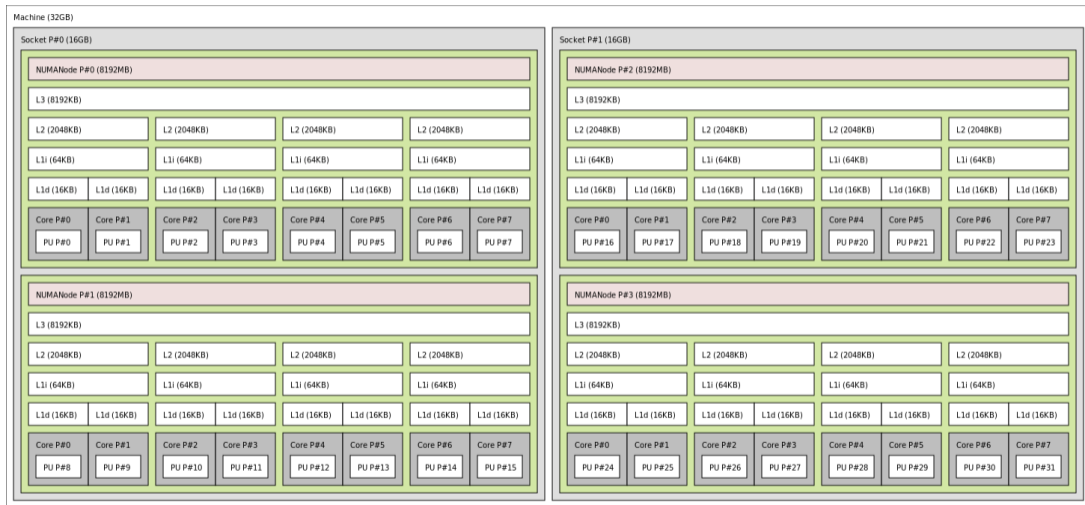
Performance

Stephen F. Siegel

Department of Computer and Information Sciences
University of Delaware

September 14, 2020

Memory hierarchy: AMD Bulldozer server



https://en.wikipedia.org/wiki/CPU_cache

Memory hierarchy: example: matrix-vector multiplication

$$\begin{bmatrix} a_{00} & a_{01} & a_{02} \\ a_{10} & a_{11} & a_{12} \\ a_{20} & a_{21} & a_{22} \\ a_{30} & a_{31} & a_{32} \end{bmatrix} \times \begin{bmatrix} x_0 \\ x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} a_{00}x_0 + a_{01}x_1 + a_{02}x_2 \\ a_{10}x_0 + a_{11}x_1 + a_{12}x_2 \\ a_{20}x_0 + a_{21}x_1 + a_{22}x_2 \\ a_{30}x_0 + a_{31}x_1 + a_{32}x_2 \end{bmatrix}$$

Memory hierarchy: example: matrix-vector multiplication

$$\begin{bmatrix} a_{00} & a_{01} & a_{02} \\ a_{10} & a_{11} & a_{12} \\ a_{20} & a_{21} & a_{22} \\ a_{30} & a_{31} & a_{32} \end{bmatrix} \times \begin{bmatrix} x_0 \\ x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} a_{00}x_0 + a_{01}x_1 + a_{02}x_2 \\ a_{10}x_0 + a_{11}x_1 + a_{12}x_2 \\ a_{20}x_0 + a_{21}x_1 + a_{22}x_2 \\ a_{30}x_0 + a_{31}x_1 + a_{32}x_2 \end{bmatrix}$$

Layout of a in memory:

a_{00}	a_{01}	a_{02}	a_{10}	a_{11}	a_{12}	a_{20}	a_{21}	a_{22}	a_{30}	a_{31}	a_{32}
----------	----------	----------	----------	----------	----------	----------	----------	----------	----------	----------	----------

Memory hierarchy: example: matrix-vector multiplication

$$\begin{bmatrix} a_{00} & a_{01} & a_{02} \\ a_{10} & a_{11} & a_{12} \\ a_{20} & a_{21} & a_{22} \\ a_{30} & a_{31} & a_{32} \end{bmatrix} \times \begin{bmatrix} x_0 \\ x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} a_{00}x_0 + a_{01}x_1 + a_{02}x_2 \\ a_{10}x_0 + a_{11}x_1 + a_{12}x_2 \\ a_{20}x_0 + a_{21}x_1 + a_{22}x_2 \\ a_{30}x_0 + a_{31}x_1 + a_{32}x_2 \end{bmatrix}$$

Layout of a in memory:

a_{00}	a_{01}	a_{02}	a_{10}	a_{11}	a_{12}	a_{20}	a_{21}	a_{22}	a_{30}	a_{31}	a_{32}
----------	----------	----------	----------	----------	----------	----------	----------	----------	----------	----------	----------

▶ see [rowmaj.c](#)

Factors that affect performance: compiler optimizations

Compilers can transform programs in myriad ways to use resources more effectively. . .

- ▶ function inlining; loop fission, loop fusion; loop interchange; loop unrolling; common subexpression elimination; constant folding, propagation . . .

Tradeoffs: more optimization generally entails. . .

- ▶ longer compile time
- ▶ larger generated code size
- ▶ program gets harder to debug
- ▶ greater sensitivity to undefined behavior (but you shouldn't use any undefined behavior!)
- ▶ generated code might actually get slower

Speedup

Let

T_{seq} = time to run sequential baseline

T_{par} = time to run parallel program

Then

$$\text{Speedup} = \frac{T_{\text{seq}}}{T_{\text{par}}}$$

- ▶ higher speedup is better

Speedup as a function of `nprocs`

- ▶ hopefully: speedup will change (increase!) with `nprocs`
- ▶ **ideal case**: speedup = `nprocs`
 - ▶ double the number of procs, cut the execution time in half
- ▶ **reality**: rarely that good
 - ▶ communication time (sending messages)
 - ▶ synchronization time (procs have to sit around waiting, e.g., at a Barrier)
 - ▶ redundant work (two procs compute the same thing)
- ▶ after some point adding more processes no longer increases speedup

Speedup as a function of `nprocs`

- ▶ hopefully: speedup will change (increase!) with `nprocs`
- ▶ **ideal case**: speedup = `nprocs`
 - ▶ double the number of procs, cut the execution time in half
- ▶ **reality**: rarely that good
 - ▶ communication time (sending messages)
 - ▶ synchronization time (procs have to sit around waiting, e.g., at a Barrier)
 - ▶ redundant work (two procs compute the same thing)
- ▶ after some point adding more processes no longer increases speedup
 - ▶ sort a list of 10^{12} elements

Amdahl's Law

Amdahl's Law

- ▶ typically some part of the code cannot be parallelized
- ▶ “inherently sequential”
- ▶ example: `diffuse1d` writes data to the screen
 - ▶ there is only one screen: no way to parallelize that part
- ▶ example
 - ▶ say in sequential program, 10% of time is spent doing “inherently sequential” work
 - ▶ the other 90% can be parallelized
 - ▶ even if a parallel program were PERFECT with unlimited resources, the best it could do is reduce the 90% to 0.
 - ▶ the 10% time would be unchanged

Weak vs. Strong Scaling

- ▶ strong scaling: baseline is constant as nprocs increases
 - ▶ always comparing against sequential run on fixed problem size
 - ▶ speedup is bounded
- ▶ strong scaling examples
 - ▶ fix list of length 10^6 ; compare sequential time to sort vs. parallel time to sort with p procs
 - ▶ fix $nx = 10^3$; compare sequential diffusion1d vs. parallel diffusion1d with p procs
 - ▶ note $nx1$, the amount of data per process, decreases as p increases

Automating performance experiments

Automating performance experiments

- ▶ see [exp/sat_strong](#) in public course repo
- ▶ a strong scaling experiment of MPI SAT solver

Data file generated by SAT performance experiment

- ▶ the `Makefile` executes `sat_mpi.exec` with 1, 2, 4, 8, 16, 32 procs
- ▶ the results are accumulated in a file `sat_mpi.dat`:

```
1 42.693483
2 29.942159
4 16.342128
8 9.844605
16 5.327622
32 2.452447
```

Data file generated by SAT performance experiment

- ▶ the `Makefile` executes `sat_mpi.exec` with 1, 2, 4, 8, 16, 32 procs
- ▶ the results are accumulated in a file `sat_mpi.dat`:

```
1 42.693483
2 29.942159
4 16.342128
8 9.844605
16 5.327622
32 2.452447
```

nprocs	time	speedup	efficiency
1	43.75	1.00	1.00
2	30.46	1.44	0.72
4	16.47	2.66	0.66
8	8.68	5.04	0.63
16	4.89	8.95	0.56
32	2.52	17.36	0.54

Graphing data with `gnuplot`

Graphing data with `gnuplot`

- ▶ free, open-source command-line tool for creating graphs
- ▶ <http://www.gnuplot.info>
- ▶ command: `gnuplot sat_mpi.gnu`

Graphing data with `gnuplot`

- ▶ free, open-source command-line tool for creating graphs
- ▶ <http://www.gnuplot.info>
- ▶ command: `gnuplot sat_mpi.gnu`

```
set terminal pdf
set output "sat_mpi.pdf"
set xlabel center "Number of processes"
set ylabel center "time (seconds)"
set xr [0:32]
set yr [0:45]
plot "sat_mpi.dat" using 1:2 title 'MPI' with linespoints
```


PDF file resulting from SAT scaling experiment

