NHR FAU

FAU Friedrich-Alexander-Universität Erlangen-Nürnberg

Programming Techniques for Supercomputers

Erlangen National High Performance Computing Center

Department of Computer Science

FAU Erlangen-Nürnberg

Sommersemester 2025

Assignment 7 – Task 1

```
$ likwid-topology -g
--------------------------------------------------------------------------------
CPU name:         AMD EPYC 7713 64-Core Processor
CPU type:         AMD K19 (Zen3) architecture
CPU stepping:     1
********************************************************************************
Hardware Thread Topology
********************************************************************************
Sockets:                  2
CPU dies:                 2
Cores per socket:         64
Threads per core:         1
--------------------------------------------------------------------------------
HWThread          Thread            Core      Die       Socket            Available
0                 0                 0         0         0                 *
1                 0                 1         0         0                 *
2                 0                 2         0         0                 *
3                 0                 3         0         0                 *
4                 0                 4         0         0                 *
[...SNIP…]
126               0                 102       0         1                 *
127               0                 103       0         1                 *
--------------------------------------------------------------------------------
[...SNIP...]
```

2 sockets,
128 cores/node

Assignment 7 – Task 1

```
$ likwid-topology -g
[...SNIP...]
********************************************************************************
Cache Topology
********************************************************************************
Level:                  1
Size:                   32 kB
Cache groups:           ( 0 ) ( 1 ) ( 2 ) ( 3 ) ( 4 ) ( 5 ) ( 6 ) ( 7 ) ( 8 ) ( 9 ) ( 10 ) ( 11 ) ( 12 ) ( 13 ) ( 14 ) (
15 ) ( 16 ) ( 17 ) ( 18 ) ( 19 ) ( 20 ) ( 21 ) ( 22 ) ( 23 ) ( 24 ) ( 25 ) ( 26 ) ( 27 ) ( 28 ) ( 29 ) ( 30 ) ( 31 ) ( 32
) ( 33 ) [...SNIP...] ( 116 ) ( 117 ) ( 118 ) ( 119 ) ( 120 ) ( 121 ) ( 122 ) ( 123 ) ( 124 ) ( 125 ) ( 126 ) ( 127 )
--------------------------------------------------------------------------------
Level:                  2
Size:                   512 kB
Cache groups:           ( 0 ) ( 1 ) ( 2 ) ( 3 ) ( 4 ) ( 5 ) ( 6 ) ( 7 ) ( 8 ) ( 9 ) ( 10 ) ( 11 ) ( 12 ) ( 13 ) ( 14 ) (
15 ) ( 16 ) ( 17 ) ( 18 ) ( 19 ) ( 20 ) ( 21 ) ( 22 ) ( 23 ) ( 24 ) ( 25 ) ( 26 ) ( 27 ) ( 28 ) ( 29 ) ( 30 ) ( 31 ) ( 32
) ( 33 ) [...SNIP...] ( 74 ) ( 75 ) ( 76 ) ( 77 ) ( 78 ) ( 79 ) ( 80 ) ( 81 ) ( 82 ) ( 83 ) ( 84 ) ( 85 ) ( 86 ) ( 87 ) (
88 ) ( 89 ) ( 90 ) ( 91 ) ( 92 ) ( 93 ) ( 94 ) ( 95 ) ( 96 ) ( 97 ) ( 98 ) ( 99 ) ( 100 ) ( 101 )
--------------------------------------------------------------------------------
Level:                  3
Size:                   32 MB
Cache groups:           ( 0 1 2 3 4 5 6 7 ) ( 8 9 10 11 12 13 14 15 ) ( 16 17 18 19 20 21 22 23 ) ( 24 25 26 27 28 29 30
31 ) ( 32 33 34 35 36 37 38 39 ) (40 41 42 43 44 45 46 47 ) ( 48 49 50 51 52 53 54 55 ) ( 56 57 58 59 60 61 62 63 ) ( 64
65 66 67 68 69 70 71 ) ( 72 73 74 75 76 77 78 79 ) ( 80 81 82 83 84 85 86 87 ) ( 88 89 90 91 92 93 94 95 ) ( 96 97 98 99
100 101 102 103 ) ( 104 105 106 107 108 109 110 111 ) ( 112 113 114 115 116 117 118 119 ) ( 120 121 122 123 124 125 126
127 )
--------------------------------------------------------------------------------
[...SNIP...]
```

Per-core L1 (32 KiB) and L2 (512 KiB) caches

CCD-shared L3 (32 MiB) cache

Assignment 7 – Task 1

```
$ likwid-topology -g
********************************************************************************
NUMA Topology
********************************************************************************
NUMA domains:             8
--------------------------------------------------------------------------------
Domain:                   0
Processors:               ( 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 )
Distances:                10 12 12 12 32 32 32 32
Free memory:              63729.5 MB
Total memory:             64319 MB
--------------------------------------------------------------------------------
[...SNIP...]
--------------------------------------------------------------------------------
Domain:                   7
Processors:               ( 112 113 114 115 116 117 118 119 120 121 122 123 124 125 126 127 )
Distances:                32 32 32 32 12 12 12 10
Free memory:              62940.4 MB
Total memory:             64505 MB
--------------------------------------------------------------------------------

--------------------------------------------------------------------------------
```

4 ccNUMA domains per socket (16 cores)

Assignment 7 – Task 2

```
double t_start,t_end;
int id, nt, k, I, NITER, N;
double t = 0;
double *a = (double*)aligned_alloc(64,N*sizeof(double));
double *b = (double*)aligned_alloc(64,N*sizeof(double));

#pragma omp parallel for
for(k=0; k<N; ++k)
  a[k]=b[k]=0.;

NITER=1;
do {
  t_start = getTimeStamp();

  for(k=0; k<NITER; ++k) {
    #pragma omp parallel for
    for(i=0; i<N; ++i) {
      a[i] = a[i] + 1.0000001 * b[i];
    }
    if(a[N/2]<0.) printf("%lf",a[N/2]);
  }
  t_end = getTimeStamp();
  NITER = NITER*2;
} while (wct_end-wct_start<1.);

NITER = NITER/2;

printf("Total walltime: %f, NITER: %d, Perf: %.3lf GF/s\n",
t_end-t_start, NITER, 2.*(double)N*NITER/(t_end-t_start)*1.e-9);
```

```
icx -O3 -xHost -qopenmp
scan_c.c timing.c
```

Assignment 7 – Task 2

```
#/bin/bash
for t in 1 2 4 8 12 18; do
    n=10
    for i in `seq 1 90`; do
        echo -n $n " "
        srun -cpu-freq=2000000-2000000:performance likwid-pin -c 0-$(( t - 1 )) -q \
                                        ./a.out $n >> daxpy_${t}.out
        n=$(( (n * 12) / 10 ))
    done
done
```

```
#/bin/bash
for t in 1 2 4 8 12 18; do
    n=10
    for i in `seq 1 90`; do
        echo -n $n " "
        OMP_NUM_THREADS=$t OMP_PLACES=cores OMP_PROC_BIND=close \
                                srun -cpu-freq=2000000-2000000:performance \
                                ./a.out $n >> daxpy_${t}.out
        n=$(( (n * 12) / 10 ))
    done
done
```

Assignment 7 – Task 2 a)

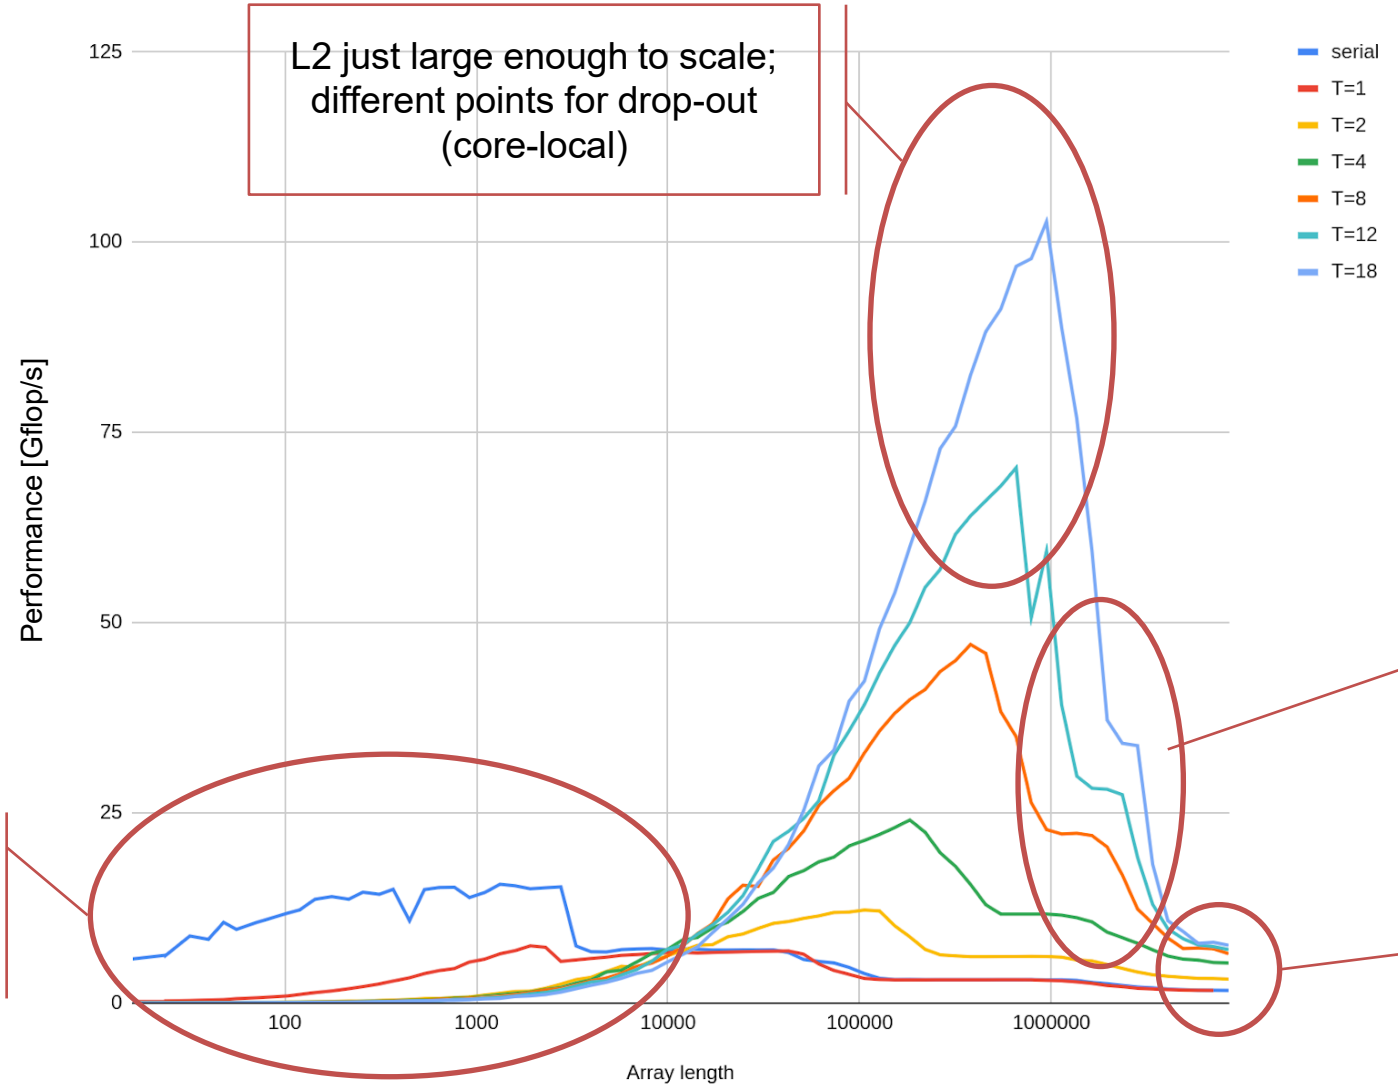

L2 just large enough to scale; different points for drop-out (core-local)

Decent scaling in L3 up to 8 cores; however, bad scaling onwards ➔ shared L3 cache

“Reverse” scaling due to OpenMP overhead

No scaling beyond 8 cores ➔ memory bandwidth exhausted

Assignment 7 – Task 2 b)

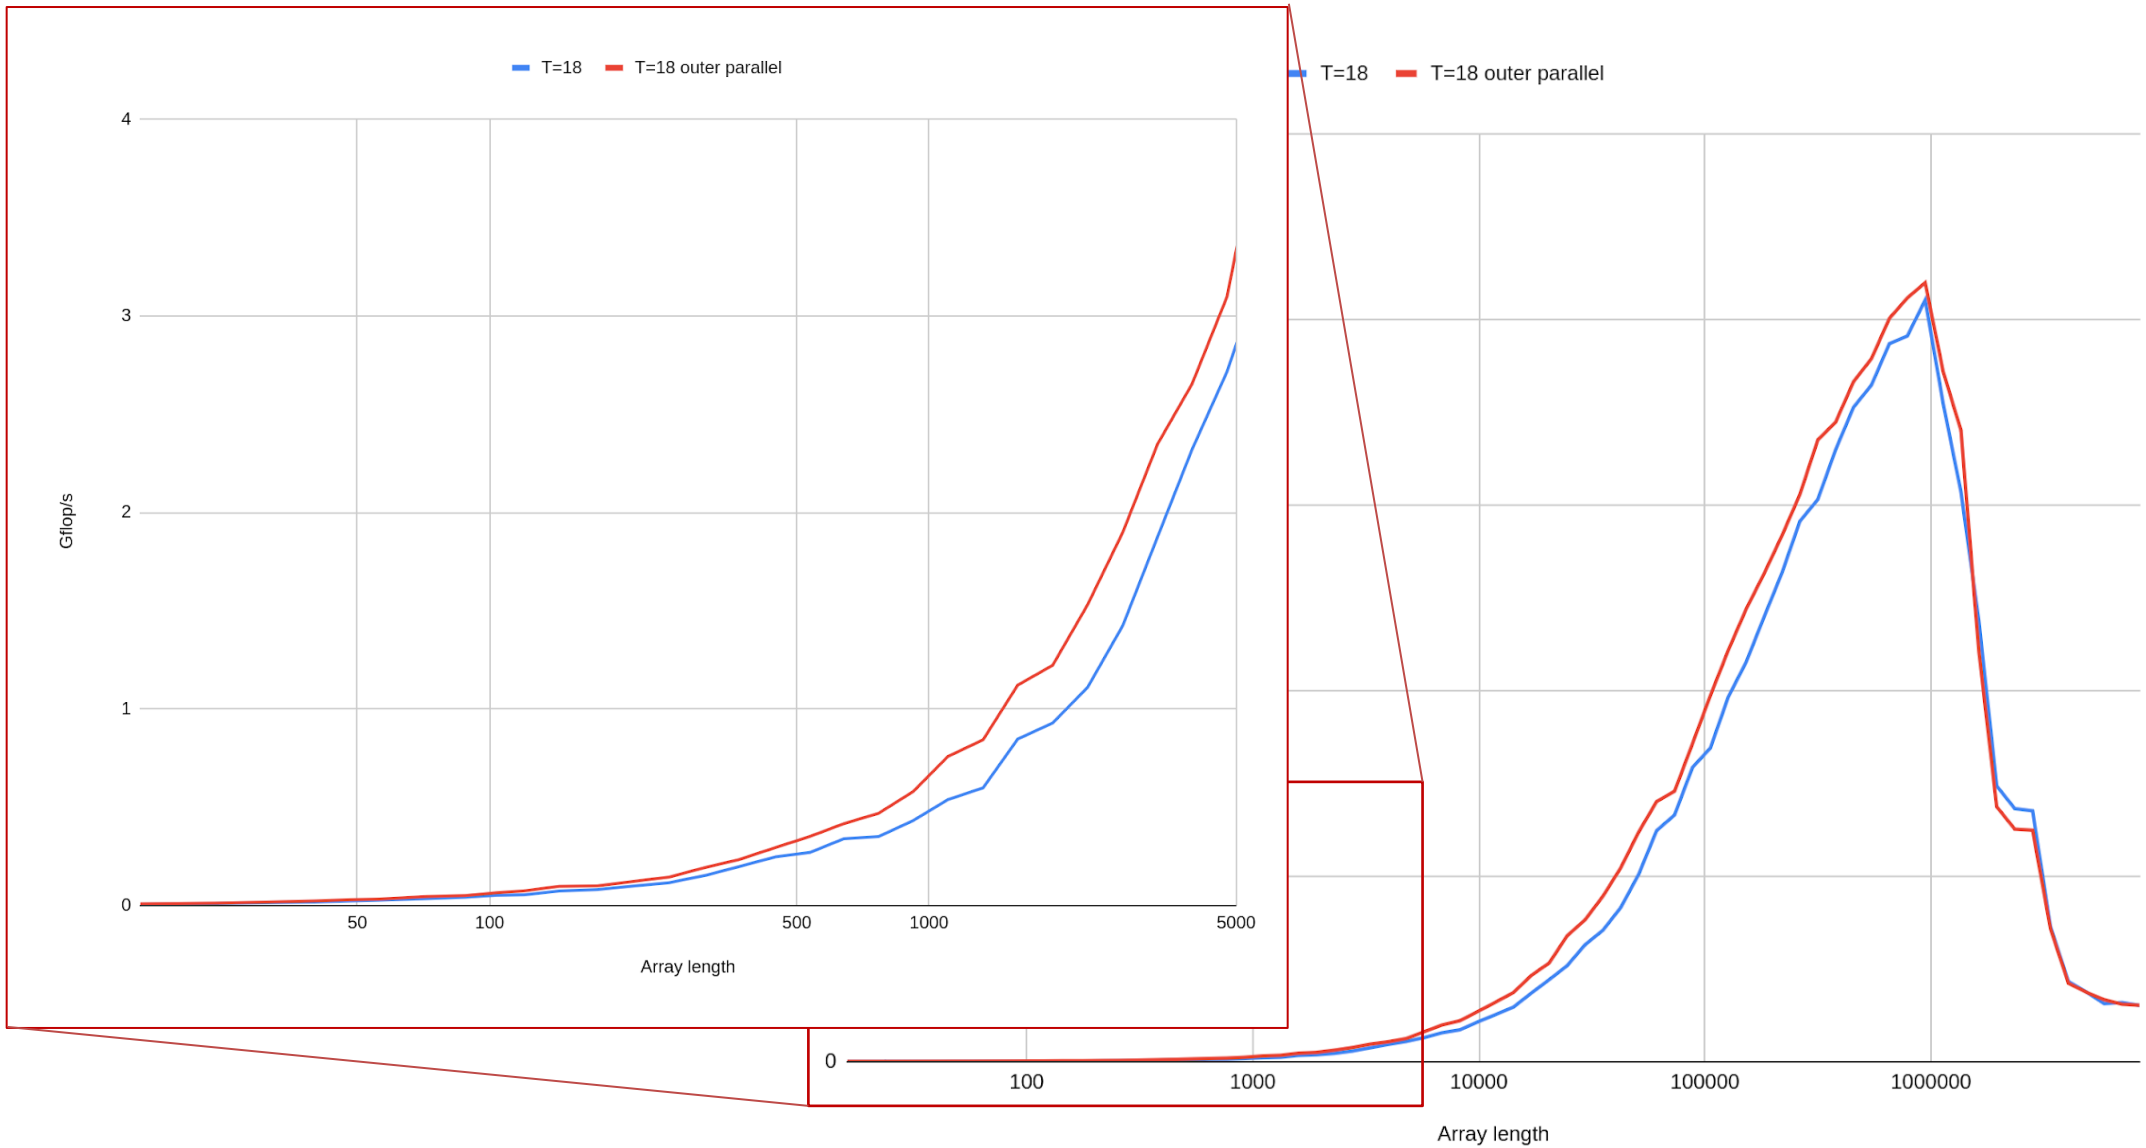

Assignment 7 – Task 2 c)

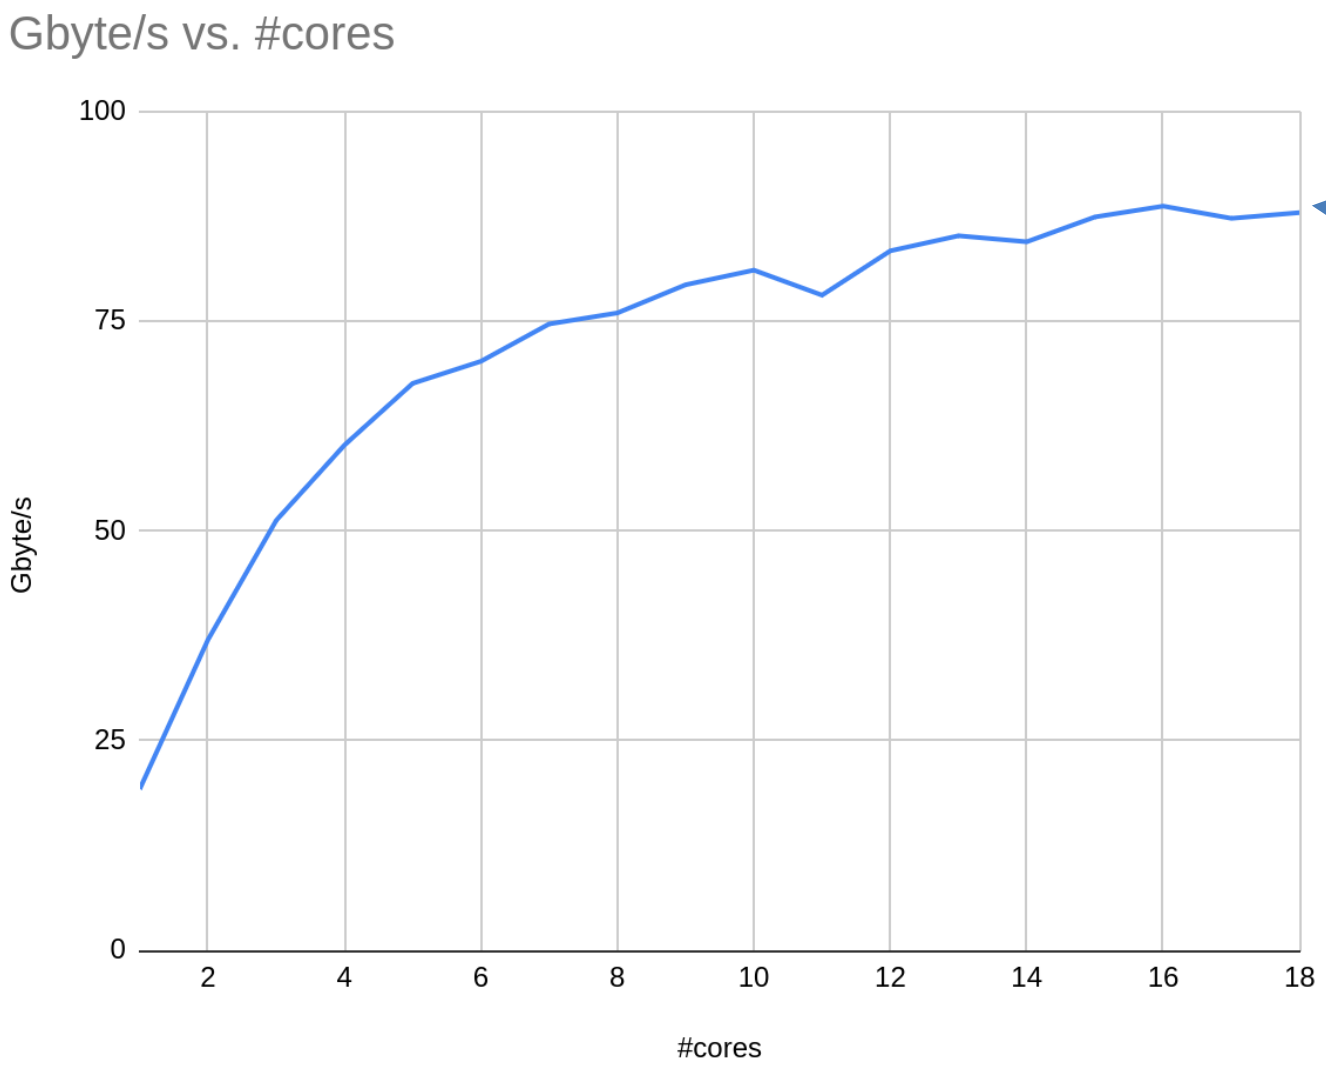


```
#pragma omp parallel for
for (int i=0; i<N; ++i)
    a[i] = a[i] + s * b[i];
```

$$b = 6300 \frac{MFLOP}{s} \times 12 \frac{B}{FLOP}$$
$$= 75.6 \frac{GB}{s}$$

Speedup: ~ 4 - 4.5x

Assignment 7 – Task 3

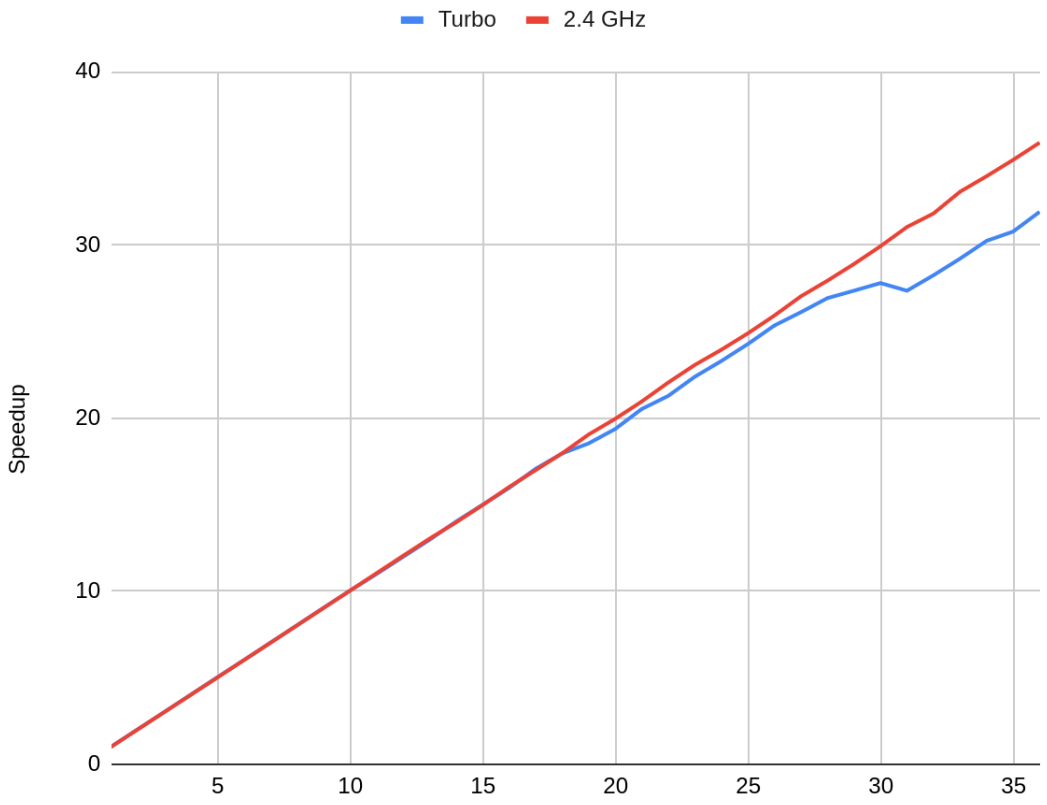


```
// bogus parallel region
#pragma omp parallel
  if(omp_get_thread_num()==0) printf("Hello");

  S = getTimeStamp();
#pragma omp parallel for reduction(+:sum) private(x)
  for(int i=0; i<N; i++) {
    x = (i + 0.5) * delta_x;
    sum += (1.0 / (1.0 + x));
  }
  double ln = sum * delta_x;
  E = getTimeStamp();
```

```
$ OMP_NUM_THREADS=$t OMP_PLACES=cores \
  OMP_PROC_BIND=close \
  srun --cpu-freq=performance ./a.out
```

- Perfect scaling with fixed frequency!
- Degraded scaling with Turbo Mode starting @19 cores → clock speed reduction!