Parallelizing Lattice Boltzmann Methods with OpenMP

Miguel Montes

July 2012

Table of Contents

Table of Contents

2D flow with a circular obstacle

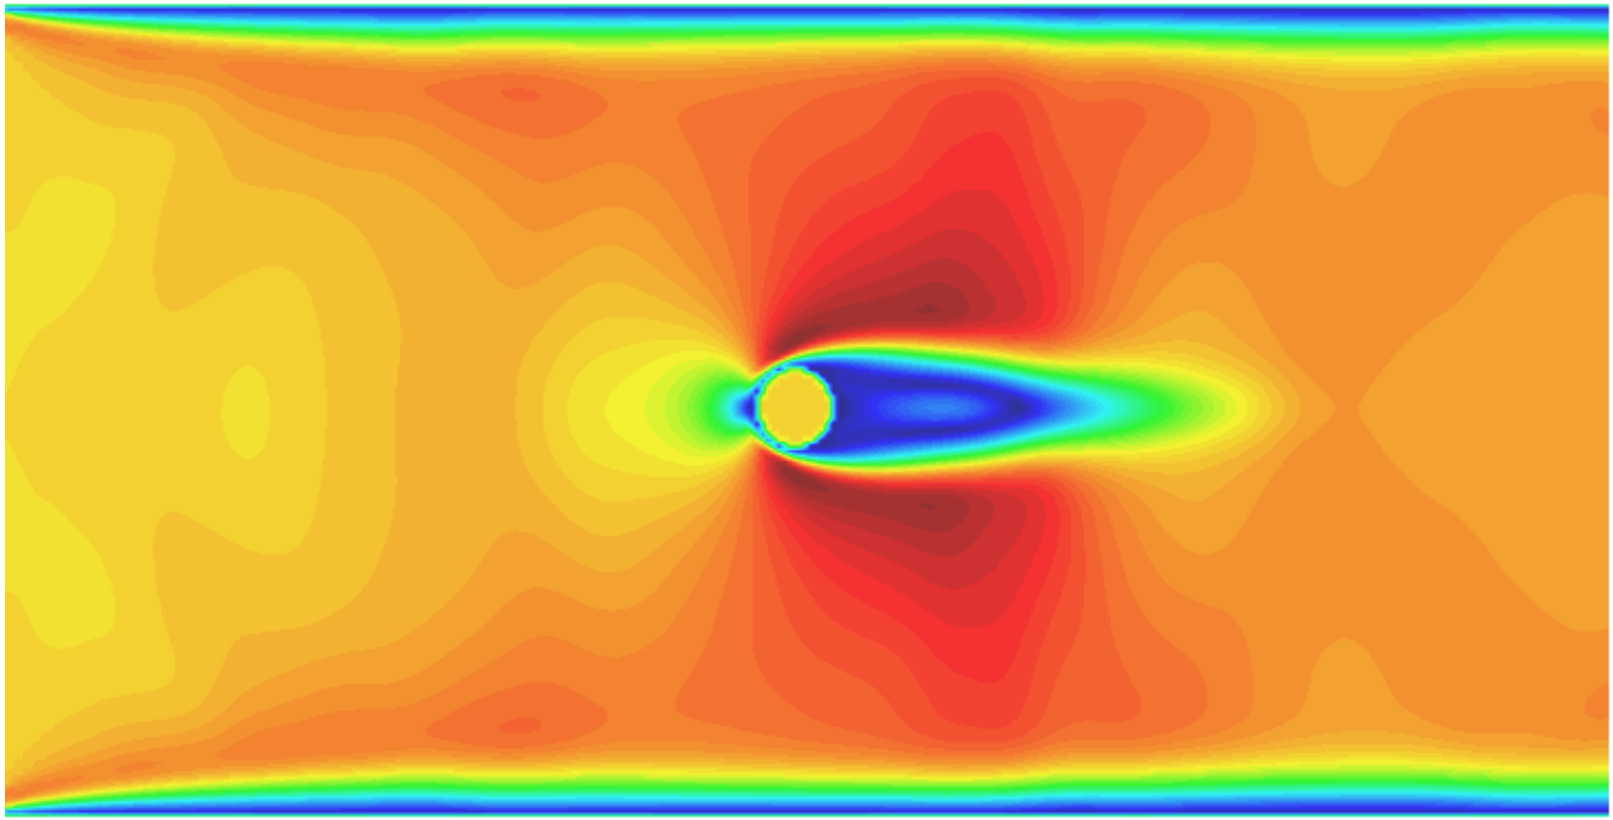

Infrastructure
Intel Manycore Testing Lab (MTL)

- 1 shared login node, for workload development.

Infrastructure
Intel Manycore Testing Lab (MTL)

- 1 shared login node, for workload development.
- 3 batch nodes for benchmarking.

Infrastructure
Intel Manycore Testing Lab (MTL)

- 1 shared login node, for workload development.
- 3 batch nodes for benchmarking.
- MPI is not a supported option.

Infrastructure
Intel Manycore Testing Lab (MTL)

- 1 shared login node, for workload development.
- 3 batch nodes for benchmarking.
- MPI is not a supported option.
- Red Hat Enterprise Linux Server release 5.4, with kernel 2.6.18-194.11.4.el5.

Infrastructure
Intel Manycore Testing Lab (MTL)

- 1 shared login node, for workload development.
- 3 batch nodes for benchmarking.
- MPI is not a supported option.
- Red Hat Enterprise Linux Server release 5.4, with kernel 2.6.18-194.11.4.el5.
- Nodes: 4 Intel® Xeon® E7-4860, 10 cores each, 24M Cache and 2.26 GHz Clock Speed.

Infrastructure
Intel Manycore Testing Lab (MTL)

- 1 shared login node, for workload development.
- 3 batch nodes for benchmarking.
- MPI is not a supported option.
- Red Hat Enterprise Linux Server release 5.4, with kernel 2.6.18-194.11.4.el5.
- Nodes: 4 Intel® Xeon® E7-4860, 10 cores each, 24M Cache and 2.26 GHz Clock Speed.
- Login node: 256Gb RAM, Hyperthreading enabled.

Infrastructure
Intel Manycore Testing Lab (MTL)

- 1 shared login node, for workload development.
- 3 batch nodes for benchmarking.
- MPI is not a supported option.
- Red Hat Enterprise Linux Server release 5.4, with kernel 2.6.18-194.11.4.el5.
- Nodes: 4 Intel® Xeon® E7-4860, 10 cores each, 24M Cache and 2.26 GHz Clock Speed.
- Login node: 256Gb RAM, Hyperthreading enabled.
- Batch nodes: 64Gb RAM, Hyperthreading disabled.

Infrastructure
Intel Manycore Testing Lab (MTL)

- 1 shared login node, for workload development.
- 3 batch nodes for benchmarking.
- MPI is not a supported option.
- Red Hat Enterprise Linux Server release 5.4, with kernel 2.6.18-194.11.4.el5.
- Nodes: 4 Intel® Xeon® E7-4860, 10 cores each, 24M Cache and 2.26 GHz Clock Speed.
- Login node: 256Gb RAM, Hyperthreading enabled.
- Batch nodes: 64Gb RAM, Hyperthreading disabled.
- Compilers: Intel 11.1, gcc (4.1.2 and 4.5.1)

Infrastructure
Intel Manycore Testing Lab (MTL)

- 1 shared login node, for workload development.
- 3 batch nodes for benchmarking.
- MPI is not a supported option.
- Red Hat Enterprise Linux Server release 5.4, with kernel 2.6.18-194.11.4.el5.
- Nodes: 4 Intel® Xeon® E7-4860, 10 cores each, 24M Cache and 2.26 GHz Clock Speed.
- Login node: 256Gb RAM, Hyperthreading enabled.
- Batch nodes: 64Gb RAM, Hyperthreading disabled.
- Compilers: Intel 11.1, gcc (4.1.2 and 4.5.1)
- PBSPro 10.2

Table of Contents

Thread affinity

- Intel: `export KMP_AFFINITY={compact|scatter|none}`

Thread affinity

- Intel: `export KMP_AFFINITY={compact|scatter|none}`
- GNU

Thread affinity

- Intel: `export KMP_AFFINITY={compact|scatter|none}`
- GNU
  - compact: `export GOMP_CPU_AFFINITY=0-39`

Thread affinity

- Intel: `export KMP_AFFINITY={compact|scatter|none}`
- GNU
  - compact: `export GOMP_CPU_AFFINITY=0-39`
  - scatter: `export GOMP_CPU_AFFINITY="0 10 20 30 1 11 21 31 41"`

Thread affinity: Compact

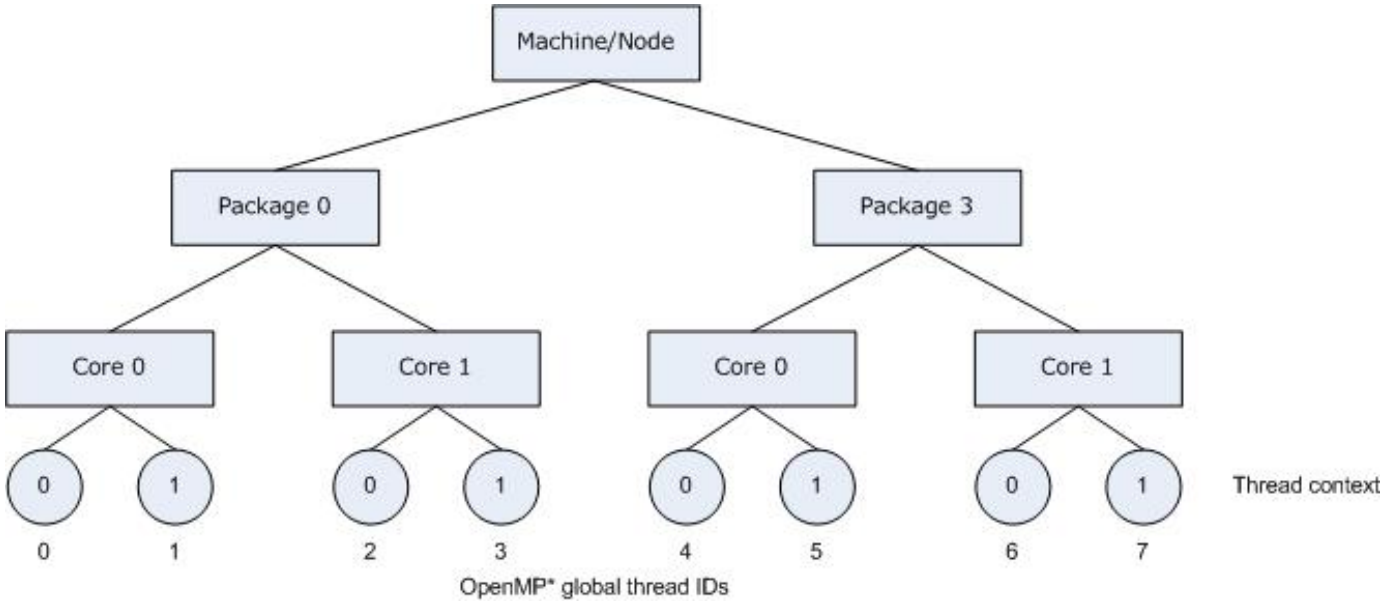

Thread affinity: Scatter

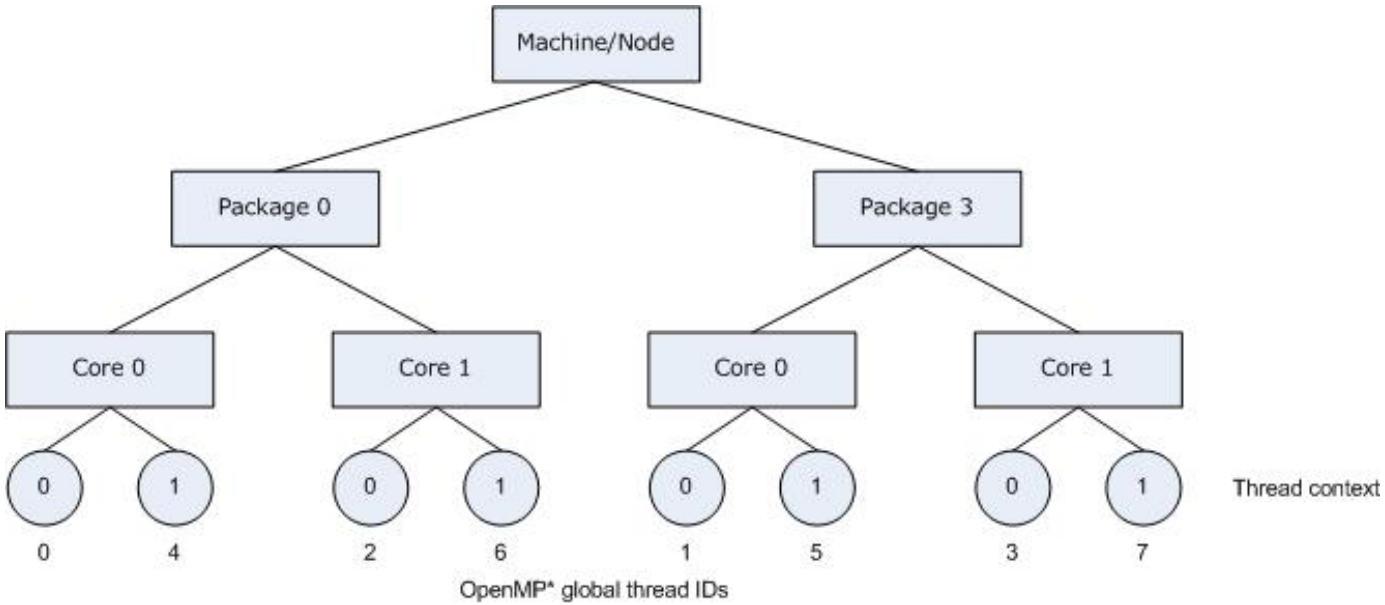

CPU affinity with gcc. Small size.

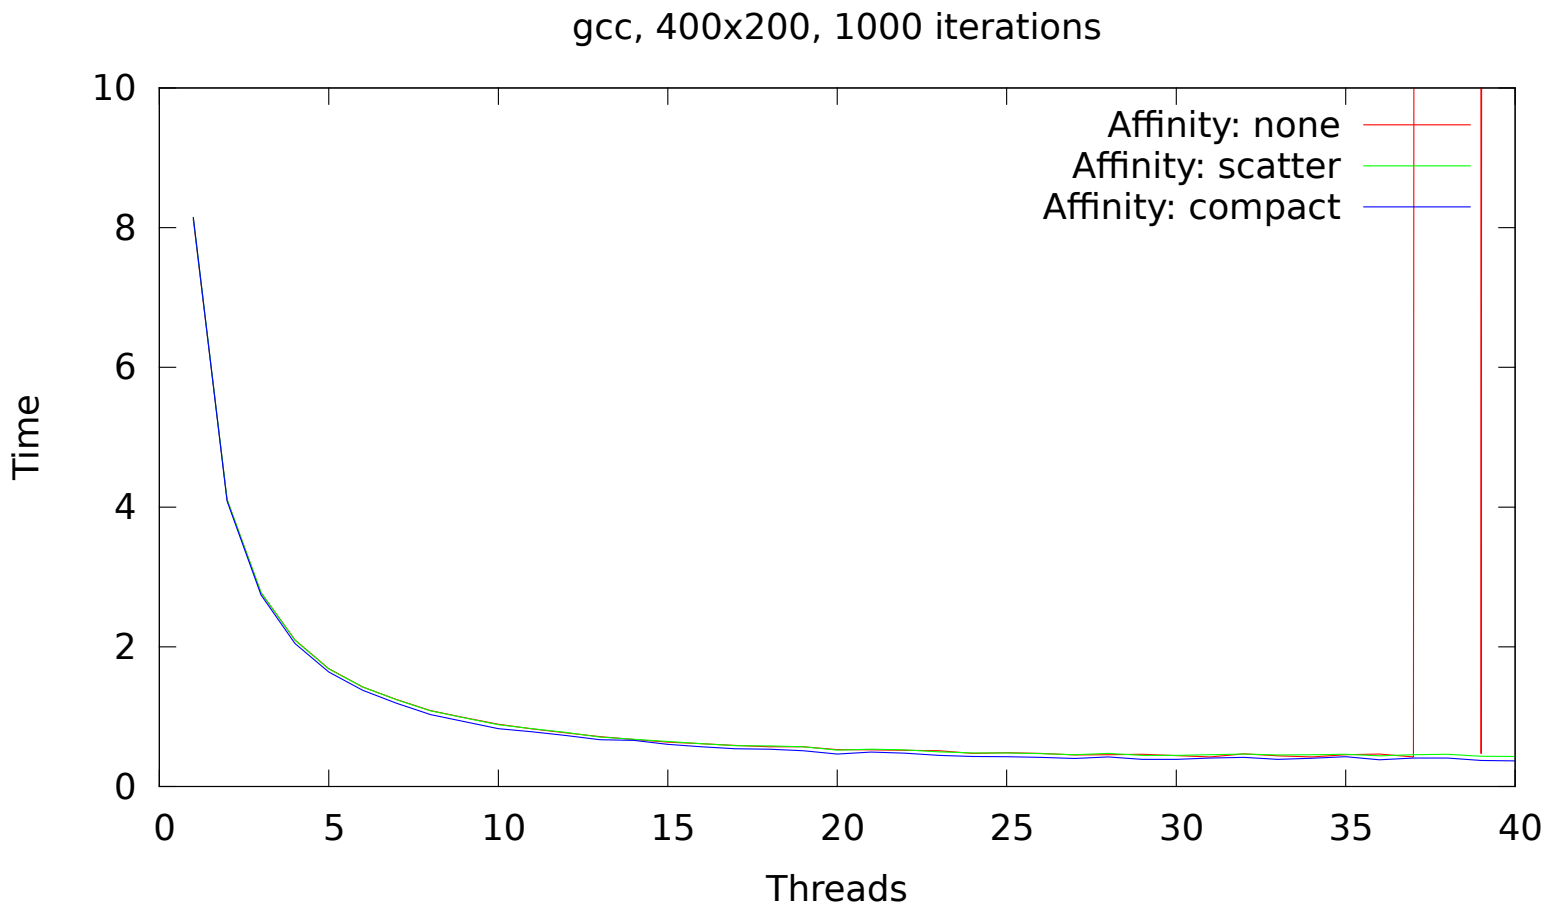

CPU affinity with gcc. Medium size.

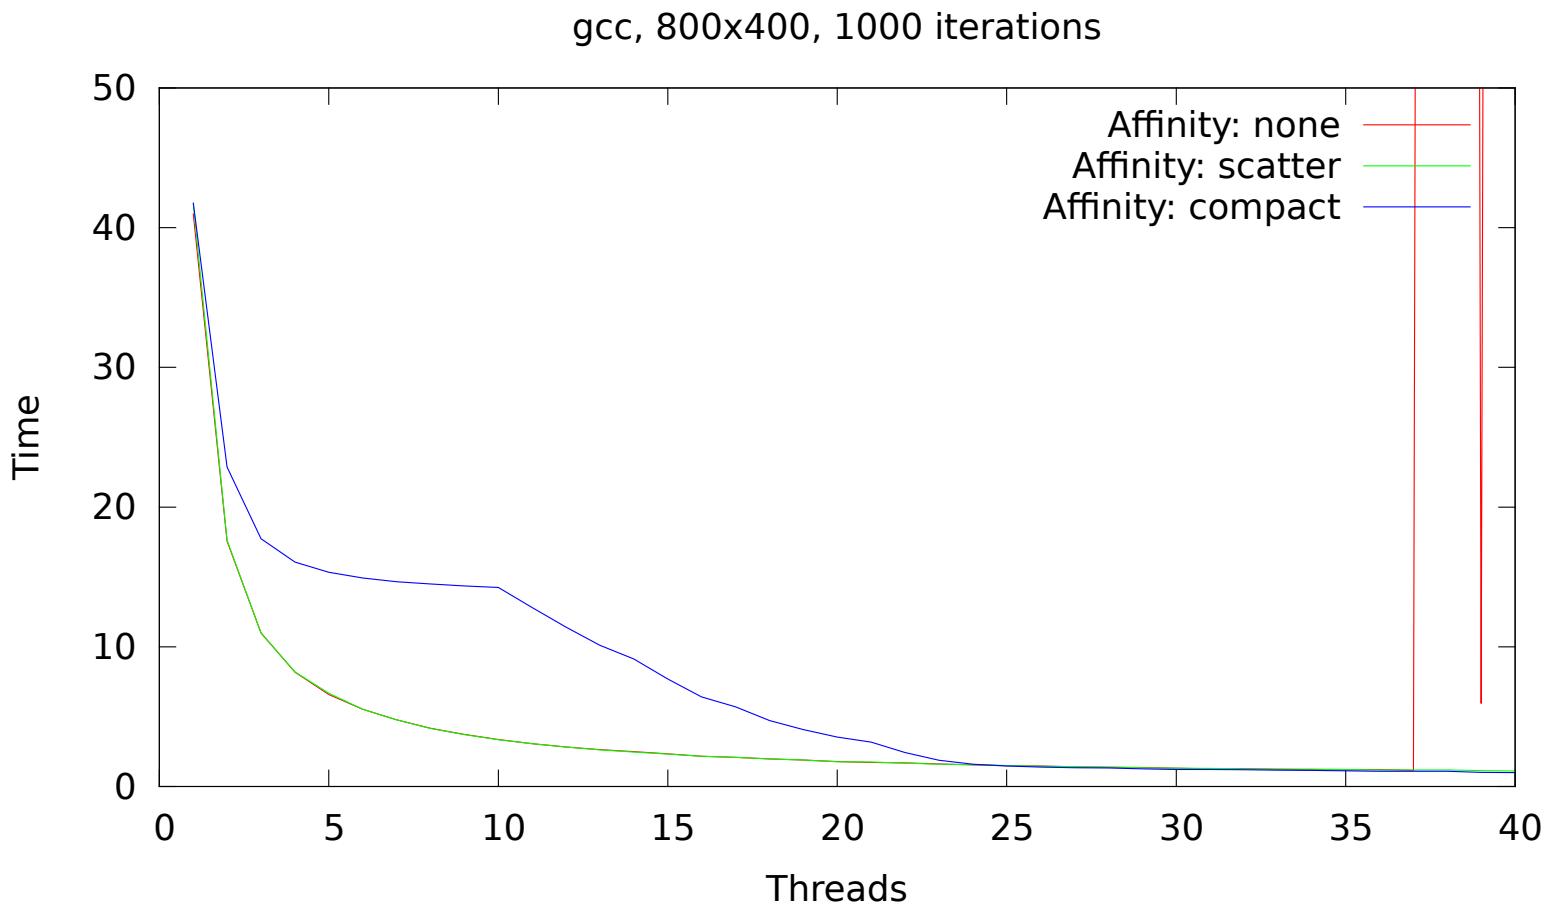

CPU affinity with gcc. Large size.

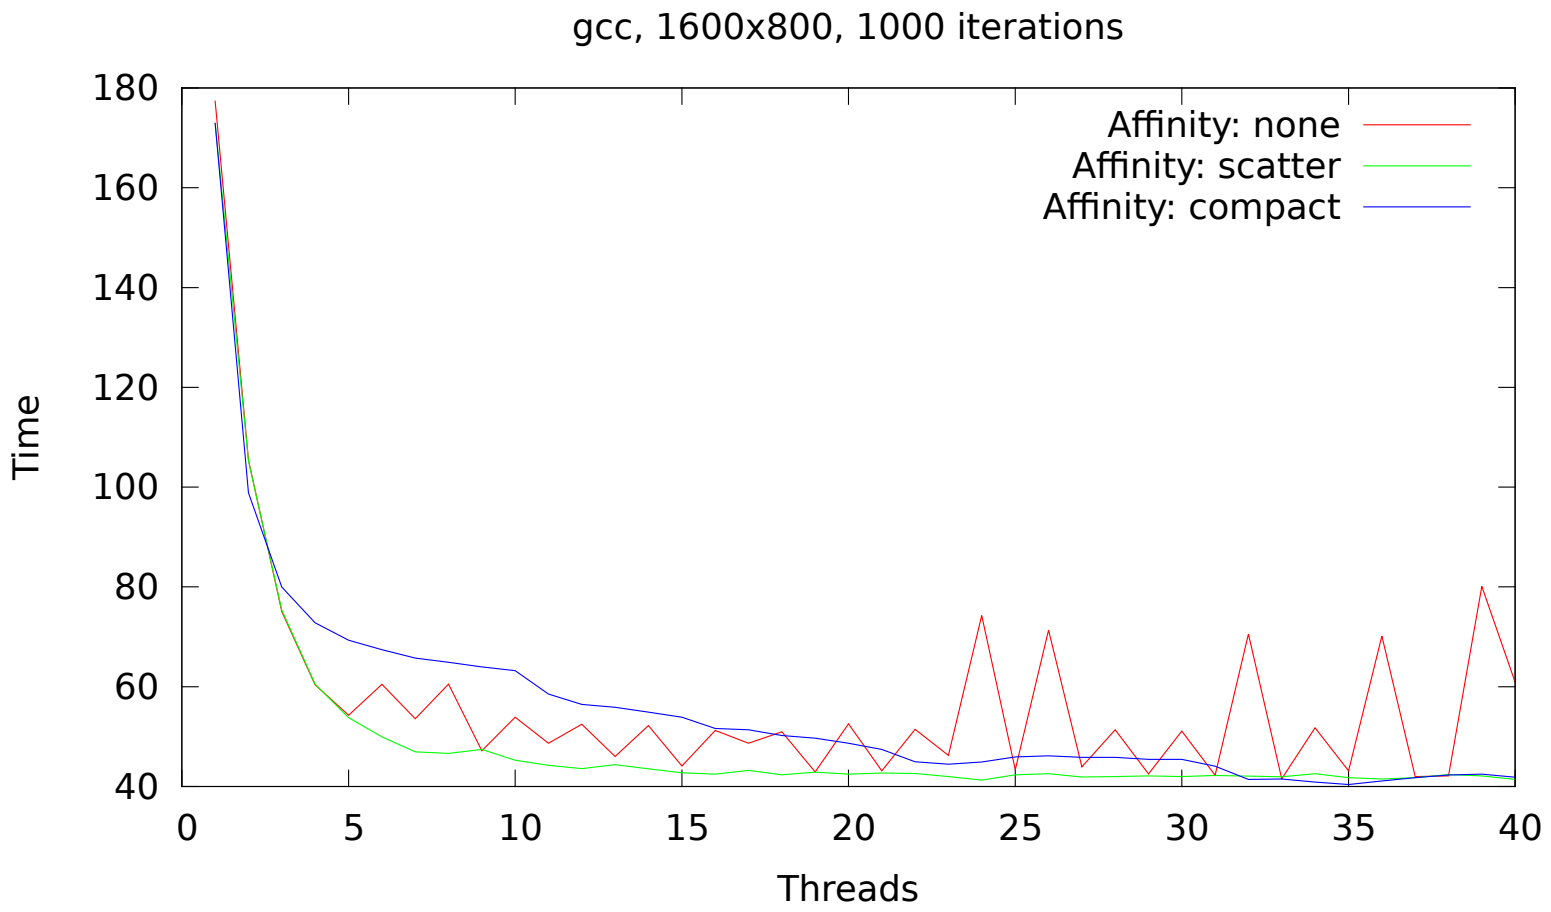

Table of Contents

Execution time of code compiled with gcc and icc

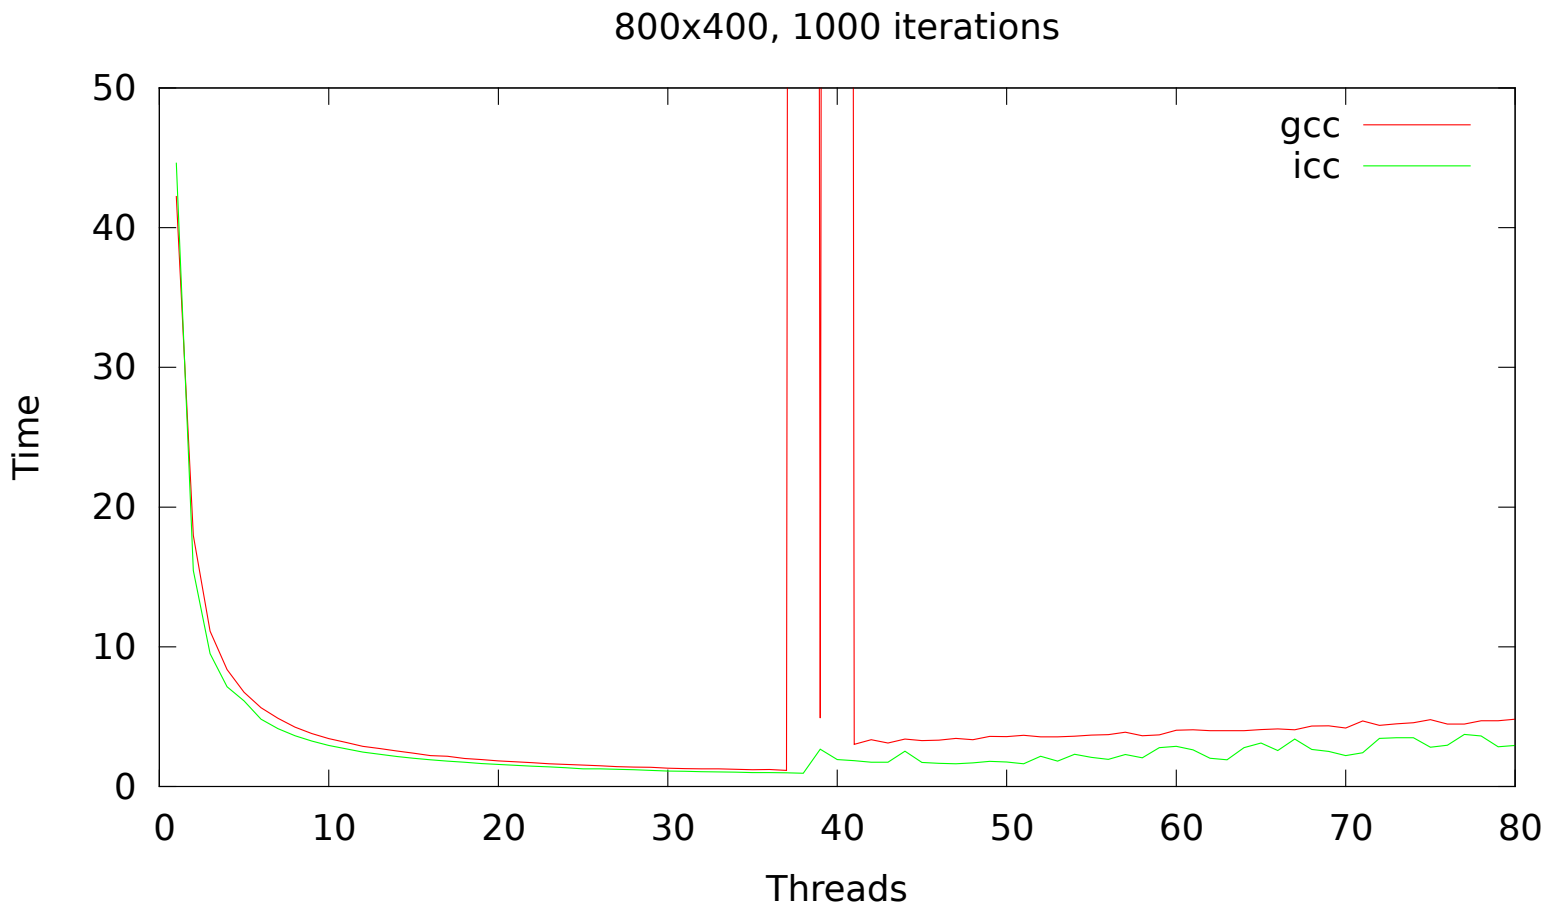

Speedup of code compiled with gcc and icc

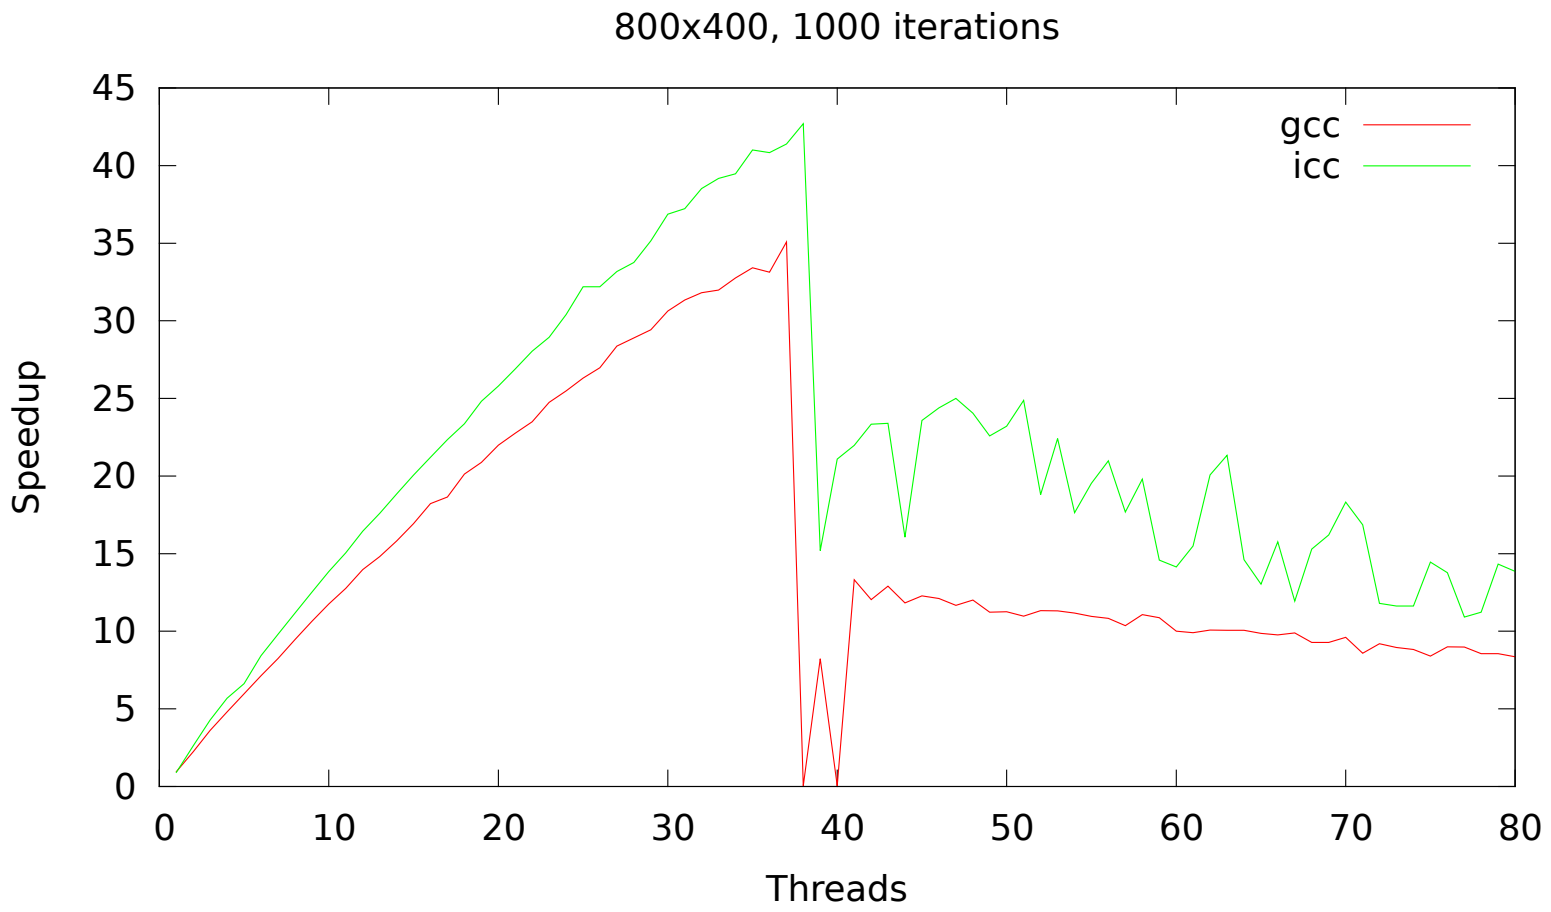

Execution time, with a subset of cores restricted via taskset

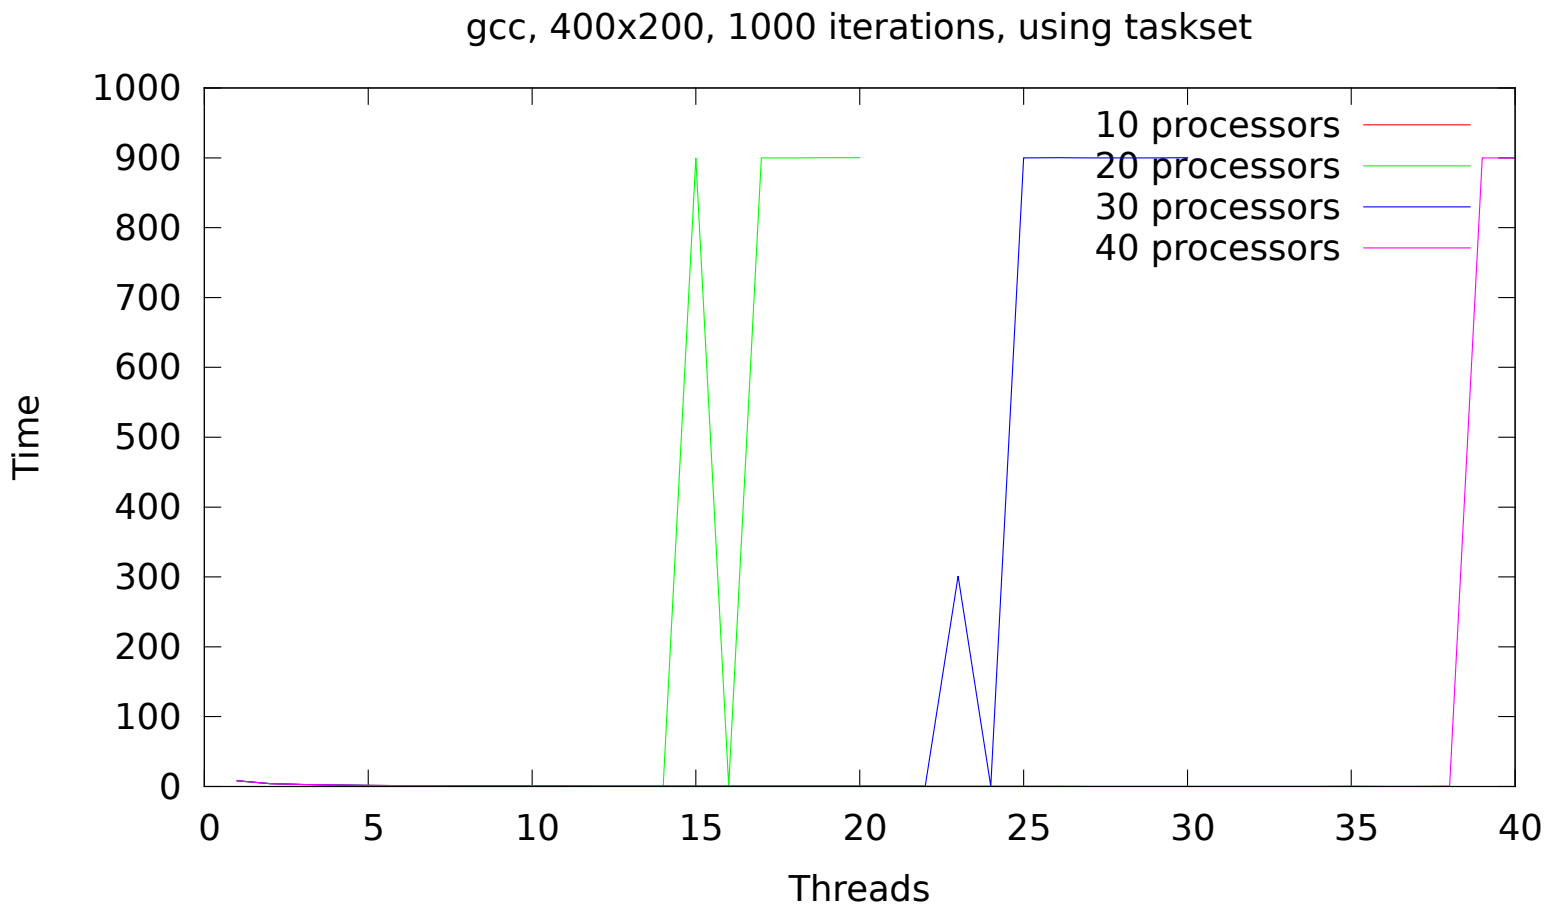

Execution times, different problem sizes

Threads	100 it. 1600x800	800x400	400x200	1000 it. 1600x800	800x400	400x200	10000 it 800x400	400x200
31	4.18	0.15	0.04	41.54	1.26	0.44	13.03	4.25
32	4.10	0.16	0.06	41.33	1.27	0.45	12.68	4.49
33	4.12	0.14	0.06	41.07	1.23	0.45	12.45	4.40
34	4.11	0.15	0.05	41.12	1.19	0.43	12.07	4.17
35	4.10	0.13	0.05	40.89	1.18	0.47	11.97	4.44
36	4.12	0.13	0.08	40.92	1.22	0.44	12.19	4.53
37	4.09	0.14	0.07	40.73	1.15	0.40	11.60	4.19
38	4.08	0.15	0.06	50.79	899.77	0.49	8999.49	4.48
39	89.87	89.90	89.97	899.86	5.94	899.82	11.87	2355.58
40	89.89	25.36	89.85	899.80	899.80	899.83	8999.38	8999.49

Table of Contents

Conclusions and future work

- GCC: There seems to be a (scheduling?) bug triggered by our code.

Conclusions and future work

- GCC: There seems to be a (scheduling?) bug triggered by our code.
- Thread Affinity

Conclusions and future work

- GCC: There seems to be a (scheduling?) bug triggered by our code.
- Thread Affinity
  - Noticeable impact on performance

Conclusions and future work

- GCC: There seems to be a (scheduling?) bug triggered by our code.
- Thread Affinity
  - Noticeable impact on performance
  - Reduces instability

Conclusions and future work

- GCC: There seems to be a (scheduling?) bug triggered by our code.
- Thread Affinity
  - Noticeable impact on performance
  - Reduces instability
  - GCC problem is not triggered when using affinity

Conclusions and future work

- GCC: There seems to be a (scheduling?) bug triggered by our code.
- Thread Affinity
  - Noticeable impact on performance
  - Reduces instability
  - GCC problem is not triggered when using affinity
- Additional analysis is needed, in particular, the effect of different scheduling algorithms.

References I

Akyil, L., Breshears, C., Corden, M., Fedorova, J., Fischer, P., Gabb, H., Gromova, V., Hoeflinger, J., Hubbard, R., Kukanov, A., O'Leary, K., Ott, D., Palmer, E., Pegushin, A., Petersen, P., Rosenquist, T., Tersteeg, A., Tsymbal, V., Voss, M., Zipplies, T.:
Intel Guide for Developing Multithreaded Application.
Intel Corporation (2011)

References II


Montes, M., Sacco, C.:
Implementación paralela de métodos de Lattice Boltzmann.
In: Primer Congreso sobre Los métodos numéricos en la enseñanza, la ingeniería y las ciencias, UTN - Facultad Regional Haedo (August 2010)


Montes, M., Sacco, C.:
Métodos de Lattice Boltzmann en equipos multicore.
In: 2do. Congreso Argentino de Ingeniería Aeronáutica, Instituto Universitario Aeronáutico (November 2010)

References III


Succi, S.:
The Lattice Boltzmann Equation for Fluid Dynamics and Beyond (Numerical Mathematics and Scientific Computation).
Numerical mathematics and scientific computation. Oxford University Press, USA (August 2001)


Wolf-Gladrow, D.A.:
Lattice-gas cellular automata and lattice Boltzmann models an introduction. 1 edn.
Springer (March 2000)