POWER CONSUMPTION OPTIMIZATION WITH PARALLEL COMPUTING

Aleksandr Igumenov, Julius Žilinskas
Vilnius University, Institute of Mathematics and Informatics

Introduction

Green computing attracts a lot of attention these days. Murugesan S. describes the field of green computing as "the study and practice of designing, manufacturing, using, and disposing of computers, servers, and associated subsystems – such as monitors, printers, storage devices, and networking and communications systems – efficiently and effectively with minimal or no impact on the environment" [3]. Green computing includes the whole arsenal of actions designed to reduce the impact of computer technology on the environment. One of the actions is optimization of electrical power consumptions.

The Green500 (www.green500.org) list provides rankings of the most energy-efficient supercomputers to raise awareness about power consumption in parallel computing. Parallel computers consume large amounts of electricity. It is important to reduce power consumption not only by design and manufacture of the efficient hardware, but also by using the hardware efficiently. Power consumption of modern multi-core computers depends on the computing load on cores. It is important to investigate this dependency and draw recommendations for the distribution of processes among computers so as to minimize the energy consumption.

Optimization is a very important task of our life. In this paper we investigate parallel optimization algorithms for two applications. One is multidimensional scaling used to analyze multidimensional data by means of visualization. Another is topology optimization of truss structures widely used in various constructions as bridges, towers and roofs. The aim of investigation is yet another application of optimization – optimization of energy consumption when solving optimization problems using parallel algorithms.

Power Consumption of Multi-core Computers

The goal of the investigation was to measure power consumption of multi-core computers with different computing load: when the computer is idle, and when one or more cores are fully loaded. We have measured power consumption of multi-core computers at clusters at Vilnius University Institute of Mathematics and Informatics. Three clusters are currently available:

- LitGrid cluster grid.mii.lt with quad-core processors Intel Quad Core Q9400 2.6 GHz running Scientific Linux and gLite middleware.
- Parallel computing cluster cluster.mii.lt with dual-core processors AMD Athlon 64 x2 Dual Core 5000 + 2.54 Ghz running Rocks cluster software.
- Parallel computing cluster hpc.mii.vu.lt with quad-core processors Intel Quad Core Q9400 2.6 GHz and Intel I5-760 2.73 GHz running Rocks cluster software.

We used electric power consumption meter PM300 to measure power consumption. Several measurements have been performed and average power consumption in Watts is shown in Table 1. It can be seen from the results that the mean power consumption per core decreases when the computing load increases. From these results recommendations can be produced to run more loaded computers than to distribute parallel tasks among separate multi-core computers.

Table 1. ***Power consumption of multi-core computers depending on load***

Load	Power consumption, W			
	Per computer	**Without idle consumption**	**Mean per core**	**Per additional core**
grid.mii.lt		Intel Quad Core Q9400		
Idle	50			
25%	80	30	30	30
50%	93	43	22	13
75%	110	60	20	17
100%	120	70	18	10
cluster.mii.lt		AMD Athlon 64 x2 Dual Core 5000+		
Idle	59			

Continued Table 1

50%	91	32	32	32
100%	115	56	28	24
hpc.mii.vu.lt		Intel Quad Core Q9400		
Idle	46			
25%	67	21	21	21
50%	81	34	17	13
75%	89	43	14	9
100%	97	51	13	7
hpc.mii.vu.lt		Intel I5-760		
Idle	50			
25%	78	28	28	28
50%	90	40	20	12
75%	105	55	18	15
100%	122	72	18	7

Energy Consumption of Parallel Computing Cluster Solving Optimization Problems Using Branch and Bound Algorithms

We have performed investigation into energy consumption of parallel computing cluster in order to compare measured energy consumption with evaluated one taking into account computing times of each process and its allocation. Two parallel branch and bound algorithms with static load balancing have been used in investigation. One is parallel branch and bound algorithm for multidimensional scaling (MDS) with city-block distances [4, 5]. Another is parallel branch and bound algorithm for topology optimization of truss structures [1, 2]. Static load balancing ensures that processes stop when they finish the work allocated to them and computing time with process allocation defines the load on each computer.

The results are presented in Table 2. We used electric power consumption meter PM300 to measure energy consumption of the whole cluster and subtracted energy consumption which would be consumed by idle cluster during the same time. We ran the algorithms on optimization problems (Cola problem in the case of MDSand 8 node structure in the case of topology optimization) and collected computing time of each process and its allocation. Using these data and measured power consumption of multi-core computers in Table 1 we have evaluated energy consumption of running cores.

Table 2. ***Energy consumption of cluster hpc.mii.vu.lt running parallel branch and bound algorithms***

Number of processes	Energy consumption, Wh		Time, s	Speed-up	Efficiency
	Measured	Evaluated			
Multidimensional scaling with city-block distances					
1	49.5	48.7	8351	1.00	1.0000
2	39.7	40.3	4239	1.97	0.9851
3	37.9	34.7	3031	2.7	0.9183
4	33.0	30.8	2235	3.74	0.9339
8	35.4	32.8	1283	6.51	0.8135
16	37.0	35.7	707	11.81	0.7383
32	42.0	39.8	398	21.01	0.6565
Topology optimization of truss structures					
1	3.7	6.0	1027	1.00	1.0000
2	3.6	5.2	577	1.78	0.8898
4	3.9	4.8	397	2.59	0.6463
8	7.9	6.2	304	3.38	0.4223
16	7.7	6.6	190	5.41	0.3379
32	7.7	8.1	111	9.29	0.2903

As it can be seen from Table 2, measured and evaluated energy consumptions are very similar and the same tendency of consumption may be seen in both numbers. We also present time of optimization, speed-up, and efficiency of parallelization. Time of optimization decreases significantly when more processes are used. Efficiency of parallelization raises some concern since it is not very high for a larger number of processes. However power consumption for these algorithms and problems is the smallest when only 4 processes are used, although it does not increase much for MDS algorithm compared to the algorithm for topology optimization.

Energy Consumption of Grid Cluster Solving Topology Optimization Problems Using Branch and Bound Algorithms

As it has been shown in the previous section, power consumption is lower when more processes are run per computer rather than distributed among more computers. However when grid middleware is used tasks may be distributed using a different strategy. In this section we investigate energy consumption of grid cluster grid.mii.lt running distributed branch and bound algorithm for topology optimization of truss structures [1]. For this purpose we ran the algorithm solving 9 node structure several times and collected computing time of each process and its allocation. Using these data and measured power consumption of multi-core computers from Table 1 we evaluated energy consumption of running cores. The results are presented in Table 3. We also present evaluated energy consumption if tasks are distributed manually packing maximum number of tasks per computer. It can be seen in the Table that middleware distribution may result in much higher energy consumption.

Table 3. ***Energy consumption of grid cluster grid.mii.lt solving problem of 9 node structure using distributed branch and bound algorithm for topology optimization of truss structures***

Number of processes	Energy consumption, kWh	
	Manual distribution	Middleware distribution
1	1.184	1.18
2	1.132	1.132-1.193
4	0.901	0.901-2.495
8	1.27	1.359-2.197
16	1.389	1.256-2.339

Conclusions

Power and energy consumptions in multi-core computers significantly depend on the load on computing cores. The measured energy consumption and evaluated energy consumption (using computing time of each process and its allocation) are very similar. The mean power consumption per core decreases when the computing load increases, therefore it can be recommended to run more loaded computers than to distribute parallel tasks among separate multi-core computers. Grid middleware distribution of tasks among computers does not take this into account which often leads to much higher energy consumption than in case of manual distribution.

Acknowledgments

The work was funded by the Research Council of Lithuania (project number MIP-108/2010).

References

1. Igumenov A., Žilinskas J., 2009, Combinatorial algorithms for topology optimization of truss structure. In: *Information technologies 2009 (IT 2009) 15th International Conference on Information and Software Technologies*. Kaunas. P. 23–24, 229–234.
2. Igumenov A., Žilinskas J., Kurowski K., Mackowiak M., 2010, Optimization of topology of truss structures using grid computing. In: *2010 IEEE Inter-national Conference on Cluster Computing Workshops and Posters (CLUS-TER WORKSHOPS)*, Heraklion, Crete, Greece, September 20–24. <doi:10.1109/CLUSTERWKSP.2010.5613101>.
3. Murugesan S., 2008, Harnessing green IT: principles and practices. *IT Professional.* No. 10 (1). P. 24–33. doi: 10.1109/MITP.2008.10.
4. Žilinskas A., Žilinskas J., 2009, Branch and bound algorithm for multidimensional scaling with city-block metric. *Journal of Global Optimization*. No. 43 (2–3). P. 357–372. <doi:10.1007/s10898-008-9306-x>.
5. Žilinskas J. Parallel branch and bound for multidimensional scaling with city-block distances. J*ournal of Global Optimization (in press)* <doi:10.1007/s10898-010-9624-7>.

POWER CONSUMPTION OPTIMIZATION WITH PARALLEL COMPUTING

Aleksandr Igumenov, Julius Žilinskas

Summary

The paper deals with power and energy consumption in multi-core systems. Power consumption of various multi-core computers with different computing load has been measured. Energy consumption of multi-core computing clusters has been evaluated taking into account computing time and allocation of each process. Measurements confirmed the evaluated estimates. Recommendations on distribution of processes among parallel multi-core computing clusters have been drawn.

Keywords: multi-core computers, power consumption, parallel branch and bound.

ELEKTROS SĄNAUDŲ OPTIMIZAVIMAS LYGIAGREČIUOSE SKAIČIAVIMUOSE

Aleksandr Igumenov, Julius Žilinskas

Santrauka

Optimizavimas šiuolaikiniame gyvenime užima labai svarbią vietą ir taikomas įvairiose srityse. Sprendžiant sudėtingus optimizavimo uždavinius, naudojami našieji superkompiuteriai, klasteriai ir skaičiuojamieji tinklai – gridai.

Pasitelkiant tokius skaičiavimo išteklius, reikia atsižvelgti į ekonominius naudojimo rodiklius. Elektros suvartojimas priklauso nuo sprendimo parametrų bei lygiagrečiųjų procesų paskirstymo. Šiame straipsnyje tiriamas elektros energijos suvartojimas daugeliu procesorių sistemose vykdant lygiagrečius skaičiavimus. Pateikiamos rekomendacijos energijos suvartojimui optimizavuoti.

Prasminiai žodžiai: našieji superkompiuteriai, energijos suvartojimas.

Įteikta 2011-11-23