7. Abbott, B. P., et al. "Einstein@ Home search for periodic gravitational waves in early S5 LIGO data." *Physical Review D* 80.4 (2009): 042003.
8. Swierczewski, Lukasz. "OProject@Home - distributed computing". Proceedings, LVEE 2013 Conference
9. Research and Infrastructure Development for Public-Resource Scientific Computing, URL: http://www.nsf.gov/awardsearch/showAward.do?AwardNumber=0221529, The National Science Foundation
10. SCI: NMI Development for Public-Resource Computing and Storage, URL: http://www.nsf.gov/awardsearch/showAward.do?AwardNumber=0438443, The National Science Foundation
11. SDCI NMI Improvement: Middleware for Volunteer Computing, URL: http://www.nsf.gov/awardsearch/showAward.do?AwardNumber=0721124, The National Science Foundation
12. Świerczewski, Łukasz. "Symulacja funkcjonalnego systemu kwantowego na równoległych komputerach klasycznych IV generacji." (2013).

Blue Gene - a parallel computer architecture

Monika Kwiatkowska and Lukasz Swierczewski

[I] *Maria Curie-Skłodowska University in Lublin,* [II] *College of Computer Science and Business Administration in Łomża, lswiercz@icm.edu.pl*

Document briefly describes the project and the Blue Gene supercomputers, Blue Gene/L , Blue Gene/P and Blue Gene/Q. In addition, performance measurement is made simple one node on the computers (Blue Gene/P and Blue Gene/Q) and other structures composed of Intel Xeon X5660 processors.

Introduction

Blue Gene is a project initiated and maintained by IBM. It aims to design a supercomputer capable of achieving speeds of PFLOPS and with low power consumption.

At present (2014) the project has created three generations of supercomputers, Blue Gene/L [1] , Blue Gene/P [2] and the Blue Gene/Q [3]. Class Blue Gene systems often occupy a leading place on the TOP500 list [4] (the fastest supercomputers) and Green500 [5] (the most energy-efficient design). Project in 2009, he received the National Medal of Technology and Innovation (National Medal of Technology and Innovation) [6].

History

In 1999, IBM announced a research project with a budget of $100,000,000 with a view to build a massively parallel computer, which was to be used to study biomolecular phenomena. The project had two objectives:

- broaden the knowledge of the mechanisms of protein folding simulations used in the context of large-scale,
- explore innovative ideas on massively parallel supercomputers and software architectures that could be implemented on them.

Originally Blue Gene project was based on an early version of the architecture Cyclops64 [7], which was designed by Monty Denneau. The initial research and development was carried out by the IBM TJ Watson

Research Center. Later. Alan Gara began work on the development of architecture QCDOC [8] to a more universal solutions .

In November 2004, the Blue Gene/L made up of 16 racks and 1024 compute nodes reaching 70.72 TFLOPS performance earned first place in the ranking of the TOP500 . Thus surpassed the NEC Earth Simulator supercomputer [9], which was the fastest in the world since 2002. Installing the Blue Gene/L in LLNL (Lawrence Livermore National Laboratory) was from 2004 to 2007 developed, and ultimately consisted of 104 cabinets. This has resulted in 478 TFLOPS performance in Linpack and 596 TFLOPS theory. Installing the Blue Gene/L in LLNL was in first place until the TOP500 ranking of 3.5 years. It was only in June 2008, this system has fallen from the first position, which took Roadrunner [10] - a solution based on Cell processors [11].

In November 2006 TOP500 list were 27 computers using the architecture of Blue Gene /L.

Blue Gene/P

In June 2007, presented the second generation supercomputer called Blue Gene/P. This solution has been designed in collaboration with LLNL IBM and Argonne National Laboratory 's Leadership Computing Facility .

Blue Gene/P is an extension of Blue Gene/L. Each Computational Blue Gene/P card consists of a quad-core PowerPC 450 processor that runs at a frequency equal to 850 MHz clock speed . The memory subsystem is composed of a small L2 cache and a shared L3 cache . The system also includes an integrated DDR2 memory controller. The system also integrates hardware logic node -to-node communication, using the same topology as the Blue Gene/L , but the use of more than double the bandwidth. A single computational card can have 2 or 4 GB of DDR2 memory. One node has a throughput of 13.6 GFLOPS. In one cabinet there are 1,024 nodes and 4,096 cores so . Blue Gene/P also had a very high energy efficiency (371 MFLOPS / W) which put him on the list of leading the construction of the TOP500 list in 2007-2008. Sample of installation Blue Gene/P Argonne National Laboratory shown in Fig. 1.

Fig. 1. Installation Blue Gene/P at Argonne National Laboratory. Source: Wikimedia Commons.

Blue Gene/Q

Third generation Blue Gene supercomputer is Blue Gene/Q. It allows a performance of 20 PFLOPS (theory) and about 17 PFLOPS in Linpack .

The Blue Gene/Q was used processors having 18 cores. PowerPC A2 uses a 64-bit architecture and runs at a clock speed of 1.6 GHz. In addition, it is supported Hyper-Threading [12] to enable the implementation on a single processor core up to four threads. In practice, however, the calculations are used for only 16 of 18 cores. Of the remaining cores one is used to support operating system functions such as interrupting, and one is used as a reserve, used to increase productivity. The system provides a theoretical yield 204.8 GFLOPS. IBM supercomputer Blue Gene/Q "Mira" is shown in Fig. 2.

Fig. 2 Blue Gene/Q at Argonne National Laboratory. Source: Wikimedia Commons.

A simple test performance application using a single node

To measure the performance of applications using a single node uses numerical calculation of the Riemann zeta function. This function is described by the following formula:

$$\zeta(z) = \sum_{n=1}^{\infty} \left(\frac{1}{n}\right)^{z}$$

Test run on a single computational node IBM Blue/P, IBM Blue Gene/Q and Intel Xeon X5660 processors. Number of iterations per second was calculated taking into account the duration of the iteration 100000000. For OpenMP parallelization technology is used [13] [14]. The results are shown in Tab. 1.

Number of threads	IBM Blue Gene/P	IBM Blue Gene/Q	Intel Xeon X5660
1	5651.95	6556.51	8459.52
2	11012.92	13102.72	16764.97
3	15985.36	19654.08	25087.54
4	21674.74	26191.72	33821.43

6	-	39308.17	50238.93
8	-	52383.44	67472.03
10	-	65316.78	84324.69
12	-	78247.26	101108.65
16	-	102669.40	-
24	-	129032.25	-
32	-	156985.87	-
48	-	200803.21	-
64	-	217864.92	-

Tab. 1. Number of iterations per second when calculating the numerical Riemann zeta function on IBM Blue Gene and Intel Xeon X5660 processors runs.

The Blue Gene/P increase in the number of iterations per second is limited to four threads. This is due to the fact that this construction has only four physical cores without Hyper-Threading.

For similar reasons, the acceleration obtained on Intel Xeon X5660 processor is limited to 12 threads.

Acknowledgment

Interdisciplinary Centre for Mathematical and Computational Modeling (ICM), Warsaw University, Poland is acknowledged for providing the computer facilities under the Grant No. G55-11.

References

1. Gara, Alan, et al. "Overview of the Blue Gene/L system architecture." IBM Journal of Research and Development 49.2.3 (2005): 195-212.
2. Kumar, Sameer, et al. "The deep computing messaging framework: generalized scalable message passing on the blue gene/p supercomputer." Proceedings of the 22nd annual international conference on Supercomputing. ACM, 2008.
3. Haring, Ruud A., et al. "The ibm blue gene/q compute chip." Micro, IEEE 32.2 (2012): 48-60.
4. Dongarra, Jack J., Hans W. Meuer, and Erich Strohmaier. "Top500 supercomputer sites." (1994).
5. Feng, Wu-chun, and Kirk W. Cameron. "The green500 list: Encouraging sustainable supercomputing." Computer 40.12 (2007): 50-55.
6. Harris, Mark (September 18, 2009). "Obama honours IBM supercomputer". Techradar
7. Del Cuvillo, Juan, et al. "FAST: A functionally accurate simulation toolset for the Cyclops64 cellular architecture." Proceedings of the Workshop on Modeling, Benchmarking and Simulation. 2005.

8. Boyle, Peter A., et al. "Overview of the QCDSP and QCDOC computers." IBM Journal of Research and Development 49.2.3 (2005): 351-365.
9. Habata, Shinichi, Mitsuo Yokokawa, and Shigemune Kitawaki. "The earth simulator system." NEC Research and Development 44.1 (2003): 21-26.
10. Barker, Kevin J., et al. "Entering the petaflop era: the architecture and performance of Roadrunner." Proceedings of the 2008 ACM/IEEE conference on Supercomputing. IEEE Press, 2008.
11. Chen, Thomas, et al. "Cell broadband engine architecture and its first implementation—a performance view." IBM Journal of Research and Development 51.5 (2007): 559-572.
12. Marr, Deborah T., et al. "Hyper-Threading Technology Architecture and Microarchitecture." Intel Technology Journal 6.1 (2002).
13. Dagum, Leonardo, and Ramesh Menon. "OpenMP: an industry standard API for shared-memory programming." *Computational Science & Engineering, IEEE* 5.1 (1998): 46-55.
14. Chandra, Rohit, ed. *Parallel programming in OpenMP*. Morgan Kaufmann, 2001.

Jade Java Agent

Kwiatkowska, M. and Swierczewski, L.

I Maria Curie-Skłodowska University in Lublin, II College of Computer Science and Business Administration in Łomża, lswiercz@icm.edu.pl

Multi-agent systems are used in situations where you have to solve the problem of a diffuse character, or a complex calculation eg search information on the web, management of telecommunications networks, air traffic control, as well as in more mundane situations, which is eg control and running appliances. Java Agent Development framework, in short JADE is an environment that supports the construction of multi-agent systems written in Java. Allows you to construct and administer agents. This publication contains basic information about agents, the criteria for their creation and standards of JADE.

Introduction

Agent is an application that runs on a device supported by the SNMP protocol. This protocol assumes the existence of two types of managed network devices, managers and managed. Application agent manages the relevant computer resources.

Foundation for Intelligent Physical Agents

An important factor in the field of agents is FIPA, Foundation for Intelligent Physical Agents registered in Geneva, Switzerland. The purpose of FIPA is to promote the success of emerging agent-based applications, services and devices, which is achieved by providing an internationally agreed specifications that maximize interoperability across applications. This is done through an open international collaboration of member organizations, which are companies and universities active in the field agent. FIPA specifications are publicly available, but these are not technology for a particular application, but