Testing the New JINR CICC Supercomputing Cluster

V.V. Korenkov, V.V. Mitsyn, E.B. Dushanov, A.S. Ayriyan, A.I. Lutsenko

Laboratory of Information Technologies, JINR

Abstract

The performance assessment for parallel computing of the new supercomputer cluster of the JINR Central Information and Computing Complex (CICC) was found by the High-Performance Linpack Benchmark. The peak performance output of 1.124 TFlops ranks the JINR CICC cluster on the 12-th place in the Top50 list of the most powerful computing systems of the Commonwealth of Independent States.

1 Introduction

The goal of this work is the performance assessment for parallel computing of the new supercomputer cluster of the JINR Central Information and Computing Complex (CICC). The peak performance obtained in the study was included in the prestigious TOP50 supercomputer list, which ranks the fastest supercomputers in the Commonwealth of Independent States (CIS) [1].

The new JINR CICC supercomputer cluster consists of three racks (two from T-Platforms and one from Hewlett-Packard), each rack holding twenty dual-processor nodes. The characteristics of each node are as follows:

Processor	Intel 2xXeon 5150
Clock rate	2660 MHz
2L cache memory per CPU	4 MB
Cores per CPU	2
CPUs per node	2
RAM per node	8 GB
Operation system	Scientific Linux 4.5
Network Interface	Gigabit Ethernet

The cluster totals:

Number of nodes	60
Number of CPUs	120
Number of cores	240
Amount of RAM	480 GB
Peak theoretical performance	2553.6 GFlops
MPI	version 1.2.7

The measurement of computer performance in solving scientific and technical floating point arithmetic problems is a question of the greatest practical interest. Its assessment by independent benchmarks allows better estimates of the available resources as well as more reliable projections on the future needs. Usually, the computational performance is measured in FLOPS (floating-point operations per second) [2].

2 High Performance LINPACK (HPL) Benchmark

LINPACK is the name of a high quality package of FORTRAN codes for solving systems of linear algebraic equations [3]. LINPACK was not specifically devised for computer performance measurements. However, the algorithms of linear algebra are widely used in various computational tasks. Therefore, the knowledge of the LINPACK processing time is quite informative for most users. The library program quality has a great impact on the efficiency of the parallel codes. That is why computer testing under LINPACK was taken as an acknowledged test for the performance measurements of the parallel computing systems. The results of LINPACK tests are used to establish the ranking of the most performing five hundred systems of the world, TOP500 [4], and the list of the fifty most performing systems of CIS, TOP50 [1]. The LINPACK benchmark test measures the processing time T needed for solving a system of linear algebraic equations using the LU-decomposition method, for which the number N_{op} of elementary arithmetic operations is known [5]. The performance is calculated with the formula:

$$P = \frac{N_{op}}{T},$$

where

$$N_{op} = \frac{2}{3} \cdot N^3 + 2 \cdot N^2,$$

with N denoting the order of the system to be solved [2].

3 Test Settings

To evaluate the system performance with the HPL Benchmark, it is necessary to build the initialization file HPL.dat [6]. The relevant fragment of this file is shown below:

```
01: HPLinpack benchmark input file
02: Innovative Computing Laboratory, University of Tennessee
03: HPL.out output file name (if any)
04: 6         device out (6=stdout,7=stderr,file)
05: 1         # of problems sizes (N)
06: 200000    Ns
07: 2         # of NBs
08: 112 120   NBs
09: 0         PMAP process mapping (0=Row-,1=Column-major)
10: 1         # of process grids (P x Q)
11: 4         Ps
12: 60        Qs
 ⋮   ⋮         ⋮
```

Here, the lines `01` through `03` are ignored. The line `04` defines the output destination. The current test asks the solution of one system of linear equations with a coefficient matrix of 200000×200000 (the lines `05` and `06`). The system of equations is solved twice (line `07`) with block decompositions of sizes 112×112 and 120×120 respectively (line `08`). This double choice allows the selection of the better partition for the final result. The task is distributed over 240 cores of the cluster as 4×60 subsets (lines `10-12` respectively). The 4×60 task option results in optimal task distribution over the existing nodes and cores. The size of the coefficient matrix is close to the total available RAM size at which peak performance is expected. The remaining parameters of the initialization files are less important and their discussion is skipped.

4 Test Results

The obtained HPL benchmark output, at a sampling of matrix sizes covering practically the whole value range at which optimal output is expected is reported in Fig. 1 and in Table 1. The dependence of both the test running time and the system performance in terms of the order N of the solved systems are given.

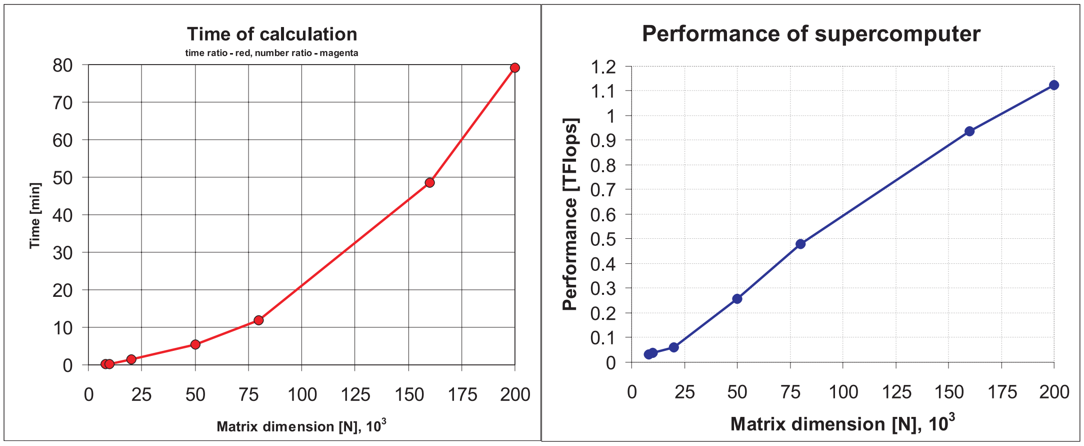

Fig. 1: Computing time and performance of JINR CICC cluster

Table 1: Running time T and performance P dependence on the matrix order N

N	T (seconds)	P (GFlops)
8000	10.59	32.24
10000	17.41	38.30
20000	87.93	60.66
50000	324.73	256.60
80000	713.75	478.20
160000	2918.21	935.70
200000	4744.29	1124.00

5 Conclusion

JINR CICC cluster was specifically acquired for distributed computing as part of the EGEE Grid infrastructure. It is worthwhile to notice that the requirements of Grid computing are very different from the low-latency requirements of the normal supercomputers [7]. Notwithstanding this fact, the test results show that the JINR CICC cluster allows effective run of parallel algorithms needed by the solution of complex scientific and technical tasks using the standard program libraries of linear algebra.

The testing results were applied to be included into the TOP50 list. The 7-th edition of the list was presented at the “Scientific Service on network INTERNET: multicore computer world” Conference on September 25, 2007 [8], and is easily available via Internet [9]. According to this list, the JINR CICC cluster ranks at the 12-th place among the most powerful computing systems within CIS (Fig. 2).

Acknowledgments. The authors thank S. Adam, Gh. Adam and V. Zhiltsov for fruitful and stimulating discussions. The work was partially supported by RFBR grant 05-01-00645-a. E. Dushanov acknowledges partial financial support from “DREAMS-ASIA”

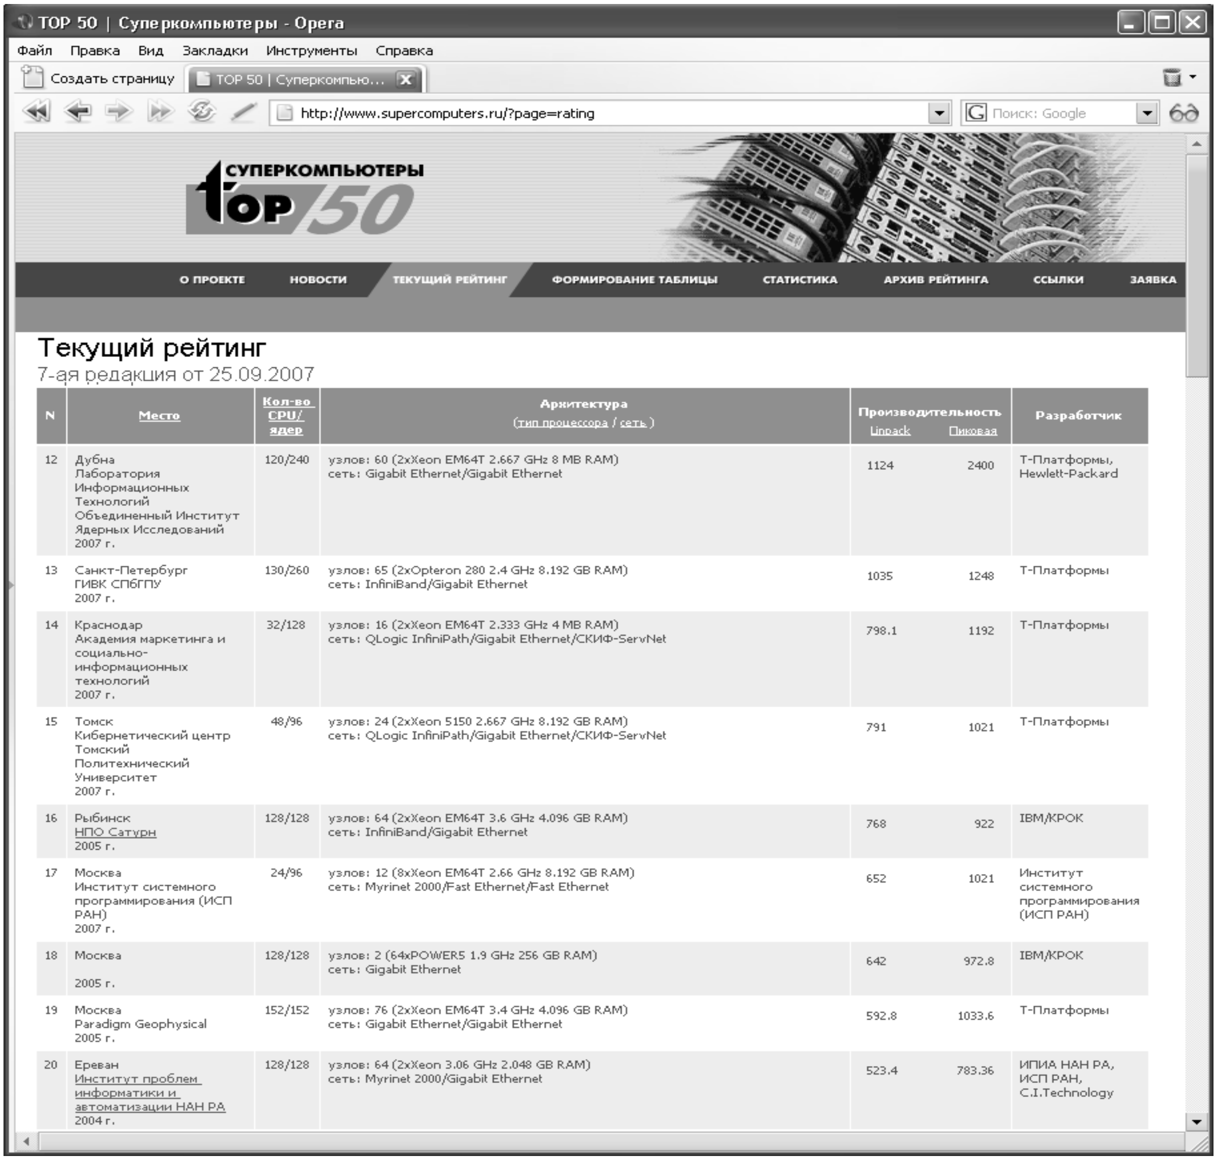

Текущий рейтинг

7-ая редакция от 25.09.2007

N	Место	Кол-во CPU/ядер	Архитектура (тип процессора / сеть)	Производительность Linpack	Производительность Пиковая	Разработчик
12	Дубна Лаборатория Информационных Технологий Объединенный Институт Ядерных Исследований 2007 г.	120/240	узлов: 60 (2xXeon EM64T 2.667 GHz 8 MB RAM) сеть: Gigabit Ethernet/Gigabit Ethernet	1124	2400	Т-Платформы, Hewlett-Packard
13	Санкт-Петербург ГИВК СПбГПУ 2007 г.	130/260	узлов: 65 (2xOpteron 280 2.4 GHz 8.192 GB RAM) сеть: InfiniBand/Gigabit Ethernet	1035	1248	Т-Платформы
14	Краснодар Академия маркетинга и социально-информационных технологий 2007 г.	32/128	узлов: 16 (2xXeon EM64T 2.333 GHz 4 MB RAM) сеть: QLogic InfiniPath/Gigabit Ethernet/СКИФ-ServNet	798.1	1192	Т-Платформы
15	Томск Кибернетический центр Томский Политехнический Университет 2007 г.	48/96	узлов: 24 (2xXeon 5150 2.667 GHz 8.192 GB RAM) сеть: QLogic InfiniPath/Gigabit Ethernet/СКИФ-ServNet	791	1021	Т-Платформы
16	Рыбинск НПО Сатурн 2005 г.	128/128	узлов: 64 (2xXeon EM64T 3.6 GHz 4.096 GB RAM) сеть: InfiniBand/Gigabit Ethernet	768	922	IBM/КРОК
17	Москва Институт системного программирования (ИСП РАН) 2007 г.	24/96	узлов: 12 (8xXeon EM64T 2.66 GHz 8.192 GB RAM) сеть: Myrinet 2000/Fast Ethernet/Fast Ethernet	652	1021	Институт системного программирования (ИСП РАН)
18	Москва 2005 г.	128/128	узлов: 2 (64xPOWER5 1.9 GHz 256 GB RAM) сеть: Gigabit Ethernet	642	972.8	IBM/КРОК
19	Москва Paradigm Geophysical 2005 г.	152/152	узлов: 76 (2xXeon EM64T 3.4 GHz 4.096 GB RAM) сеть: Gigabit Ethernet/Gigabit Ethernet	592.8	1033.6	Т-Платформы
20	Ереван Институт проблем информатики и автоматизации НАН РА 2004 г.	128/128	узлов: 64 (2xXeon 3.06 GHz 2.048 GB RAM) сеть: Myrinet 2000/Gigabit Ethernet	523.4	783.36	ИПИА НАН РА, ИСП РАН, C.I.Technology

Fig. 2: The JINR CICC cluster occupies 12-th place at the TOP50 list

(Development of gRid EnAbling technology in Medicine&Science for Central ASIA) NATO Grant EAP.NIG 982956.

References

[1] http://www.supercomputers.ru/

[2] Voevodin V.V., Voevodin Vl.V. Parallel computations. - SPb.: BHB-Peterburg, 2002, pages 162-178.

[3] http://www.netlib.org/linpack/

[4] http://www.top500.org/

[5] http://www.netlib.org/benchmark/hpl/

[6] http://www.netlib.org/benchmark/hpl/tuning.html/

[7] H. Hämmerle, N. Crémel, Computing News and Feature: CERN makes it into supercomputing TOP500, CERN COURIER, volume 47, Number 7, September 2007, p. 16.

[8] Voevodin Vl.V. Seventh reduction of high-performance computers of CISs TOP50 list: quickly, more, but whether be better? All-Russian scientific Conference "Scientific Service on network INTERNET: multicore computer world. 15 years of RFBR" Abrau-Durso, September 24-29, 2007.

[9] http://www.supercomputers.ru/?page=rating/