Performance assessment of the 2008 configuration of the JINR CICC cluster

A. Ayriyan[1,*], Gh. Adam[1,2], S. Adam[1,2], E. Dushanov[1], V. Korenkov[1], A. Lutsenko[1], V. Mitsyn[1]

[1] *Laboratory of Information Technologies, Joint Institute for Nuclear Research, 141980, Dubna, Russia*
[2] *Horia Hulubei National Institute for Physics and Nuclear Engineering, IFIN-HH, Bucharest, Romania*
* *E-mail:* ***ayriyan@jinr.ru***

Work is partly supported by grants **RFBR #08-01-00800** and **RFBR #08-07-09240**

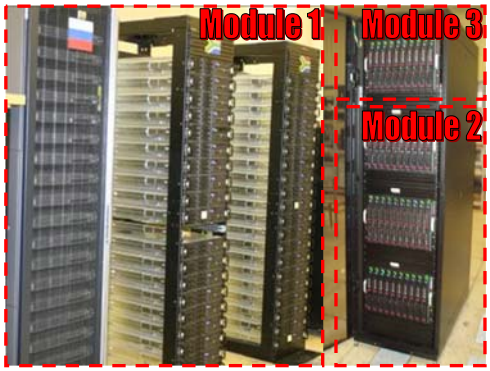

Fig. 1. The three modules of JINR CICC cluster.

INTRODUCTION

JINR covers a broad multidisciplinary spectrum of research topics and LIT does mathematical modeling and software development for that.
Among them there are two kind of problems
~ very large single problems (e.g. lattice chromodynamics),
~ very large number of different problems of similar nature (like in LHC experiment).
Central Information and Computing Complex (CICC) has to optimally implement the acquired hardware and open software sources.

MOTIVATION

~ Research of efficiency of the software implementation and optimization of hardware modules connection.
~ Finding and shooting weak points.
~ Finding way for next upgrade.
~ To make proposal for Top50 (www.supercomputers.ru) and to make comparison with Top500 (www.top500.org) data.

Provider	[Totals]	T-Platfor.&HP	SuperMicro	SuperMicro	
Structure	Heterogeneous	Homogeneous	Homogeneous	Homogeneous	
Year	2007-2008	2007	2008	2008	
CPU/cores	200/560	120/240	60/240	20/80	
Intel Processor	–	Xeon 5150 (dual-core)	Xeon E5430 (quad-core)	Xeon X5450 (quad-core)	
Frequency, MHz	–	2660	2660	3000	
RAM, Gb	1120	480	480	160	
Network	Gigabit Ethernet	Gigabit Ethernet	Gigabit Ethernet	Gigabit Ethernet & InfiniBand	
Performance Linpack/Peak, GFlops	2982 / 6067.2	1325 / 2511.04	1414 / 2553.6	GbE	IfB
				598.4/960	757.5/960
Efficiency Linpack/Peak	0.491	0.528	0.554	GbE	IfB
				0.623	0.797
Coefficient matrix order	274400	220000	220000	GbE	IfB
				120000	120000
RAM occupancy by coefficient matrix	0.501	0.751	0.751	GbE	IfB
				0.671	0.671
Basic aim of computing	Distributed	Distributed	Distributed	Distributed	
	Parallel			Parallel	

Table 1. Summary of system and performance characterization.

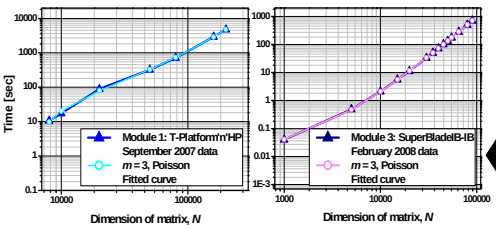

Fig. 2. Example of the data fitting, fitted curve (polynomial with degree *m*) to February 2008 Module 1 (left) data and Module 3 (right) data.

CONSISTENCY TIME DATA CHECK

The Top500 High-Performance Linpack benchmark (HPL) was used for performance assessment. The consistency of the measured times with the computational complexity of the *LU*-decomposition involved in the HPL was checked by solving a least squares problem. This analysis yields a third power law of the increase of the computing time with *N*, the order of the solved linear algebraic system.

Detailed description of the *lsq*-problem was given in:
Gh.Adam, S.Adam, A.Ayriyan, E.Dushanov, E.Hayryan, V.Korenkov, A.Lutsenko, V.Mitsyn, T.Sapozhnikova, A.Sapozhnikov,O.Streltsova, F.Buzatu, M.Dulea, I.Vasile, A.Sima, C.Visan, J.Busa, I.Pokorny, "Performance assessment of the SIMFAP parallel cluster at IFIN-HH Bucharest", Romanian Journal of Physics, Vol. 53, Nos. 5-6, P. 665-677, 2008.

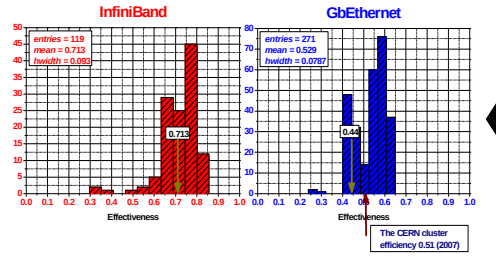

Fig. 3. Comparison of former efficiencies of Modules 1 and 3 (arrows) with NOV'07 TOP500 data (histograms)

FINDING AND SHOOTING WEAK POINTS

The information transfer is secured by one Gigabit Ethernet (GbE) in each machine, while every module connect to the main Backbone Ethernet switch via four-ports GbE trunk, so aggregated in-between modules rate was upgraded up to 4 Gbps.

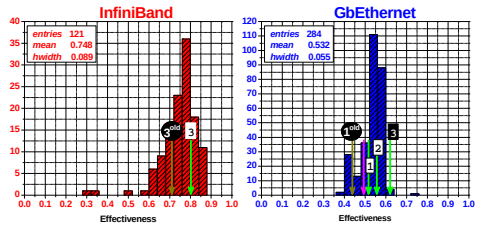

Fig. 4. Comparison of the last and present efficiencies of Modules 1 to 3 (arrows) with JUN'08 TOP500 data (histograms)

Modules	Network	R_{peak} [Gflops]	R_{max} [Gflops]	Efficiency	Past Eff.
Module 1	Gigabit Ethernet	2553.6	1325.0	0.528	0.44
Module 2	Gigabit Ethernet	2553.6	1414.0	0.554	–
Module 3	Gigabit Ethernet	960	598.4	0.623	–
	InfiniBand	960	757.5	0.797	0.713
Sum	Gigabit Ethernet	6067.2	3337.4	–	–
CICC Cluster	Gigabit Ethernet	6067.2	2982.0	0.491	0.44

Table 2. Table of the current results.

EFFICIENCY IN RDIG

JINR CICC shows roughly 40 percent of the active resources involved in Tier2 RDIG.

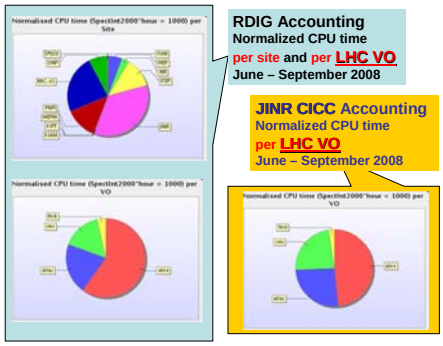

Fig. 5. Normalized CPU time per site and per LHC VO

EGEE and OSG SITES:
All VOs, Normalised CPU time by SITE
June 2008 - September 2008

1.	USCMS-FNAL-WC1-CE	5 367 192
2.	FZK-LCG2	3 795 800
3.	GRIF	3 020 511
4.	NIKHEF-ELPROD	2 956 977
5.	IN2P3-CC	2 586 393
6.	IN2P3-CC-T2	2 577 241
7.	TRIUMF-LCG2	2 408 040
8.	GLOW	2 317 754
9.	BNL_ATLAS_1	2 172 275
10.	CERN-PROD	1 949 913
11.	JINR-LCG2	1 769 925

TOTAL SITES: 263 EGEE and OSG sites

Statistics obtained from the EGEE Accounting Portal:
http://www3.egee.cesga.es/gridsite/accounting/CESGA/egee_view.html

Fig. 6. List of top sites sorted by normalized CPU time.

CONCLUSIONS

The highest measured performance of the CICC cluster provided entries to the Top50:
- 12th rank within the 7th edition SEP'07,
- 23rd rank within the 8th edition MAR'08,
- 24th rank within the 9th edition SEP'08.

We have found that low cost improvements of the hardware connections resulted in substantial gains concerning the latency decrease and the consistent data handling to/from the associated mass storage peripherals. The increase of the RAM/core also resulted in sensible performance gains too.

With these improvements, the efficiency of the high performance computing done on each of the three homogeneous modules of the LIT-JINR cluster stays at the upper end of the reported statistics on the TOP500 [4] data. Therefore, the home made free software installation is consistent with the similar works worldwide.

As it concerns the efficiency of the distributed computing of interest in Grid data analyses, it also benefited from these additions since the improved rack interconnect is straightforwardly reflected in the speed of this kind of computations.

Fig. 7. Entries of the CICC cluster to the Top50 (in Russian)