Письма в ЭЧАЯ. 2009. Т. 6, № 3(152). С. 407–417

JINR CICC IN COMPUTATIONAL CHEMISTRY AND NANOTECHNOLOGY PROBLEMS: DL_POLY PERFORMANCE FOR DIFFERENT COMMUNICATION ARCHITECTURES

E. Dushanov [a,b], *Kh. Kholmurodov* [a], *G. Aru* [a], *V. Korenkov* [a], *W. Smith* [c], *Y. Ohno* [d], *T. Narumi* [d], *G. Morimoto* [d], *M. Taiji* [d], *K. Yasuoka* [e]

[a] Joint Institute for Nuclear Research, Dubna
[b] Institute of Nuclear Physics, Tashkent, Uzbekistan
[c] Daresbury Laboratory, Warrington, United Kingdom
[d] High-Performance Computing Team, Integrated Simulation of Living Matter Group, Computational Science Research Program, RIKEN (The Institute of Physical and Chemical Research), Yokohama, Kanagawa, Japan
[e] Keio University, Yokohama, Japan

This report compares the performance of the DL_POLY general-purpose molecular dynamics simulation package on the LIT JINR computing cluster CICC with various communication systems. The comparison involved two cluster architectures: Gigabit Ethernet and InfiniBand technologies, respectively. The code performance tests include some comparison of the CICC cluster with the special-purpose computer MDGRAPE-3 developed at RIKEN for a high-speed acceleration of the MD (molecular dynamics) without a fixed cutoff. The DL_POLY benchmark covers a set of typical MD system simulations detailed below.

В работе сравнивается производительность пакета многоцелевого использования для молекулярно-динамического моделирования DL_POLY на вычислительном кластере ЦИВК ЛИТ ОИЯИ с различными сетевыми технологиями. Сравнение кодов DL_POLY проводилось для двух видов кластерных архитектур на основе сетевых технологий Gigabit Ethernet и InfiniBand. Тестирование кода включает сравнение кластера ЦИВК с компьютером специального назначения MDGRAPE-3, разработанным в RIKEN для высокопроизводительных МД (молекулярно-динамических) вычислений без использования радиуса обрезания дальнодействующих межмолекулярных сил. Тестирование производительности пакета DL_POLY охватывает типичные молекулярно-динамические вычисления, описанные ниже.

PACS: 31.15.A-; 31.15.at

INTRODUCTION

A modern state-of-art computational study of a molecular system is generally undertaken using classical molecular simulation or quantum chemistry methods, or a hybrid of these two methods. The methods of molecular simulation (conventional or hybrid molecular dynamics (MD), Monte-Carlo (MC), *ab initio* quantum-chemistry, etc.) of large molecular systems, first proposed more than 50 years ago with the advent of computers have shown a surge

of development in the last decade. With the creation of new parallel/vector supercomputers and special-purpose computer clusters, the molecular simulation methods are now a powerful tool in bioengineering, nanotechnology and material science, capable of revealing details of processes at the atomic scale and describing technologically important phenomena. Molecular simulation represents a practical tool in the development of new materials and new drugs, performing large-scale simulations of molecular complexes composed of hundreds, thousands or multi-millions of particle systems. Thus, computer molecular simulations play a fascinating role in fundamental physics, biochemical and life sciences, having an increasingly significant impact on many applied industries and modern nanotechnology. Molecular simulations (conventional and hybrid MD) exploit classical Newtonian physics laws to model the particle interaction in molecules via force fields that are defined in advance, given empirically, or calculated by other methods. In MD studies the molecular systems are modeled deterministically by the integration of classical equations of motion and in MC by stochastic processes linked to various ensembles. The MD methods are capable of modeling atomic molecular systems up to thousands and even millions of particles, and simulating many system properties and environmental configurations. The MD simulation enables to efficient prediction of ensemble properties and behaviors, such as P–V–T relations, phase equilibria, transport properties, structures of synthetic and biological macromolecules, docking of one molecule against another, etc.

Other aspects of molecular science make use of computational quantum chemistry methods. The quantum chemistry studies — *ab initio* density functional theories (DFT) and others, in contrast to the conventional molecular simulations, are based on quantum mechanics. The computational quantum chemistry methods primarily serve to calculate electronic structures of atoms or molecules, yielding wave functions or probability density functionals of electron states. The quantum chemistry methods provide greater accuracy but are restricted to smaller molecular sizes because of the complexity of the models and computational costs. The quantum chemistry simulation is essential when chemical bonds are created or broken. It is also used when force parameters are unknown or not applicable. DFT methods are well established and used with increasing accuracy; high-level wave function methods with large atomic-orbital basis sets are currently standard. As a result of quantum chemistry studies, we have at hand force-field data, and are able to calculate thermochemistry, kinetics, optical properties, NMR shifts, etc. [1–3].

The aim of the present work is to perform benchmark simulations on Central Information and Computing Complex (CICC), two LIT JINR computing clusters with the DL_POLY code. DL_POLY is a general-purpose serial and parallel molecular dynamics simulation package developed at Daresbury Laboratory by W. Smith et al. [4, 5]. The Molecular Simulation Group (MSG) at Daresbury Laboratory developed the original package DL_POLY under the auspices of the Engineering and Physical Sciences Research Council (EPSRC) for the EPSRC's Collaborative Computational Project for the Computer Simulation of Condensed Phases (CCP5). The Natural Environment Research Council through the eMinerals project also supported later developments. The package is the property of the Central Laboratory of the Research Councils.

We have used the DL_POLY package at LIT JINR under a joint collaboration agreement aimed to compare the code performance in CICC (lxpub01 – lxpub04) for different communication environments. At the time of testing, the CICC consisted of 70 nodes (60 HP and T-platform architecture nodes and 10 SuperBlade nodes) that employ different commu-

nication tools. The communication networks that operate within computing cluster CICC are represented by the Gigabit Ethernet and InfiniBand Switch modules [6].

Workstations that have been benchmarked include those from Hewlett Packard with the Intel 2xXeon 5150 CPUs, the former in the C2660 (2.66 GHz) [6, 7] in the first case (cluster I) and SuperBlade with the Intel Xeon 5300 CPUs [8] in the second case (cluster II).

The characteristics of the computing nodes of each cluster are given in Table 1:

Table 1. **The characteristics of computing nodes of each cluster**

	Cluster I	Cluster II
Processor	Intel 2xXeon 5150	Intel Xeon 5300
Clock rate, MHz	2660	3000
2L cache memory per CPU, MB	4	8
Cores per CPU	2	4
CPUs per node	2	2
RAM per node	8	16
Operation system	Scientific Linux 4.5	Scientific Linux 4.5
Network Interface	Gigabit Ethernet	Gigabit Ethernet, InfiniBand
Totals:		
Number of nodes	60	10
Number of CPUs	120	20
Number of cores	240	80
Amount of RAM, GB	480	80
Peak theoretical performance, Gflops	2553.6	960
MPI	Version 1.2.7	Version 1.2.7, Open MPI 1.2.5

We should stress that our access to part of the hardware resources at CICC was not optimized to the full. The job run often involved the temporary loan or donation of the machine units as part of the hardware evaluation exercises run at the LIT cluster. In many cases the CICC machines were not optimally configured in terms of either job queue or access to the high-speed disk memories. Thus, the results presented in this report should be overestimated. The following Secs.1 and 2 present results of DL_POLY performance.

The DL_POLY simulation programs we used are represented by the versions 2.17 and 3.07 of the code. DL_POLY_2 is the original version which has been parallelized using the Replicated Data strategy and is useful for simulations of up to 30,000 atoms on 100 processors. DL_POLY_3 is a version which uses Domain Decomposition to achieve parallelism and is suitable for simulations of order 1 million atoms on 8-1024 processors [4, 5].

1. THE DL_POLY BENCHMARK

The benchmark summarized below is designed to reflect the typical range of simulations undertaken by the molecular dynamics.

The benchmark simulations with DL_POLY_2.17 were:

— Benchmark 3 (Test3): Simulation of valinomycin in 1223 water molecules (3837 atoms, 100 time steps);

— Benchmark 5 (Test5): Dynamic Shell model $MgCl_2$ structure (768 atoms, 1280 sites, 1000 time steps);

— Benchmark 9 (Test9): Simulation of a model membrane with 2 membrane chains, 202 solute molecules and 2746 solvent molecules (3148 atoms, 1000 time steps).

The benchmark simulations with DL_POLY_3.07 were:

— Benchmark 1 (Test1): Simulation of sodium chloride with Ewald sum (27 000 ions);

— Benchmark 2 (Test2): Simulation of sodium chloride with Ewald sum (216 000 ions);

— Benchmark 4 (Test4): DMPC in water (413 896 atoms).

Below in Table 2 we present the basis parameters of these molecular systems.

Table 2. **The basis parameters of testing molecular systems**

Parameters	Test1	Test2	Test3	Test4	Test5	Test9
Temperature	500.0	500.0	310.0	300.0	1200.0	300.0
Equilibration steps	200	100	100	30	100	100
Steps	200	100	500	30	500	500
Time step	0.001	0.001	0.002	0.001	0.0005	0.001
Cutoff	12.0	12.0	12.0	10.0	12.0	9.0
Ewald precision	10^{-6}	10^{-6}	—	10^{-5}	10^{-5}	—
Job time	100	400	6000	600	3600	43 200

In Figs. 1–5 the snapshots of each molecular system are shown, respectively.

Gigabit Ethernet Type Network. In Table 3 we present the simulation results for the code DL_POLY_2.17 on cluster I. It is seen that the efficiency of DL_POLY_2.17 in cluster I is linear at 8 CPUs with Leap-Frog algorithm and 4 CPUs with Velocity-Verlet algorithm.

In Table 4 we present the simulation results for DL_POLY_2.17 molecular dynamics code on cluster II. It is seen that the efficiency of DL_POLY_2.17 in cluster I behaves linearly at 8 CPUs with Leap-Frog algorithm and 4 CPUs with Velocity-Verlet algorithm. The simulation time of cluster II is a factor of two less than the simulation time of cluster I with Benchmark 3.

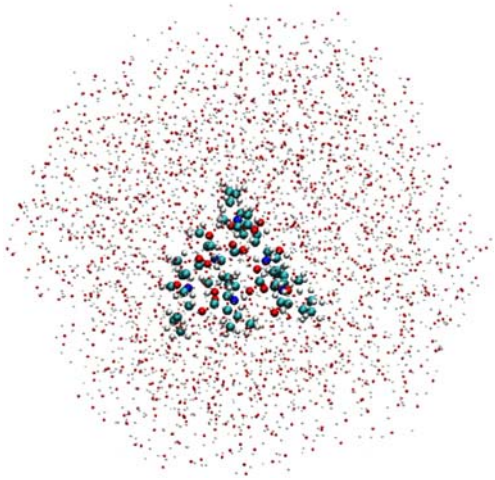

Fig. 1. The snapshot of initial configuration of valinomycin (central chain) surrounded by water molecules (molecular system of Benchmark 3)

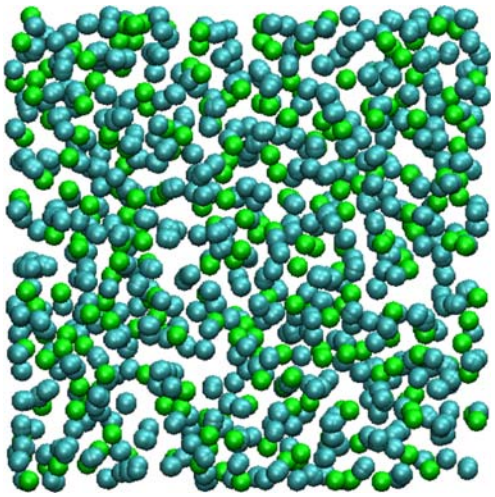

Fig. 2. The snapshot of initial configuration of $MgCl_2$ model structure (molecular system of Benchmark 5)

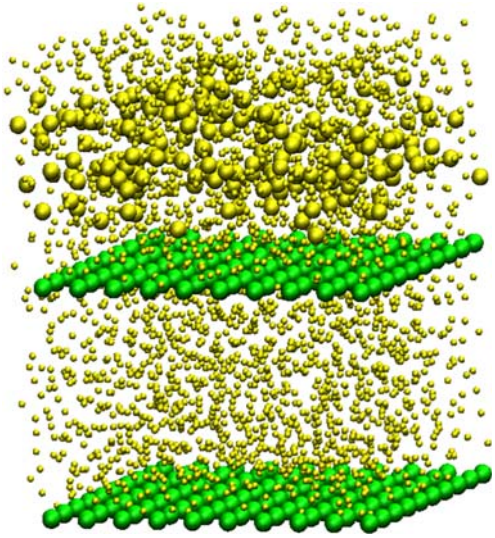

Fig. 3. The snapshot of a model membrane with 2 membrane chains, solute and solvent molecules (molecular system of Benchmark 9)

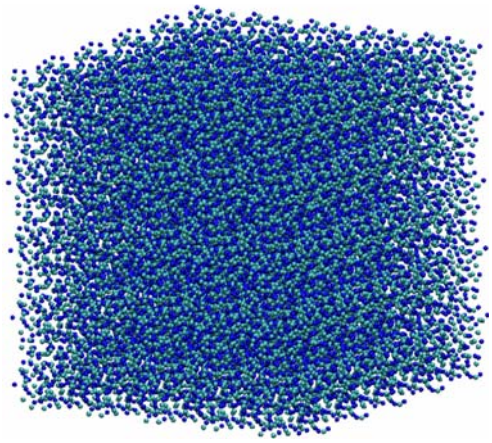

Fig. 4. The snapshot of initial configuration of sodium chloride structure (molecular system of Benchmark 1)

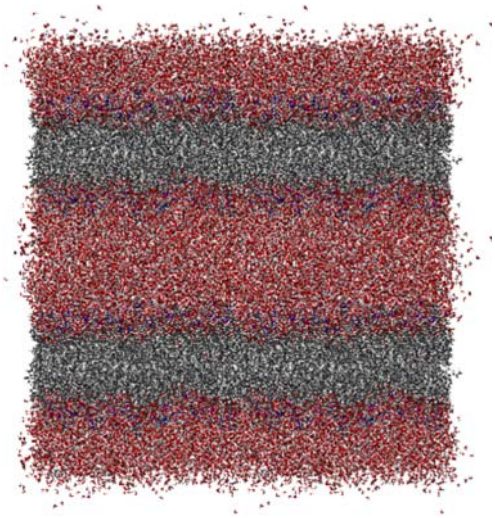

Fig. 5. The snapshot of initial configuration of DMPC in between water slabs (molecular system of Benchmark 4)

In Fig. 6 we present the simulation results for DL_POLY_2.17 code on clusters I and II, respectively.

Table 3. **Results of simulations with DL_POLY_2.17 molecular dynamics code on cluster I**

	Leap-Frog algorithm			Velocity-Verlet algorithm		
Proc.	Test3	Test5	Test9	Test3	Test5	Test9
1	138.104	16.884	56.597	140.137	17.105	57.729
2	80.651	10.028	34.800	91.131	10.188	41.056
4	63.556	8.400	24.664	79.881	7.950	30.195
8	59.639	8.791	34.277	87.545	18.731	40.985
16	82.990	42.255	82.891	151.694	21.654	48.105

Table 4. **Results of simulations with DL_POLY_2.17 molecular dynamics code on cluster II**

	Leap-Frog algorithm			Velocity-Verlet algorithm		
Proc.	Test3	Test5	Test9	Test3	Test5	Test9
1	95.246	11.352	39.399	96.866	11.410	40.216
2	49.849	5.945	23.037	53.382	5.974	24.235
4	29.122	3.514	13.091	33.310	3.490	15.051
8	20.933	2.678	9.244	26.246	2.488	11.846
16	27.817	4.247	13.064	43.330	4.148	19.116
32	36.404	5.465	18.452	59.903	5.175	28.177

In Table 5 we present the simulation results for DL_POLY_3.07 code on clusters I and II, respectively. It is seen that the efficiency of DL_POLY_3.07 in cluster II behaves linearly with the growing number of processors.

Table 5. **Results of simulations with DL_POLY_3.07 molecular dynamics code**

	Cluster I			Cluster II		
Proc.	Test1	Test2	Test4	Test1	Test2	Test4
1	92.375	398.996	622.133	92.043	398.539	621.152
2	92.180	394.902	631.488	91.551	395.320	621.473
4	91.816	409.262	546.945	65.840	361.363	429.930
8	77.445	412.891	410.809	54.387	299.953	356.973
16	91.504	410.324	389.113	46.805	254.516	267.062
32	95.516	400.723	418.082	41.930	207.062	235.129
64	94.859	521.516	723.336	37.930	173.664	260.148

In Fig. 7 we present the simulation results for DL_POLY_3.07 code on clusters I and II, respectively.

InfiniBand Type Network. In Table 6 we present the simulation results for the code DL_POLY_2.17 on cluster II. It is seen that the efficiency of the code behaves linearly for 32 CPUs.

In Table 7 we present the simulation results for DL_POLY_3.07 code on cluster II.

In Fig. 8 we present the simulation results for DL_POLY_2.17 and DL_POLY_3.07 codes on cluster II with InfiniBand type network. It is seen that the efficiency of the codes behaves linearly for 32 CPUs.

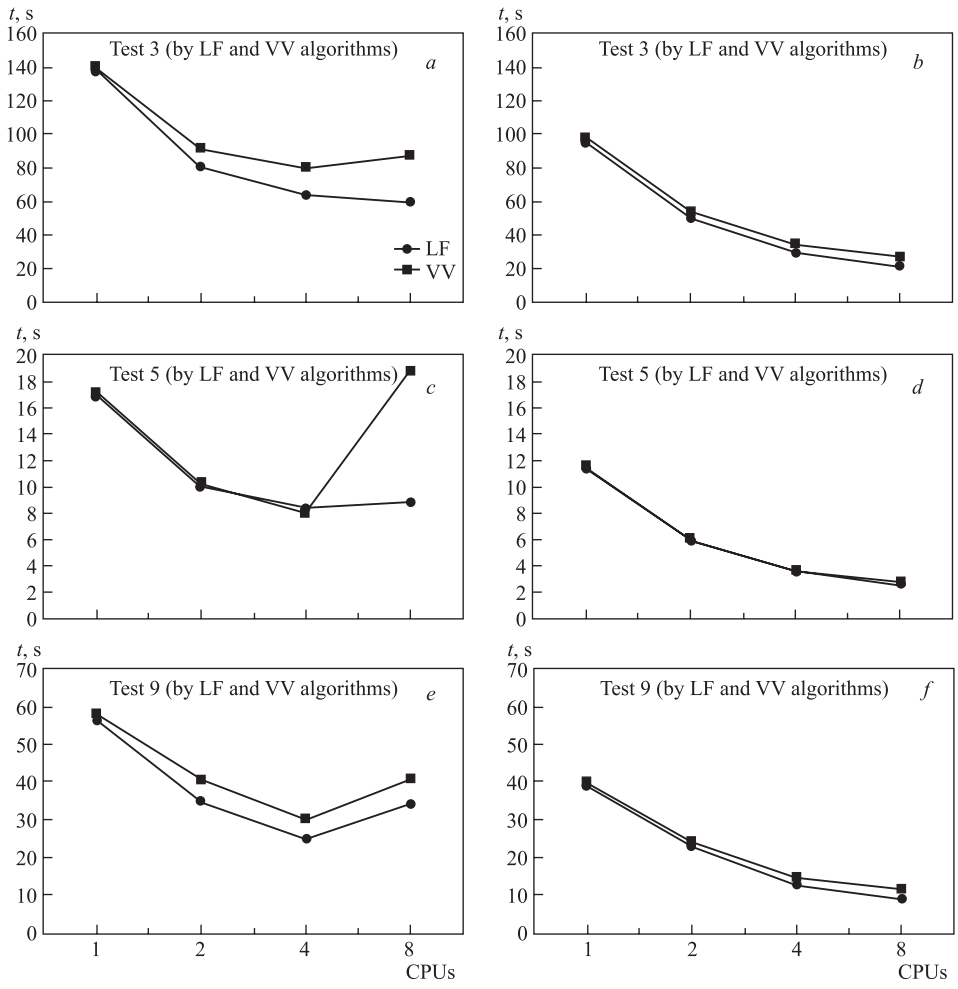

Fig. 6. Comparison results of simulations on cluster I (a, c, e) and cluster II (b, d, f)

Table 6. **Results of simulations with DL_POLY_2.17 molecular dynamics code on cluster II**

	Leap-Frog algorithm			Velocity-Verlet algorithm		
Proc.	Test3	Test5	Test9	Test3	Test5	Test9
1	95.246	11.352	39.399	96.886	11.410	40.216
2	49.377	5.705	22.293	51.146	5.755	23.206
4	25.958	3.076	12.039	27.666	3.087	12.762
8	15.596	1.863	6.978	17.164	1.940	8.043
16	10.134	1.402	5.040	12.845	1.370	5.997
32	8.942	1.365	4.594	12.672	1.369	6.346

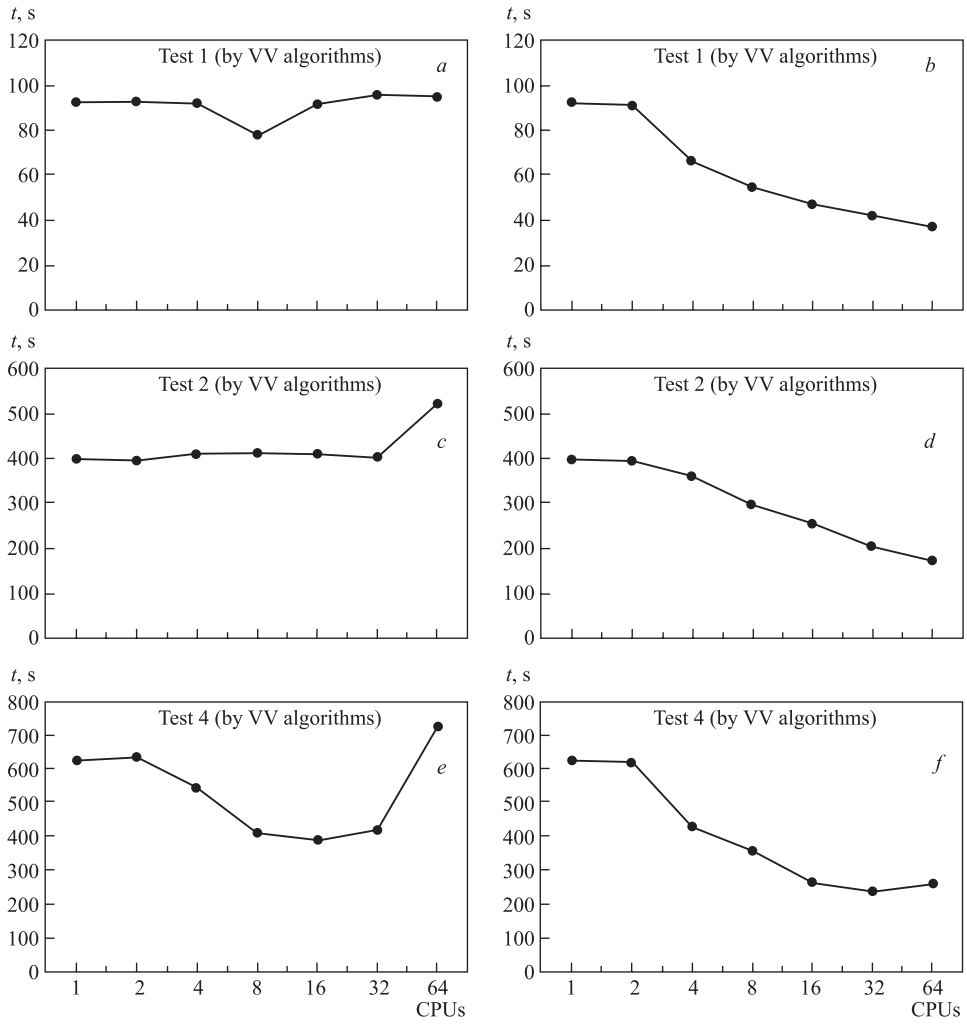

Fig. 7. Comparison results of simulations on cluster I (a, c, e) and cluster II (b, d, f)

Table 7. **Results of simulations with DL_POLY_3.07 molecular dynamics code**

Proc.	Test1	Test2	Test4
1	92.043	398.539	621.152
2	91.492	395.017	602.887
4	69.789	321.717	368.254
8	56.672	263.068	316.193
16	21.991	149.218	182.501
32	36.774	83.925	129.092
64	36.900	66.379	108.286

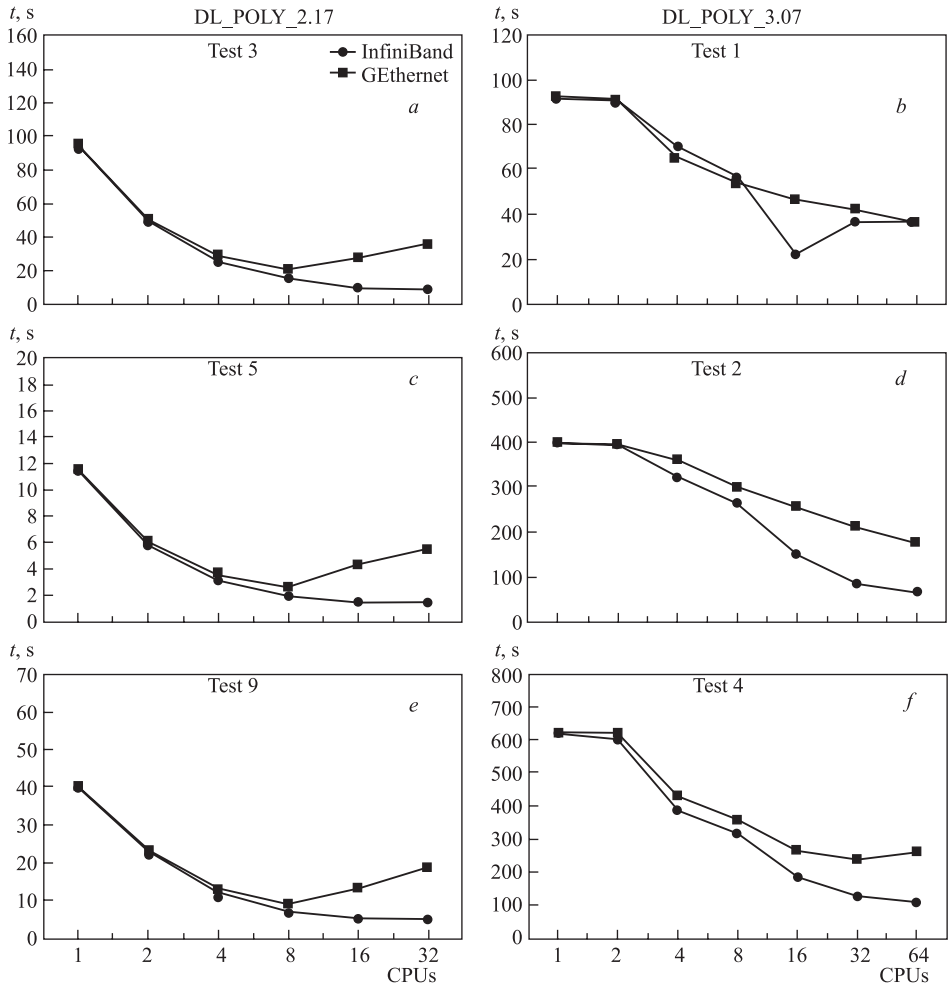

Fig. 8. Results of simulations on cluster II with InfiniBand type of network

2. PERFORMANCE OF MDGRAPE ARCHITECTURE

The architecture of MDGRAPE-3 is quite similar to its predecessors, the GRAPE (GRAvity PipE) systems. The GRAPE systems are special-purpose computers for gravitational N-body simulations and molecular dynamics simulations developed at University of Tokyo. Its direct predecessor MDM (Molecular Dynamics Machine) developed by RIKEN achieved 78 Tflops peak performance in 2001 and MDGRAPE-3 reached 1 Pflops in 2006 (see Fig. 9). The GRAPE systems won eight Gordon Bell prizes in total (http://sc08.supercomputing.org /?pg=awards.html).

The MDGRAPE-3 system consists of a host computer and special-purpose engines. The special-purpose engine calculates non-bonding forces (i.e., electrostatic and intermolecular forces), which dominates computational load. All other calculations are performed by a

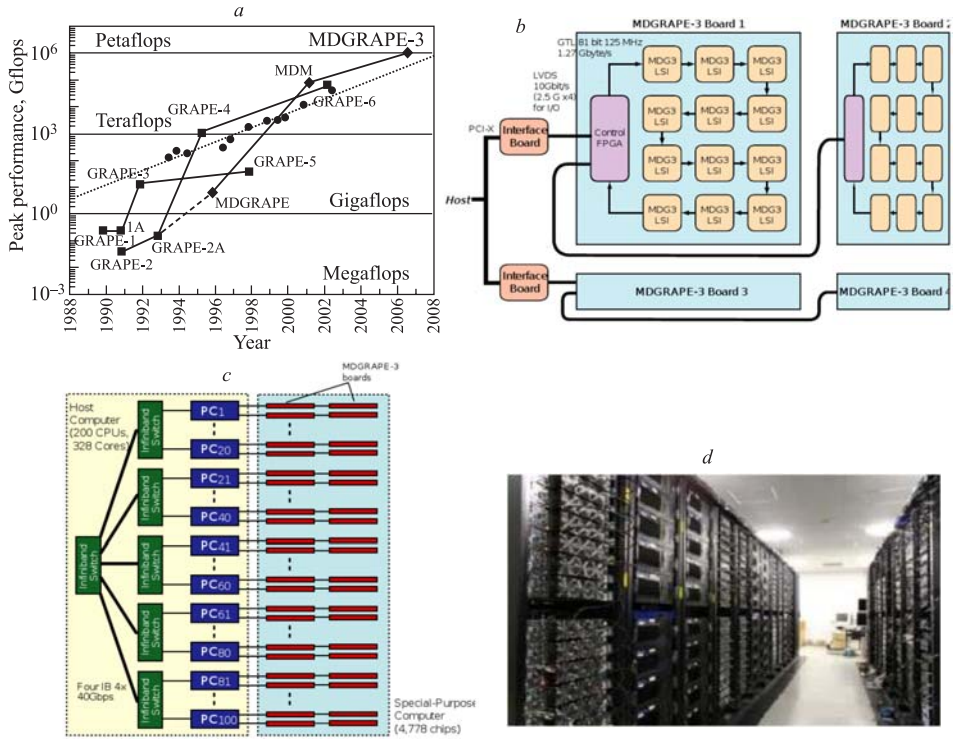

Fig. 9. Peak performance of the MDGRAPE systems (a); block diagram of the MDGRAPE-3 board (b); that of the system (c) and photograph (d) for the final PFLOPS MDGRAPE-3 system

host computer. The calculation of the non-bonding forces is accelerated by dedicated LSI «MDGRAPE-3 chip», which has 165 Gflops peak performance at 250 MHz and 200 Gflops at 300 MHz. It has 20 force calculation pipelines. A host computer will be a PC cluster with 200 CPU/100 nodes. Each node will have a special-purpose engine with 48 MDGRAPE-3 chips in total.

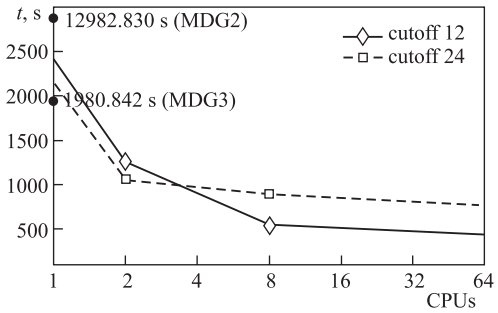

Fig. 10. Test system simulation for the sodium chloride with Ewald sum (64 000 ions, 1000 steps). The simulation results are shown for the MDGRAPE-2 (MDG2: $t_1 = 12982.830$ s), MDGRAPE-3 (MDG3: $t_2 = 1980.842$ s), and JINR CICC cluster II (by using cutoff = 12 and 24 Å), respectively

We have tested the DL_POLY_2.17 program on MDGRAPE-2 (4 chips) and MDGRAPE-3 (2 chips) machines to compare the results with JINR CICC cluster II. In Fig. 10 the test simulations are presented for the sodium chloride with Ewald sum (number of ions = 64 000; time steps = 1000). The results show the following simulation times for each machine, $t_1 = 12982.830$ s (MDGRAPE-2) and $t_2 = 1980.842$ s (MDGRAPE-3), respectively. The JINR CICC simulations for the same number of ions (64 000) and time steps (1000) give us the simulation time $t_3 = 2424.590$ s for one processor. We should stress, however, that for the JINR CICC cluster II the simulations were restricted to a cutoff radii for non-bonded and Coulombian interactions ($R_{\mathrm{cutoff}} = 12$ and 24). In the sense of comparison, simulations with MDGRAPE-2 and MDGRAPE-3 computers are more exact. Using larger cutoff radii for the ionic system under consideration requires more simulation time with the JINR CICC cluster.

CONCLUSION

In summary we conclude that InfiniBand network is very efficient for both versions DL_POLY_2.17 and DL_POLY_3.07 even considering a large number of processors. The code DL_POLY_2 is the original version which has been parallelized using the Replicated Data strategy and is useful for simulations of up to 30,000 atoms on 100 processors. DL_POLY_3 is a version which uses Domain Decomposition to achieve parallelism and is suitable for simulations of order 1 million atoms on 8-1024 processors. The simulation time with Gigabit Ethernet varies with increasing number of processors. The obtained results show some comparison of the CICC cluster and MDGRAPE-2 and 3 special-purpose computers. While using a large cutoff, the performance of MDGRAPE-3 is significantly better than CICC and MDGRAPE-2 for simulations of large ionic system.

REFERENCES

1. *Car R., Parrinello M.* Unified Approach for Molecular Dynamics and Density Functional Theory // Phys. Rev. Lett. 1985. V. 55. P. 2471–2474.
2. *Skolnick J., Kolinski A.* Computer Simulations of the Folding of a Globular Protein // Science. 1990. V. 250. P. 1121–1125.
3. Molecular Simulation Studies in Material and Biological Sciences / Ed. Kh. Kholmurodov. N.Y.: Nova Sci. Publ., 2007.
4. *Todorov I. T., Smith W., Trachenko K.* DL_POLY_3: New Dimensions in Molecular Dynamics Simulations via Massive Parallelism // J. Mater. Chem. 2006. V. 16. P. 1911.
5. *Smith W., Forester T. R., Todorov I. T.* The DL_POLY User Manual (STFC Daresbury Laboratory, Version 2.18. 2007); http://www.cse.scitech.ac.uk/ccg/software/DL_POLY/
6. *Korenkov V. V. et al.* Testing the New JINR CICC Supercomputing Cluster // LIT Sci. Rep. 2007. P. 43–46.
7. TOP50 Computers. http://www.supercomputers.ru/?page=rating/
8. SuperBlade User's Manual (Revision 1.0b). Super Micro Comp. Inc., 2007; http://www.supermicro.com/products/superblade/

Received on July 26, 2008.