

UDC 519.6

## СУПЕРКОМПЬЮТЕРНОЕ ОБРАЗОВАНИЕ: ТРЕТЬЯ СОСТАВЛЯЮЩАЯ СУПЕРКОМПЬЮТЕРНЫХ ТЕХНОЛОГИЙ

Вл. В. Воеводин<sup>1</sup>, В. П. Гергель<sup>2</sup>

Ставится вопрос о необходимости внедрения суперкомпьютерного образования в практику ведущих естественнонаучных и технологических университетов. Представлены первые шаги по формированию национальной системы подготовки кадров в области суперкомпьютерных технологий. Этот же опыт может быть учтен и при организации масштабных образовательных программ по другим актуальным направлениям информационных технологий.

**Ключевые слова:** суперкомпьютерные технологии, параллельные вычисления, высокопроизводительные вычисления, высшее образование, университетское образование, национальный проект.

**1. Supercomputing education: why?** The supercomputing technologies are referred today to the most important directions of scientific and technological development in many leading countries of the world, including Russia. The potential of the supercomputer industry makes it possible to solve many fundamental and applied scientific and technical problems that require large-scale computations.

Supercomputers are being developed incredibly fast, which impacts the performance and scalability of modern computing systems. Let us refer to the Top500 list of the most powerful computer systems of the world ([www.top500.org](http://www.top500.org)), the newest version of the list was released at the ISC-2010 conference at the end of May. The Cray Jaguar supercomputer being the leader of the list has the performance of 1.76 Pflops (1 Pflops =  $10^{15}$  floating point operations per second) and has 224 162 cores working in parallel. The performance is vast, but it will be available for a user only if all the 224 162 cores are used efficiently. How to create a parallel algorithm for problem solving that could scale up to such a number? How to distribute the data among more than 26 000 compute nodes of Cray Jaguar to minimize the data exchange? Is it possible to create an algorithm with no data exchange at all? The questions are not simple, and without a special training it is difficult to find the answers to these questions.

The second position of Top500 is given to the China's Nebulae computer with a performance of 1.27 Pflops. In contrast to Cray Jaguar, Nebulae uses the NVidia Tesla C2050 graphics processors, which are the co-processors to the traditional Intel Xeon X5650 processors. How to formulate an algorithm for problem solving in such a way that it would be possible to use the SIMD nature of the graphics processors and at the same time to minimize the data transmission time between the major processors and co-processors? If this is not done, the user will get only small percentage share of the performance. The Nebulae supercomputer performance shown on the Linpack benchmark (only 43% of the maximum possible performance) is a very convincing confirmation that this is not an easy question. It should also be noted that these results were shown on the benchmark which is extremely simple, regular and has a high computational power. During 30 years of using this benchmark, we have learned everything about it. What results will be shown on real applications?

We can continue this analysis for all the systems of Top500 and for all the representatives of the supercomputer industry. The number of processors and cores in computing systems is steadily growing, which in combination with the heterogeneity raises many nontrivial issues of how to use them. If we extrapolate the tendency of supercomputer development shown in Top500, then in 2018 we can expect the emergence of computing systems of exaflop performance that will have hundreds of millions cores working in parallel. Realizing the complexity of working with such large-scale systems, the computational community is making attempts to understand what software will be able to work on the systems of the future, where the main bottlenecks are and how to overcome them. It is exactly the purpose of J. Dongarra and P. Beckman who initiated the International Exascale Software Project (IESP, [www.exascale.org](http://www.exascale.org)) in the framework of which the expert community discusses and tries to understand what changes in software should be introduced and what software components should become the point of applying the efforts of professionals.

<sup>1</sup> Научно-исследовательский вычислительный центр, Московский государственный университет им. М. В. Ломоносова, 119991, Москва; зам. директора, чл.-корр. РАН, e-mail: [vovodina@parallel.ru](mailto:vovodina@parallel.ru)

<sup>2</sup> Нижегородский государственный университет им. Н. И. Лобачевского, факультет вычислительной математики и кибернетики, просп. Гагарина, 23, 603950, г. Нижний Новгород; декан, профессор, e-mail: [gergel@unn.ru](mailto:gergel@unn.ru)

Supercomputers are rapidly developing. IESP discusses the software design for the future computers. However, there arises the question: who will work on these computers? To be more exact, who will be able to work on them? Who will be able to develop a method, to create an algorithm, and to write an efficient program for these computers? Very few will. And we should not have any illusions that it is relevant only for the systems of the highest performance range containing hundreds of millions cores. The situation will be similar for the “ordinary” desktop systems whose performance is tens of thousands times lower and whose hardware includes “only” about  $10^4$ – $10^5$  cores. It is necessary to start working purposefully at developing the system of supercomputing education in order to prepare the specialists for the realities of the future “superparallel” computer world. Only the three components together — hardware, software, and supercomputing education — will create a steady basis for the development of the entire high performance computing area.

**2. Supercomputing education: when?** In recent years we can observe a paradoxical situation in the supercomputing area. On the one hand, we see the tremendous development of the computing systems. Computer capability has been developing in accordance with the Moor’s law for quite a while and the computer performance doubles every 18 months. That is why every 11–12 years the supercomputer performance increases by three orders of magnitude, which leads to the transfer from Mega to Giga, from Giga to Tera, from Tera to Peta and so on and so forth. As a result, today we have the incredible characteristics of computing systems that are stated in the current version of the Top500 List of the most powerful computers in the world.

On the other hand, what is the progress in the development of parallel programming technologies or in the methods of problem solving? What happened in these areas that would be comparable to the happenings in the hardware area? Almost nothing. The pace of computing machinery development remains invariably fast while no adequate actions on developing the methods of working on such computers have been taken. Many years of inactivity in this field of knowledge led to a huge gap between the possibilities of the modern supercomputers and the practice of their real usage. As a result, we spend a lot of efforts on the struggle with MPI, we try to scale nonscalable methods, we rewrite programs constantly moving them from one computing platform to another.

And the reason for the current situation is the absence of the corresponding supercomputing education. We should have started to actively implement it 10–15 years ago. This delay should be urgently compensated by intensive actions.

**3. Supercomputing education: how?** Having realized the significance of the supercomputing education, it is extremely important to solve the issue of what should be taught in the area of supercomputing technologies. An analysis of the state of things shows that unfortunately in many universities this education is narrowed to studying only several simple technological subjects. However, it is absolutely clear that a specialist in the area of supercomputing technologies must have the knowledge and skills in a wide range of computer science issues. We consider it essential to work on the information structure of programs and algorithms, the architecture of parallel computing systems, the models of computations and the methods of complexity analysis, the parallel numerical methods, the parallel programming (languages, development environments, libraries) and many other areas [1, 2]. Moreover, it should be clear that creating the parallel software is impossible without the professional knowledge and skills in “usual” sequential programming. It is also very important to understand that in most cases the development of efficient supercomputer programs for solving most complicated problems is impossible without having the highest qualification in mathematics.

How to provide the training in such a wide range of required areas? Moreover, it is important to understand that the activity in the area of supercomputing technologies will vary for various specialists. It can be creating new efficient parallel algorithms, or developing parallel programs, or solving complex computation-intensive applied problems, etc. As a result, it becomes clear that we should identify various directions in the supercomputing technology education. Among them there could be the following:

- 1) designing, developing and using supercomputing systems,
- 2) administration of supercomputing systems,
- 3) system programming,
- 4) applied programming,
- 5) advanced training for the university teachers that are giving the training in various aspects of parallel programming (system programming, programming in various subject areas: physics, mechanics, etc.),
- 6) advanced training for specialists in various applied areas where supercomputing systems can be used for problem solving.

It should also be clear that the training in a certain selected area should significantly depend on the target category of the trainees. Thus, the programs for teaching the students should undoubtedly differ from the programs of the advanced training for IT specialists or university teachers.

Forming a certain “complete set” of supercomputing knowledge is still ahead and the significance of this task will apparently attract the attention of many scientists, teachers, and specialists. A possible approach [3] may involve the creative usage of the curriculum recommendations approved by the international associations ACM and IEEE.

Based on this approach, it can be possible to propose an initial version of the “complete set” of supercomputing knowledge.

*Mathematical foundations of supercomputing technologies:*

- 1) mathematical foundations of parallel computing,
- 2) algorithms and analysis of parallel computing complexity,
- 3) computational mathematics and methods of parallel computing.

*Information technologies and high performance computing:*

- 4) architecture and organization of high performance computing systems,
- 5) operating systems and management systems for supercomputers,
- 6) distributed computing and grid technologies,
- 7) organization of the human–computer interaction in conducting numerical supercomputing experiments,
- 8) computer graphics and visualization of high performance computing results,
- 9) information management in conducting high performance computing experiments.

*Methods and technologies of supercomputing:*

- 10) fundamentals of parallel programming,
- 11) languages and parallel programming technologies,
- 12) software engineering of supercomputing software,
- 13) software systems for supercomputing experiments,
- 14) social and professional issues of supercomputing technologies.

**4. Supercomputing education: where?** The importance of supercomputing education requires a consolidation of the efforts of the university community alongside with the close collaboration with the world scientific centers and leading IT companies. The creation of the Supercomputing Consortium of Russian Universities (<http://www.hpc-russia.ru>) became one of the most significant events in this area in 2008.

The purpose of the Supercomputing Consortium is to develop and to ensure holding a set of events and measures aimed at the efficient usage of the higher education system potential for developing and implementing the supercomputing technologies in the Russian education, science, and industry.

Within the framework of the purpose in view, the Consortium is oriented to solving the following major issues:

- coordination and organization of interaction among the higher education institutions in Russia with the aim of using and developing the modern supercomputing technologies in education, science, and industry;
- development of education programs for training the highly qualified specialists, advanced training and continuing education in the area of supercomputing technologies and high performance computing;
- contribution to the development of the network of highly efficient supercomputing centers in Russia oriented to implementing the high performance computing technologies in all the areas of national economics;
- creation of the advanced experience exchange system and distribution of the latest achievements in the area of high performance computing, and others.

Another important direction in the activities of the Consortium is the active promotion of works on the practical usage of supercomputing technologies for solving the today’s problems of science, technology, industry, and business that could not have been analyzed earlier because of the limitations in the existing computing tools and technologies.

At the beginning of 2010, more than 25 leading universities of Russia became the members of the Consortium. A number of institutions of the Russian Academy of Sciences and IT companies participate in this activity as associated members.

**5. Supercomputing education — a national program in Russia.** The significance of the development and practical use of supercomputing technologies elevates the problem of supercomputing education to the state level. The National Project of Supercomputing Education has been created in Russia to fulfill this task in 2010–2012. It was initiated by the Supercomputing Consortium and was supported by the President and the government. Now this Project has a status of the Presidential Project.

Creation of the system of training highly qualified specialists in the area of supercomputing technologies (SCT) must become the major result of the Project. The basis of this scheme will be the creation of scientific educational centers (SEC) in the SCT. The main task for these centers will be the efficient organization of universities’ activities aimed at teaching, continuing and advanced training in the SCT area. The first result

of their activity will be the development of the Set of Knowledge and Skills (professional competencies) in the SCT area according to the activity forms and target groups. Based on this Set, there will be implemented the modernization (update) of the Federal State Educational Standards of the third generation. The modernization will be aimed both at the in-depth training in the SCT in mastership and at the SCT implementation in the basic education (bachelorship) starting from the first years of university studies. The following educational standards will be updated first and foremost: Mathematics, Mathematics and Computer Science, Fundamental Informatics and Information Technologies, Applied Mathematics and Informatics. It is planned to train no less than 500 highly qualified specialists in the area of supercomputing technologies and applications and to provide the advanced and continuing training of 150 teachers on the basis of the updated educational standards.

To implement the new educational programs, a number of advanced courses and continuing training courses for teachers in the SCT area will be organized on the basis of the SEC SCT. As a result of the Project, no less than 25 universities will be involved in the training of specialists in the SCT.

Another important activity for the SEC SCT will be forming special groups of the trainees for various categories of participants (students, masters, postgraduate students, teachers, specialists) for intensive target-oriented training in the SCT area. Such a kind of training takes into account the cross-disciplinary nature of the SCT and will make it possible to train highly qualified specialists that will be able to develop and to efficiently use the SCT in conducting the fundamental and applied research and to implement them in industry and economics.

The national system of SCT education quality monitoring will be established on the basis of the SEC SCT. Under the SEC SCT control, there will be developed new educational courses and programs in the SCT and their applications. In the course of the project execution process, there will be prepared a series of textbooks and manuals in the SCT.

Another way to teach the fundamentals of the SCT will be the Internet University of Supercomputing Technologies. On the basis of this university, the trainees will be able to study distantly.

The SEC SCT will have another important function: to maintain a close relationship with the higher education institutions, the Russian Academy of Sciences, industry and business. Within the implementation of the Supercomputing Education Project, they will cooperate to train specialists and to conduct research in the SCT area.

The Project implementation will be aimed at the in-depth integration of the Russian universities into the world education process in the SCT area. In the framework of this direction, it will be possible to cooperate with the foreign scientific educational partner organizations, to develop cooperative educational programs, to attract foreign scientists for training Russian highly qualified specialists in the SCT area.

Series of popular scientific articles, printed and electronic publications, other presentations devoted to the project results will help to inform the society about the achievements in the development of the SCT and to attract talented young people to it. There will be arranged a periodical edition "Supercomputing Education in the World", which is an approved and efficient way to promote the SCT. The presentations at significant events will also be arranged to inform the society and to present the achievements in the SCT area.

The major objectives of the Supercomputing Education Project are:

*Objective 1.* Creating a network of scientific educational centers of supercomputing technologies. The objective includes creating the infrastructure, developing the regulatory framework and methodological support, purchasing hardware and software, maintaining and staffing the SEC. The implementation of this objective is aimed at the concentration of the scientific educational potential on the basis of the universities having the resources and significant experience in the development of supercomputing technologies.

*Objective 2.* Developing the methodology and methods of teaching, advanced and continuing training for the human resources in the area of supercomputing technologies. The following activities will be realized in the framework of this objective: developing and/or updating the educational standards of the new generation, preparing the new set of knowledge and skills containing the requirements to the professional competencies of the specialists in the SCT area, developing the curriculum and programs, producing textbooks, manuals, and other literature on the SCT, preparing and publishing the consolidated list of the main scientific conferences, seminars, contests, schools on the SCT, developing the strategy of the supercomputing education quality monitoring. This objective is aimed at the creation of the scientific methodical basis built on the educational standards of the new generation for training specialists in the SCT area.

*Objective 3.* Implementing the programs for teaching, advanced and continuing training for the human resources in the area of supercomputing technologies. The objective includes the following aims:

- training 500 specialists in the area of supercomputing technologies;
- advanced and continuing training for the teaching staff of universities;

- advanced and continuing training for specialists having the professional higher education with the use of the distance learning technologies;

- creating Internet centers of the system of educational resources in the SCT area.

This objective is aimed at creating the scientific educational potential that ensures the high quality of training provided to specialists in the SCT area.

*Objective 4.* Integration of fundamental and applied research and education in the area of supercomputing technologies. Ensuring the cooperation with the Russian Academy of Sciences, industry and business.

The objective is aimed at the development of the integration of fundamental and applied research and education in the area of supercomputing technologies, the cooperation with the Russian Academy of Sciences, industry and business. Reaching this objective presupposes establishing the partner relationships among the SEC SCT, the Russian Academy of Sciences, industry and business. There will be published series of reviews, articles, analytical reviews, works in the area of supercomputing technologies and their application in science and business. In the framework of this objective, there will be held meetings with invited representatives of industry, business and science, which will make it possible to integrate the leading hi-tech areas of economy and to staff these areas with the specialists in the area of supercomputing technologies.

*Objective 5.* Enhancing the international collaboration in creating a system of supercomputing education.

This objective includes attracting the leading foreign specialists in the SCT area to participate in the training of specialists. The objective presupposes establishing the partner relationships with foreign scientific educational partner organizations, arranging traineeship for Russian teachers, participation of foreign professors in the development and realization of collaborative educational programs. There will be a periodic release of the anthology “Supercomputer Education in the World”. The objective is aimed at increasing the academic mobility and enhancing the quality of the education in the SCT area.

*Objective 6.* Developing and implementing a system of information support of the society on the achievements in the SCT area, which presupposes publishing popular science articles, preparing and broadcasting TV programs, publishing materials in electronic media and other presentations of the results of the SCT education system development. The objective is aimed at the popularization of scientific knowledge in the area of supercomputing technologies and at attracting talented young people to science and to the development of the strategic economics directions.

The Supercomputing Education Project implementation is planned for a period of 2010–2012. The following events can be considered as the first prominent results of the Project implementation:

*All-Russian Youth School.* The Supercomputing Consortium of Russian Universities conducted a unique All-Russia Youth School “Supercomputing Technologies and High Performance Computing in Education, Science and Industry”.

All-Russia Youth School made it possible to move to a new quality level of conducting scientific educational schools. The peculiarities of this school are as follows:

- active support of the school from the educational community: the educational program was realized by the leading scientists and well-known specialists of the country;

- comprehensive coverage of the supercomputing technologies and high performance computing issues: there were seven parallel sections covering a wide range of topics;

- the variety of forms and methods of teaching: the school program included intensive classes with enlarged laboratory courses, survey lectures from the leading scientists and specialists, trainings of various complexity levels;

- a significant number of the participants: 75 participants from 25 cities of Russia, young specialists, postgraduate students, students of various qualification levels underwent the school training.

The participation in the school provided a unique opportunity of the additional professional training in the area of supercomputing technologies. The feedback shows that this school was well-organized and excellently conducted; it proved to be exceptionally useful. It was decided to include this school in the plan of annual events of the Supercomputing Education Project.

*Target educational programs.* The training for the first special group consisting of 40 senior students of Computational Mathematics and Cybernetics, Mechanics and Mathematics and Physics Faculties of Moscow State University was conducted according to the specially developed educational program “Supercomputing Technologies”. Leading specialists of MSU and institutions of Russian Academy of Sciences were engaged for lectures and seminars. A computational laboratory course for these special groups was conducted on the basis of the high performance systems IBM Blue Gene/P and SKIF MSU “Chebyshev” of the MSU Supercomputing Center.

*The Supercomputing Consortium Meeting.* The annual meeting of the Supercomputing Consortium of

Russian Universities was held at Moscow State University and covered the topic “The Realization of the State Decisions on Education in the Area of Supercomputing Technologies and High Performance Computing”. Apart from the rectors of the Universities, among the participants there were the leading scientists of the Russian Academy of Sciences and the representatives of a number of the industrial companies. The meeting covered the tasks of developing the training and the issues of using the supercomputing technologies in the fields of engine construction, rocket engineering, ship construction, oil and gas industry, biotechnology and pharmacology.

*The Internet University of Supercomputing Technologies* (<http://www.hpcu.ru>) is oriented at teaching a wide range of trainees (students, specialists, teachers) and presupposes the availability of various areas of the training to meet various professional requirements in the area of supercomputing technologies (users, programmers, engineers). The Project was supported and became an integral part of the Supercomputing Education Project activities.

During the two academic years 2009–2010, more than 200 trainees had the training in the University, which can be considered as the first results of the Internet University of the Supercomputing Technologies. The results of the training show that the chosen direction is correct. The ultimate success will depend a lot on several factors: how the Project will be perceived by the supercomputing community of the country, if it will be possible to involve the best teachers and specialists into giving classes and how interested the Russian universities will be in this form of education.

**Conclusion.** The selected number of key issues of the national system of supercomputing education shows the importance of the efforts taken in this direction in Russia. This experience can be useful in solving the problems of supercomputing education in other countries and can also be taken into account in organizing the large-scale educational programs in other areas of information technologies.

It should be noted that the supercomputing education system, which is being formed in Russia now, is oriented at the deep integration of the Russian universities into the world educational process in the area of supercomputing technologies. With this objective in view, the widest possible collaboration with foreign scientific educational partner organizations is planned. This collaboration can also contain the development of the joint educational programs, the participation of foreign scientists in training the highly qualified Russian specialists in the SCT area, the student and teacher exchange, etc. Attracting the attention of all the interested parties to such a kind of collaboration is one of the major objectives of this paper.

We are grateful to the Rector of Lomonosov Moscow State University V. Sadovnichy for his enormous efforts in the support of supercomputing technologies and continuative attention to the supercomputing education in Moscow State University.

This paper was initially presented at the European Computer Science Summit in Prague, October 2010. This work was partially supported by the project “Supercomputing Education” and by the Government project “Development of high-tech production of complex software solutions for domain specific cloud computing”, No. 2010–218–01–209.

## References

1. *Voevodin V.V., Voevodin Vl.V.* Parallel computing. St. Petersburg: BHV-Petersburg, 2002 (In Russian).
2. *Gergel V.P.* Theory and practical use of parallel computations. Moscow: Binom, 2007 (In Russian).
3. Computing Curricula 2010 (CC2010). Association for Computing Machinery and Computer Society of IEEE.

Received October 15, 2010

---