The State-of-the-art Trends in Education Strategy for Sustainable Development of the High Performance Computing Ecosystem

Sergey Mosin[0000-0003-1389-2602]

Kazan Federal University, Kazan, Russia
smosin@ieee.org

Abstract. High-performance computing (HPC) plays very important role in the sphere of information technology as well as defines the strategic direction for inter- and trans-disciplinary breakthroughs ensuring the essential influence on local and global markets. The current status of HPC systems development in different countries is analyzed in the paper. The constraints for an active involvement the HPC in many business processes of different industrial, academic and research partners deal with low competence of the regular users and lack of the HPC proficient personnel. Both the technical infrastructure development and training the competent staff with wide range of the HPC related knowledge and skills are the strategic tasks of the national level. The second task is principal for stable development of the HPC ecosystems especially forwarding to the exascale era. The features of curricula focused on education in the HPC area are considered. The experience of implementation the education strategy of Kazan Federal University in the HPC field based on skills-driven model and partnership with IT-companies is discussed.

Keywords: High Performance Computing · Education · Trends · Sustainable Development.

1 Introduction

The comprehensive informatization of the state-of-the-art society and the active introduction of information technology into the business processes of all sectors of the economy determine the intensive development of hardware and software platforms and the IT sphere in a whole.

Nowadays there is an expansion of the range and complexity of the tasks demanded by the business community. Against this background, the demand of the labor market in IT specialists of different levels and qualifications increases: from the project managers of high level (leaders, architects, project managers, etc.) to rank-and-file executors (programmer, tester, technician, etc.). The efficiency of IT companies in many ways relays with the easiness of integration the university graduates into the processes of hardware-software co-design and implementation the systems of automated data processing in accordance with the requirements and specification of the customer.

The reform of higher education in the Russian Federation is a consequence of the economic transformation of society associated with the transition from industrial orientation to the market. In this regard, the training of a specialist for a specific sector of industry gives way to a competence model of education, when theoretical knowledge is backed by the technological skills and the ability to use it in practice according to employers' requests. The application of the concept of practice-oriented education into the implementation of higher vocational study programs opens the possibility for training IT professionals with a combination of competencies that are in demand on the labor market and specified in the state and professional standards. In this case, universities and the corresponding educational programs acquire competitive advantages, providing greater appeal for entrants.

High-performance computing (HPC) plays very important role in the sphere of information technology (IT) as well as defines the strategic direction for inter- and trans-discipline breakthroughs ensuring the essential influence on local and global markets. Many countries which pretend on the global leadership have developed and implement national strategy for progress in the HPC and corresponding growing of industry, science and economy [1]. The HPC is very specific and key area in the IT sphere with the following distinctive features:

1. Complex infrastructure, which needs high quality specialists for use and support.
2. High direct and indirect cost on high-performance computer design, implementation, use and maintenance.
3. The unique architecture depending on the class of tasks or even specific task which should be solved. Each HPC system is designed for specific task.
4. Essential gap between hardware performance and available software possibilities.
5. Non-equal involvement of the HPC into interdisciplinary R&D.
6. A lot of skills and knowledge in different areas such as computer science, telecommunications, program engineering, power supply and consumption, applied and computational mathematics, management, etc., are required for efficient HPC user and computational scientists.

The constraints for active involvement the HPC in many business processes of different industrial, academic and research partners deal with low competence of the regular users, the problem originators and the general IT specialists in the HPC topics. The current situation can be changed by the complex modernization of the vocational study programs emphasizing the wide range of HPC applications. The common trends and experience of the Institute of Computational Mathematics and Information Technologies at Kazan Federal University in education specialists for sustainable development of the high-performance computing ecosystem are considered in the paper.

The rest of the paper is organized as the following: Section 2 provides an introduction into the HPC strategies of the main world players. The principal features of ACM Curricula for the HPC and NSF/IEEE-TCPP Curriculum Initiative on Parallel and Distributed Computing are considered in Section 3. Section 4 highlights the education strategy of Kazan Federal University for sustainable development of the HPC ecosystem. Final section ensures the concluding remarks.

2 The HPC strategies of different countries

The advantage in both the development and application of high performance computing is a vital for countries' economic competitiveness and innovation potential [2]. Accordingly, many countries have made significant investments and fulfilled holistic strategies to position themselves at the forefront of the rivalry for the global HPC leadership. Many countries have created national programs that are investing large sums of money to wide use the existent HPC systems and develop exascale supercomputers. Mastering and active use the HPC systems, tools and technologies open the ways forward to generation of new significant technologies for overtaken the global grand challenges and improve both the innovative character of national economy and supremacy on the global market.

The state-of-the-art high performance computers are very complex systems, efficient exploiting of which requires a huge amount of financial, the power supply and human resources. There are a limited set of countries that can be able to design, manufacture and use the HPC systems. The influence and contribution of the high performance computers into scientific progress, industrial competitiveness, national security, and quality of life are significant. The open results of the world competitions for the HPC leadership are provided twice per year in the Top 500 ranking. The dynamics of changing the number of the most powerful supercomputers for some leading countries during the last three editions in Top 500 is shown in Fig. 1 [1].

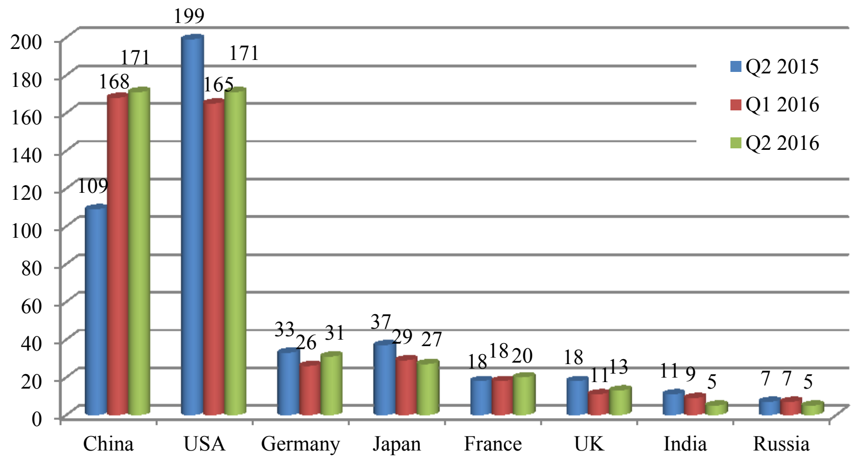

Fig. 1. The dynamic of changes the number of HPC systems in the leading countries at last three TOP500 editions

According to the last 48th Top 500 edition two countries, China and USA, dominate nowadays in both number and total performance of HPC systems (Table 1). The total performance of all 500 computers on the list is now 672 petaflops. The USA holds the narrowest of leads in the aggregate Linpack performance with 33.9 percent of the total and China is the second with 33.3 percent. The retention and especially augmentation the number of HPC systems in the list are very complex efforts for each country because all participants of Top 500 not stand still, but reinforce the processes to develop new systems with orientation onto the exascale architectures. Such stable development in many ways deals with strong progress in the integrated technologies [3], as well as the official state policy in the area of HPC technologies, granted the state programs and investment by government, industrial and business partners. The national strategies and financial resources for realization are represented in Table 2.

Table 1. HPC specification for top 10 countries in the 48th edition of Top 500

Country	Count	System Share (%)	Rmax (GFlops)	Rpeak (GFlops)	Cores
USA	171	34,2	228,032,809	327,303,955	11,660,816
China	171	34,2	223,571,136	394,013,392	21,546,512
Germany	31	6,2	36,501,435	45,628,388	1,600,240
Japan	27	5,4	54,486,820	77,371,577	3,946,560
France	20	4	25,398,803	31,727,765	1,158,428
United Kingdom	13	2,6	27,602,596	31,682,369	1,148,968
Poland	7	1,4	6,162,214	8,157,370	208,284
Italy	6	1,2	14,062,113	21,140,514	606,312
India	5	1	3,092,368	4,456,051	133,172
Russia	5	1	4,411,812	6,515,928	181,070

Table 2. Short specification of the national HPC strategies

Country	HPC Strategy / Program	Investment, $
USA	National Strategic Computing Initiative (NSCI)	320 million/year
China	13th Five-Year Development Plan (Develop Multiple Exascale Systems)	200 million/year (for next five years)
European Union	ETP4HPC; PRACE; ExaNeSt	1.1 in billion total allocated through 2020
Japan	Flagship2020 Program	@$200 million/year (for next five years)
India	National Supercomputing Mission	140 million/year (for 2016-2020)
South Korea	National Supercomputing Act	20 million/year (for 2016-2020)

China has made the HPC leadership a national priority. The leadership in high-performance computing for China is central to the country's goal of transitioning away from reliance on foreign technology to using home-made technology. For instance, the Sunway TaihuLight system placed the first place in the Top 500 list was developed by Chinese National Research Center of Parallel Computer Engineering & Technology using the state-of-the-art 260-core manycore processors ShenWei SW26010 designed by the National High Performance Integrated Circuit Design Center in Shanghai.

2.1 National Strategic Computing Initiative in USA

The National Strategic Computing Initiative (NSCI) was launched to advance the USA leadership in the HPC [4]. The NSCI is a whole-of-nation effort designed to create a cohesive, multi-agency strategic vision and Federal investment strategy, executed in collaboration with industry and academia, to maximize the benefits of HPC for the United States.

The NSCI seeks to accomplish five strategic objectives in the government collaboration with industry and academia [5]:

1. Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.
2. Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.
3. Establishing, over the next 15 years, a viable path forward for future HPC systems even after the limits of current semiconductor technology are reached (the "post-Moore's Law era").
4. Increasing the capacity and capability of an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development.
5. Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the United States Government and industrial and academic sectors.

The NSCI is supported and realized by many national agencies, which also have to develop an ambitious workforce development plan to educate the current generation and train the next generation of scientists and engineers to adopt HPC as an effective approach to solving problems of societal importance.

National Science Foundation (NSF) plays a central role in scientific discovery advances, the broader HPC ecosystem for scientific discovery, and workforce development. According to NSCI the NSF should [5]:

- Provide leadership in learning and workforce development to encompass support of basic HPC training for a broad user community as well as support for career path development for computational and data scientists;
- Increase engagement with industry and academia through existing programs;

- Support the broad deployment of NSCI technologies to increase the capacity and capability of the HPC ecosystem, enabling fundamental understanding across frontiers consistent with NSF scientific and engineering priorities;
- Lead the development of domestic and international collaborations that will advance transformative computational science and engineering with an integrated approach to high-end computing, data, networking, facilities, software, and multidisciplinary expertise, consistent with NSCI strategic objectives.

2.2 European HPC Strategy

HPC is considered as a high strategic importance for European society, competitiveness and innovation. The use of HPC has contributed significantly and increasingly to scientific progress, industrial competitiveness, national and regional security, and the quality of human life. HPC-enabled simulation is widely recognized as the third branch of the scientific method, complementing traditional theory and experimentation.

The European HPC strategy has three pillars [6]:

1. Developing the next generation of HPC technologies, applications and systems towards exascale;
2. Providing access to the best supercomputing facilities and services for the industry (including SMEs) and academia (Partnership for Advanced Computing in Europe – PRACE);
3. Achieving excellence in HPC application delivery and use through establishment of Centers of Excellence in HPC applications.

These pillars are complemented with awareness raising, training, education and skills development in HPC.

The European Technology Platform for High Performance Computing (ETP4HPC) is an industry-led think tank and advisory group of companies and research centers involved in the HPC technology research in Europe, which was formed in 2011 with the aim to build a world-class HPC Technology Supply Chain in Europe, increase the global share of European HPC and HPC technology vendors as well as maximize the benefit of HPC technology for the European HPC user community [7].

The PRACE ensures the wide availability of HPC resources on equal access terms, in order to strengthen the position of European industry and academia in the use, development and manufacturing of advanced computing products, services and technologies. The training an adequate number of professional personnel, including computational scientists, programmers, system administrators, technologists, etc. is considered as one of the key factor for successful development HPC ecosystem in the EU [8, 9].

The PRACE has an extensive education and training effort for effective use of the research infrastructure through seasonal schools, workshops and scientific and industrial seminars throughout Europe. Seasonal schools target broad HPC audiences, whereas workshops are focused on particular technologies, tools or disciplines or research areas.

All EU state programs in the field of HPC are oriented onto strengthening the position of European industry and academia in the use, development and manufacturing of advanced computing products, services and technologies [8].

3 Features of the curricula focused on education in HPC

The major part of national strategies oriented onto development HPC as the key tasks defines generation and development of: 1) de-facto HPC systems emphasing in the short- and middle-terms onto exascale architectures; 2) infrastructure accumulating the public structures, private industry and business as well as academia involved into complex processes of development, support, maintenance and use the HPC systems; 3) educational platform ensuring the training a huge number of required high qualified personnel responsible for effective use of existent hardware and software tools as well as generating a new knowledges and technologies in the HPC area, training the next generation of scientists, designers, engineers, users and task managers.

The implementation of the third task requires development of new or adaptation of existent curricula taking into accounts both the global tendencies of evolution the HPC tools and technologies and the local demands of the state sector, industry and business.

According to exhaustive analysis the efficient training of HPC professionals should be realized in the framework of computer science / computer engineering curricula by including the new specialized courses [10]. End users represent academic, research and industrial organizations and communities. Their applications in engineering, human, social and natural sciences are typically compute and/or data intensive. In some of the areas there are long traditions in using HPC, but in some areas computational science is just entering the domain. Therefore, delivering the introductory basic courses on HPC for master students especially in technical and technological fields in order to increase general awareness and knowledges about HPC possibilities and prepare qualified task managers is very important.

Many universities develop curricula on HPC based on ACM CS/CE Curricula [11], [12] and/or NSF/IEEE-TCPP Curriculum Initiative on Parallel and Distributed Computing [13]. The experience of different universities in adaptation and implementation the curricula focused on HPC and parallel and distributed computing is actively publicized and discussed [14]-[17].

3.1 ACM/IEEE-CS Joint Task Force: Computer Engineering and Computer Science Curricula

Main focus on the HPC is localized through the parallel and distributed computing techniques. The ACM/IEEE Computer Engineering curriculum [12] considers the following main aspects of HPC: 1) Computer architecture and organization with instruction-level and processor-level parallelism (multicore processor and multiprocessor system); 2) Distributed system architectures, high performance computing and networks, memory hierarchy architecture for single core and multicore systems; 3) Parallel algorithms and multi-threading; 4) Introduction to High Performance Computing, which covers the organization of high performance computer, design methods of parallel programming, performance model of programs, performance evaluation and optimization techniques, programming in MPI and OpenMP and algorithms in high performance computing.

The latest ACM/IEEE-CS Joint Task Force: Computer Science Curricula [11] proposal vastly upgraded the coverage of parallel thinking proposing topics such as: 1) Parallel and Distributed Computing; 2) Parallelism Fundamentals; 3) Parallel Decomposition; 4) Parallel Algorithms, Analysis, and Programming; 5) Parallel Architecture; 6) Parallel Performance; 7) Distributed Systems; and 8) Cloud Computing.

Parallel and distributed computing builds on foundations in many areas, including an understanding of fundamental systems concepts such as concurrency and parallel execution, consistency in state/memory manipulation, and latency. Communication and coordination among processes is rooted in the message-passing and shared-memory models of computing and such algorithmic concepts as atomicity, consensus, and conditional waiting.

Special attention is paid to software engineering, which considers different technologies, techniques and tools for software development with orientation on wide range of systems, such as real time systems; client-server systems; distributed systems; parallel systems; web-based systems; high integrity systems, etc. and specifics of parallel programming vs. concurrent programming.

3.2 NSF/IEEE-TCPP Curriculum Initiative on Parallel and Distributed Computing

The draft of parallel and distributed computing (PDC) curricula was designed by IEEE Computer Society Technical Committee on Parallel Processing (TCPP) with support of National Science Foundation (NSF) [13]. This document provides guidance and support for departments looking to expand the coverage of parallel and distributed topics in their undergraduate programs. According to the recommendations the problems of parallel and distributed computing fall into the following four knowledge areas:

1) Architecture.
2) Programming.
3) Algorithms.
4) Cross Cutting and Advanced Topics.

A primary goal of proposed curriculum is the definition for the computer science (CS) / computer engineering (CE) students and their instructors to receive periodic guidelines that identify aspects of PDC that are important to be covered, and suggest specific core courses in which their coverage might find an appropriate context. The proposed curriculum enables students to be fully prepared for their future careers in light of the technological shifts and mass marketing of parallelism through multicores, GPUs, and corresponding software environments, and to make a real impact with respect to all of the stakeholders for PDC, including employers, authors, and educators.

4 Education strategy of Kazan Federal University in the HPC field

Kazan Federal University (KFU) founded in 1804 nowadays is the biggest research and educational center in the Volga region federal district of Russia. The main priorities in the R&D area as well as innovations are organized and developed in the form of the following Strategic Academic Units (SAU): 1) 7P Translational Medicine; 2) Ecooil – global energy and resources for the materials of the future; 3) Astrochallenge: cosmology, monitoring, navigation, applications and 4) The quadrature of transforming teacher education – 4T. There are more than 150 OpenLabs and research centers involved in the state-of-the-art R&D projects as a part of SAU. The major part of research works uses numerical simulation, intellectual data analysis based on data mining and machine learning algorithms. The effective solution of many tasks may be obtained only using the up-to-date hardware and software tools oriented on parallel and distributed computations. Each SAU implements several tens trans- and interdisciplinary research works combining specialists from different scientific fields. IT professionals especially with strong experience in the HPC area play vital important role in the research groups.

The Institute of Computational Mathematics and Information Technologies at KFU (ICMIT) trains the IT specialists competent in the HPC technologies and tools on three levels of study: bachelor programs, master programs and Ph.D. programs. All curricula in the ICMIT are based on the ACM/IEEE CS Curriculum and NSF/IEEE-TCPP Curriculum Initiative on Parallel and Distributed Computing.

The educational process uses the practical skills-driven model. The professional courses combine theoretical knowledge and practical skills. The laboratory works are constructed in such way to master different technologies and tools of parallel and distributed programming for SMP, NUMA, MPP and Cluster architectures, using CUDA programming, OpenMP, OpenCL, MPI, threads programming, etc. The KFU HPC cluster system is used in education process as well. The cluster has the hybrid architecture and combines HPC subsystem, GPU-based cluster subsystem and Big Data processing subsystem. The total peak performance of KFU cluster consists of 39 TFLOPS.

Access of students to real HPC systems plays important role at training specialists adapted to real conditions and studying not only theory of parallel programming but also rules and processes specific in the HPC and data centers.

Additional workshops and short courses delivered by well experienced professionals from the partner's IT companies are important part of training process. The ICMIT regularly organizes such courses in partnership with Intel and NVidia, as well as some academic organization in the framework of Computer Science Club initiative. The KFU has close cooperation with the Supercomputing Consortium of Russian Universities [18].

Bachelor degree students receive basic competences in parallel and distributed computing. The master and Ph.D. students study advanced courses and combine training with R&D. Such multilayer education system allows generating different specialists for the local and global HPC ecosystems.

Conclusions

The global problems and tendencies in development of the HPC ecosystems were discussed. The necessity for continued collaboration and innovative initiatives is obvious and permanent grows. The number of required personnel competent in the HPC system design, implementation and maintenance; parallel programming and application development; numerical and computational modelling as well as task management is increased regularly due to developing the HPC centers and wide using the HPC systems by public organizations, private industry and business. The increased use of computational and information technologies brings innovation and efficiency in many production and business processes, generates products and services favoring the growing of industry, science and economy. The preparation and implementation the professional courses to train the new generation of specialists with knowledge and skills in mathematical simulation and modelling, intellectual data analysis and HPC using, administrating and management are very important tasks for future development the HPC ecosystems. The experience of Kazan Federal University in training IT specialists on three layers of study based on skills-driven model was described as well as concept of trans- and interdisciplinary collaboration in the project of Strategic Academic Units. The realization of introductory basic courses on the HPC for non-IT specialties can provide conditions for active use the HPC systems and technologies at interdisciplinary R&D in the short- and middle-terms.

Acknowledgement

The work is performed according to the Russian Government Program of Competitive Growth of Kazan Federal University.

References

1. Top 500 Homepage, https://www.top500.org, last accessed 2017/04/30.
2. Ezell S. J., Atkinson R. D. The Vital Importance of High-Performance Computing to U.S. Competitiveness. ITIF, April (2016).
3. Mosin S. G.: The Features of Integrated Technologies Development in Area of ASIC Design, In Proc. of 9th Conf. the Experience of Designing and Application of CAD System in Microelectronics (CADSM'07), pp. 292–295, Lviv – Polyana, Ukraine (2007).
4. The National Strategic Computing Initiative (NSCI) Homepage, https://www.nitrd.gov/nsci/index.aspx, last accessed 2017/04/30.
5. National Strategic Computing Initiative Strategic Plan, 2016.
6. HORIZON 2020. High-Performance Computing Homepage, https://ec.europa.eu/programmes/horizon2020/en/h2020-section/high-performance-computing-hpc, last accessed 2017/04/30.
7. ETP4HPC strategic research agenda: 2015 update (2015).
8. High Performance Computing in the EU: Progress on the Implementation of the European HPC Strategy. IDC. 2015. (DOI: 10.2759/034719)

9. Margetts L.: Survey of Computing Platforms for Engineering Simulation, Supplementary Report, European Exascale Software Initiative (2015).
10. Monismith D. R.: Incorporating parallelism and high performance computing into computer science courses. In IEEE Frontiers in Education Conference (FIE), 9 p., IEEE, El Paso, TX, USA (2015).
11. ACM/IEEE-CS Joint Task Force, Computer Science Curricula 2013: Curriculum Guidelines for Undergraduate Degree Programs in Computer Science, Online: http://www.acm.org/education/CS2013-finalreport.pdf, last accessed 2017/04/30.
12. ACM/IEEE Computer Engineering Curricula 2016, Online: http://www.acm.org/binaries/content/assets/education/ce2016-final-report.pdf, last accessed 2017/04/30.
13. Prasad, S. K., Chtchelkanova, A., Dehne, F., Gouda, M., Gupta, A., Jaja, J., Kant, K., La Salle, A., LeBlanc, R., Lumsdaine, A., Padua, D., Parashar, M., Prasanna, V., Robert, Y., Rosenberg, A., Sahni, S., Shirazi, B., Sussman, A., Weems, C., and Wu, J. 2012. NSF/IEEE-TCPP Curriculum Initiative on Parallel and Distributed Computing - Core Topics for Undergraduates, Version I, Online: http://www.cs.gsu.edu/~tcpp/curriculum/, last accessed 2017/04/30.
14. Connor C., Bonnie A., Grider G., Jacobson A.: Next generation HPC workforce development: the computer system, cluster, and networking summer institute. In Proc. of the Workshop on Education for High Performance Computing (EduHPC '16), pp. 32-39, IEEE Press, Salt Lake City, Utah, USA (2016).
15. Gergel V., Liniov A., Meyerov I., Sysoyev A.: NSF/IEEE-TCPP Curriculum Implementation at the State University of Nizhni Novgorod. In IEEE 28th International Parallel & Distributed Processing Symposium Workshops, pp. 1079-1084, IEEE, Phoenix, AZ, USA (2014).
16. Škrinárová J., Vesel E.: Model of education and training strategy for the high performance computing. In International Conference on Emerging eLearning Technologies and Applications (ICETA), 6 p., IEEE, Vysoke Tatry, Slovakia (2016).
17. Wilson L. A., Dey S. C.: Computational Science Education Focused on Future Domain Scientists. In Proc. of the Workshop on Education for High Performance Computing (EduHPC '16), pp. 19-24, IEEE Press, Salt Lake City, Utah, USA (2016).
18. Supercomputing Consortium of Russian Universities Homepage, Online: http://hpc.msu.ru/?q=node/136, last accessed 2017/04/30.