Developing GPU HPC Infrastructure at Moscow State University and Beyond

Vladimir Voevodin, Sergey Sobolev, Dmitry Andreev,
CUDA Center of Excellence, Moscow State University

Moscow State University (MSU) is one of the oldest (since 1755) and largest Russian educational and research centers. The history of high performance computing in MSU goes back to the mid-fifties of the 20th century when MSU Research Computing Center (RCC) was founded in 1955 and equipped with up-to-date computing hardware. Nowadays MSU is the leading Russian supercomputing center serving more than 600 scientific groups from MSU, other Russian universities and institutes of Russian Academy of Sciences. At the end of 2011 MSU was awarded a highly honored CUDA Center of Excellence status from NVIDIA Corporation. The award inspired us to the challenge of creating a complete GPU HPC infrastructure which would include all necessary components for the wide dissemination of GPU-technologies in Moscow State University and far beyond. In this paper we describe briefly what has been done within this project.

Hardware Platforms. The flagship of MSU HPC facilities is "Lomonosov" supercomputer. Being installed in 2009 with peak performance of 420 TFlops, it had a number of successive upgrades. The most significant upgrades were made on the basis of NVIDIA X2070 graphical processors in 2011-2012. "Lomonosov's" 2130 GPUs brought 1.2 PFlops and led to 1.7 PFlops of the total peak performance. It is worth to mention that all necessary preliminary experiments and tests were performed on a special "GraphIT!" GPU testbed. The testbed was a one-rack 16-node supercomputer with 48 NVIDIA M2050 cards purchased from Hewlett-Packard, installed at MSU a year before GPU part of "Lomonosov" was deployed. There were two goals which we achieved with "GraphIT!". If first one was to refine GPU structure of the upcoming "Lomonosov" giant, the second goal was about users. It was necessary not only to port some of the major applications to a new platform, but rather to make our users more experienced with GPU technologies. With the use of "GraphIT!" it was done extremely well: a subsequent transition of many users to "Lomonosov" was painless and fast.

Software Infrastructure. Software is the question of primary importance. That is why "Lomonosov" GPU partition, the core of CCOE MSU computing power, is properly accompanied with modern and actual software to satisfy broad users needs. Now, software infrastructure of "Lomonosov" contains a wide set of commercial and/or public tools, libraries and applications. Developers may use nvcc compiler with CUDA 4.2 as well as PGI compiler 12.9 with OpenACC support. Packages compiled with multi-GPU and hybrid (CPU+GPU) architecture are available for researchers including Amber, Lammps, and NAMD for molecular dynamics simulation, Abinit for materials sciences research, Quantum Espresso for computational chemistry and Magma library for dense linear algebra. Software infrastructure is constantly evolving. Once a year, in January, we interview our users to compose a list of software they need for their research and use this list as a guideline during a year when we decide to renew licenses or acquire new products.

GPU Applications. Applications are the central point for all supercomputing centers and this is a primary question for us. Here are just several examples of scientific projects of MSU and institutes of the Russian academy of sciences which use intensively GPU potential of MSU's "Lomonosov". Modeling of graphene electronic properties using quantum field theory was performed by specialists from the Institute for Theoretical and Experimental Physics with their own CUDA-written application. MSU Biological faculty also works on their own code for creating GPU-version of ProKSim application for investigation of intermolecular interaction in protein-

protein complexes. New results were obtained in MSU Physics faculty using Lammps package for atomistic modeling of polymer composites. Another group from this department developed a method for computing of a liquid crystal drops structure using CUDA and C programming.

Getting Ready for Exascale. Thinking about the future, CCOE MSU pays a special attention to "superscalable" GPU-based applications which are able to utilize efficiently the most of "Lomonosov" GPU partition. Once a quarter we block a queue to provide selected applications with the total power of "Lomonosov". Here are a few examples of this kind. A 75-round reduced SHA-1 collision was found using 22 days and 445 GPUs on "Lomonosov" two years ago. It was and still remains world's the best result by number of rounds for SHA-1. The results of this work were presented at GTC-2012 as a poster and at EuroPar-2012 conference as a full paper. Another interesting project is aimed to developing a hybrid (CPU+GPU) approach for multidimensional integration using Monte Carlo methods. The application shows high scalability in a wide range of GPU-based nodes up to all 1 065 nodes of "Lomonosov" GPU partition. Very perspective results are obtained on "Lomonosov" by a group of scientists from Keldysh Institute of Applied Mathematics, RAS, on aeroacoustics (noise reduction in aircrafts).

GPU programming technologies. Straight GPU programming and code debugging could be a bit sophisticated procedure for users. So not only applications but new GPU programming technologies, languages and benchmarking suites are created with CCOE MSU support. NUDA (Nemerle Unified Device Architecture) is a project aimed at creating a high-level GPU programming system based on Nemerle, an extensible programming language. Another project is intended for efficiency improving and bottlenecks finding of GPU programs. This project involves developing a sophisticated set of tests to measure various performance characteristics of various GPUs accompanied with monitoring and performance analysis tools of GPU applications.

GPU Education. Education is extremely important in the world of intensively progressing technologies. No education, no results, no perspectives. This applies both to the higher education, and advanced trainings for qualified specialists. CCOE MSU has arranged a number of trainings on GPU programming technologies. The most popular trainings set is devoted to the basics of CUDA and OpenACC technologies. Following topics are covered in common: GPGPU evolution, architecture of NVIDIA GPUs, CUDA programming model, interaction between CPU and GPU, CUDA memory types, CUDA application libraries (CUFFT, Thrust etc.), asynchronous execution, CUDA streams, multi-GPU technologies and interaction methods, program profiling and performance analysis, debugging and optimization, fast GPU programs development with directives, OpenACC and PGI compiler. These trainings are carried out approximately once in a quarter and accompanied with practical sessions using "Lomonosov" GPU facilities.

From the number of trainings we would like to highlight the one arranged in the frame of Summer Supercomputing Academy at MSU (June 25 – July 7, 2012). The Academy is a unique annual two-week event aimed for advanced education of participants with a wide range of skills and knowledge. In 2012, 78 graduate and postgraduate students, 33 PhDs even 2 full doctors were selected to be participants of the Academy and 25 participants gave preference to NVIDIA HPC educational track. NVIDIA Fellow David Kirk was invited to the Academy and gave a bright plenary talk "GPU is revolutionizing science".

Another notable event was NVIDIA HPC Day at MSU (October 22, 2012). Aimed to extremely wide MSU auditorium, its key speakers were professor Wen-mei W. Hwu and Mark Ebersole. They were accompanied by Russian scientists and researchers who use GPUs to

accelerate their HPC applications in biology, climatology, oil and gas, information security etc. NVIDIA HPC Day gave a start for an advanced GPU technologies one-week training. Its program includes such topics as Kepler architecture, CUDA 5 basics, multi-GPU programming and profiling, common scientific libraries with CUDA support. A one-day DevTech Kitchen event performed by Thomas Bradley took place during the training. It allowed researchers to talk directly with high qualified CUDA engineers on optimization of their real scientific applications. General methods of optimizations were shown by examples of aero- and gas dynamics and biophysics tasks.

As a good tradition, a contest was arranged between participants during practical sessions of later GPU trainings. The aim of contests was to build an optimized specific CUDA kernel function for a given application to achieve the highest performance. Contest winners were awarded with "GPU Computing Gems" book personally signed by the editor, prof. Wen-mei W. Hwu.

Short intensive trainings are not the only form of GPU education at MSU. Now, a number of courses at the Computational Mathematics and Cybernetics faculty of MSU are carried out every semester on a regular basis. Some examples of the courses are "Programming for the CUDA architecture", "Computing on graphics processors" and "Advanced GPU programming techniques". Moreover, as a result of negotiations with the vice-dean of the faculty upcoming 3-day training (February 20-22, 2013) on GPU programming technologies was approved as an official faculty course.

National Project "Supercomputing Education". During 2010-2012 MSU was a leader of the national "Supercomputing Education" project. The project was fulfilled by the Supercomputing Consortium of Russian Universities under the Russian Federation Presidential Commission on the modernization and technological development of Russian economy. The strategic goal of the "Supercomputing Education" project was to create a national system for training of highly skilled supercomputing professionals. GPU and hybrid HPC architectures were made an integral part of the Body of Knowledge in supercomputing technologies which is used as a basis for development of new curriculum in Russian universities. MSU proposed an idea to include a book on GPU technologies in a series "Supercomputing Education" prepared and published in the frame of the project. 1500 copies of the book "Parallel Computing on GPU: Architecture and Programming Models" were disseminated among 43 Russian universities.

As a result, according to the last statistics, more than 35 universities - project participants have regular courses with practical sessions on GPU programming technologies. In total, more than 1400 students attend to these courses every year.

Awarding the Strongest. We deeply believe that HPC and GPU technologies should be learned and used by students from the very beginning of their research activities. To encourage this, in the autumn of 2012 CCOE MSU announced a scholarship contest for MSU students who uses GPU computing in their research projects. The contest has two stages, November-January and February-May. On each stage, up to 10 students could be selected for scholarship of $1000. Students' projects are selected by its scientific and practical value, role of GPU technologies, implementation level, innovative approaches, author's contribution and quality of project description. A compulsory condition of the full funding is a ten-minute presentation on results achieved during the first stage at a student scientific conference. On the 1st stage of 2012 contest, 9 students were awarded by scholarships. Next contest stage begins February 14, 2013.

CCOE MSU proposed a special status "CCOE MSU Project" for bright research projects that run on "Lomonosov" supercomputer. CCOE MSU projects are provided with a starting grant of

$1000, high priority access to “Lomonosov” supercomputer and personal pages on CCOE MSU web site. A number of CCOE MSU projects will be announced in the middle of February, 2013.

Expanding the Boundaries. One of the most important and perspective direction of HPC development is a wide collaboration with industry. On this way, CCOE MSU, NVIDIA and T-Platforms arrange a round table with a wide participation of aircraft building, space, automotive and construction industry on February 8th, 2013. The main challenge of the round table is to find an answer to the question: “What is the best way to marry HPC and the Russian industry and how to make this alliance happy and efficient?”.

For broader engagement of GPU computing into different area of research we connected “Lomonosov” supercomputer with the National Nanotechnological Grid Network and made GPU packages installed on “Lomonosov” available to this community. Now these packages are available not only for MSU supercomputing center users but also for external scientists.

Dissemination and Popularization. This is a very important part of our activities targeted to the broad audience in education, science and industry. We will not list all the things that have been done, here is just one example. CCOE MSU initiated a series of articles in the Russian HPC magazine “Supercomputers” devoted to solving computationally intensive problems on GPUs. “Using GPUs for resisting force computing in sea ice movements”, “Designing of ultrasonic tomography scanners and GPUs”, “Cryptographic records of Lomonosov” are some good examples of this articles written with easy language and aimed to public at large.