HPC Ecosystem and Competences in Bulgaria

Emanouil Atanassov[0000-0002-7442-7096], Aneta Karaivanova[0000-0002-6493-7981], Todor Gurov[0000-0003-4900-0899]

Institute of Information and Communication Technologies, Bulgarian Academy of Sciences, Sofia, Bulgaria
{emanouil,anet,gurov}@parallel.bas.bg

Abstract. In short, the current status of the HPC ecosystem in Bulgaria includes diverse modern hardware (including two supercomputers from the TOP500 list – Discoverer and HEMUS, as well as multiple HPC clusters), and extensive expertise in HPC and the adjacent fields. The competencies and expertise were acquired during a multitude of successful EU projects, as well as leadership and participation in high-profile national programmes and projects with industry/SMEs. Experience in negotiating and managing contracts with different sources of funding and for the use of both equipment and expertise has been gathered. In this paper, we discuss the new developments in the HPC Ecosystem, supported by the acquisition of the HEMUS supercomputer in 2023, as well as our design choices and optimization approaches.

Keywords: HPC+, Supercomputer, Competences, EuroCC2.

1 Introduction

High performance computing (HPC) plays a pivotal role in stimulating world's economic growth allowing industry and academia to develop world-class products, services, and innovations. HPC can process data and perform complex calculations at high speeds (Sterling et al., 2017).

Bulgaria has a long history of deploying and operating state-of-the-art supercomputers. Even during the years of restrictions on acquiring high-tech computing equipment, HPC systems were acquired or even developed within the country. The first supercomputer acquired for serving the research community was the Blue Gene/P deployed at the Bulgarian State Agency for Information Technology and Communications (SAITC). It is ranked, 332nd in the TOP500 (June 2008) for 1st time, as for the last time it appeared 379th in the TOP500 (November 2009) (Blue Gene/P Solution, 2024). It had excellent efficiency in terms of power consumption and floor space usage, while its main drawback was the limited memory.

Right when the Blue Gene P was retired, the new supercomputer Avitohol was deployed at IICT-BAS (Fig. 1). Avitohol is ranked, 332nd in the TOP500 (June 2015), as for the last time it appeared 389th in the TOP500 (November 2015), but it still operates (Avitohol supercomputer, 2024).

Fig. 1. Avitohol supercomputer at IICT-BAS.

It featured Intel Xeon Phi 7120P as a coprocessor/accelerator on 150 identical dual-CPU servers. Due to being based on the x86-64 instruction set, it was able to support diverse sets of applications and use cases from many areas of science. Certain research groups used it also for running services and deploying their middleware on virtual machines. A notable drawback of the system that could be observed was that many applications did not make use of the Xeon Phi coprocessors, even though they provided most of their computational performance.

Currently, Bulgaria has two supercomputers in the TOP500 list: the Discoverer supercomputer, which is ranked 188th in the TOP500 list (June 2024) (Discoverer system, 2024) and the HEMUS supercomputer ranked 395th in the TOP500 list (June 2024) (HEMUS, 2024).

Discoverer supercomputer, installed at Sofia Tech Park in 2021, has 1128 identical servers equipped with AMD EPYC 7H12 64C 2.6GHz processors. In this way, it is again suitable for diverse workloads due to its compatibility with the x86-64 instruction set. However, it cannot address the soaring interest in training Machine Learning models especially Large Language Models using GPUs, even though it has a Petascale level of Linpack performance. As it is financed by the EuroHPC Joint Undertaking (EuroHPC JU, 2024) it serves both Bulgarian and European research groups.

HEMUS supercomputer (**He**terogeneous **MU**lti-functional **S**upercomputer) (Fig. 7) installed at IICT in 2023 builds upon the experience and feedback obtained through the years and provides a heterogeneous environment, consisting of 128 CPU-only and 20 GPU-enabled servers, to match better the needs of its projected user base.

In section 2 we will provide more information about the usage of the HPC infrastructure and then, in section 3, we will describe our motivation in making the various design choices for HEMUS during its planning and implementation. The HPC competencies, necessary for optimal supercomputer use, are considered in section 4. The last section concludes this paper.

2 Feedback on HPC Infrastructure Usage

Recently, in the framework of the European project EuroCC (NCC Bulgaria , 2024), a survey and analysis of the use of the infrastructure for high-performance computing by Bulgarian researchers, as well as their training needs, was carried out. The survey was carried out with the main aim of categorizing specific topics and forms of training that are in demand by Bulgarian researchers. Questions were added to the survey to determine the respondents' experience in using HPC/AI / HPDA. The survey was completed by 85 participants.

The most important conclusions and recommendations resulting from the analysis are:

- A considerable part of the participants has tasks that would benefit from the use of HPC/HPDA/AI, however, the remaining 25% of them have a tangible need to get better acquainted with the capabilities of the technologies in question.
- The majority of the users have significant experience, so the emphasis can be on training for more advanced users and developers.
- Some participants have started using HPC / HPDA / AI very recently or have yet to be trained. There is a tangible need for entry-level training as well.
- A significant part of users are interested in long-term – more than 5 years, data storage. This should be reflected by dedicating separate resources and maintaining protocols and by including training material accordingly.
- When working on artificial intelligence applications, the focus should be on the use of CPU and GPU, but it seems appropriate to provide training on using FPGA as well.
- Introducing users to the FAIR principles for data storage– a large part of the users from academia seem to be unfamiliar with them.
- Attention should be paid to the use of higher level languages for high-performance computing and data processing, such as Python, R, and Matlab.
- Need to respond to the significant interest in using HPC in the cloud by deploying relevant services and providing appropriate training

Some of the quantitative results are summarized below.

A significant proportion of the participants had started using HPC very recently or had yet to be trained (Fig. 2). There is a visible need for entry-level training.

To achieve the goals of NCC Bulgaria, it is important to connect the activities of the Competence Centre to the already existing and prospective national and international activities to benefit from the potential synergies. The HPC competence map for Bulgaria shall be continuously upgraded to ensure a clear picture of available competencies. The interactions with the industry will be continuously monitored and analysed to improve the processes and increase the overall impact.

The Competence Centre will be collaborating in a better way with researchers, public administration, and industry in the HPC+ area, exposing a comprehensive portfolio of services. Following the developed methodology for knowledge transfer, the consultancy work will be supported by streamlined procedures and templates, and its progress will be monitored and evaluated. Potential users could gain ideas and learn from the good practices and success stories and will be directed to appropriate consultants from the teams formed. Various dissemination activities will contribute to further HPC awareness in Bulgaria. Regular dissemination and marketing events will be organized to integrate the HPC+ stakeholders in Bulgaria and to complement the knowledge available in the NCC.

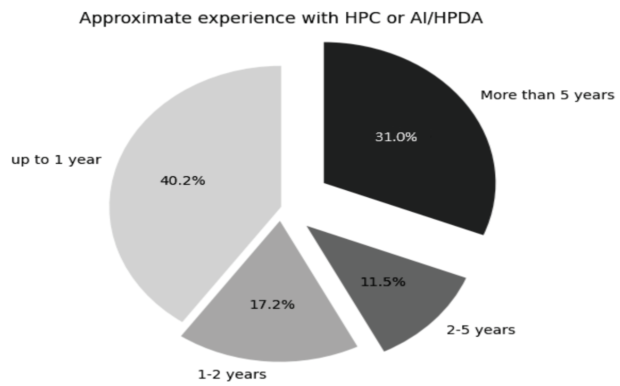

Fig. 2. Distribution of users according to their work experience.

Predictably, the use of existing HPC software (in most cases probably open source, although there are also those using licensed software) is expected, but also the pool of those developing entirely proprietary software or a combination of proprietary and off-the-shelf software is significant in size and therefore there is a need for developer-oriented courses (Fig. 3).

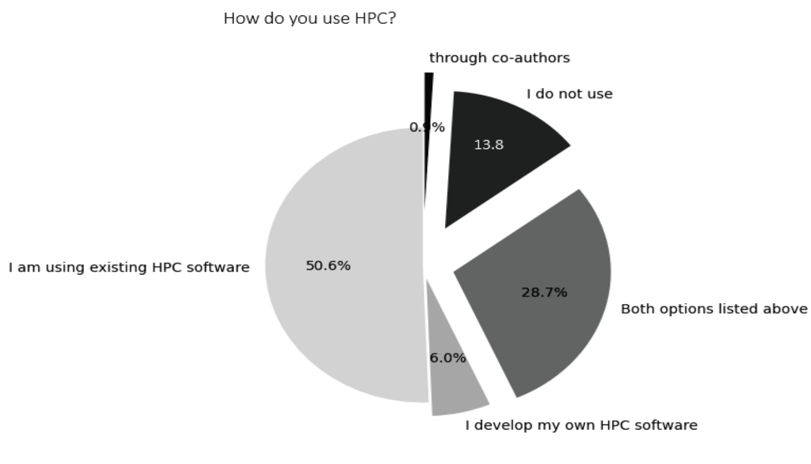

Fig. 3. Distribution of users by software used.

A small percentage of users can use a significant number of cores and are therefore at the minimum level necessary to efficiently use HPC resources. Here, there is a significant shear between the current state and the desired state, as 128 cores is not a lot at all

(Fig. 4). Note that this question has been answered some years ago and thus the current state has improved.

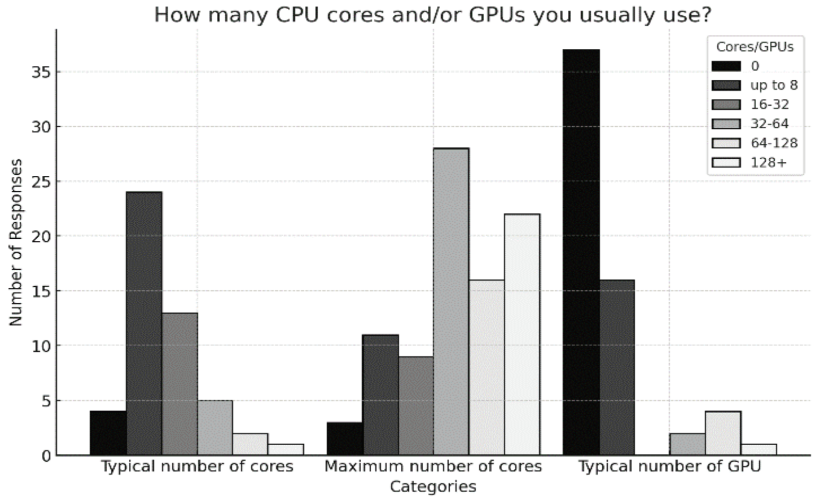

Fig. 4. Distribution of users by number of processor cores used.

Although not overwhelming, the amount of users of GPU resources is significant and there is a need for training to enable users to move to much larger-scale usage (Fig. 5).

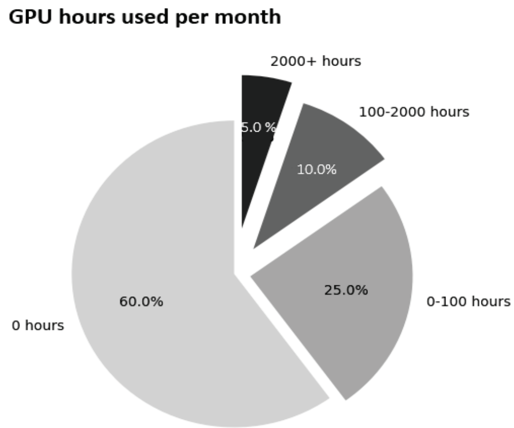

Fig. 5. Distribution of users by GPU hours used per month.

The percentage of those who answered "Yes" to the question of whether the software developed by the respondents can work on heterogeneous architectures (Fig. 6.) is surprisingly high and should be taken into account in the preparation of courses that also address hybrid types of use.

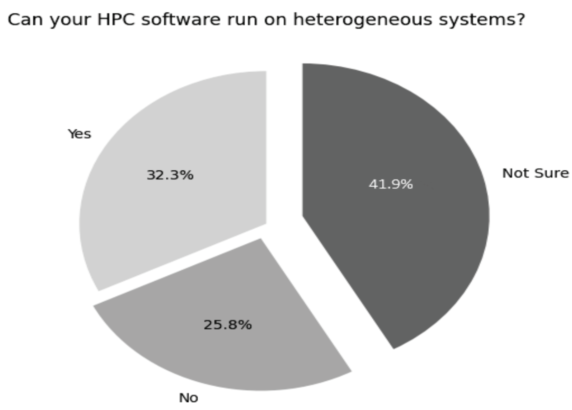

Fig. 6. Use of heterogeneous computing systems.

3 HEMUS – the New Heterogeneous Multifunctional Supercomputer

Our motivation in acquiring the supercomputer and building all the necessary support infrastructure for it has been to serve the needs of the Bulgarian academic community and more specifically the needs of the partnership of the Centre of Excellence in Informatics and ICT, which was funded under the Operational programme "Science and education for smart growth". That is why we evaluated the usage patterns observed at our previous supercomputer Avitohol and undertook a series of surveys and interviews with researchers, and concluded that we have basically two types of usage. Some parallel applications are not yet optimized for the use of accelerators and thus cannot benefit from the availability of advanced accelerators like GPUs, although their parallel efficiency is well understood, and they can scale well on regular servers. This is the case for many applications in climatology, air pollution modeling, statistical mechanics, and astronomy. Other applications have already been ported to GPUs and can effectively utilize servers equipped with multiple GPUs, with most of the computational work outsourced to these GPUs and with their scalability dependent on fast links between GPUs on the same server as well as between servers. Although currently the focus in this direction is in Machine Learning/AI, many of the applications used in the area of Quantum Chemistry/Molecular Dynamics are already fully GPU-aware. Many of the participating institutions in the consortium have capable developers that develop parallel applications using mainly some combination of OpenMP/MPI/CUDA, or simply relying on the use of optimized libraries/frameworks.

Taking into account also the current trends in the development of future versions of the popular applications, we arrived at the conclusion that we need a hybrid system, with balanced division between CPU-only and GPU-enabled servers. That is how HEMUS was conceived as consisting of two subsystems:

- CPU-only subsystem with 128 HPE ProLiant XL220n Gen10 Plus servers, equipped with two 2xIntel Xeon-Platinum 8352Y 2.2GHz 32-cores CPUs, 256 GB RAM;
- GPU-enabled subsystem with 20 HPE ProLiant XL675d Gen10 Plus servers, equipped with two AMD EPYC 7742 (2.25GHz/64-core/225W) 64 cores CPUs, eight NVIDIA A100 TENSOR CORE 40GB SXM GPUs, 512 GB RAM.

Taking into account the available power and cooling capacity, we put an effective limit on the maximum power draw of HEMUS at 200kW (without considering the power usage for cooling, which forms a separate subsystem).

In real usage, we observed a maximum of about 120 kW power usage by the GPU-enabled servers, while the CPU-only subsystem uses considerably less.

Although few applications can readily utilize a hybrid system like that, we have observed that many workloads that perform training of Machine Learning/AI models have distinct phases, some of which are mostly CPU- or IO-intensive. As the GPU-enabled servers are the most contested resource, we expected users to be able to distribute such workloads in a way that makes use of the CPU-only servers where advantageous, thus speeding up the training and improving the utilization of the GPUs. That is why we not only provided InfiniBand interconnection for all server nodes, but in the case of the GPU-enabled servers we provide one InfiniBand port for each GPU. We also ensured that the connection is non-blocking, i.e., any two servers can communicate at the maximum possible speed, without interference from any other communications.

Due to the limited floor space and from considerations of energy efficiency, the use of water cooling was an obvious choice. The provision of direct water cooling for the processors and GPUs was a design choice of the supplier, which improved the energy efficiency parameters even more, while the serviceability of the system was not impacted negatively, as it is still possible to easily service individual servers when necessary.

The operating system installed is SUSE HPC Linux, which enhances the regular SUSE distribution with some HPC-specific packages. As the main users of the system are researchers, who are used to submitting their workloads as batch jobs, we manage the utilization of HEMUS via the SLURM batch system (Slurm manager, 2024).

A diverse set of compilers is in use, consisting of GNU Compiler Collection, and compilers from Intel and NVIDIA, with the possibility to add more if requested.

Although HEMUS is designed to support mainly research, the rules of the funding programme allow up to 20% of its capacity to be used by SMEs/industry. For such users, it is important to ensure not only security, but also utmost reliability and availability. We were able to upgrade the power supply infrastructure by including a diesel generator and obtaining a dedicated transformer, which goes a long way towards decreasing downtimes. A Building Management System (BMS) has also been deployed, enabling proactive monitoring of the environmental conditions at the data centre.

Fig. 7. HEMUS supercomputer at IICT-BAS.

The networking is facilitated through an optical link with the nearly located GÉANT Point-of-Presence, so that the intermediate devices are not a limiting factor when connecting to the GÉANT pan-European network (GEANT Network, 2024).

The capabilities of HEMUS were tested with a set of well-known benchmarks, of which the Linpack benchmark is the most widely recognized. During the testing, some problems in the configuration were discovered and fixed so that we could achieve the best possible performance with our available hardware. We found out that most of the computations were performed by the A100 cards, while the CPUs were used for only 2-3% of the floating-point operations. The cards were drawing 400W of power each one, which we consider to be typical for well-optimized workloads.

As certain research groups deal regularly with large datasets, sometimes consisting of speech/video-recordings or high-resolution images, especially in the case of dealing with digitized cultural heritage artifacts, we ensured that data can be stored and processed at high speed, using the data storage system with a capacity of more than 6 Petabytes. IBM Storage Scale and HPE Ezmeral software have been deployed and made available to users, providing a rich set of options for data storage and processing. Overall the new HEMUS system has been supplanted with an extensive spectrum of middleware, software, and infrastructure support options to ensure it can serve the widest possible audience.

4 HPC Competences

Competences are an important pillar for a successful implementation of the European HPC (EuroHPC) strategy (Atanassov et al., 2023). There are two main categories of HPC competences:

Technical Competences: These competences focus on the underlying hardware and software of HPC systems. They include:

- Hardware knowledge: understanding of the different components of an HPC system, such as processors, memory, and storage.

- Software knowledge: familiarity with HPC operating systems, job schedulers, and programming models.
- System administration: the ability to install, configure, and maintain HPC systems.

Scientific Competences: These competencies focus on using HPC systems to solve scientific problems. They include:

- Problem-solving skills: the ability to identify and define problems that can be addressed using HPC.
- Algorithmic thinking: the ability to design and implement algorithms that can be effectively run on HPC systems.
- Domain knowledge: expertise in a specific scientific field, such as physics, chemistry, or engineering.

Having a strong foundation in both technical and scientific competencies is essential for success in the field of HPC.

The National Competence Centre in HPC (NCC Bulgaria , 2024) was established within the framework of the EuroCC project which is now in its second phase EuroCC2. It is concentrated on concentrating skills in the area of HPC and then offering access to this expertise by offering a portfolio of related services, mainly consulting and training. The portfolio is based on the already existing expertise of the partnership (consisting of the Institute of Information and Communication Technologies – Bulgarian Academy of Sciences (IICT-BAS), leading organization, Sofia University "St. Kliment Ohridski" (SU) and the University of National and World Economics (UNWE). Within the project that has been an ongoing effort in classifying and adjusting the competencies in the area of HPC in order to have a good view of the European ecosystem. Many workshops, webinars and other training events were devoted to improving the technical competencies in the area of HPC and its adjacent fields. The Bulgarian team was among the most active participant in those events, (EuroCC ACCESS, 2024). In this way, we improved the overall technical level of the participant institutions (as these events were open) and acquired knowledge in areas where we had little coverage, e.g., quantum computing. The project does not deal with the scientific competencies, that are highly domain specific, but they are an obvious basis for interactions with industry, as there is little incentive for SMEs/big industry to collaborate with academia on purely technical issues. The ability to provide a full solution to an industrial problem is a cornerstone of our strategy for industry engagement.

5 Conclusions

In the last year, there have been substantial new developments in the HPC ecosystem in Bulgaria. An agile approach has been chosen in the acquisition of HEMUS, considering the lessons learned from the previous supercomputers as well as the users' feedback. The fast interconnection within the heterogeneous HEMUS system and with the high-capacity data storage system provides a balanced HPC environment enable diverse set of applications. The Bulgarian National Competence Centre in HPC has developed a wide range of skills and knowledge, combining technical, scientific, and administrative expertise and providing open access for external users from academia, public administration, or private companies, so that they can directly connect with our experts and work on concrete projects. Some of the results of these collaborations have been described in success stories widely disseminated within the European HPC community. Through the vibrant collaboration with other national centres we hope to expand further our capacity and capabilities so that we can better address the important problems that stand before Bulgarian science, society and economy and obtain the benefits that supercomputing and HPC technologies provide.

Acknowledgements.

This work of authors has been supported by the EuroCC2 National Competence Center Bulgaria; project No 101101903, which has received funding from the European Union through the European High Performance Computing Join Undertaking (JU) and the Ministry of Education and Science of the Republic of Bulgaria. The research that led to these results was carried out using the infrastructure purchased under the National Roadmap for RI, financially coordinated by the MES of the Republic of Bulgaria (Grant No. DO1-325/01.12.2023).

References

Atanassov, E., Gurov, T., & Karaivanova, A. (2023). Towards Effective Bulgarian Competence Centre in High Performance Computing – Service Portfolio and Competences. *Digital Presentation and Preservation of Cultural and Scientific Heritage*, *13*, 333–340. https://doi.org/10.55630/dipp.2023.13.35

Avitohol supercomputer. (202). https://top500.org/system/178609

Blue Gene/P Solution. (2024). https://www.top500.org/system/176516

Discoverer system. (2024). https://www.top500.org/system/179948

EuroCC ACCESS. (2024). https://www.eurocc-access.eu/

EuroHPC JU. (2024). https://eurohpc-ju.europa.eu/supercomputers/our-supercomputers_en

GEANT Network. (2024). https://network.geant.org

HEMUS. (202). https://www.top500.org/system/180208

NCC Bulgaria . (2024). http://eurocc-bulgaria.bg/

Slurm manager. (2024). https://slurm.schedmd.com

Sterling, T., Brodowicz, M., & Anderson, M. (2017). *High Performance Computing: Modern Systems and Practices.* Elsevier.

Received: March 15, 2024
Reviewed: May 01, 2024
Finally Accepted: May 14, 2024