HIERARCHICAL MODEL OF ARCHITECTURE OF SUPERCOMPUTER SYSTEMS FOR COMPARISON AND RANKING

© 2022 D.A. Nikitenko

Lomonosov Moscow State University, Research Computing Center
(Leninskie Gory 1 bld. 4, Moscow, 119991 Russia)
E-mail: dan@parallel.ru

The task of comparing the capabilities of computing systems with each other and forming various ratings has many possible goals. Here, there is the identification of trends, the promotion of proven general-purpose architectures, and the demonstration of superiority in a certain class of tasks, etc. It is, of course, not enough to describe the achieved performance for all these purposes, various rankings and comparisons use different levels of abstraction and generalization up to that level, which would allow to associate the identified performance indicators with certain features of the system. In practice, descriptions of the architectural peculiarities of systems in ratings are rather scarce, and the authors of the work solve the problem of development a formal description of computer systems of a relatively high level, which, at the same time, would allow to increase the required level of detail, corresponding to the goals of applied research. Such a hierarchical system description model has been proposed and tested on well-known systems from the Top50 and Top500 lists.

Keywords: model of a supercomputer system, description of the architecture of supercomputer systems, comparison of performance of computing systems, performance ratings.

FOR CITATION

Introduction

The problem of a detailed description of the architectures of computing systems for their comparative analysis is sufficiently manifested when considering the current ratings of computing systems. Indeed, the desire and need to highlight significant characteristics that affect achievable performance are understandable, but, the key obstacles to the high detail of such descriptions are: firstly, the available information about systems is largely limited, and for advanced systems there is know-how and trade secrets, and, secondly, the real goals of the ratings, which are, to a greater extent, in a competitive way than in a research one. The logical result is that the data on the configuration of top supercomputer systems, presented in the world ratings, are limited to a description of an administrative and marketing nature and the most basic description of the configuration [1], in fact, consisting in a description of the scale of systems and the generations of components used. The Russian Top50 rating demonstrates a significantly larger backlog, but in its form it is not suitable for comparing two different architectures, but only allows you to get more rating slices on the results of using certain components on the HPL.

A common weakness of the existing ratings is a narrow focus on ranking according to a certain criterion. Importantly, within the framework of the problem being solved, it is also necessary to take into account the principles of co-design, according to which the choice of a software and hardware platform is carried out based on the characteristics of the problem being

solved, and vice versa, for a certain configuration, recommendations can be made on what kind of applications can potentially work on it with higher efficiency. From here, a fundamentally important conclusion is made — the developed methods of comparative analysis should be able to operate with various benchmarking results and should be described in a way for joint and complex analysis, this is their fundamental difference from traditional ratings with ranking by a fixed parameter.

In the Section 1 we give a brief overview of the state o f the art. Section 2 introduces the new method of system description and model. Section 3 provides details on the approbation of the proposed approach on well-known systems from Top50 and Top500 lists. In the final section we summarize the paper.

1. Related work

World ratings are characterized by the fact that, on the one hand, they actually allow to see trends in the HPC systems market, on the other hand, in view of the large number of participants and applicants, rating compilers often deliberately limit the details in the participant's application so as not to scare participants away, and on the other hand, to cope with the flow of heterogeneous information.

Regional ratings allow much more ordinary systems to get into the area of attention. In view of the relatively low rate of renewal, it is necessary to speak about global trends based on regional ratings with caution, but, at the same time, such ratings reflect the distribution of computing resources in the region by industry.

One possible way to look at supercomputer ratings is with regard to their specialization. There is no rating that accurately reflects the capabilities of systems in relation to solving specific applied problems. Therefore, in order to determine the properties of systems that are important and outline characteristic of systems that perform best on the class of problems of interest, one should, first of all, evaluate the systems participating in the ratings according to the closest benchmarks. In particular, regarding I/O — IO500, regarding working with memory — Graph500, regarding machine learning — MLPerf. Regarding the descriptive part of the architectures, only basic information about the system and its components is available in the Top500 rating. This is enough to build sections and analyze trends in the use of certain technologies, however, the lack of information about the structure of the system, the structure of its nodes and memory subsystems makes it impossible to use it to compare systems. A formal description of the model is not available in the literature. So, we evaluate its capabilities and parameters based on the results of discussions with Top500 authors at supercomputer conferences of previous years.

In the Top50 rating, the situation is significantly different: the concepts of nodes are highlighted, their basic characteristics are identified by processor models, the categories of networks and their topologies are highlighted. So, there is a good foundation for analyzing architectures. The model is described in detail in [1]. However, the model used has a number of significant drawbacks — a weak description of the memory hierarchy, no description of intranodal connections, etc.

The description model, that is being developed within the Algo500 project [2, 3], should also be mentioned. At the moment, the project uses a hierarchical description model based on XML [4], while it is designed in such a way that it allows to get by with a minimum of descriptions of components and their properties, and does not involve the introduction of new nested entities (for example, to describe a computational group or an interconnect between

GPU). Unfortunately, in our opinion, this significantly limits the prospects for the development and application of such a model, as well as its processing in the event of a significant increase in the number of attributes and the described systems. The questions of system modelling are also touched upon in many adjacent areas like simulation of Database Systems [5] and others, but they target different of the project aims. Some other models are discussed in [6], which more correspond to the models of computing. Thus, regarding the research goals, the model used in the Top50 rating seems to be the most developed of the existing ones, however, it is not suitable for solving the problem of comparing architectures.

2. Proposed method

To describe the configuration of computing systems, it is proposed to use a hierarchical graph shown (Fig. 1). Green shades of vertices correspond to compute items on various layers. Shades of grey depict various levels of interconnects and networks, blues are for accelerators, including GPUs, yellow stands for RAM, and orange is for the local on-node storage.
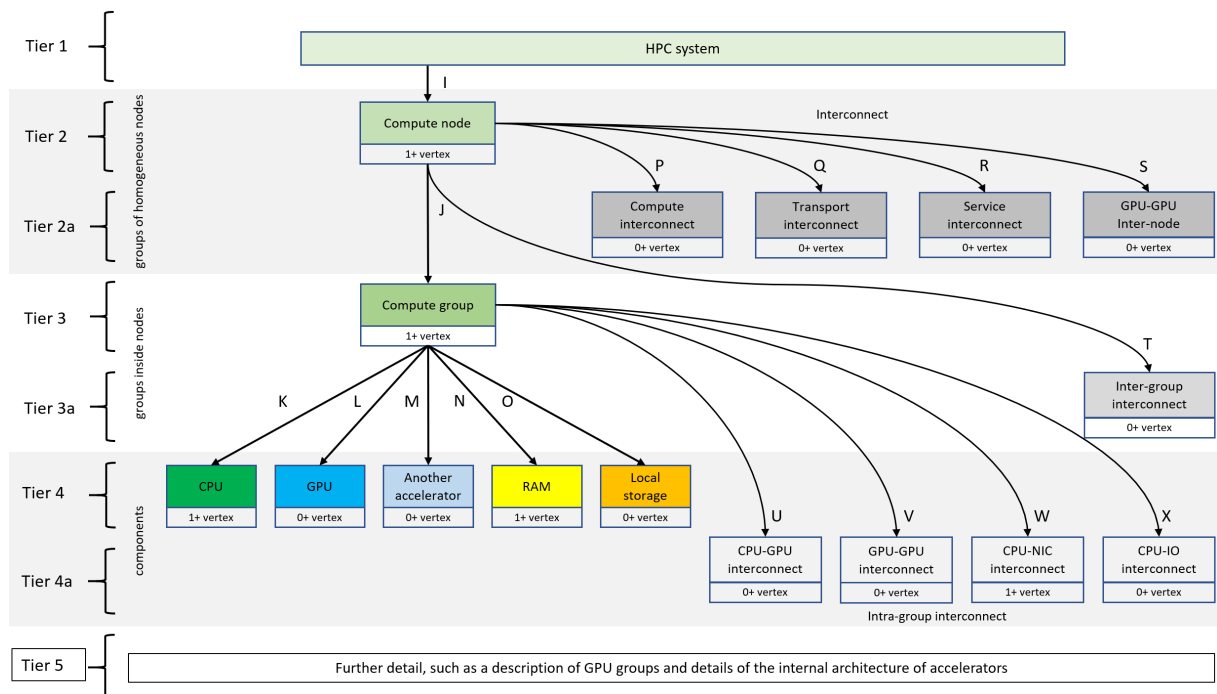

Fig. 1. Graph tiered hierarchical model for describing supercomputers
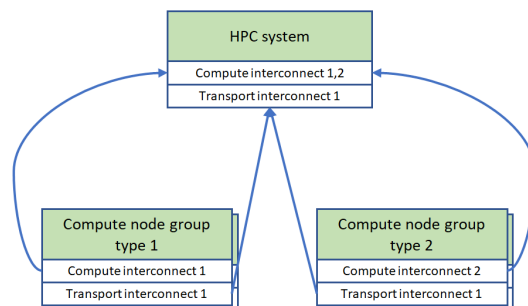

Fig. 2. Inheritance of properties from groups of computing nodes to the system as a whole

Such graphs are called models because they are modelling the general features of acrhitecture. Of course, we should not consider this as a model of the system itself. The model of architecture consists of objects located at the vertices of the graph, connected by links — edges. Each object and relationship has attributes. The vertices are divided into the tiers. The examples of attributes for different vertex objects are given in Tab. 1.

- **Tier 1.** The level of the computing system is the root vertex of the graph, which describes the system itself and its characteristics as a whole.
- **Tier 2.** The level of description of groups of nodes. Each vertex corresponds to a group of nodes of the same configuration. The edge parameter I corresponds to the number of nodes of the type it belongs to.
- **Tier 2a.** The level of description of inter-node connections at the level of groups of nodes and, transitively, above (Fig. 2):
  - computer network (for MPI transfers, etc.);

Table 1. Preliminary attributes of vertex objects for the objects of top tiers and interconnect

Tier 1. System description	
– Name	
– Place of installation	
– Date of described configuration	
– Date when the configuration is expired	
– Peak performance (FP64, FP32, FP16, ...)	
– Achieved performance (FP64, FP32, FP16, ...) on the relevant tests/criteria	
– Power consumption	
– Resource manager and version	
– Links to other system descriptions	
– Link to previous configuration	
Tier 2. Node type description	**Tier 2a. Inter-node interconnect**
– Manufacturer of the base server	– Name
– Base server model	– Family
– Communication network	– Bandwidth
– Transport network	– Latency
– Service network	– Topology
– Intergroup interconnect	– Carrier (copper, fiber)
– Operating system (family, type, kernel, ...)	– Number of interfaces per node
– File system (family, version)	
Tier 3. Compute group configuration	**Tier 3a&4a. Intra-node interconnect**
– CPU model and number	– Name
– GPU model and number	– Version
– RAM	– Bandwidth
– Local storage	– Latency
– Memory access model (SMP/NUMA)	– Topology
– Interconnects:	– Number of interfaces per unit
CPU-CPU, CPU-GPU, GPU-GPU,	– Developer of the standard
CPU-NIC, CPU-IO, GPU-I/O, GPU-NIC	– Type

- transport (network file system);
- service network (image management, monitoring, etc.);
- internodal connection between accelerators.
- **Tier 3.** The level of the computing group. Each vertex corresponds to a computational group of a certain type. Computing group refers to the central processor and controlled or accelerators, directly addressable memory, local storage on the node. The arc parameter J corresponds to the number of computational groups on the node that it connects.
- **Tier 3a.** The intra-node communication level on the node. For example, QPI/Infinity/X-BUS for two or more CPUs per node, or NVLink/Infinity for GPU-to-GPU communication.
- **Tier 4.** The level of description of the computational group components. At the moment, the central processor, graphics accelerator, other type of accelerator, RAM, local storage (top) are allocated.
- **Tier 4a.** Level of description of intra-group relations:
  - communication between the central processor and the graphic accelerator;
  - communication between graphics accelerators;
  - communication between the CPU/GPU and network adapters;
  - communication between the CPU/GPU and local storage.
- **Tier 5.** Provided for a more detailed description of the components. For example, a graphics accelerator chip or a memory hierarchy on an accelerator can be described separately.

It is supposed that the level of detail is such as is required for the analysis of the model, that is, in principle, it is allowed to use an incompletely described model. For example, if there is no information about intra-site connections, their description can be omitted, while the description of the system will remain correct, but with a lower level of detail.

The storage of performance data is intended to be in a format such that:

- saving a lot of measurement results for each configuration;
- selection of subsets of the system components that were involved in performance measurements;
- construction of various ratings, slices according to customizable criteria.

3. Experimental evaluation

We have tested the model on 20 leading systems from the top of Top50 and Top500 lists, with 10 from each. Let us consider several of these well-known systems in terms of the model.

3.1. Frontier, Oak Ridge National Laboratory, USA

Frontier [7, 8] is the most productive system in the world according to Top500 edition No. 59 dated 06.2022. The system is based on HPE Cray EX235a, and consists of 9,408 nodes (Fig. 3).

Each node (Fig. 4) contains one 64-core AMD "3rd generation optimized" processor, 4 AMD MI250X graphics accelerators and 512Gb of DDR4 memory. Logically, the cores are divided into 4 NUMA groups, each of which works with a separate accelerator, consisting of two integrated GPUs. Each node has two local NVMe SSDs of 1.92Tb each.

The nodes are interconnected by a proprietary Slingshot-11 interconnect, 4 interfaces per node. CPU-GPU and GPU-GPU are interconnected by AMD Infinity Fabric.

3.2. Summit, Oak Ridge National Laboratory, USA

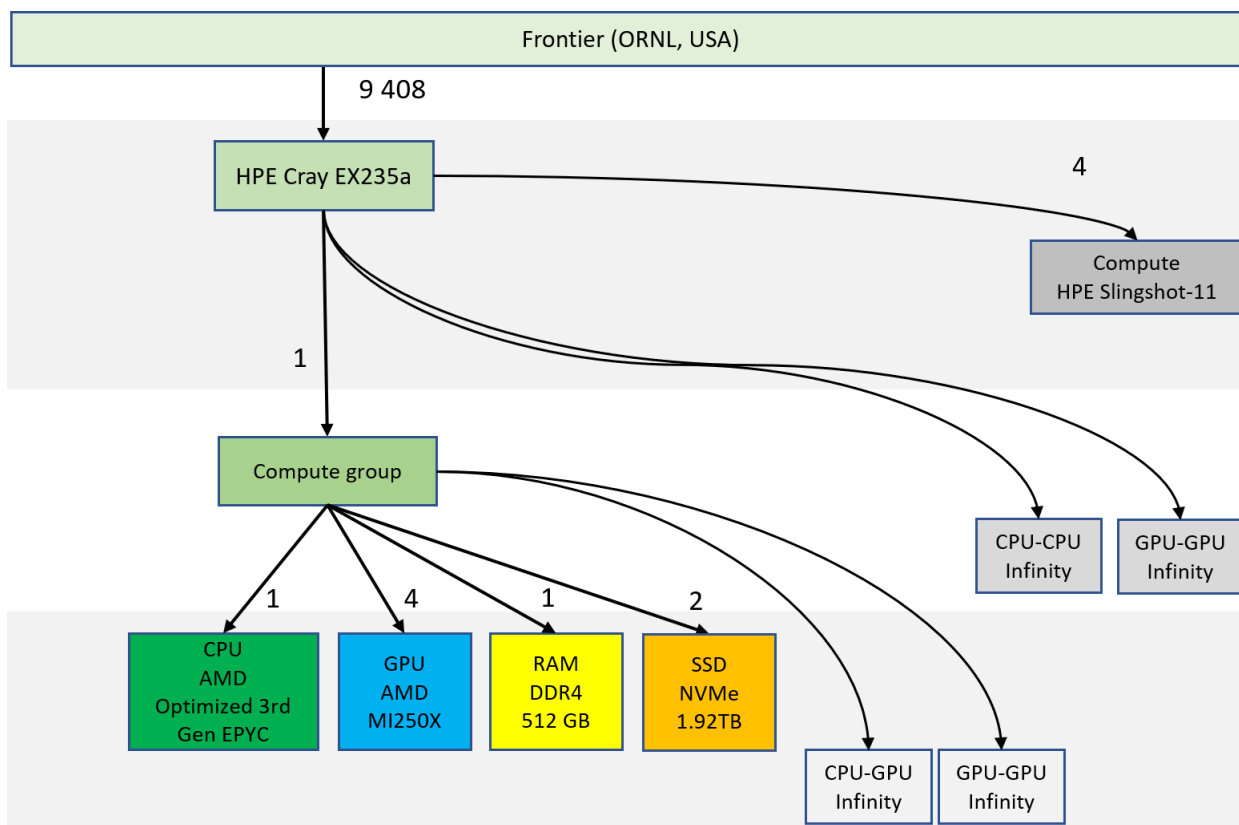The Summit system [10, 11] ranks fourth in the Top500 list of edition No. 59 dated 06.2022.

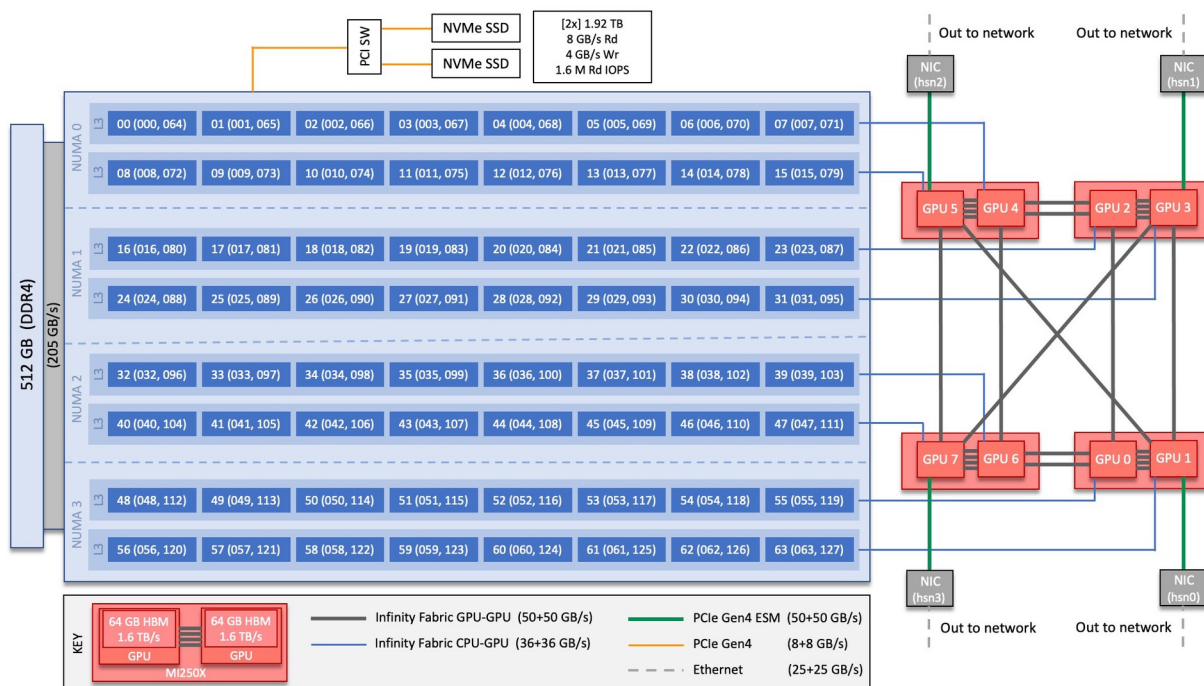**Fig. 3.** Description of the Frontier system model

Fig. 4. Frontier system node architecture [9]

The system is based on IBM POWER SYSTEM AC922, contains 4608 computing nodes, each of which contains 2 IBM POWER9 CPUs and 6 NVIDIA GV100 GPUs (Fig. 5). The structure of the node is remarkable (Fig. 6), the computing groups are pronounced, however, the
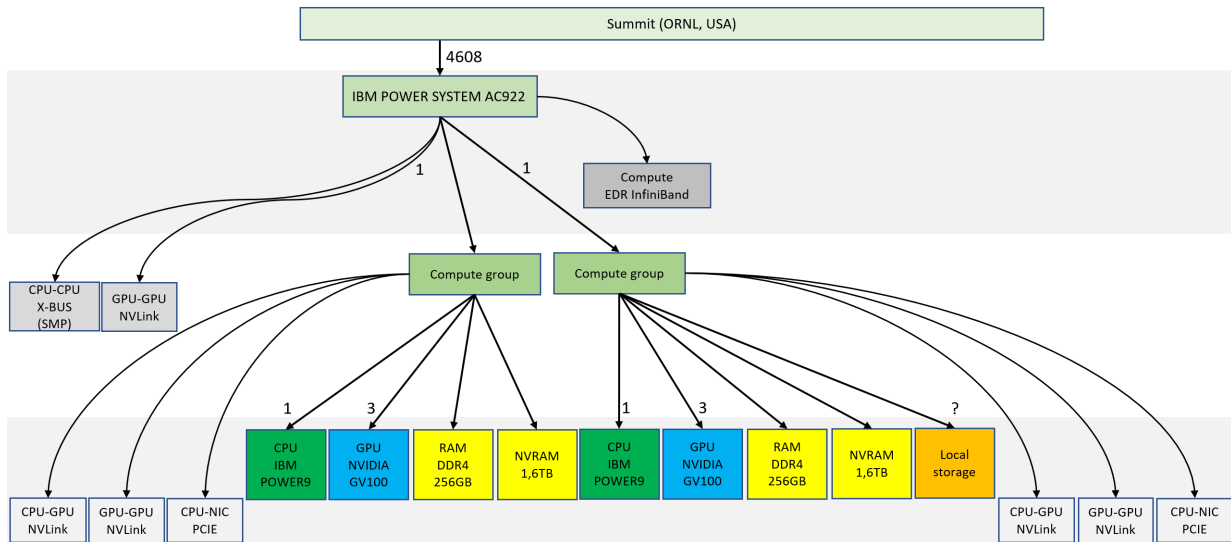
Fig. 5. Description of the Summit system model

X-BUS bus that unites them is SMP. The 22-core processors used are interesting, they already have NVLink support built in them, which avoids the traditionally used PCIe.
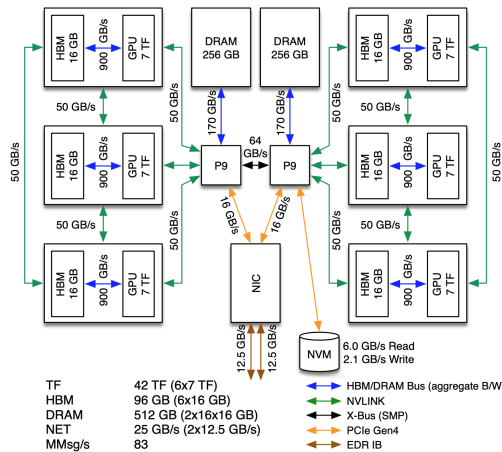

Fig. 6. Summit supercomputer node [12]

In addition to DDR4 memory, non-volatile NVRAM is also available on the node. The computing nodes of this system are interconnected by an EDR InfiniBand interconnect.

3.3. Selene, NVIDIA, USA

The Selene supercomputer [13, 14] was built by NVIDIA in accordance with the DGX SuperPOD principle [15] from DGX A100 nodes. By the way, one of the leading Russian Christofari Neo systems is built using the same SuperPOD technology.

The Selene system consists of 560 nodes, each of them with 8 high-performance NVIDIA Ampere A100 GPU accelerators allowed this system to take the 8th place in the TOP500 edition of June 2022.

Each node (NVIDIA DGX A100) is hybrid, and contains 2 AMD Epyc Rome 7742 CPUs and 8 NVIDIA A100 GPUs connected by an NVLink interconnect (Fig. 7). The AMD Epyc Rome 7742 processor contains 48 cores. The NVIDIA Ampere A100 GPU accelerator consists of

108 multiprocessors (SM), each of which can execute 64 CUDA cores or 4 "tensor" cores. Each node is equipped with 1024Gb of DDR4 RAM in NUMA mode, 2×1.92Tb local NVMe SSDs and 4×3.84Tb NVMe SSDs. The computing nodes of this system are interconnected by an HDR InfiniBand interconnect using the Fat Tree topology.
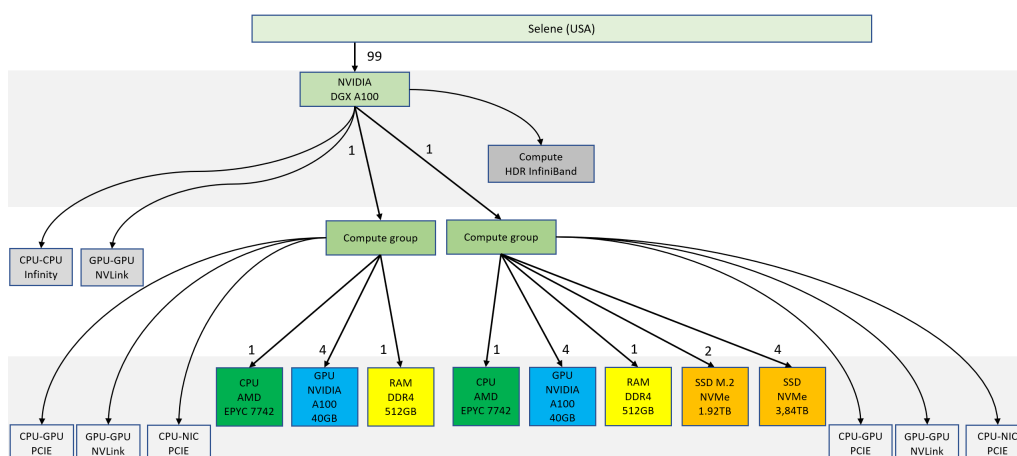
Fig. 7. Description of the Selene system model

Let us consider more spreading architectures that have developed historically as a result of multi-stage development. Both described systems can be found in the top of the Top50 list of the most powerful installations in Russia. Good news is that most of system holder from the Top50 list are open for discussions and are interested in developing the tools aimed at improving the efficiency of systems at all levels [15].

3.4. Lomonosov-2, Moscow State University, Russia

Lomonosov-2 with a performance of FP64 2.478 PFLOP/s (FP64 4.95 PFLOP/s peak) achieved on the HPL, the Lomonosov-2 system [16], installed at the HPC center of Moscow State University [17, 18], ranks the sixth in the Top50 list of the current edition (No. 37 dated 09.2022) [19, 20].
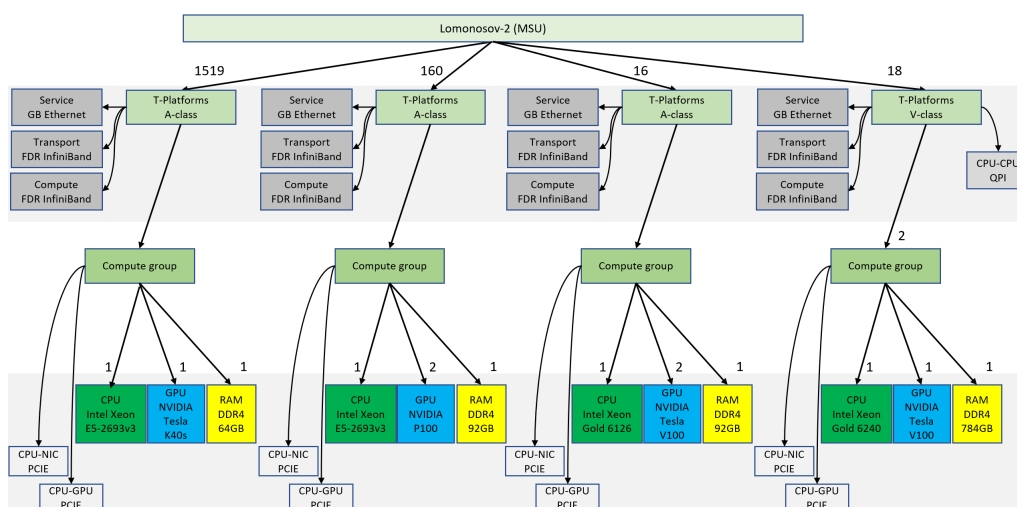
Fig. 8. Description of the Lomonosov-2 system model

The system has been the leader in the list of Russian supercomputers for many years and consists of a number of heterogeneous nodes based on the A-Class and V-Class of the T-Platform company. The nodes are predominantly uniprocessor, except for one segment (Fig. 8). There are no local disks on the nodes. The computing nodes of the Lomonosov-2 supercomputer are connected by an FDR InfiniBand interconnect, the transport network is FDR InfiniBand, and the service network is GB Ethernet.

3.5. cHARISMa, Higher School of Economics, Russia

In the tenth place in terms of performance achieved on HPL FP64 927.4 PFLOP/s (FP64 2027.27 PFLOP/s peak) in the Top50 list of the current edition (No. 37 dated 09.2022) is the cHARISMa system [21], installed at the HSE [22]. The system is heterogeneous, consisting of
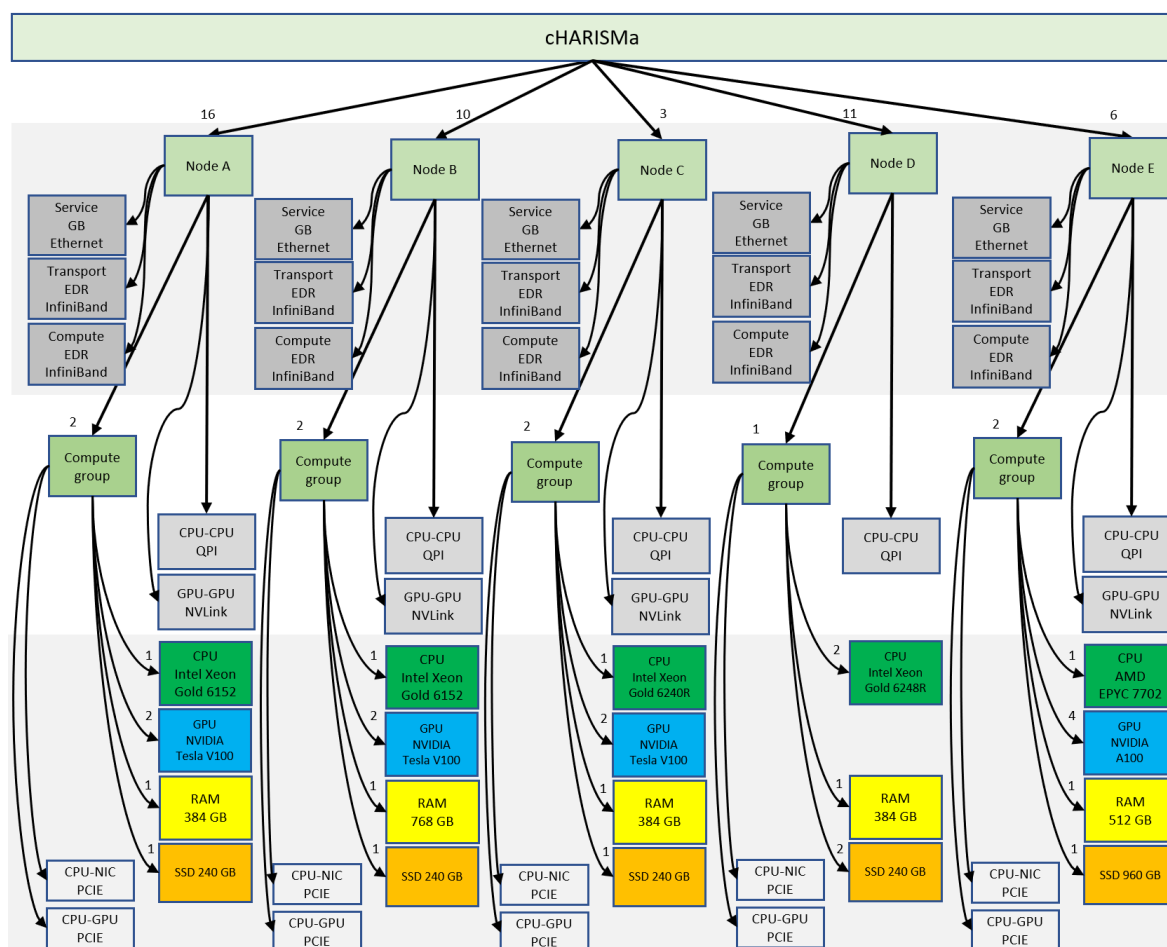

Fig. 9. Description of the cHARISMa system model

five types of nodes, both based on Intel and AMD processors (Fig. 9). Most part of the nodes are equipped with GPU like NVIDIA A100 and V100, small part are classic dual-processor servers. Heterogeneity is due to different classes of computational problems that are solved on the supercomputer. Each node type has two SSD drives in RAID 1. Computing nodes of the cHARISMa are connected by two aggregated EDR InfiniBand NICs (2*100Gb/s), and also has Gigabit Ethernet service and monitoring networks.

Conclusion

A basic version of the model of supercomputer systems has been developed for their comparative analysis. The developed model is hierarchical and supports the introduction of additional levels and relationships between levels.

Basic versions of models of supercomputer systems architectures make it possible to calculate the characteristics of supercomputer systems, in particular, higher-level parameters from the parameters of lower levels. It also makes it possible to calculate the characteristics of the components of supercomputer systems according to given general characteristics, in particular, by limiting the sample of systems by a certain parameter, as a result of screening, only records with a finite number of components remain, all of whose characteristics are known. The proposed approach allows processing the results of calculations: all performance results are stored in accordance with the model, which allows a comprehensive comparison of systems.

The model was tested on the leading and latest promising systems in Russia and the world: the top 10 Russian computing systems from the Top50 rating of the current edition were considered and described; reviewed and described the top 10 computing systems of the world from the Top500 rating of the current edition. Some of the models are given in this article. At this stage, the number of systems considered, does not allow for meaningful analysis of described systems but for some general conclusions. The description of the systems is carried out in order to test the approach and to provide its further adjustment. We will gladly report the analysis results next year after analysing considerable number of other outstanding systems.

From the point of view of the development of the proposed model, testing showed the feasibility of a detailed consideration at the next stage of work of the following aspects: description of the hierarchy of memory and the means of exchange with memory; ability to describe the nesting of the GPU (revealed on the example of AMD MI250X in the Frontier system, etc.).

From the point of view of the development of comparative analysis methods, it is necessary to ensure work with the promising criteria identified in the previous sections: for computing nodes, groups of nodes and systems as a whole — the specific amount of memory per computing core; for processors and accelerators — the specific amount of cache memory per computing core, the number of channels for working with memory.

In addition, it is necessary to ensure the possibility of ranking in the generated slices and ratings simultaneously by several characteristics.

References

1. Nikitenko D.A., Zheltkov A.A. The Top50 list vivification in the evolution of HPC rankings. Parallel Computational Technologies. Vol. 753 / ed. by L. Sokolinsky, M. Zymbler. Cham: Springer, 2017. P. 14–26. Communications in Computer and Information Science. DOI: 10.1007/978-3-319-67035-5_2.

2. Antonov A., Dongarra J., Voevodin V. AlgoWiki Project as an Extension of the Top500 Methodology. Supercomputing Frontiers and Innovations. 2018. Vol. 5, no. 1. P. 4–10. DOI: 10.14529/jsfi180101.

3. Antonov A.S., Nikitenko D.A., Voevodin V.V. Algo500 – A New Approach to the Joint Analysis of Algorithms and Computers. Lobachevskii J Math. 2020. Vol. 41, no. 6. P. 1435–1443. DOI: 10.1134/S1995080220080041.

4. Antonov A.S., Maier R.V. Development and Implementation of the Algo500 Scalable Digital Platform Architecture. Lobachevskii J Math. 2022. Vol. 43, no. 7. P. 837–847. DOI: 10.1134/S1995080222070058.

5. Kostenetskii P.S., Sokolinsky L.B. Simulation of Hierarchical Multiprocessor Database Systems. Programming and Computer Software. 2013. Vol. 39, no. 1. P. 10–24. DOI: 10.1134/S0361768813010040.

6. Zhang Y., Chen G., Sun G., Miao Q. Models of Parallel Computation: A Survey and Classification, Frontiers Comput. Sci. China. 2007. Vol. 1, no. 2. P. 156–165. DOI: 10.1007/s11704-007-0016-1.

7. Official Frontier website at ORNL. URL: `https://www.olcf.ornl.gov/frontier` (accessed: 07.11.2022).

8. Frontier system at Top500 rating. URL: `https://www.top500.org/system/180047` (accessed: 07.11.2022).

9. Frontier system User Guide. URL: `https://docs.olcf.ornl.gov/systems/frontier_user_guide.html` (accessed: 07.11.2022).

10. Official Summit website at ORNL. URL: `https://www.olcf.ornl.gov/frontier` (accessed: 07.11.2022).

11. Summit system at Top500 rating. URL: `https://www.top500.org/system/179397` (accessed: 07.11.2022).

12. Summit system User Guide. URL: `https://docs.olcf.ornl.gov/systems/frontier_user_guide.html` (accessed: 07.11.2022).

13. Official Selene website. URL: `https://www.nvidia.com/en-us/on-demand/session/gtcspring21-s31700/` (accessed: 07.11.2022).

14. Selene system at Top500 rating. URL: `https://www.top500.org/system/179842/` (accessed: 07.11.2022).

15. NVIDIA SuperPOD. URL: `https://www.nvidia.com/en-us/data-center/dgx-superpod/` (accessed: 07.11.2022).

16. Voevodin V.V., Chulkevich R.A., Kostenetskiy P.S., *et al.* Administration, Monitoring and Analysis of Supercomputers in Russia: a Survey of 10 HPC Centers. Supercomputing Frontiers and Innovations. 2021. Vol. 8, no. 3. P. 82–103. DOI: 10.14529/jsfi210305.

17. Voevodin V.V., Antonov A.S., Nikitenko D.A., *et al.* Supercomputer Lomonosov-2: Large scale, deep monitoring and fine analytics for the user community. Supercomputing Frontiers and Innovations. 2019. Vol. 6, no. 2. P. 4–11. DOI: 10.14529/jsfi190201.

18. Lomonosov-2 User's Guide. URL: https://parallel.ru/cluster/lomonosov2.html (accessed: 07.11.2022). (in Russian)

19. Voevodin V., Antonov A., Nikitenko D., *et al.* Lomonosov-2: Petascale supercomputing at Lomonosov Moscow State University. Contemporary High Performance Computing: From Petascale toward Exascale. Vol. 3. Boca Raton, United States: CRC Press, 2019. P. 305–330. DOI: 10.1201/9781351036863-12.

20. Lomonosov-2 system at Top50 rating. URL: http://top50.supercomputers.ru/systems/4568 (accessed: 07.11.2022). (in Russian)

21. Lomonosov-2 system at Top500 rating. URL: https://www.top500.org/system/178444/ (accessed: 07.11.2022).

22. HSE cHARISMa system at Top50 rating. URL: http://top50.supercomputers.ru/systems/6294 (accessed: 07.11.2022). (in Russian)

23. Kostenetskiy P.S., Chulkevich R.A., Kozyrev V.I. HPC Resources of the Higher School of Economics. Journal of Physics: Conference Series. 2021. Vol. 1740, no. 1. P. 012050. DOI: 10.1088/1742-6596/1740/1/012050.

ИЕРАРХИЧЕСКАЯ МОДЕЛЬ АРХИТЕКТУРЫ СУПЕРКОМПЬЮТЕРНЫХ СИСТЕМ ДЛЯ СРАВНЕНИЯ И РАНЖИРОВАНИЯ

© 2022 Д.А. Никитенко

Научно-исследовательский вычислительный центр
Московского государственного университета имени М.В. Ломоносова
(119991 Москва, Ленинские горы, д. 1, стр. 4)
E-mail: dan@parallel.ru
Поступила в редакцию: 07.11.2022

Задача сравнения возможностей вычислительных систем между собой и формирования различных рейтингов преследует множество возможных целей. Здесь и выявление трендов, и продвижение отработанных архитектур общего назначения, и демонстрация превосходства на определенном классе задач и др. Одного лишь описания достигнутой производительности для всех этих целей, конечно, недостаточно, в различных рейтингах и сравнения используются различные уровни абстракции и обобщения до того уровня, который бы позволил связать выявленные показатели производительности с теми или иными особенностями системы. На практике описания архитектурных особенностей систем в рейтингах достаточно скудны, и авторами работы решается задача по составлению формального описания вычислительных систем относительно высокого уровня, которое при этом позволило бы увеличивать требуемый уровень детализации, соответствующий целям прикладных исследований. Такая иерархическая модель описания систем предложена и апробирована на известных системах из списков Топ50 и Top500.

Ключевые слова: модель суперкомпьютерной системы, описание архитектуры суперкомпьютерных систем, сравнение производительности вычислительных систем, рейтинги производительности.

ОБРАЗЕЦ ЦИТИРОВАНИЯ

Литература

1. Nikitenko D.A., Zheltkov A.A. The Top50 list vivification in the evolution of HPC rankings // Parallel Computational Technologies. Vol. 753 / ed. by L. Sokolinsky, M. Zymbler. Cham: Springer, 2017. P. 14–26. Communications in Computer and Information Science. DOI: 10.1007/978-3-319-67035-5_2.

2. Antonov A., Dongarra J., Voevodin V. AlgoWiki Project as an Extension of the Top500 Methodology // Supercomputing Frontiers and Innovations. 2018. Vol. 5, no. 1. P. 4–10. DOI: 10.14529/jsfi180101.

3. Antonov A.S., Nikitenko D.A., Voevodin V.V. Algo500 — A New Approach to the Joint Analysis of Algorithms and Computers // Lobachevskii J Math. 2020. Vol. 41, no. 6. P. 1435–1443. DOI: 10.1134/S1995080220080041.

4. Antonov A.S., Maier R.V. Development and Implementation of the Algo500 Scalable Digital Platform Architecture // Lobachevskii J Math. 2022. Vol. 43, no. 7. P. 837–847. DOI: 10.1134/S1995080222070058.

5. Kostenetskii P.S., Sokolinsky L.B. Simulation of Hierarchical Multiprocessor Database Systems // Programming and Computer Software. 2013. Vol. 39, no. 1. P. 10–24. DOI: 10.1134/S0361768813010040.

6. Zhang Y., Chen G., Sun G., Miao Q. Models of Parallel Computation: A Survey and Classification // Frontiers Comput. Sci. China. 2007. Vol. 1, no. 2. P. 156–165. DOI: 10.1007/s11704-007-0016-1.

7. Official Frontier website at ORNL. URL: `https://www.olcf.ornl.gov/frontier` (дата обращения: 07.11.2022).

8. Frontier system at Top500 rating. URL: `https://www.top500.org/system/180047` (дата обращения: 07.11.2022).

9. Frontier system User Guide. URL: `https://docs.olcf.ornl.gov/systems/frontier_user_guide.html` (дата обращения: 07.11.2022).

10. Official Summit website at ORNL. URL: `https://www.olcf.ornl.gov/frontier` (дата обращения: 07.11.2022).

11. Summit system at Top500 rating. URL: `https://www.top500.org/system/179397` (дата обращения: 07.11.2022).

12. Summit system User Guide. URL: `https://docs.olcf.ornl.gov/systems/frontier_user_guide.html` (дата обращения: 07.11.2022).

13. Official Selene website. URL: `https://www.nvidia.com/en-us/on-demand/session/gtcspring21-s31700/` (дата обращения: 07.11.2022).

14. Selene system at Top500 rating. URL: `https://www.top500.org/system/179842/` (дата обращения: 07.11.2022).

15. Voevodin V.V., Chulkevich R.A., Kostenetskiy P.S., *et al.* Administration, Monitoring and Analysis of Supercomputers in Russia: a Survey of 10 HPC Centers // Supercomputing Frontiers and Innovations. 2021. Vol. 8, no. 3. P. 82–103. DOI: 10.14529/jsfi210305.

16. Voevodin V.V., Antonov A.S., Nikitenko D.A., *et al.* Supercomputer Lomonosov-2: Large scale, deep monitoring and fine analytics for the user community // Supercomputing Frontiers and Innovations. 2019. Vol. 6, no. 2. P. 4–11. DOI: 10.14529/jsfi190201.

17. Lomonosov-2 User's Guide. URL: `https://parallel.ru/cluster/lomonosov2.html` (дата обращения: 07.11.2022).

18. Voevodin V., Antonov A., Nikitenko D., *et al.* Lomonosov-2: Petascale supercomputing at Lomonosov Moscow State University // Contemporary High Performance Computing: From Petascale toward Exascale. Vol. 3. Boca Raton, United States: CRC Press, 2019. P. 305–330. DOI: 10.1201/9781351036863-12.

19. Lomonosov-2 system at Top50 rating. URL: `http://top50.supercomputers.ru/systems/4568` (дата обращения: 07.11.2022).

20. Lomonosov-2 system at Top500 rating. URL: `https://www.top500.org/system/178444/` (дата обращения: 07.11.2022).

21. HSE cHARISMa system at Top50 rating. URL: `http://top50.supercomputers.ru/systems/6294` (дата обращения: 07.11.2022).

22. Kostenetskiy P.S., Chulkevich R.A., Kozyrev V.I. HPC Resources of the Higher School of Economics // Journal of Physics: Conference Series. 2021. Vol. 1740, no. 1. P. 012050. DOI: 10.1088/1742-6596/1740/1/012050.

Никитенко Дмитрий Александрович, к.ф.-м.н., Научно-исследовательский вычислительный центр Московского государственного университета имени М.В. Ломоносова (Москва, Российская Федерация)