

## HyperQueue facilitates full utilisation of the computational resources of supercomputers

Ostrava, 9<sup>th</sup> January 2023 – Development is moving forward inexorably, and supercomputers that ruled the TOP 500 ten years ago are now forgotten; many of them are even literally on the scrap heap. New supercomputers are much more powerful and can contain over a hundred cores within a single compute node. When computing certain types of massively parallel tasks, running them can be complicated, leading to the potential of new supercomputers not being fully exploited. The solution is HyperQueue, a tool developed by scientists at IT4Innovations National Supercomputing Center for efficient organisation of computations to solve scientific problems.

IT4Innovations is involved in [the LIGATE project](#), which aims to use European supercomputers for drug design. This involves Computer-Aided Drug Design (CADD) using cutting-edge supercomputers, and in the future, exascale systems are also envisaged. Specifically, the LIGATE project deals with the screening of huge quantities of molecules, which is the kind of task that can benefit from exploiting the full potential of compute nodes. Thus, scientists from IT4Innovations at VSB – Technical University of Ostrava have delivered a unique solution in the form of HyperQueue, which is a tool that allows a large number of computational tasks to be run efficiently and easily on modern heterogeneous supercomputers. What exactly does this mean?

*“Modern supercomputers are characterised by combining different computer architectures and containing a large number of heterogeneous resources. Using them efficiently with traditional computational tools can be difficult,”* explains Jan Martinovič from the Advanced Data and Simulation Lab at IT4Innovations, adding: *“That’s why we created the HyperQueue tool, which simplifies the use of supercomputers with complex resources and provides a simple interface for entering computational tasks. At the same time, it can efficiently use the available computational resources of a supercomputer.”*

Additionally, Branislav Jansík, IT4Innovations' Supercomputing Services Director states that *“HyperQueue has been successfully deployed and tested on several supercomputers with different hardware architectures. These include the EuroHPC supercomputer Karolina, operated by IT4Innovations, the Czech National Supercomputing Center, and Europe’s most powerful supercomputer, LUMI, operated by the Finnish CSC.”*

*“HyperQueue has been immensely useful in scaling up existing workflows run by researchers, requiring little to no changes to their code. Working with HyperQueue is very straightforward, and we have even added native support for HyperQueue to the bioinformatics workflow manager Nextflow, so that users can benefit from HyperQueue when running their genomics workflows without having to even learn HyperQueue. It’s nice for once to encounter a tool which plays nicely with the system scheduler and does not negatively affect system stability when doing very large workflows,”* says Henrik Nortamo, Applications Specialist at [CSC](#).

HyperQueue is also being deployed on supercomputers at the Italian CINECA as part of the LIGATE project. It has the potential to become a primary tool to help efficiently schedule a large number of jobs that would not be otherwise able to individually use the full capacity of a supercomputer's compute node. *“We found HyperQueue an easy-to-use tool that simplifies deployment on novel HPC machines, making the use of resources more efficient for workloads composed of many small tasks.*

*This is exactly the case for the in-silico virtual screening application we are developing in the context of the LIGATE project, and it has already been used to help fight the COVID-19 pandemic,”* added Gianluca Palermo from Politecnico di Milano, who is also the Technical Manager of the LIGATE project.

### Technical details

Modern HPC clusters contain a large number of heterogeneous resources that provide vast amounts of computational power. It is challenging to design monolithic programs that can leverage that performance potential effectively (e.g., by scaling to hundreds of cores); HPC users often design their computational workflows as a set of smaller, interdependent tasks that use only a fraction of the resources of a single cluster node. Yet executing these workflows on HPC clusters in the presence of job managers such as Torque/PBS or Slurm can be challenging. They can impose limits on the concurrent execution of multiple tasks on a single node, thus hampering node utilisation, and their design in general is not accustomed to an enormous amount of smaller, less resource-intensive tasks, which can lead to the manager being overloaded.

HyperQueue is an HPC task execution framework that offers a solution to this problem. It allows users to submit tasks in a simple way outside of a computational job. HyperQueue then takes care of asking for computational resources from the job manager and executing the tasks on all available computing nodes. It uses a sophisticated scheduler to load balance the tasks while taking into account arbitrary resource specifications and current node utilisation. As an example use-case, it is trivial to define a computation with many tasks that use a small number of cores and execute it on a cluster with very powerful nodes (with e.g. 128 cores) whilst achieving very high node utilisation out of the box.

For more information, see: <https://github.com/it4innovations/hyperqueue>



*The LIGATE project receives funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 956137. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Italy, Sweden, Austria, the Czech Republic, and Switzerland. This project receives funding from the Ministry of Education, Youth and Sports of the Czech Republic (ID: MC2102).*

**Project contact person**

Jan Martinovič, Design and integration of AI, Big Data and workflow components work package leader of the LIGATE project and Head of the Advanced Data Analysis and Simulation Lab of IT4Innovations National Supercomputing Center, VSB – Technical University of Ostrava, [jan.martinovic@vsb.cz](mailto:jan.martinovic@vsb.cz)

**Media contact**

Zuzana Červenková, Spokesperson for IT4Innovations National Supercomputing Center  
[zuzana.cervenkova@vsb.cz](mailto:zuzana.cervenkova@vsb.cz)

Phone: +420 602 593 335

---

**IT4Innovations National Supercomputing Center** at VSB – Technical University of Ostrava is a leading research, development, and innovation centre active in the fields of High-Performance Computing (HPC), Data Analysis (HPDA), and Artificial Intelligence (AI) and their application to other scientific fields, industry, and society. Since 2013, IT4Innovations has been operating the most powerful supercomputing systems in the Czech Republic, which are provided to Czech and foreign research teams from both academia and industry. Together with the CESNET and CERIT-SC institutions, IT4Innovations constitutes [e-INFRA\\_CZ](#), a strategic research infrastructure of the Czech Republic. At present, IT4Innovations runs three supercomputers: [Karolina](#) (15.7 PFlop/s, installed in the summer of 2021), [Barbora](#) (849 TFlop/s, installed in the autumn of 2019), and the specialised [NVIDIA DGX-2](#) system for artificial intelligence computation, (130 TFlop/s and 2 PFlop/s in AI, installed in the spring of 2019). The Karolina supercomputer, acquired as part of the EuroHPC Joint Undertaking, was the most powerful supercomputer in the Czech Republic on the day of its installation, 69<sup>th</sup> in the world (on the TOP500 list), and ranked 8<sup>th</sup> in the Green500 ranking of the most energy-efficient supercomputers. The first supercomputers installed at IT4Innovations were Anselm (94 TFlop/s, installed in the summer of 2013), and Salomon (2 PFlop/s, installed in the summer of 2015 and decommissioned in 2021). The key research areas of IT4Innovations include big data processing and analysis, machine learning, development of parallel scalable algorithms, solution of computationally demanding engineering problems, advanced visualisation, virtual reality, modelling for nanotechnologies, and material design. For more information, see [www.it4i.cz](http://www.it4i.cz).