

JUWELS Booster – Early User Experiences

Andreas Herten

a.herten@fz-juelich.de

Jülich Supercomputing Centre, Forschungszentrum Jülich GmbH
Jülich, Germany

ABSTRACT

Over the last few years, GPUs became ubiquitous in HPC installations around the world. Today, they provide the main source of performance in a number of Top500 machines – for example Summit, Sierra, and JUWELS Booster. Also for the upcoming Exascale era, GPUs are selected as key enablers and will be installed numerously. While individual GPU devices already offer plenty of performance ($O(10)$ TFLOP/ s_{FP64}), current and next-generation supercomputers employ them in the thousands. Using these machines to the fullest extend means not only utilizing individual devices efficiently, but using the entire interconnected system of devices thoroughly.

JUWELS Booster is a recently installed Tier-0/1 system at Jülich Supercomputing Centre (JSC), currently the 7th-fastest supercomputer in the world, and the fastest in Europe. JUWELS Booster features 936 nodes, each equipped with 4 NVIDIA A100 Tensor Core GPUs and 4 Mellanox HDR200 InfiniBand HCAs. The peak performance of all GPUs together sums up to 73 PFLOP/s and it features a DragonFly+ network topology with 800 Gbit/s network injection bandwidth per node.

During installation of JUWELS Booster, a selected set of applications were given access to the system as part of the *JUWELS Booster Early Access Program*. To prepare for their first compute time allocation, scientific users were able to gain first experiences on the machine. They gave direct feedback to the system operations team during installation and beyond. Close collaboration was facilitated with the application support staff of JSC, giving unique insights into the individual processes of utilizing a brand-new large-sale system for a first time. Likewise, performance profiles of applications could be studied and collaboratively analyzed, employing available tools and methods. Performance limiters of the specific application on the platform were identified and proposals for improvement developed.

This talk will present first experiences with JUWELS Booster and the applications utilizing the system during its first months. Applied methods for onboarding, analysis, and optimization will be shown and assessed. Highlights of the state of the art of performance analysis and modeling for GPUs will be presented with concrete examples from the JUWELS Booster Early Access Program.

CCS CONCEPTS

• **Computing methodologies** → *Parallel programming languages*; **Graphics processors**; • **Computer systems organization** → *n-tier architectures*.

KEYWORDS

GPU, GPU Computing, Profiling, Parallel Computing, Distributed Computing, Exascale

ACM Reference Format:

Andreas Herten. 2021. JUWELS Booster – Early User Experiences. In *Proceedings of the 2021 Performance Engineering, Modelling, Analysis, and Visualization On Strategy (PERMAVOST '21), June 25, 2021, Virtual Event, Sweden*. ACM, New York, NY, USA, 1 page. <https://doi.org/10.1145/3452412.3462752>

BIOGRAPHY



Andreas Herten is the group leader of the *Accelerating Devices* Lab at Jülich Supercomputing Centre in which he investigates usage of accelerators by HPC applications. Within the *NVIDIA Application Lab at Jülich* he supports scientific users in analyzing, porting, and optimizing their workloads to GPUs. He was the coordinator of the *JUWELS Booster Early Access Program* and has extensive knowledge of the machine. He is involved with various international collaborations for improving performance and scalability on GPUs.

Before working in HPC, Andreas did his PhD in experimental particle physics, developing methods for real-time data suppression on GPUs.

Apart from massively parallel computing with GPUs, Andreas is interested in making HPC more accessible. He teaches different programming courses on GPU computing and beyond.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

PERMAVOST '21, June 25, 2021, Virtual Event, Sweden

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8387-5/21/06.

<https://doi.org/10.1145/3452412.3462752>