
JÜLICH
Forschungszentrum

JUWELS BOOSTER EARLY ACCESS PROGRAM
ENABLING 1ST APPLICATIONS FOR EUROPE'S #1 SYSTEM

17 December 2020 | Andreas Herten | JSC EoYC 2020; Jülich Supercomputing Centre, Forschungszentrum Jülich

JÜLICH Forschungszentrum
JÜLICH SUPERCOMPUTING CENTRE

JUWELS Overall Architecture

JUWELS Cluster (2018)

- 2511 compute nodes (2× Skylake)
- 48 GPU nodes (4× V100 w/ NVLink2)
- Mellanox EDR 100 Gbit/s network, fat-tree topology (1:2@L1)
- 12 PFLOP/s

JUWELS Booster (2020)

- 936 compute nodes (2× AMD Rome, 4× A100 w/ NVLink3)
- Mellanox HDR 200 Gbit/s network, DragonFly+ topology
- 73 PFLOP/s

JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE

JUWELS Overall Architecture

JUWELS Cluster (2018)

- 251 [illegible]ke)
- 48 [illegible]nk2)
- Me[illegible]rk, fat-
- 12 [illegible]

TOP 500 The List.

Top500 Nov-2020:

- #1 Europe
- #7 World
- #3* Green500

JUWELS Booster (2020)

- 936 compute nodes (2× AMD Rome, 4× A100 w/ NVLink3)
- Mellanox HDR 200 Gbit/s network, DragonFly+ topology
- 73 PFLOP/s

JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE

JUWELS Overall Architecture

More System Details:
⇒ Damian's Talk

- 936 compute nodes (2× AMD Rome, 4× A100 w/ NVLink3)
- Mellanox HDR 200 Gbit/s network, DragonFly+ topology
- 73 PFLOP/s


Early Access Program

Overview

- Started in early 2020
- Selected 14 applications from various scientific domains
  - Aimed for applications that could use JUWELS Booster at scale
  - Some teams already use JUWELS Cluster, others new
- **Offer**: Use JUWELS Booster before general access; **Request**: Help improve system
- Endeavor of many parts in **JSC** and beyond
  - NVIDIA Application Lab: Steering, GPU optimization, application support, system support
  - Application support, Simulation Labs
  - Performance Optimisation and Productivity team
  - System operations team
  - Vendors: NVIDIA, ParTec, Atos

Timeline to Booster

- Preparation Timeline

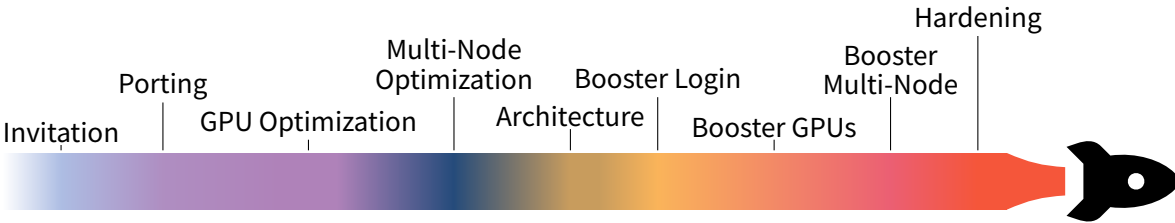

Timeline to Booster

- Preparation Timeline
- Additionally: events

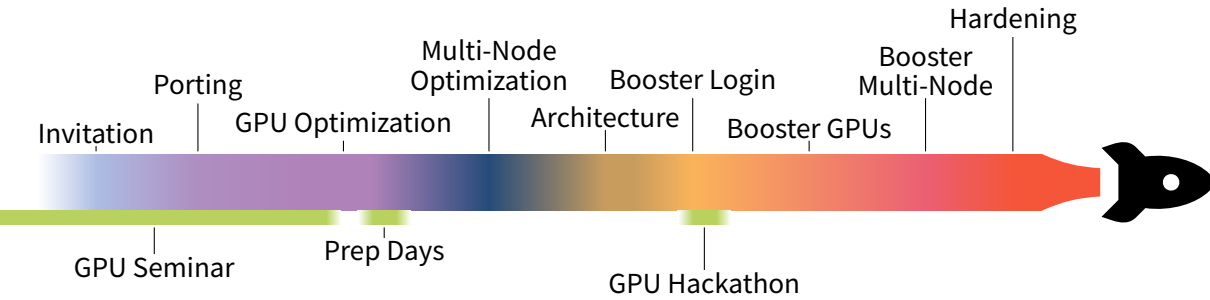

Timeline to Booster

- Preparation Timeline
- Additionally: events

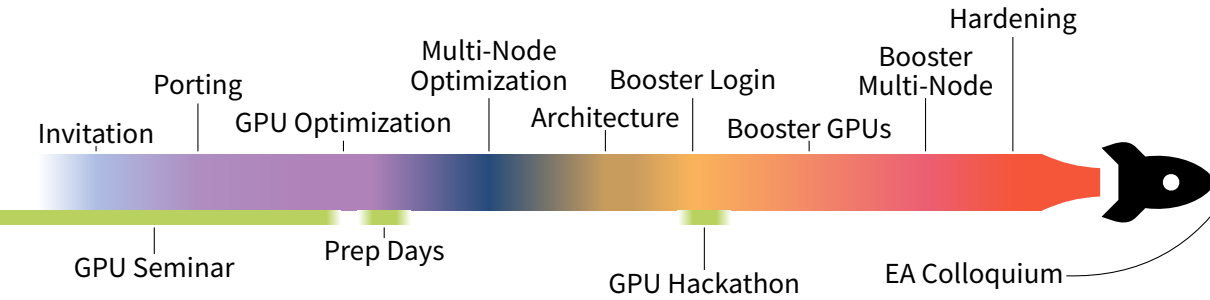

Applications I

Climate/Meteo/Hydro (*ESM*)

DeepACF
- High-resolution Weather Forecast Based on Deep Learning
- JSC: Bing Gong, Michael Langguth, Amirpasha Mozaffari, Martin Schultz, Scarlet Stadtler

ICON
- Next-Generation Physical Weather and Climate Models
- MPI Met: Luis Kornblueh; NVIDIA: Dmitry Alexeev

MPTRAC
- Massive Parallel Trajectory Calculations of Volcanic Emissions
- JSC: Sabine Grießbach, Lars Hoffmann

ParFlow
- Surface, Soil, Ground Water Flow
- IBG-3: Jaro Hokkanen, Stefan Kollet

Biological Matter

Amber
- Drug Binding over Biologically Relevant Timescales (MD)
- JSC/HHU: Holger Gohlke, Christopher Pfleger, Michele Bonus

SOMA
- Kinetics of Nanomaterial Formation (Soft Matter)
- U Göttingen: Ludwig Schneider, Niklas Blagojevic

Applications I

Climate/Meteo/Hydro (*ESM*)

DeepACF High-resolution Weather Forecast Based on Deep Learning
JSC: Bing Gong, Michael Langguth, Amirpasha Mozaffari, Martin Schultz, Scarlet Stadtler

ICON Next-Generation Physical Weather and Climate Models
MPI Met: Luis Kornblueh; NVIDIA: Dmitry Alexeev

MPTRAC Massive Parallel Trajectory Calculations of Volcanic Emissions
JSC: Sabine Grießbach, Lars Hoffmann

ParFlow Surface, Soil, Ground Water Flow
IBG-3: Jaro Hokkanen, Stefan Kollet

Biological Matter

Amber Drug Binding over Biologically Relevant Timescales (MD)
JSC/HHU: Holger Gohlke, Christopher Pfleger, Michele Bonus

SOMA Kinetics of Nanomaterial Formation (Soft Matter)
U Göttingen: Ludwig Schneider, Niklas Blagojevic

Applications II

PIConGPU ⚛ Plasma Simulations for Next Generation Particle Accelerators (Plasma)
HZDR: Alexander Debus, Anton Lebedev, Rene Widera, Michael Bussmann

JUQCS-G ⚛ Simulating Universal Quantum Computer (Quantum)
JSC: Hans De Raedt, Kristel Michielsen, Dennis Willsch

E-train ⚛ Understanding Learning Processes in Brain (Neuro)
U Graz: Franz Scherr, Wolfgang Maass; U Sussex: James Knight; INM-6: Sacha van Albada

NBODY6++GPU ⚛ Dense Star Clusters and Gravitational Waves (Astro)
U Heidelberg: Rainer Spurzem

Lattice QCD

Bonn ⚛ Flavour Singlet Structure of Hadrons
U Bonn: Simone Bacchio, Bartosz Kostrzewa, Carsten Urbach

Wuppertal ⚛ SignQCD – Studying the Hottest Man-made Liquid
U Wuppertal: Szabolcs Borsányi, Kalman Szabo

Bielefeld ⚛ HotQCD – Studying Extreme States of Matter
U Bielefeld: Christian Schmit, Dennis Bollweg, Frithjof Karsch

Regensburg ⚛ Baryons with Charm
Peter Boyle, Christoph Lehner, Gunnar Bali, Sara Collins

Applications II

PIConGPU Plasma Simulations for Next Generation Particle Accelerators (Plasma)
HZDR: Alexander Debus, Anton Lebedev, Rene Widera, Michael Bussmann

JUQCS-G Simulating Universal Quantum Computer (Quantum)
JSC: Hans De Raedt, K
Dennis Willsch

E-train Understanding Learning Processes in Brain (Neuro)
U Graz: Franz Scherr, Wolfgang Maass; U Sussex: James Knight; INM-6: Sacha van Albada

NBODY6++GPU Dense Star Clusters and Gravitational Waves (Astro)
U Heidelberg: Rainer Spurzem

→ Details on each app online

Lattice QCD

Bonn Flavour Singlet Structure of Hadrons
U Bonn: Simone Bacchio, Bartosz Kostrzewa, Carsten Urbach

Wuppertal SignQCD – Studying the Hottest Man-made Liquid
Wuppertal: Szabolcs Borsányi, an Szabo

Bielefeld HotQCD – Studying Extreme States of Matter
U Bielefeld: Christian Schmit, Dennis Bollweg, Frithjof Karsch

Regensburg Baryons with Charm
Peter Boyle, Christoph Lehner, Gunnar Bali, Sara Collins

Application Status

To my knowledge and opinion

ESM:DeepACF

ESM:ICON

ESM:MPTRAC

ESM:ParFlow

Bio:Amber

Bio:SOMA

NBODY6++GPU

JUQCS-G

E-train

PIConGPU

LQCD:Bonn

LQCD:Wuppertal

LQCD:Bielefeld

LQCD:Regensburg

Application Status

To my knowledge and opinion

ESM:DeepACF

ESM:ICON

ESM:MPTRAC

ESM:ParFlow

Bio:Amber

Bio:SOMA

NBODY6++GPU

JUQCS-G

E-train

PIConGPU

LQCD:Bonn

LQCD:Wuppertal

LQCD:Bielefeld

LQCD:Regensburg

Programming Model

Application Status

To my knowledge and opinion

ESM:DeepACF	Lib:DL
ESM:ICON	OpenACC
ESM:MPTRAC	OpenACC
ESM:ParFlow	CUDA C
Bio:Amber	Lib
Bio:SOMA	OpenACC
NBODY6++GPU	CUDA Fortran
JUQCS-G	CUDA Fortran
E-train	Lib:DL
PIConGPU	CUDA C++
LQCD:Bonn	Lib:QUDA
LQCD:Wuppertal	Lib:QUDA
LQCD:Bielefeld	CUDA C++
LQCD:Regensburg	Lib:Grid

Programming Model

Application Status

To my knowledge and opinion

ESM:DeepACF	Lib:DL	1		JUQCS-G	CUDA Fortran	2	
ESM:ICON	OpenACC	2		E-train	Lib:DL	2	
ESM:MPTRAC	OpenACC	1		PIConGPU	CUDA C++	3	
ESM:ParFlow	CUDA C	1		LQCD:Bonn	Lib:QUDA	3	
Bio:Amber	Lib	3		LQCD:Wuppertal	Lib:QUDA	2	
Bio:SOMA	OpenACC	3		LQCD:Bielefeld	CUDA C++	3	
NBODY6++GPU	CUDA Fortran	2		LQCD:Regensburg	Lib:Grid	3	

Programming Model | **GPU level:** | **1** – Just started | **2** – Started in months before JWB | **3** – Already around for long | **Lab Involvement:** | – Hands-on | – Meetings (regular) | – Meetings (rare) | – Hackathons | – Support (mostly chat) | – Bug reporters

Lessons Learned

- EA Program had to be tailored individually around each application
- Regular/fast updates of development environment was necessary and helpful
- Knowledge dissemination programs well-liked (talks, newsletter, overview documentations, chat)
- Challenging to schedule EA runs, bug hunts, and low-level system tests at same time
- EA Program very well received (large number of teams, 380 000 h_{node} consumed)


Early Performance Results

Disclaimer

- Following results obtained on fresh JUWELS Booster
- …while system integration work was done at same time
- System is being be tuned and improved
- …also due to results obtained by EA applications!

- Software used
  - GCC 9.3.0
  - CUDA 11.0 (with CUDA Driver 450.80.02)
  - NVHPC 20.7
  - ParaStationMPI 5.4.7 (with UCX 1.8.1)

Soft Matter: SOMA

- **SOMA**: Soft, coarse-grained Monte-Carlo Acceleration
  L. Schneider and M. Müller, Comput. Phys. Commun. 235C 463–476 (2019) and GPU Seminar Talk
- Kinetics of nanomaterial formation; multi-component polymer systems (battery materials, membranes, …)
- Unique: Resolve details of polymer, but study lengths relevant to engineering
- Team: L. Schneider, N. Blagojevic, L. Pigard, M. Müller, et al

→ `gitlab.com/InnocentBug/SOMA/`

- C, OpenACC, MPI
- Frequent JUWELS user

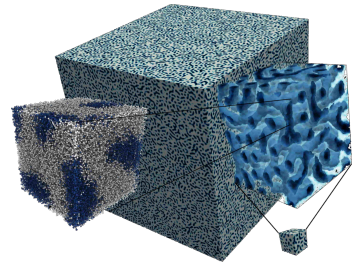

GEORG-AUGUST-UNIVERSITÄT GÖTTINGEN

SOMA
yes, soft matters

JÜLICH Forschungszentrum
JÜLICH SUPERCOMPUTING CENTRE

SOMA Performance Results

Comparison of GPU Generations

- Long experience with various GPU architectures

→ **Update to new generations early!**

- Some algorithmic changes between generations; also feature additions
- *PTPS: Particle Timesteps Per Second*

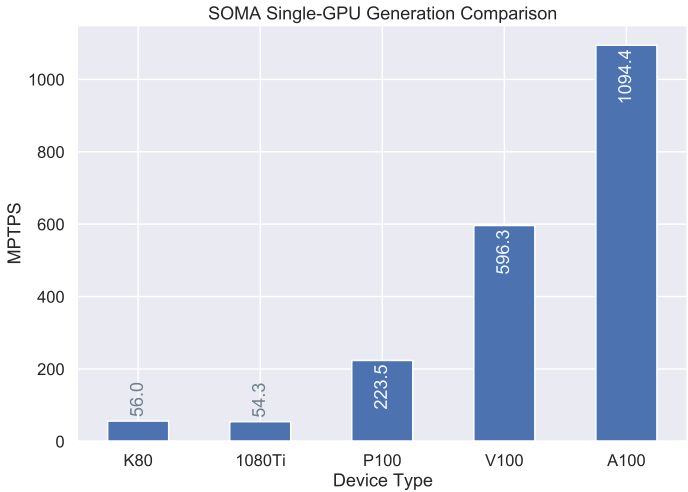

Data provided by SOMA / L. Schneider

SOMA Performance Results

New Method for Scaling

- **Scale of Booster: New algorithms, implementations with more scalability!**
- New project for Booster: *String* Method
- String-coupled SOMA ensemble simulation
- Master thesis of N. Blagojevic

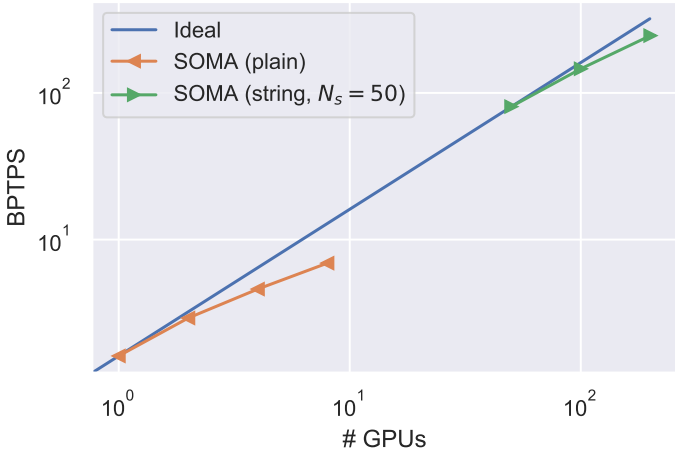

Data provided by SOMA / N. Blagojevic

Earth-system modelling: ParFlow

- **ParFlow**: Numerical model for groundwater and surface water flow

  J. Hokkanen, S. Kollet, et al, EGU General Assembly 2020, 4–8 May 2020, EGU2020-12904, and GPU Seminar Talk

- Model hydrologic processes, hill-slope to continental scale; forecasting, water cycle research, climate change; since 1990s
- Finite-difference scheme with implicit time integration
- Team: J. Hokkanen, S. Kollet

→ `parflow.org`

- C, C++, CUDA, MPI
- Fresh GPU port in prepartion for Booster

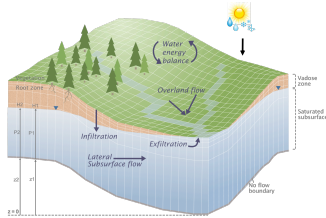

JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE

ParFlow Performance Results

Single-Node Performance

- Comparing CPU of Booster node with GPUs
- **Good speed-up, max. 29×**
- Memory pool (*RMM*) gives extra boost
- Larger problem sizes solvable per node

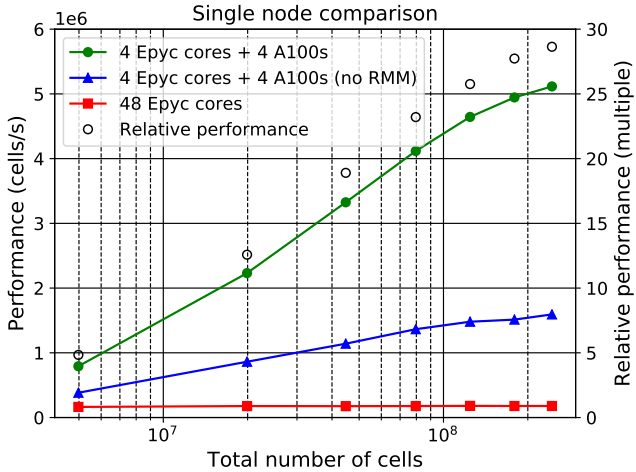

Plot provided by ParFlow / J. Hokkanen

ParFlow Performance Results

Weak Scaling

- Fixed problem size per node
- **26× speed-up achieved over $\mathcal{O}(100)$ nodes**

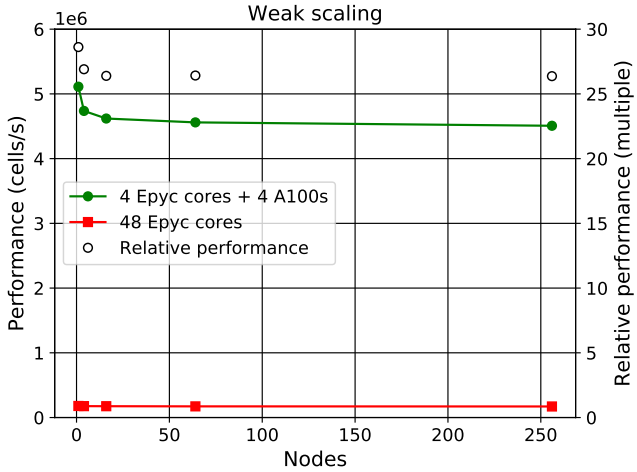

Plot provided by ParFlow / J. Hokkanen

Quantum Computing: JUQCS

- **JUQCS**: Jülich Universal Quantum Computer Simulator
  De Raedt et al., Comp. Phys. Comm. 237 47–61 (2019)
- Universal quantum computing on digital computer
- Network-, memory-intensive computations
- Team: Research group Quantum Information Processing
- Fortran, CUDA Fortran
- Frequent JUWELS user

JÜLICH SUPERCOMPUTING CENTRE

JUQCS

- **40 qubits:**
  - $>$ 16 TiB memory needed
    $\rightarrow$ 512 A100s
  - Each quantum operation: Update states, 8 TB transfer
- Weak scaling: Compute constant, MPI as expected
- Strong scaling: Still investigate DragonFly+ topology

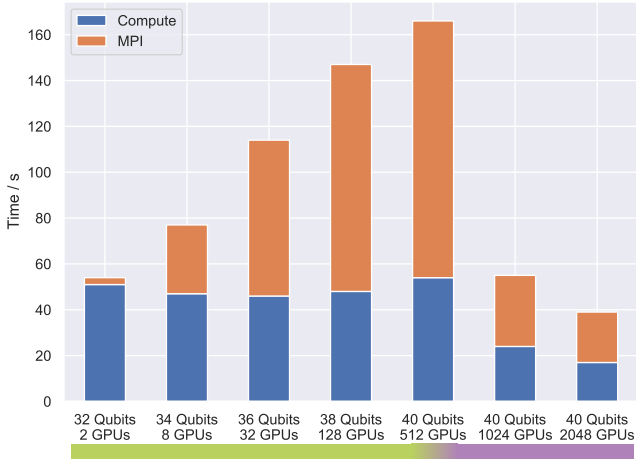

Data provided by JuQCS / H. De Raedt

Summary and Conclusions

Summary

- JUWELS Booster: European flagship system based on A100 GPUs
- In production since end of November, applications prepared through Early Access Program
- Early performance results are very encouraging

Summary

- JUWELS Booster: European flagship system based on A100 GPUs
- In production since end of November, applications prepared through Early Access Program
- Early performance results are very encouraging

- Early Access Colloquium 20 January 2021:
  `https://indico-jsc.fz-juelich.de/event/166/`

Summary

- JUWELS Booster: European flagship system based on A100 GPUs
- In production since end of November, applications prepared through Early Access Program
- Early performance results are very encouraging

- Early Access Colloquium 20 January 2021:
  https://indico-jsc.fz-juelich.de/event/166/

Acknowledgments

- JSC Application-Oriented Technology Development: Kaveh Haghighi-Mood
- JSC High Performance Systems: Dorian Krause, Damian Alvarez, Benedikt von St. Vieth
- NVIDIA Collaborators: Markus Hrywniak, Jiri Kraus, Mathias Wagner
- Participants of Early Access Program, especially

SOMA	Ludwig Schneider, Louis Pigard, Niklas Blagojevic
ParFlow	Jaro Hokkanen
JuQCS	Hans De Raedt

Summary

- JUWELS Booster: European flagship system based on A100 GPUs
- In production since end of November, applications prepared through Early Access Program
- Early performance results are very encouraging
- Early Access Colloquium 20 January 2021: `https://indico-jsc.fz-juelich.de/event/166/`

Thank you for your attention!
a.herten@fz-juelich.de

Acknowledgments

- JSC Application-Oriented Technology Development: Kaveh Haghighi-Mood
- JSC High Performance Systems: Dorian Krause, Damian Alvarez, Benedikt von St. Vieth
- NVIDIA Collaborators: Markus Hrywniak, Jiri Kraus, Mathias Wagner
- Participants of Early Access Program, especially
  - SOMA Ludwig Schneider, Louis Pigard, Niklas Blagojevic
  - ParFlow Jaro Hokkanen
  - JuQCS Hans De Raedt

Appendix

Appendix

- Network Performance
- SOMA Supplementary
- LQCD: Bonn
- References

Appendix

Network Performance

Network Performance

OSU Micro-Benchmarks: Bandwidth

- OSU Microbenchmarks: device-device bandwidth (`osu_bw D D`)
- Good results, expected limiters
- Intra-node: NVLink3 bandwidth
- Inter-node: HDR200 bandwidth
- Model fits show 2 regimes (---/ ---)

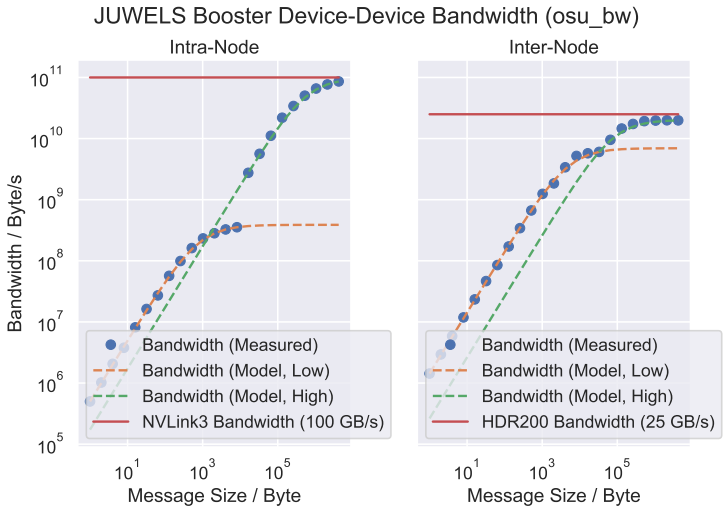

Appendix

SOMA Supplementary

SOMA Supplementary

Kernel Comparison: Memory Chart

- Many random accesses

→ Benefit from larger L1, L2 caches

→ More FP64 throughput

- Knock-on effect: less memory traffic
- **Kernel runtime:**

V100	25.8 ms
A100	21.5 ms
A100*	18.9 ms

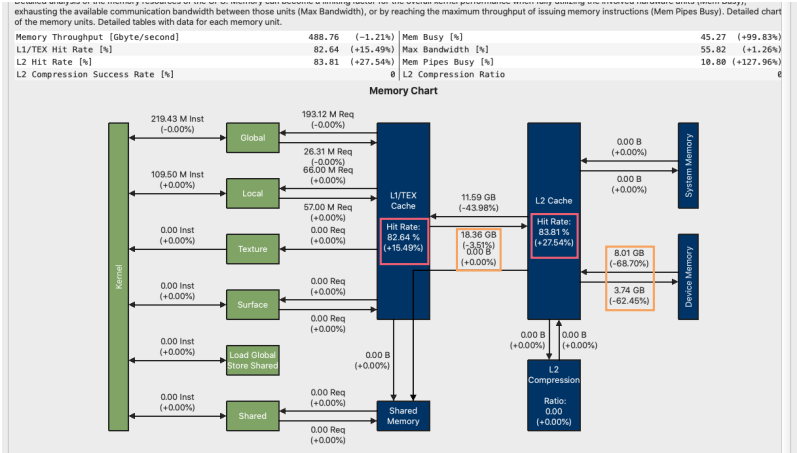

Nsight Compute profile provided by SOMA

Appendix

LQCD: Bonn

LQCD: Bonn

- **ETMC**: Extended Twisted Mass Collaboration
  C. Alexandrou and S. Bacchio et al, Phys. Rev. D 101 094513 (2020)
- Study of the Flavour Singlet Structure of Hadrons
- Team: S. Bacchio, B. Kostrzewa, et al; Uni Bonn, Uni Cyprus, Cyprus Institute, Uni Rome, …

→ `github.com/etmc`, PLEGMA, QUDA, tmLQCD

- C/C++, CUDA, MPI, OpenMP
- Frequent JUWELS user

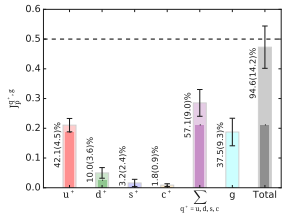

JÜLICH Forschungszentrum
JÜLICH SUPERCOMPUTING CENTRE

LQCD: Bonn

Comparison of GPU HPC Machines

- Multigrid inversion
- Mean time-to-solution, spread
- Systems

  Piz Daint Haswell, P100; DragonFly

  Marconi100 POWER9, V100; DragonFly+
- JUWELS Booster: Low time to solution; but large spread (being investigated)

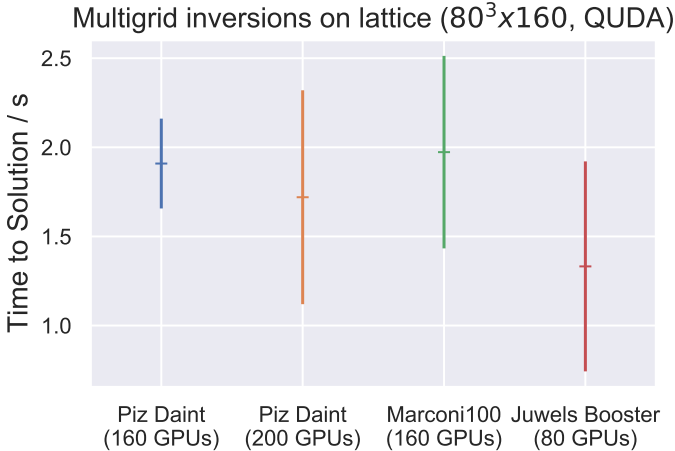

Data provided by ETMC / S. Bacchio

Appendix

References

References I

[6] Hans De Raedt et al. “Massively parallel quantum computer simulator, eleven years later.” In: *Computer Physics Communications* 237 (2019), pp. 47–61. ISSN: 0010-4655. DOI: https://doi.org/10.1016/j.cpc.2018.11.005. URL: http://www.sciencedirect.com/science/article/pii/S0010465518303977 (page 28).

References: Images, Graphics I

[1] Forschungszentrum Jülich GmbH (Ralf-Uwe Limbach). *JUWELS Booster*. Christmasified by Michael Bresser.

[2] Forschungszentrum Jülich GmbH (Ralf-Uwe Limbach). *JUWELS Cluster*.

[3] Forschungszentrum Jülich GmbH (Ralf-Uwe Limbach). *JUWELS Booster*.

[4] SpaceX. *Falcon Rocket in Sunrise*. Freely available at Unsplash. URL: https://unsplash.com/photos/PdeP2ZxooVo.

[5] NASA. *Nov. 14, 1969 - Apollo 12 Liftoff*. The Apollo 12 space vehicle is launched from Kennedy Space Center (KSC), at 11:22 a.m. (EST), Nov. 14, 1969. Apollo 12 was the United States' second lunar landing mission. URL: https://www.nasa.gov/image-feature/apollo-12-liftoff.

References: Images, Graphics II

[7] NASA/Carl Winebarger. *Atlantis Breaks Through the Clouds*. Space shuttle Atlantis emerges through the clouds. Liftoff on its STS-129 mission on Nov. 16, 2009. URL: https://www.nasa.gov/mission_pages/shuttle/shuttlemissions/sts129/multimedia/gallery/09-11-16-5.html.