ParCIO Publications

November 1, 2024

References

[1] Anna Fuchs, Jannek Squar, and Michael Kuhn. Ensemble-Based System Benchmarking for HPC. In *23rd International Symposium on Parallel and Distributed Computing, ISPDC 2024, Chur, Switzerland, July 8-10, 2024*, pages 1–8. IEEE, 2024.

[2] Anna Fuchs, Jannek Squar, and Michael Kuhn. Towards End-to-End Compression in Lustre. In *23rd International Symposium on Parallel and Distributed Computing, ISPDC 2024, Chur, Switzerland, July 8-10, 2024*, pages 1–8. IEEE, 2024.

[3] Anna-Lena Roth, David James, Michael Kuhn, and Johannes Konert. Enhancing Parallel Programming Education with High-Performance Clusters Utilizing Performance Analysis. In Sandra Schulz and Natalie Kiesler, editors, *DELFI 2024 - Die 22. Fachtagung Bildungstechnologien der Gesellschaft für Informatik e.V., DELFI 2024, Fulda, Germany, September 9-11, 2024*, volume P-356 of *LNI*. Gesellschaft für Informatik e.V., 2024.

[4] Kira Duwe and Michael Kuhn. DAI: How Pre-computation Speeds up Data Analysis. In Leonardo Franco, Clélia de Mulatier, Maciej Paszynski, Valeria V. Krzhizhanovskaya, Jack J. Dongarra, and Peter M. A. Sloot, editors, *Computational Science - ICCS 2024 - 24th International Conference, Malaga, Spain, July 2-4, 2024, Proceedings, Part II*, volume 14833 of *Lecture Notes in Computer Science*, pages 116–130. Springer, 2024.

[5] Sajad Karim, Johannes Wünsche, David Broneske, Michael Kuhn, and Gunter Saake. A Design Proposal for a Unified B-epsilon-tree: Embracing NVM in Memory Hierarchies. In Uta Störl, editor, *Proceedings of the 35th GI-Workshop Grundlagen von Datenbanken, Herdecke, Germany, May 22-24, 2024*, volume 3710 of *CEUR Workshop Proceedings*, pages 43–50. CEUR-WS.org, 2024.

[6] Georgiana Mania, Nicholas Styles, Michael Kuhn, Andreas Salzburger, Beomki Yeo, and Thomas Ludwig. Vecpar - A Framework for Portability and Parallelization. In Jirí Mikyska, Clélia de Mulatier, Maciej Paszynski, Valeria V. Krzhizhanovskaya, Jack J. Dongarra, and Peter M. A. Sloot, editors, *Computational Science - ICCS 2023 - 23rd International Conference, Prague, Czech Republic, July 3-5, 2023, Proceedings, Part I*, volume 14073 of *Lecture Notes in Computer Science*, pages 253–267. Springer, 2023.

[7] Sajad Karim, Johannes Wünsche, David Broneske, Michael Kuhn, and Gunter Saake. Assessing Non-volatile Memory in Modern Heterogeneous Storage Landscape using a Write-optimized Storage Stack. In Holger Schwarz, editor, *Proceedings of the 34th GI-Workshop on Foundations of Databases (Grundlagen von Datenbanken), Hirsau, Germany, June 7-9, 2023*, volume 3714 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2023.

[8] Johannes Wünsche, Sajad Karim, Michael Kuhn, David Broneske, and Gunter Saake. Intelligent Data Migration Policies in a Write-Optimized Copy-on-Write Tiered Storage Stack. In Jean-Thomas Acquaviva, Shadi Ibrahim, and Suren Byna, editors, *Proceedings of the 3rd Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems, CHEOPS 2023, Rome, Italy, 8 May 2023*, pages 17–26. ACM, 2023.

[9] Jannek Squar, Niclas Schroeter, Anna Fuchs, Michael Kuhn, and Thomas Ludwig. Content queries and in-depth analysis on version-controlled software. In Matteo Cristani, Carlos Toro, Cecilia Zanni-Merk, Robert J. Howlett, and Lakhmi C. Jain, editors, *Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 26th International Conference KES-2022, Verona, Italy and Virtual Event, 7-9 September 2022*, volume 207 of *Procedia Computer Science*, pages 1261–1270. Elsevier, 2022.

[10] Maximilian Keiff, Frederic Voigt, Anna Fuchs, Michael Kuhn, Jannek Squar, and Thomas Ludwig. Automated performance analysis tools framework for HPC programs. In Matteo Cristani, Carlos Toro, Cecilia Zanni-Merk, Robert J. Howlett, and Lakhmi C. Jain, editors, *Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 26th International Conference KES-2022, Verona, Italy and Virtual Event, 7-9 September 2022*, volume 207 of *Procedia Computer Science*, pages 1067–1076. Elsevier, 2022.

[11] Julius Plehn, Anna Fuchs, Michael Kuhn, Jakob Lüttgau, and Thomas Ludwig. Data-Aware Compression for HPC using Machine Learning. *ACM SIGOPS Oper. Syst. Rev.*, 56(1):62–69, 2022.

[12] Julius Plehn, Anna Fuchs, Michael Kuhn, Jakob Lüttgau, and Thomas Ludwig. Data-Aware Compression for HPC using Machine Learning. In Michael Kuhn, Kira Duwe, Jean-Thomas Acquaviva, Konstantinos Chasapis, and Jalil Boukhobza, editors, *CHEOPS@EuroSys 2022: Proceedings of the Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems, Rennes, France, 5 April 2022*, pages 8–15. ACM, 2022.

[13] Michael Kuhn, Kira Duwe, Jean-Thomas Acquaviva, Konstantinos Chasapis, and Jalil Boukhobza, editors. *CHEOPS@EuroSys 2022: Proceedings of the Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems, Rennes, France, 5 April 2022*. ACM, 2022.

[14] Michael Blesel, Michael Kuhn, and Jannek Squar. heimdallr: Improving Compile Time Correctness Checking for Message Passing with Rust. In Heike Jagode, Hartwig Anzt, Hatem Ltaief, and Piotr Luszczek, editors, *High Performance Computing - ISC High Performance Digital 2021 International Workshops, Frankfurt am Main, Germany, June 24 - July 2, 2021, Revised Selected Papers*, volume 12761 of *Lecture Notes in Computer Science*, pages 199–211. Springer, 2021.

[15] Xiaocong Ai, Georgiana Mania, Heather M. Gray, Michael Kuhn, and Nicholas Styles. A GPU-Based Kalman Filter for Track Fitting. *Comput. Softw. Big Sci.*, 5(1), 2021.

[16] Bartosz Balis, Dora B. Heras, Laura Antonelli, Andrea Bracciali, Thomas Gruber, Jin Hyun-Wook, Michael Kuhn, Stephen L. Scott, Didem Unat, and Roman Wyrzykowski, editors. *Euro-Par 2020: Parallel Processing Workshops - Euro-Par 2020 International Workshops, Warsaw, Poland, August 24-25, 2020, Revised Selected Papers*, volume 12480 of *Lecture Notes in Computer Science*. Springer, 2021.

[17] Michael Kuhn, Kira Duwe, Jean-Thomas Acquaviva, Konstantinos Chasapis, and Jalil Boukhobza, editors. *CHEOPS '21: Proceedings of the Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems, In Conjunction with EuroSys 2021, Online Event, United Kingdom, April, 2021*. ACM, 2021.

[18] Kira Duwe and Michael Kuhn. Dissecting Self-Describing Data Formats to Enable Advanced Querying of File Metadata. In Bruno Wassermann, Michal Malka, Vijay Chidambaram, and Danny Raz, editors, *SYSTOR '21: The 14th ACM International Systems and Storage Conference, Haifa, Israel, June 14-16, 2021*, pages 12:1–12:7. ACM, 2021.

[19] Kira Duwe and Michael Kuhn. Using Ceph's BlueStore as Object Storage in HPC Storage Framework. In Michael Kuhn, Kira Duwe, Jean-Thomas Acquaviva, Konstantinos Chasapis, and Jalil Boukhobza, editors, *CHEOPS '21: Proceedings of the Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems, In Conjunction with EuroSys 2021, Online Event, United Kingdom, April, 2021*, pages 3:1–3:6. ACM, 2021.

[20] Jannek Squar, Tim Jammer, Michael Blesel, Michael Kuhn, and Thomas Ludwig. Compiler Assisted Source Transformation of OpenMP Kernels. In *19th International Symposium on Parallel and Distributed Computing, ISPDC 2020, Warsaw, Poland, July 5-8, 2020*, pages 44–51. IEEE, 2020.

[21] Kira Duwe, Jakob Lüttgau, Georgiana Mania, Jannek Squar, Anna Fuchs, Michael Kuhn, Eugen Betke, and Thomas Ludwig. State of the Art and Future Trends in Data Reduction for High-Performance Computing. *Supercomput. Front. Innov.*, 7(1):4–36, 2020.

[22] Michael Kuhn. Parallele Dateisysteme. *Inform. Spektrum*, 42(5):360–364, 2019.

[23] Yevhen Alforov, Thomas Ludwig, Anastasiia Novikova, Michael Kuhn, and Julian M. Kunkel. Towards Green Scientific Data Compression Through High-Level I/O Interfaces. In *30th International Symposium on Computer Architecture and High Performance Computing, SBAC-PAD 2018, Lyon, France, September 24-27, 2018*, pages 209–216. IEEE, 2018.

[24] Pierre Matri, Yevhen Alforov, Álvaro Brandón, María S. Pérez, Alexandru Costan, Gabriel Antoniu, Michael Kuhn, Philip H. Carns, and Thomas Ludwig. Mission possible: Unify HPC and Big Data stacks towards application-defined blobs at the storage layer. *Future Gener. Comput. Syst.*, 109:668–677, 2020.

[25] Jakob Lüttgau, Michael Kuhn, Kira Duwe, Yevhen Alforov, Eugen Betke, Julian M. Kunkel, and Thomas Ludwig. Survey of Storage Systems for High-Performance Computing. *Supercomput. Front. Innov.*, 5(1):31–58, 2018.

[26] Michael Kuhn. JULEA: A Flexible Storage Framework for HPC. In Julian M. Kunkel, Rio Yokota, Michela Taufer, and John Shalf, editors, *High Performance Computing - ISC High Performance 2017 International Workshops, DRBSD, ExaComm, HCPM, HPC-IODC, IWOPH, IXPUG, P^3MA, VHPC, Visualization at Scale, WOPSSS, Frankfurt, Germany, June 18-22, 2017, Revised Selected Papers*, volume 10524 of *Lecture Notes in Computer Science*, pages 712–723. Springer, 2017.

[27] Pierre Matri, Yevhen Alforov, Álvaro Brandón, Michael Kuhn, Philip H. Carns, and Thomas Ludwig. Could Blobs Fuel Storage-Based Convergence Between HPC and Big Data? In *2017 IEEE International Conference on Cluster Computing, CLUSTER 2017, Honolulu, HI, USA, September 5-8, 2017*, pages 81–86. IEEE Computer Society, 2017.

[28] Michael Kuhn, Julian M. Kunkel, and Thomas Ludwig. Data Compression for Climate Data. *Supercomput. Front. Innov.*, 3(1):75–94, 2016.

[29] Pablo Llopis, Manuel F. Dolz, Francisco Javier García Blas, Florin Isaila, Mohammad Reza Heidari, and Michael Kuhn. Analyzing the energy consumption of the storage data path. *J. Supercomput.*, 72(11):4089–4106, 2016.

[30] Manuel F. Dolz, Mohammad Reza Heidari, Michael Kuhn, Thomas Ludwig, and Germán Fabregat. ArduPower: A Low-cost Wattmeter to improve Energy Efficiency of HPC Applications. In *Sixth International Green and Sustainable Computing Conference, IGSC 2015, Las Vegas, NV, USA, December 14-16, 2015*, pages 1–8. IEEE Computer Society, 2015.

[31] Alexander Droste, Michael Kuhn, and Thomas Ludwig. MPI-Checker: Static Analysis for MPI. In Hal Finkel, editor, *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC, LLVM 2015, Austin, Texas, USA, November 15, 2015*, pages 3:1–3:10. ACM, 2015.

[32] Christopher Bartz, Konstantinos Chasapis, Michael Kuhn, Petra Nerge, and Thomas Ludwig. A Best Practice Analysis of HDF5 and NetCDF-4 Using Lustre. In Julian M. Kunkel and Thomas Ludwig, editors, *High Performance Computing - 30th International Conference, ISC High Performance 2015, Frankfurt, Germany, July 12-16, 2015, Proceedings*, volume 9137 of *Lecture Notes in Computer Science*, pages 274–281. Springer, 2015.

[33] Michael Kuhn. Dynamically Adaptable I/O Semantics for High Performance Computing. In Julian M. Kunkel and Thomas Ludwig, editors, *High Performance Computing - 30th International Conference, ISC High Performance 2015, Frankfurt, Germany, July 12-16, 2015, Proceedings*, volume 9137 of *Lecture Notes in Computer Science*, pages 240–256. Springer, 2015.

[34] Michael Lautenschlager, Panagiotis Adamidis, and Michael Kuhn. Big Data Research at DKRZ – Climate Model Data Production Workflow. In Lucio Grandinetti, Gerhard R. Joubert, Marcel Kunze, and Valerio Pascucci, editors, *Big Data and High Performance Computing - Selected Papers from the High Performance Computing Workshop, Cetraro, Italy, July 7-11, 2014*, volume 26 of *Advances in Parallel Computing*, pages 133–155. IOS Press, 2014.

[35] Julian Martin Kunkel, Michael Kuhn, and Thomas Ludwig. Exascale Storage Systems – An Analytical Study of Expenses. *Supercomput. Front. Innov.*, 1(1):116–134, 2014.

[36] Konstantinos Chasapis, Manuel F. Dolz, Michael Kuhn, and Thomas Ludwig. Evaluating Lustre's Metadata Server on a Multi-Socket Platform. In Dean Hildebrand, Dries Kimpe, and Xiaosong Ma, editors, *Proceedings of the 9th Parallel Data Storage Workshop, PDSW '14, New Orleans, Louisiana, USA, November 16-21, 2014*, pages 13–18. IEEE, 2014.

[37] Michael Kuhn. A Semantics-Aware I/O Interface for High Performance Computing. In Julian M. Kunkel, Thomas Ludwig, and Hans Werner Meuer, editors, *Supercomputing - 28th International Supercomputing Conference, ISC 2013, Leipzig, Germany, June 16-20, 2013.*

Proceedings, volume 7905 of *Lecture Notes in Computer Science*, pages 408–421. Springer, 2013.

[38] Dirk Meister, Jürgen Kaiser, André Brinkmann, Toni Cortes, Michael Kuhn, and Julian M. Kunkel. A Study on Data Deduplication in HPC Storage Systems. In Jeffrey K. Hollingsworth, editor, *SC Conference on High Performance Computing Networking, Storage and Analysis, SC '12, Salt Lake City, UT, USA - November 11 - 15, 2012*, page 7. IEEE/ACM, 2012.

[39] Michael Kuhn, Julian M. Kunkel, and Thomas Ludwig. Simulation-Aided Performance Evaluation of Server-Side Input/Output Optimizations. In Rainer Stotzka, Michael Schiffers, and Yannis Cotronis, editors, *Proceedings of the 20th Euromicro International Conference on Parallel, Distributed and Network-Based Processing, PDP 2012, Munich, Germany, February 15-17, 2012*, pages 562–566. IEEE, 2012.

[40] Michael Kuhn, Julian M. Kunkel, Yuichi Tsujita, Hidetaka Muguruma, and Thomas Ludwig. Optimizations for Two-Phase Collective I/O. In Koen De Bosschere, Erik H. D'Hollander, Gerhard R. Joubert, David A. Padua, Frans J. Peters, and Mark Sawyer, editors, *Applications, Tools and Techniques on the Road to Exascale Computing, Proceedings of the conference ParCo 2011, 31 August - 3 September 2011, Ghent, Belgium*, volume 22 of *Advances in Parallel Computing*, pages 455–462. IOS Press, 2011.

[41] Timo Minartz, Daniel Molka, Julian M. Kunkel, Michael Knobloch, Michael Kuhn, and Thomas Ludwig. Tool Environments to Measure Power Consumption and Computational Performance. In Ishfaq Ahmad and Sanjay Ranka, editors, *Handbook of Energy-Aware and Green Computing - Two Volume Set*, pages 709–744. Chapman and Hall/CRC, 2012.

[42] Julian M. Kunkel, Timo Minartz, Michael Kuhn, and Thomas Ludwig. Towards an Energy-Aware Scientific I/O Interface – Stretching the ADIOS Interface to Foster Performance Analysis and Energy Awareness. *Comput. Sci. Res. Dev.*, 27(4):337–345, 2012.

[43] Julian M. Kunkel, Olga Mordvinova, Michael Kuhn, and Thomas Ludwig. Collecting Energy Consumption of Scientific Data – Energy Demands for Files During Their Life Cycle. *Comput. Sci. Res. Dev.*, 25(3-4):197–205, 2010.

[44] Michael Kuhn, Julian M. Kunkel, and Thomas Ludwig. Dynamic file system semantics to enable metadata optimizations in PVFS. *Concurr. Comput. Pract. Exp.*, 21(14):1775–1788, 2009.

[45] Michael Kuhn, Julian M. Kunkel, and Thomas Ludwig. Directory-Based Metadata Optimizations for Small Files in PVFS. In Emilio Luque, Tomàs Margalef, and Domingo Benitez, editors, *Euro-Par 2008 - Parallel Processing, 14th International Euro-Par Conference, Las Palmas de Gran Canaria, Spain, August 26-29, 2008, Proceedings*, volume 5168 of *Lecture Notes in Computer Science*, pages 90–99. Springer, 2008.

[46] Thomas Ludwig, Stephan Krempel, Michael Kuhn, Julian M. Kunkel, and Christian Lohse. Analysis of the MPI-IO Optimization Levels with the PIOViz Jumpshot Enhancement. In Franck Cappello, Thomas Hérault, and Jack J. Dongarra, editors, *Recent Advances in Parallel Virtual Machine and Message Passing Interface, 14th European PVM/MPI User's Group Meeting, Paris, France, September 30 - October 3, 2007, Proceedings*, volume 4757 of *Lecture Notes in Computer Science*, pages 213–222. Springer, 2007.