


Breaking Limits: Scaling HPC Performance Engineering Horizons to Maximize Potential

Sarah M. Neuwirth
Johannes Gutenberg University Mainz, Germany
neuwirth@uni-mainz.de

PASC 2024, Zurich, Switzerland, June 2024

Motivation
Heterogeneous and Complex HPC Infrastructures

- HPC infrastructure *too complex*, humans are *overwhelmed*
- Complexity and scope increase the *urgency*
  - *New computational paradigms* (AI/ML apps vs. BSP-style HPC)
  - *New architectural directions* (e.g., IPU, RISC-V, data flow)
  - *Heterogeneity overall*: node architectures, within the system, storage and parallel file system during application design (e.g., ML within HPC applications)
  - *New operations paradigms* (e.g., cloud, container)
  - Simplistic approaches to increasing compute demand result in *unacceptable power costs*
- Difficult for humans to optimally adapt applications to systems and to detect and diagnose vulnerabilities

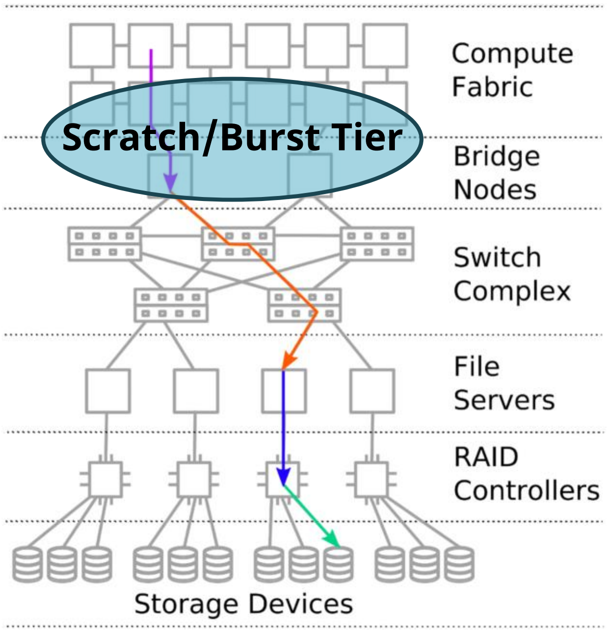

Carns, P., 2023. *HPC Storage: Adapting to Change*. Keynote at REX-IO'23 Workshop.

Ciorba, F., 2023. *Revolutionizing HPC Operations and Research*. Keynote at HPCMASPA'23 Workshop.

B. Settlemyer, G. Amvrosiadis, P. Carns and R. Ross, 2021. *It's Time to Talk About HPC Storage: Perspectives on the Past and Future*, in Computing in Science & Engineering, vol. 23, no. 6, pp. 63-68.

Motivation
Emerging HPC Workloads

- **<u>Traditional HPC:</u>**
  - Dominated by bulk-synchronous I/O phases
  - Simulation workloads
  - Checkpoint / Restart Files
  - Data Input / Data Output (bulky reads or writes)
- **<u>Emerging HPC workloads also encompass:</u>**
  - Artificial Intelligence
  - Data Analytics and Big Data
  - Deep Learning
  - Multi-step Workflows
  - In-situ analysis

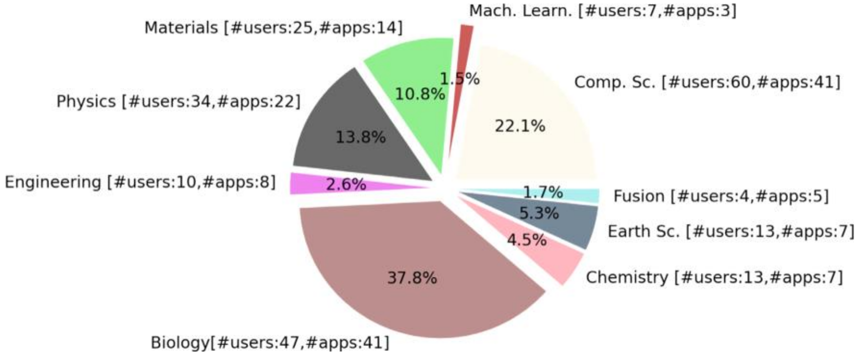

Classification of 23,389 ML jobs on Summit by science domains.

Karimi, A.M., Paul, A.K. and Wang, F., 2022. *I/O performance analysis of machine learning workloads on leadership scale supercomputer*. Performance Evaluation, 157, p.102318.

- ***<u>Vastly different I/O and performance characteristics</u>*** (random small file accesses, non-sequential, metadata-intensive, and small-transaction reads and writes)

=> HPC I/O is more than just checkpointing and bulk-synchronous parallel I/O phases

Performance Engineering and I/O

Performance Engineering and I/O
The Performance Optimization Cycle

<u>**Performance engineering typically is a cyclic process:**</u>

- ***Instrumentation:*** common term for preparing the performance measurement
- ***Measurement:*** During measurement, raw performance data is collected
  - **Profiles:** hold aggregated data (e.g. total time spent in function foo())
  - **Traces:** consist of a sorted list of timed application events/samples (e.g. enter function foo() at 0.11 s)
- ***Analysis:*** Well defined performance metrics are derived from raw performance data during analysis
- ***Presentation:*** Presenting performance metrics graphically fosters human intuition
- ***Evaluation (and Code Optimization):*** Requires tools and lots of thinking

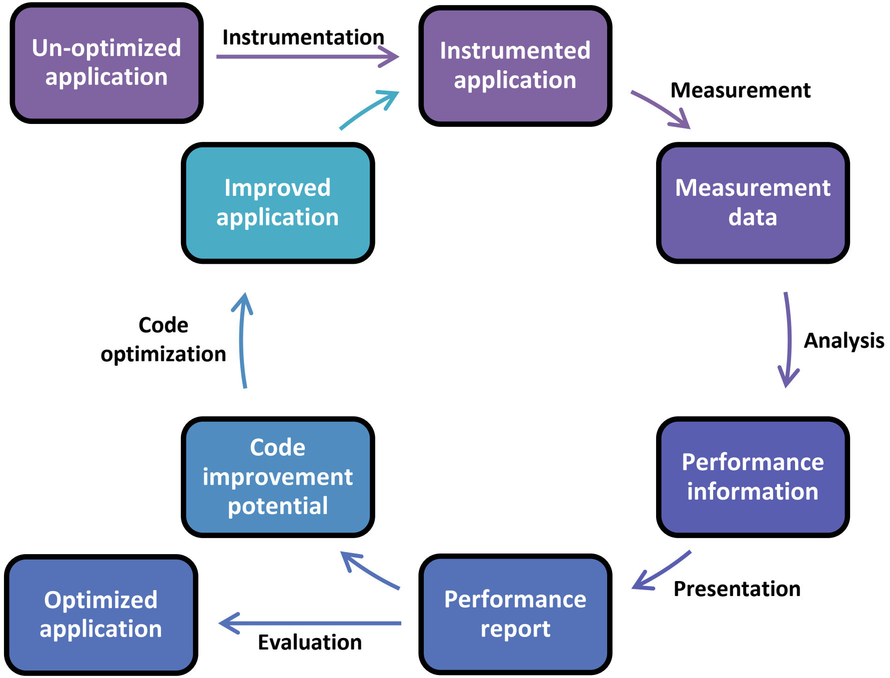

Performance Engineering Overview, https://doc.zih.tu-dresden.de/software/performance_engineering_overview/

Performance Engineering and I/O
Example: Overview of the VI-HPS Tools

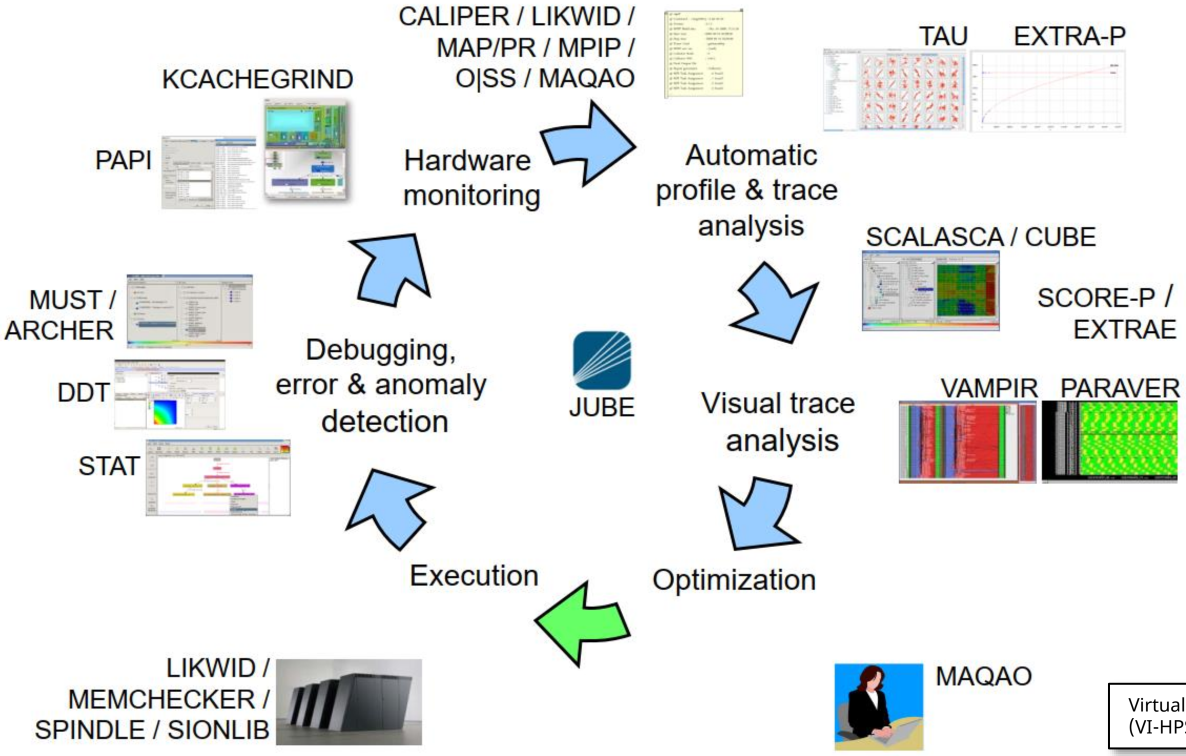

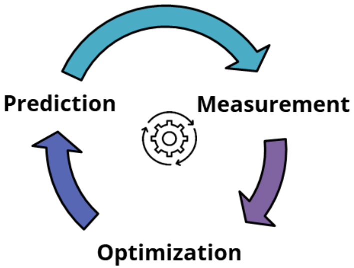

Virtual Institute – High Productivity Supercomputing (VI-HPS), https://www.vi-hps.org/tools/tools.html

Performance Engineering and I/O
Exemplary I/O Performance Analysis Tools

PIKA: Center-Wide and Job-Aware Cluster Monitoring

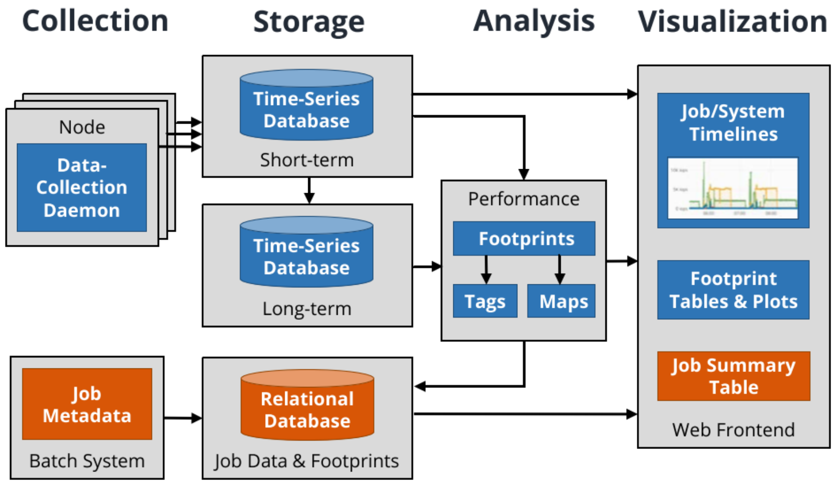

Oeste, S., 2022. *Introduction on parallel I/O and distributed file systems*. NHR Lecture.

Darshan I/O-Characterization Tool

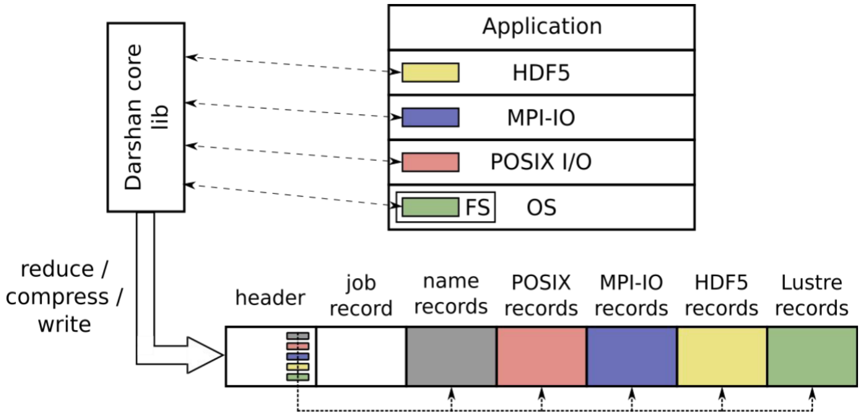

Snyder, S., 2022. *Darshan: Enabling Insights into HPC I/O Behavior*. ECP Community BoF Days.

Performance Engineering and I/O
Changing Landscape of Parallel I/O Architectures

- Characterizing and understanding I/O behavior is critical => *increasingly complex I/O stack*
  - More diverse applications, computational frameworks, etc.
  - Emerging hardware and storage paradigms
- Understanding and re-envisioning I/O stands to benefit numerous HPC stakeholders:
  - Application scientists: Improved I/O performance ⇒ decreased time to scientific discovery
  - Admins: Inform decisions related to procuring new systems
  - Researchers: Optimizing storage system and I/O library designs

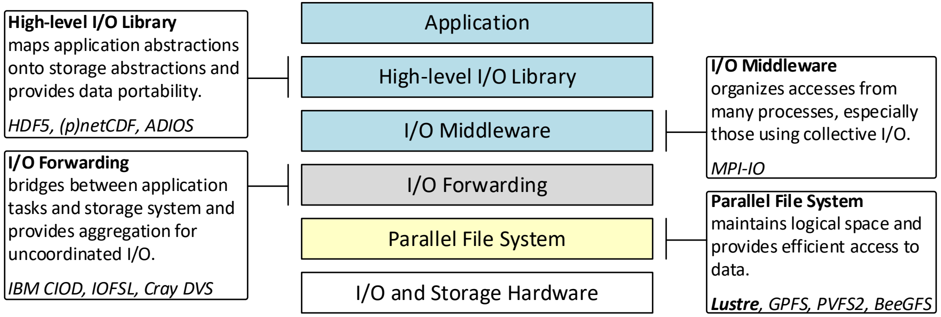

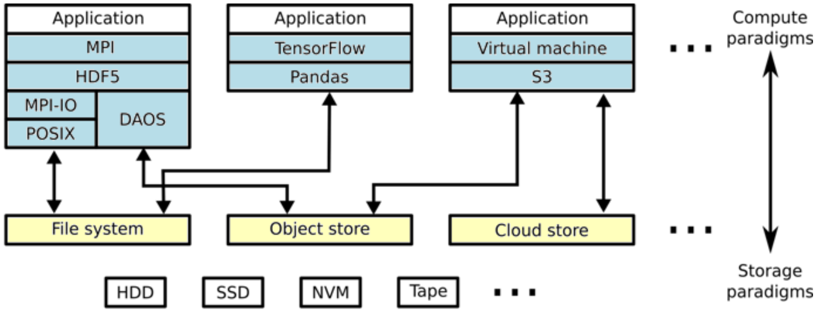

Performance Engineering and I/O

Status of I/O Characterization Tools

Blue Waters, Mira, and Theta popular Darshan log sources used for research:

- https://bluewaters.ncsa.illinois.edu/data-sets
- https://reports.alcf.anl.gov/data/
- ftp://ftp.mcs.anl.gov/pub/darshan/data

Open questions:

- How relevant are the logs to current systems?
- How do we know the integrity of the logs?

Community comments:

- “Darshan is one of the first tools to be deactivated in the event of I/O problems.”
- “Darshan cannot grasp the complexity of state-of-the-art parallel storage systems.”

Darshan I/O Characterization Tool

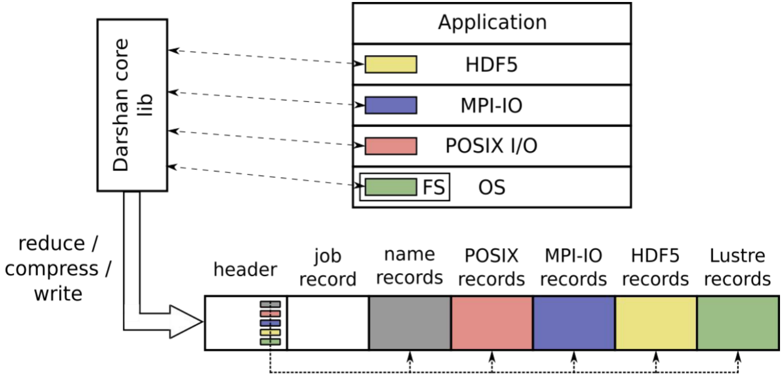

Snyder, S., 2022. *Darshan: Enabling Insights into HPC I/O Behavior*. ECP Community BoF Days.

What are the implications of these questions and observations?

Performance Engineering and I/O

I/O Performance Factors and Metrics

Factors Potentially Affecting I/O Performance:

Application

- Number of processes
- Request sizes
- Access patterns
- I/O operation
- Data volume

Network

- Message sizes
- Network topology
- Network paths
- Network type

File System

- Type of file system
- Disk types
- Stripe sizes
- File hierarchy
- Shared access

Multiple Tools for I/O Performance Analysis:

- May be a problem when users need to change the tool and want to ensure the measurement continuity and comparability
- There is no easy way to verify metrics consistency between tools

=> *Mango-IO is first attempt to provide tools-agnostic metrics calculation*

Liem, Radita, Sebastian Oeste, Jay Lofstead, and Julian Kunkel. *Mango-IO: I/O Metrics Consistency Analysis*. In 2023 IEEE International Conference on Cluster Computing Workshops (CLUSTER Workshops), pp. 18-24. IEEE, 2023.

Tracking the Data Trail

Tracking the Data Trail

Key Questions and Goals of the Project

❶ How does the HPC compute and storage ecosystem look like?

- Overview of modern HPC infrastructures and architectures from tier-1 to tier-3
- Identification of new storage tiers and all possible data paths

❷ Which monitoring infrastructures for I/O and data are deployed globally?

- Identification of the world's most popular monitoring software and toolchains

❸ To what extent are current monitoring infrastructures capable of supporting heterogeneous I/O and storage architectures?

- Identification of shortcomings in widely used monitoring software
- Development of concepts to better support complex parallel I/O and storage systems

Tracking the Data Trail

Goal 1: HPC Compute and Storage Ecosystem

Joint work-in-progress with <u>Hariharan Devarajan</u> (LLNL) and <u>Jay Lofstead</u> (SNL)

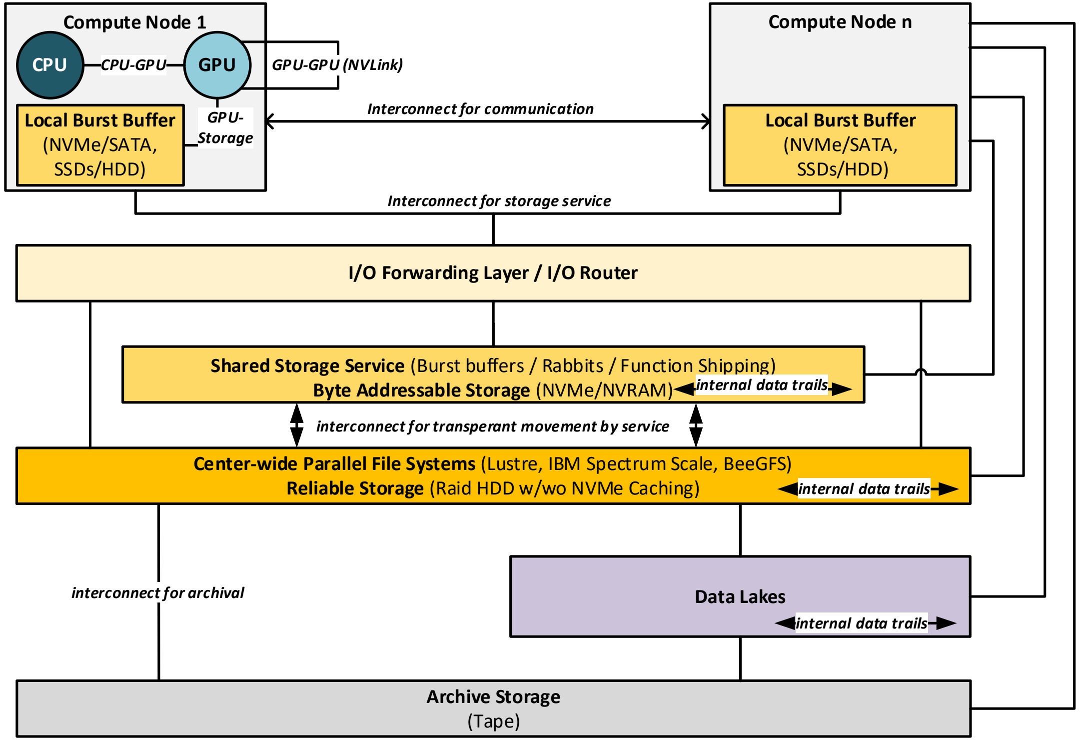

Tracking the Data Trail

Goal 1: Gathering System Insights from IO500 and Top500

Insights from the Top500:

- Unfortunately, very limited (little information about architecture)

Insights from the IO500:

- Test conditions (e.g., #nodes, #processes/node)
- File system information (e.g., type, vendor, software, mount options, capacity, etc.)
- Distinguishes between client-side storage and data servers (i.e., STORAGESYSTEM, SUPERCOMPUTER)
- Effort to collect:
  - Different interconnects for storage (metadata and data) and clients => *not mandatory*
  - Storage media deployed (storage tiers not always clear, optional)

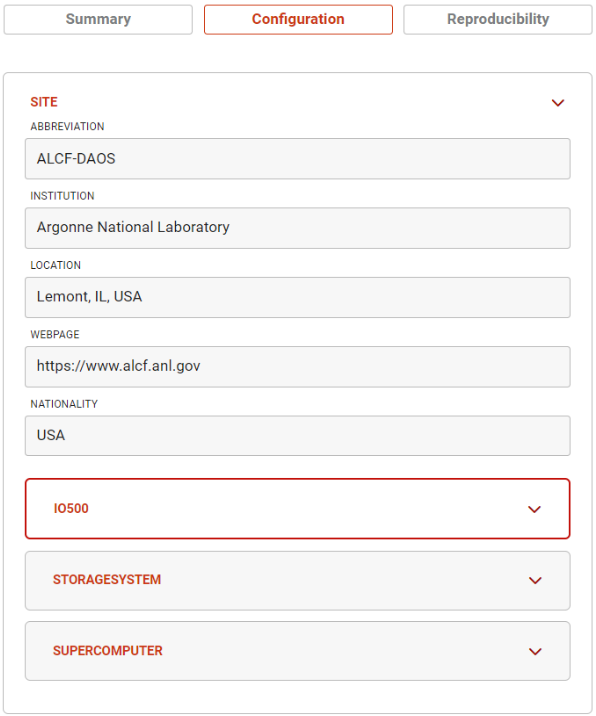

Tracking the Data Trail
Goal 2: Data Trails Community Survey

- Please help us with complementing the information collected by IO500 and CDCL!
- Which systems should take part?
  - **Tier-1:** Top-tier supercomputers with the highest performance for large-scale, national or international projects.
  - **Tier-2:** Mid-level systems for regional or institutional use, balancing cost and performance.
  - **Tier-3:** Entry-level clusters for smaller scale, departmental use, and less intensive computational tasks.

=> Basically everyone! We want a diverse and global mix! ☺

>> https://forms.gle/uPUQLYaciT41q7of6 <<

SCAN ME

Tracking the Data Trail

Goal 3: Holistic Performance Engineering and Analysis

- ***Idea:*** Develop and implement standardized and tool-independent approach for HPC workload and application analysis
- Identify potential limitations and bottlenecks of popular monitoring tools (e.g., unsupported hardware and libraries)
- Integration of various community tools, increasing the compatibility and coverage of different use cases
- ***Goal:*** Establish a ***performance history database*** to categorize systems, workload behaviors, and characteristic patterns for different science domains

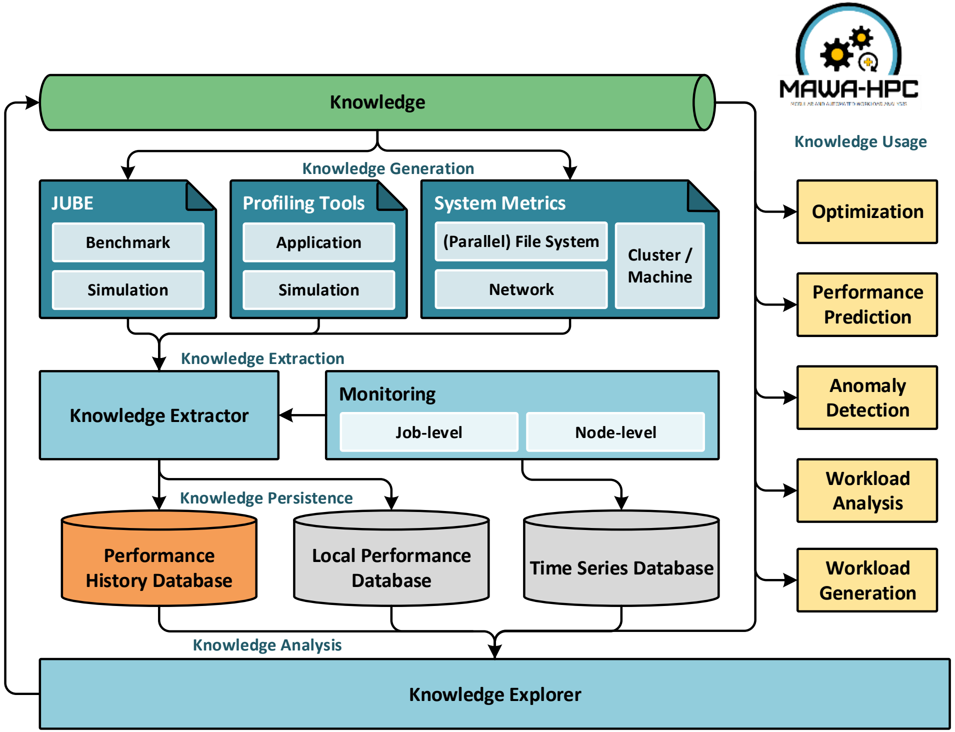

Tracking the Data Trail

Goal 3: Multi-dimensional Performance Modeling

- *Goal:* provide a comprehensive view of application and system performance ⇨ *emerging workloads*
- Multi-dimensional performance models, for example Roofline model, to account for multiple performance factors (e.g. network, compute power, and parallel I/O)
- Including time as an additional dimension, the Roofline model can provide insight into an application's performance over time, enabling the identification of performance anomalies

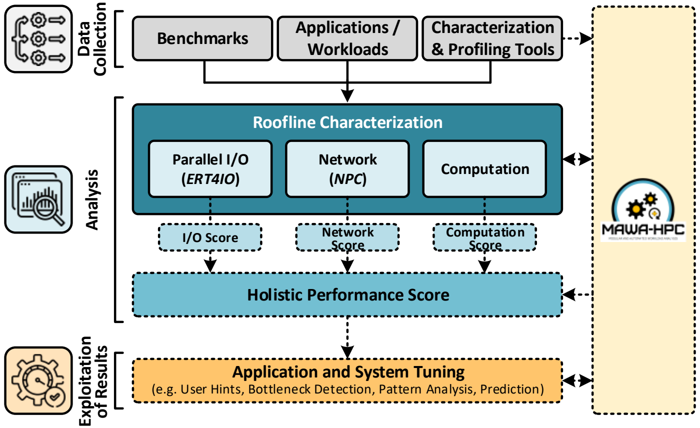

Zhu, Z., Bartelheimer, N. and Neuwirth, S., 2023. *MAWA-HPC: Modular and Automated Workload Analysis for HPC Systems*. ISC'23.

Zhu, Z., Bartelheimer, N. and Neuwirth, S., 2023. *An Empirical Roofline Model for Extreme-Scale I/O Workload Analysis*. IPDPSW'23.

Bartelheimer, N., Zhu, Z., and Neuwirth, S., 2023. *Toward a Modular Workflow for Network Performance Characterization*. IPDPSW'23.

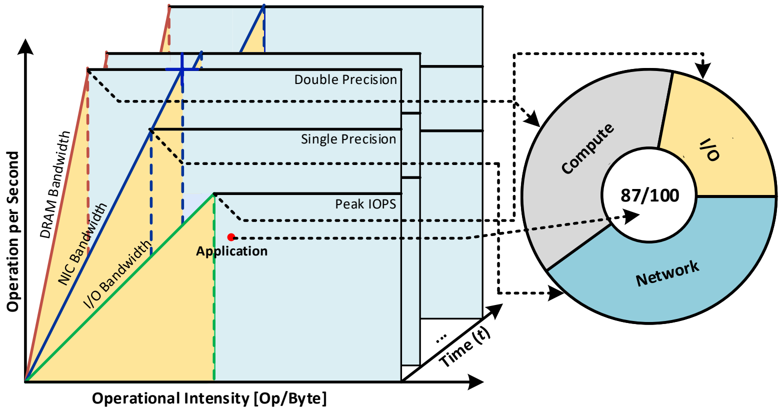

Tracking the Data Trail

Goal 3: Reproducible Evaluation – Design Discussion

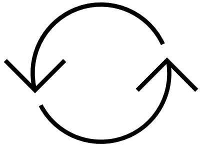

Repeatability

Runs should be able to be repeated with little configuration effort

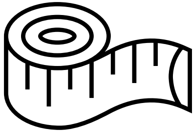

Comparability

A simplified structure amplifies comparability of results between runs or even machines

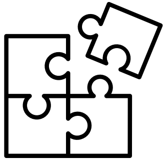

Modularity

Enhanced modularity increases possiblities for further extensions

Schifrin, A., 2023. *Automated Performance Characterization of HPC Systems.* Bachelor thesis, Goethe University Frankfurt.

Bartelheimer, N. and Neuwirth, S., 2023. *Toward Reproducible Benchmarking of PGAS and MPI Communication Schemes.* ICPADS'23.

Bartelheimer, N., Zhu, Z., and Neuwirth, S., 2024. *Automated Network Performance Characterization for HPC Systems*. International Journal of Networking and Computing.

Tracking the Data Trail

Goal 3: Reproducible Evaluation – Example Results

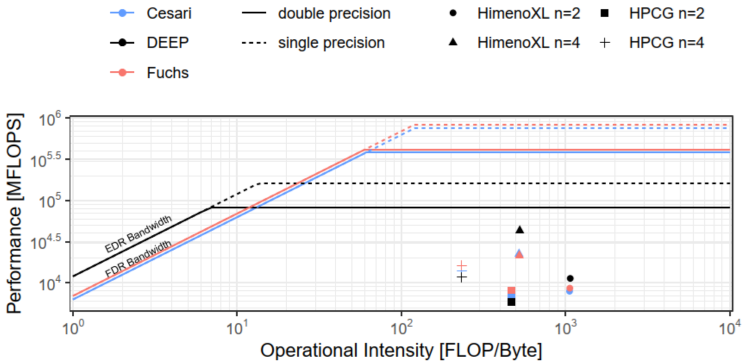

Roofline model with Himeno and HPCG results.

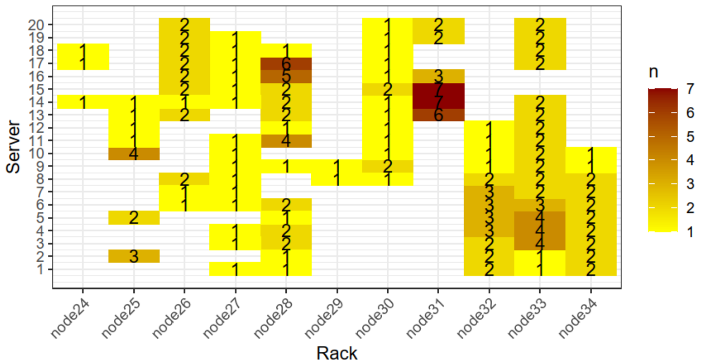

Heat map of the allocated nodes (overall benchmark runs).

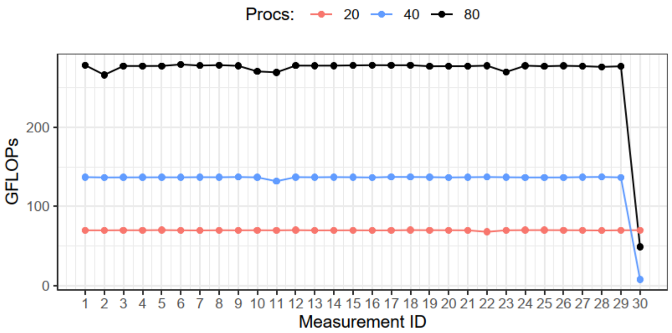

Himeno benchmark over 15 days / 2 measurements per day.

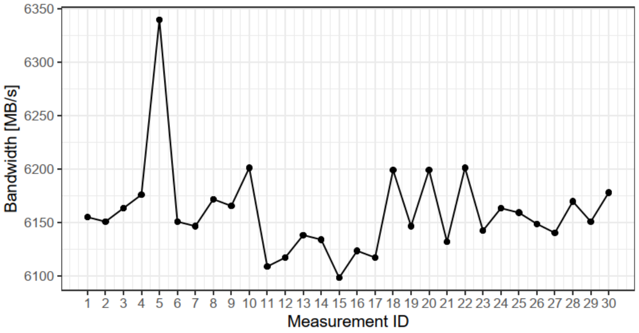

RDMA point-to-point performance over 15 days / 2 measurements per day.

Challenges and Vision

Challenges and Vision

Importance of Reproducibility in Scientific Research

Verification and Trust

- Allows other researchers to verify the original findings
- If findings cannot be reproduced, it casts doubt on the entire study and its conclusions

Building on Knowledge

- Science is a collaborative effort. By reproducing research, scientists can build upon existing knowledge. They can confirm findings, explore them further, or even identify inconsistencies that lead to new discoveries.

Reducing Errors

- Not all irreproducible research is due to misconduct. Mistakes happen.
- Reproducibility helps identify these errors in methodology, data collection, or analysis

Transparency and Openness

- Reproducibility enables transparency in research. When researchers make their data and methods openly available, it allows others to see how they reached their conclusions. This openness is essential for scientific integrity.

Challenges and Vision

Toward traceroute for Data Trails?

How can the differences between modern monitoring infrastructures and the actual data trails be reconciled?

What if there would be *traceroute* for parallel storage and I/O architectures?

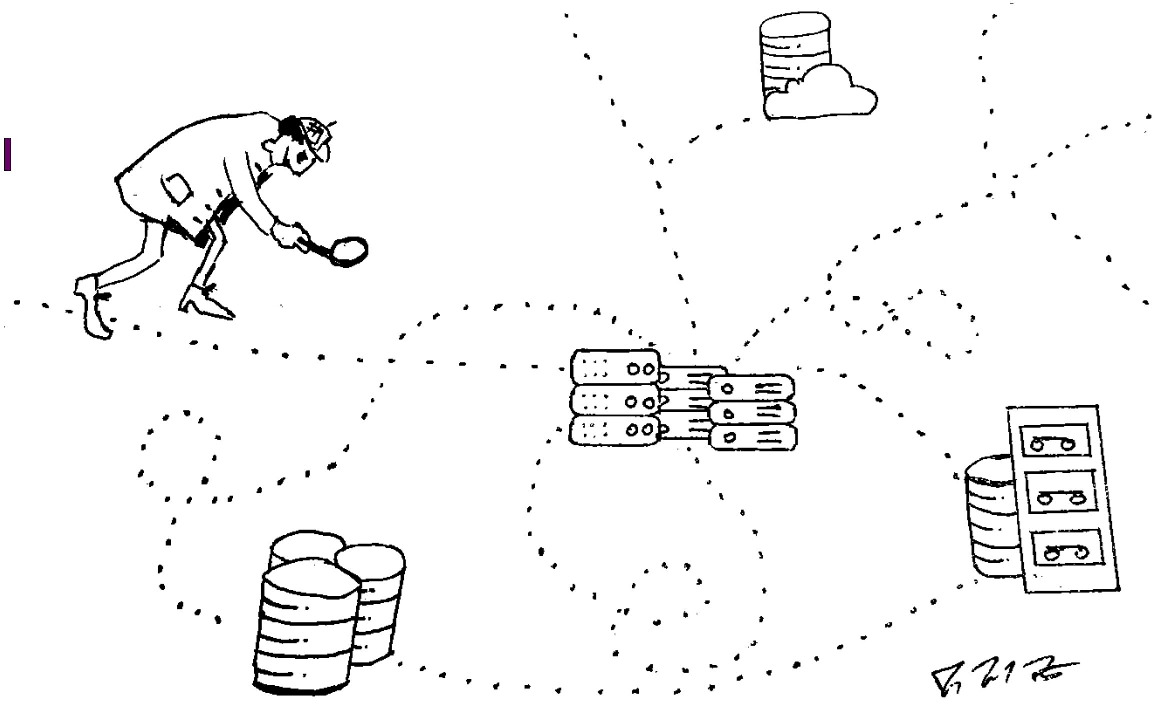

© Joey White-Swift

Challenges and Vision

Holistic Monitoring and Operational Data Analytics

- Continuous and holistic *monitoring*, *archiving*, and *analysis* of operational and performance data open up interactivity with applications, system software, and hardware through
  - Automated feedback
  - Dynamic analysis of workloads and application demands, architecture and resource state
  - Actionable analytics and adaptive response
- Enable *efficient HPC operations* by revolutionizing HPC operations & research through *autonomous* monitor-analyze-feedback-response *loops*

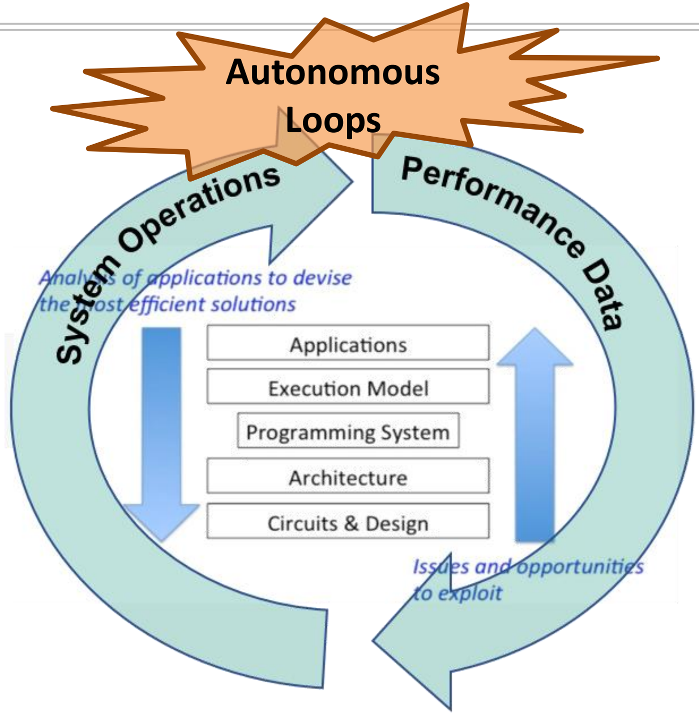

Gentile, A., 2021. *Enabling Application and System Data Fusion*. Keynote at MODA'21 Workshop.

Ciorba, F., 2023. *Revolutionizing HPC Operations and Research*. Keynote at HPCMASPA'23 Workshop.

Dagstuhl Seminar 23171, 2023. *Driving HPC Operations With Holistic Monitoring and Operational Data Analytics*. https://www.dagstuhl.de/23171

Thank you for your Attention!