

Driving HPC Operations With Holistic Monitoring and Operational Data Analytics

Dagstuhl Seminar 23171

Apr 23 – Apr 28, 2023

<https://www.dagstuhl.de/23171>

Schedule (as of Apr 25 — subject to revision)

Sunday - Technical Dinner at Dagstuhl (6-9pm)

Monday (9am-5:30pm):

- **Block A (9am -12:15pm)** Welcome and Keynote speakers *on visions for HPC system and facility monitoring and operational data analytics* (Chair Florina)
 - Welcome and Meeting Expectations (Michael and Florina) (10 min):
 - Instructions:
 - Crazy ideas. Experience-driven aspirations.
 - **How can we move forward and reach the vision of fully automated monitoring, analysis, feedback, and response?** Move past the well-explored topics of data collection systems and “drowning in data” discussions. Positive movement forward, not problem listing.
 - Outcomes:
 - **High level diagram about how capabilities would have to interact, latency, possibilities for analysis and response, what that would enable** etc. Not something that looks like anyone’s current architecture, nor APIs, diagram but rather is notional vision with what areas need revolution and funding for progress.
 - **Community establishment**
 - **Wonderful and Fundable Vision Report** to be shopped around as a community as the basis for national and international funding. Seed agency calls.
 - Participants Introductions (50 min)
 - Name
 - Affiliation
 - 1 sentence about the topic that you work on/are most interested in related to this seminar
 - 9:30 Keynote 1 Jim Brandt (Sandia) (45 min incl. q.)
 - Coffee break (30min) ~10:15am
 - 11:00 Keynote 2 Utz-Uwe Haus (HPE) (45 min incl. q.)
- Lunch - 12:15
- **Block B (2:00pm - 3:30pm): Session 1 Use Cases and Data: Use cases to determine analyses and information desired by humans and software. Challenges and Opportunities** (Chair Ann)
 - 20 min (incl. q.) - **Michael Ott** (LRZ) - Continuous performance counter sampling to assess the characteristics of user workloads. Since user applications are a black box for many HPC

data centers, continuous monitoring and analysis can provide insight to help users improve performance and guide future procurements,

- 20 min - **Kevin Huck** (U Oregon) - Started with application monitoring and is bringing in system monitoring, so inverse of Jim. All aspects of application monitoring through to response. Given your deep knowledge of application tools, how can we have this type of monitoring always-on w/o too much data or overhead? Given your current response work, what are/should be the targets for the community? For example, What is missing in the data, our understanding of the data, hooks to get data or respond, etc?
 - 20 min - **Esa Heiskanen** (LUMI) - Presents the operations view. In addition to your visions, provide some context so we can understand the operations center drivers. What are the challenges in operations? What are the most urgent issues for your center? What are the controls you have and those you would like to have? What would you want to know about the applications running on your system? How would analysis and insight help your operations? How are you getting the infrastructure data and what are the challenges in that data (calibration, lack of instrumentation, ensuring data quality)?
 - Group **Discussion** (30 min) (Killer Use Case of Tomorrow?)
- **Coffee Break (3:30pm-4:00pm)**
 - **Block C (4:00pm - 5:30pm): Position Talks from Attendees** (10 min incl. q / talk) (**Chair Ann**)
 - Taylor Groves, NERSC: "Workflows & Patchwork: Building a Big Picture Out of Scraps of Data"
 - Terry Jones, ORNL: "Challenges for Holistic Monitoring When Attempting Codesign Support"
 - Tapasya Patki, LLNL: "Large-scale Vendor-Neutral Power Monitoring"
 - Francieli Boito, University of Bordeaux, Inria: "Towards an efficient and concise characterization of temporal I/O behavior"
 - Jeff Hanson, HPE: "Metadata management: making sense of mounds of data"
 - Mike Showerman, NCSA: "Containerized real time job anomaly detection"
 - Group **Discussion** (30 min)

Tuesday (9am-5:30pm):

- **Block A (9am -12:15pm) Session 2 Response: Feedback-Driven Response Presentations on Actionable Considerations/Architectural Gaps for Application and System Feedback (Chair Torsten)**
 - 20 min incl. q. - Phil Carns (ANL) How people have used or been thwarted in using results of his tool (Darshan/IO). Issues and questions in portability of solutions to other systems, workflows, etc.
 - 20 min incl. q. - Frederic Suter (ORNL) - Potential interactions of applications with system data. How would that be used for scheduling decisions? Considerations in an infrastructure to share and respond to this type of information with minimal application impact/alteration.
 - 20 min incl. q.- Valeria Cardellini (Uni Roma) - Hierarchical control of distributed systems. Run time control of applications using Reinforcement Learning. In RL used in this space, what are the parameters to optimize on? How should those parameters be configured?
 - 20 min incl. q. - Ann Gentile (SNL) - What has to change to get to fully autonomous, self-organizing systems? (e.g., vendors moving toward more encapsulated view, but we need exposure and changes to the stack)
 - Coffee Break (30 min) ~10:30am
 - Breakout Discussions (55 min) in two groups:
 - What are convincing use cases for autonomic computing?
 - What are the enablers (components, hooks, and functionalities) feedback-driven response for these use cases?
 - What is missing to realize these use cases?
 - *TODO: Ensure that we come out with report-worthy, constructive solution-oriented, collaborative insights*
 - Group Discussion (30 min) (One or more compelling responses and enablers?)
 - Report from each of the two groups on motivating use cases, feedback needed to drive response, entities to respond, response with biggest impact, biggest enabler for this response.
 - Draw the roadmap together.
 - *TODO: Ensure that we come out with report-worthy, constructive solution-oriented, collaborative insights*
- Lunch - 12:15
- **Block B (1:30pm - 3:30pm): Session 3A: Analytic and Informatics Approaches Applications of Analytic Methods (Analysis talks receive an additional 5 min to include background on the analysis technique(s) necessary an audience of non-experts) (Chair Jim)**
 - Group Discussion (30 min): Community discussion on techniques and challenges for analysis
 - What types of analysis techniques need to be developed to address the complex use cases from the morning session?
 - How can we define these in a way that we can significantly move these forward?
 - *TODO: ensure that we come out with report-worthy answers to these questions*
 - 25 min - Abdullah Mueen (UNM) - Time-series analysis applied to this domain. Generative AI as a possible approach?
 - 25 min - Devesh Tiwari (Northeastern) Insights and revelations on applications and their behavior based on analysis. There is some conventional wisdom that is wrong - what are the implications of that for how we would operate systems? Background on some of the techniques, their application, and challenges in this domain.

- 25 min - Krishnan Raghavan (ANL) - Application of GNN to anomaly detection in workflows. Challenges are amplified when applying continuous learning mechanisms: volume of data, convergence, etc. Comparison of offline vs online training and learning. How do you assess that the model is no longer applicable? What is the stability of the model and of the system? What is the latency for analysis? How does the desire to take response drive requirements for your analysis?
- Group Discussion (15 min) (Q&A)
- **Coffee Break (3:30pm-4:00pm)**
- **Block C (4pm - 5:40pm): Session 3B Analytic and Informatics Approaches** *Applications of Analytic Methods* (Analysis talks receive an additional 5 min to include background on the analysis technique(s) necessary an audience of non-experts) **(Chair Jim)**
 - 25 min - Thaleia Doudali (IMDEA (ES))- Models for memory access patterns. What can we deploy and make them practical? Opportunity to use visualization and images and ML-techniques for those in the system learning pipeline. Generative AI for viz.
 - 25 min - Hilary Egan (NREL) - Modularity in applying models developed for one system to apply to many in the data center. Estimate of impact/benefit of anomaly detection and response. Different work with both HPE and NREL – what were the motivations in each case? What did you try and how well did it work?
 - Group Discussion (45 min)
 - What existing analysis techniques are promising, directly applicable or adaptable to our area (and data)?
 - What new techniques do we need to develop?
 - How to constructively move forward in analysis?
 - *TODO: ensure that we come out with report-worthy, constructive, solution-oriented, collaborative insights*
 - Instructions (5 min) for Wed Artifact Working Session

Wednesday (9am start with late excursion)

- **Block A (9am - 12:15pm): Session 4 Group Working Session Artifact Working Session**
 - *Form 3 breakout groups and assign scribes (15 min) (Chair 1-5)*
 - **(80 min) 3 Breakouts Groups** - Planning and development of mind map and report topics and funding path
 - 3 highest use cases - walk through timescales, space scales, and control. At the end of each hour what have we identified as the key research areas to work on.
 - Mind map of the future:
 - Autonomic computing for Scientific discovery and solving humanity's problems.....How does your work impact society?
 - Goal of the mind map:
 - How can we get analysis and response on the necessary timescales while still maintaining scalability?
 - Control: What entities are responsible for making decisions and communicating state and response? Who has authority for making those decisions? (e.g., who has the necessary data, insight to build models, to make decisions, etc)
 - MAPE-K?
 - We do this for the top 2-3. Identified possible monitoring, analysis, feedback, and response.
 - WAFR - Identify open questions for how important and impactful are these areas and what does that mean for driving research areas for funding?
 - 10 min each breakout group to prepare 1 coherent slide: answers to above + research recommendations
 - Coffee Break (30 min) ~10:30 - 11:00
 - 5 minutes each — **Working Group Reports**: Each group presents their mind map and research directions for the WAFR
 - 45 min group discussion and questions
- Lunch - 12:15
- **After lunch (~ 1:00pm): Group Excursion**

Planning in progress....
- Dinner - 18:00 back at Dagstuhl

Thursday (9am - 5:30pm):

- **Block A (9am - 12:15pm): Session 5 Response: Actionable Analytics and Responses (Chair Michael)**
 - 20 min Jeff Hanson and Torsten Wilde (HPE) - Have been doing analysis/prediction, however a unique vendor perspective that they bring is: What hooks could the vendors expose? If the community defined aspirations, how could this influence design down to even the HW level? How do we get hooks for response in (that could be triggered by analyses of Oriol and Mariam)? What would vendor(s) be comfortable with having the SW touch the HW controls?
 - 20 min Oriol Vida (BSC (ES))- Work in the context of power and energy optimization: How to let the users and admins enact the optimization positions that they want? Are there insights that can be extrapolated to optimizations in considerations beyond power and energy?
 - 20 min Mariam Kiran (NERSC) - AI in the networks. You offered to share techniques — what are these? How portable? How to extend to other areas of the problem space? How do we get techniques deployed?
 - **Parallel Working Group Breakouts**
 - Topics including, but not limited to: How can we work as a community, including the vendors, and move these things forward? We are gated by what vendors enable us to do.
 - Funding avenues
 - Working Group Reports
 - Group Discussion
 - Coffee Break (30 min) ~10:30
- **Block B (1:30pm - 3:30pm) Session 5a: Followup Topic(s) Additional coverage of any questions/thesis statement from the sessions**
 - Discussions
 - Parallel Working Group Breakouts
 - Working Group Reports
 - Group Discussion
- **Coffee Break (3:30pm-4:00pm)**
- **Block C (4pm - 5:30pm): Session 5b Followup Topic(s) Additional coverage of any questions/thesis statement from the sessions (cont'd)**

Friday (9am-1:30pm)

- **Block A (9am - 12:15pm): Session 6: Report Outline and Planning**
 - Summary Working Group Presentations
 - Community Path forward
 - Group identification of key takeaways for inclusion in the seminar report
 - Group creation of report outline
-

Meals:

- Dinner options from 6pm (Sun)
- Breakfast - 7:30 am - 8:45 am (M-F)
- Coffee - 10:30am-11am (M-F)
- Lunch - 12:15 (M-F)
- Coffee and Cake - 3:30pm-4pm (M-F)
- Dinner - 6pm (M-Th)