



**HAL**  
open science

## Autonomy Loops for Monitoring, Operational Data Analytics, Feedback, and Response in HPC Operations

Francieli Zanon Boito, Jim Brandt, Valeria Cardellini, Philip Carns, Florina Ciorba, Hilary Egan, Ahmed Eleliemy, Ann Gentile, Thomas Gruber, Jeff Hanson, et al.

► **To cite this version:**

Francieli Zanon Boito, Jim Brandt, Valeria Cardellini, Philip Carns, Florina Ciorba, et al.. Autonomy Loops for Monitoring, Operational Data Analytics, Feedback, and Response in HPC Operations. HPCMASPA 2023 - Workshop on Monitoring and Analysis for HPC Systems Plus Applications, Oct 2023, Santa Fe, United States. pp.7, 10.1109/CLUSTERWorkshops61457.2023.00016 . hal-04382088

**HAL Id: hal-04382088**

**<https://inria.hal.science/hal-04382088>**

Submitted on 9 Jan 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# Autonomy Loops for Monitoring, Operational Data Analytics, Feedback, and Response in HPC Operations

Francieli Boito<sup>||</sup>, Jim Brandt<sup>†</sup>, Valeria Cardellini<sup>x</sup>, Philip Carns<sup>\*</sup>, Florina M. Ciorba<sup>‡</sup>, Hilary Egan<sup>xviii</sup>, Ahmed Eleliemy<sup>‡</sup>, Ann Gentile<sup>†</sup>, Thomas Gruber<sup>xvi</sup>, Jeff Hanson<sup>xiv</sup>, Utz-Uwe Haus<sup>xxi</sup>, Kevin Huck<sup>§</sup>, Thomas Ilsche<sup>††</sup>, Thomas Jakobsche<sup>‡</sup>, Terry Jones<sup>xiii</sup>, Sven Karlsson<sup>\*\*</sup>, Abdullah Mueen<sup>xvii</sup>, Michael Ott<sup>xii</sup>, Tapasya Patki<sup>xv</sup>, Ivy Peng<sup>xx</sup>, Krishnan Raghavan<sup>\*</sup>, Stephen Simms<sup>¶</sup>, Kathleen Shoga<sup>xix</sup>, Michael Showerman<sup>xv</sup>, Devesh Tiwari<sup>xi</sup>, Torsten Wilde<sup>xiv</sup>, and Keiji Yamamoto<sup>‡‡</sup>

<sup>||</sup>University of Bordeaux, CNRS, Bordeaux INP, INRIA, LaBRI, Talence, FR;

<sup>†</sup>Sandia National Laboratories, CA and NM, US; <sup>x</sup>University of Rome Tor Vergata, IT;

<sup>\*</sup>Argonne National Laboratory, IL, US; <sup>‡</sup>University of Basel, Basel, CH;

<sup>xviii</sup>National Renewable Energy Laboratory, CO, US;

<sup>xvi</sup>Erlangen National High Performance Computing Center (NHR@FAU), DE;

<sup>xiv</sup>Hewlett Packard Enterprise, DE and US; <sup>xxi</sup>Hewlett Packard Labs, EMEA Research Lab, CH;

<sup>§</sup>University of Oregon, Eugene, OR, US; <sup>††</sup>Technische Universität Dresden, DE;

<sup>xiii</sup>Oak Ridge National Laboratory, TN, US; <sup>\*\*</sup>Technical University of Denmark, DK;

<sup>xvii</sup>University of New Mexico, NM, US; <sup>xii</sup>Leibniz Supercomputing Centre, DE;

<sup>¶</sup>Lawrence Berkeley National Laboratory, CA, US; <sup>xx</sup>KTH Royal Institute of Technology, SE;

<sup>xix</sup>Lawrence Livermore National Laboratory, CA, US; <sup>xv</sup>University of Illinois, IL, US;

<sup>xi</sup>Northeastern University, MA, US; and <sup>‡‡</sup>RIKEN R-CCS, JP

Email: <sup>||</sup> francieli.zanon-boito@u-bordeaux.fr, <sup>†</sup> {brandt, gentile}@sandia.gov, <sup>x</sup> cardellini@ing.uniroma2.it,

<sup>\*</sup> {carns, kraghavan}@mcs.anl.gov, <sup>‡</sup> {florina.ciorba, ahmed.eleliemy, thomas.jakobsche}@unibas.ch,

<sup>xviii</sup> hilary.egan@nrel.gov, <sup>xvi</sup> thomas.gruber@fau.de, <sup>xiv</sup> {jeff.hanson, torsten.wilde}@hpe.com, <sup>xxi</sup> uhaus@hpe.com,

<sup>§</sup> khuck@cs.uoregon.edu, <sup>††</sup> thomas.ilsche@tu-dresden.de, <sup>xiii</sup> trj@ornl.gov, <sup>\*\*</sup> svea@dtu.dk, <sup>xvii</sup> mueen@unm.edu,

<sup>xii</sup> ott@lrz.de, <sup>xix</sup> {patki1,shoga1}@llnl.gov, <sup>xx</sup> ivybobpeng@kth.se, <sup>¶</sup> ssimms@lbl.gov, <sup>xv</sup> mung@illinois.edu,

<sup>xi</sup> d.tiwari@northeastern.edu, and <sup>‡‡</sup> keiji.yamamoto@riken.jp

**Abstract**—Many High Performance Computing (HPC) facilities have developed and deployed frameworks in support of continuous monitoring and operational data analytics (MODA) to help improve efficiency and throughput. Because of the complexity and scale of systems and workflows and the need for low-latency response to address dynamic circumstances, automated feedback and response have the potential to be more effective than current human-in-the-loop approaches which are laborious and error prone. Progress has been limited, however, by factors such as the lack of infrastructure and feedback hooks, and successful deployment is often site- and case-specific. In this position paper we report on the outcomes and plans from a recent Dagstuhl Seminar, seeking to carve a path for community progress in the development of autonomous feedback loops for MODA, based on the established formalism of similar (*MAPE-K*) loops in autonomous computing and self-adaptive systems. By defining and developing such loops for significant cases experienced across HPC sites, we seek to extract commonalities and develop conventions that will facilitate interoperability and interchangeability with system hardware, software, and applications across different sites, and will motivate vendors and others to provide telemetry interfaces and feedback hooks to enable community development and pervasive deployment of MODA autonomy loops.

**Index Terms**—high performance computing, monitoring and operational data analytics, autonomy loops, *MAPE-K*

## I. INTRODUCTION

Using monitoring data to extract actionable insights on system behaviors regarding facilities and building infrastructure, system hardware, system software, and applications has been referred to as monitoring and operational data analytics (MODA) [1]. Many high performance computing (HPC) and data centers have developed and/or deployed data collection frameworks that facilitate continuous and holistic monitoring and analysis (e.g., [2]–[14]) in support of MODA goals. While the *monitoring* component of MODA is well established and deployed, the *analysis* part is still performed mostly by visual inspection, and feedback typically involves a *human in the loop* [4], [5], [15] to make analysis-based responses. As HPC systems are growing larger and ever more complex, manually analyzing high-dimensional time-series operational data becomes intractable without *automation*. Moreover, having a human in the loop limits the speed of response and consequently, the opportunities for feedback-driven improvements.

Some HPC and data centers have started to work on automated feedback and response functionalities (e.g., [6]–[8], [16]–[24]). However, these are often one-off approaches that

are limited in the scope of data and responses and dependent on the specifics of particular use cases, architectures, available actuators for responses, and so forth. It has become evident that closer collaboration would enable the community to better leverage aggregate work and experience to maximize impact.

With this in mind, a cross-section of the HPC ecosystem gathered at Schloss Dagstuhl in April 2023 for a seminar on “Driving HPC Operations With Holistic Monitoring and Operational Data Analytics” [25], with the goal of advancing the field and establishing a community path forward. During the seminar, we converged on the high-level functionalities of *monitoring*, *analysis*, *feedback*, and *response* to comprise MODA-specific *autonomy loops*, as illustrated in Fig. 1. Then, a set of important use cases, across multiple sites, for autonomy loops were identified, which the group plans to develop as prototypes. Through the prototypes, we intend to extract commonalities and develop conventions that will facilitate interoperability with system hardware, software, applications, etc., across different sites. The use cases and established conventions will then be used to motivate vendors and others to provide the necessary telemetry and feedback hooks.

In this paper we posit that bringing formalism to MODA autonomy loops will help the community build more generalized interfaces, infrastructure, and interactivity approaches enabling reusable and more comprehensive (in terms of system, applications, and facility components and response opportunities) approaches to self-adaptivity in HPC systems. We intend to leverage the formalism defined in the fields of autonomous computing and self-adaptive systems [26], in the form of *MAPE-K* loops: *Monitor*, *Analyze*, *Plan*, and *Execute* over some *Knowledge* [27], because of the similarity to our concepts of monitoring, analysis, feedback, and response. This would both enable a common context for thinking about the problem and allow us to take advantage of the *MAPE-K* architectural design patterns. We provide background on the autonomy loop approach and considerations to help identify opportunities for developing generalizable *MAPE-K* loops-based infrastructure and approaches for MODA. We highlight details for a foundational case that we intend to tackle as a community.

## II. AUTONOMY LOOP CASE STUDY METHODOLOGY

*MAPE-K* refers to the architectural pattern using *Monitor*, *Analyze*, *Plan*, and *Execute* loops over some *Knowledge*, introduced in the autonomic computing (AC) initiative by IBM [27] and used in self-adaptive systems [26]. The AC concept is based on constant checking and optimization of system status through adaptive decisions. The AC reference model includes a managed system, sensors, actuators, a managing system (utilizing *MAPE*), and *Knowledge* about the managed system and its environment (where the latter is not under control). In our context a *managed system* is any HPC hardware or software system in which *sensors* provide data about the HPC system, and *actuators* are response hooks in building infrastructure, system hardware, system software, or applications. *Monitor* refers to the process of collecting data

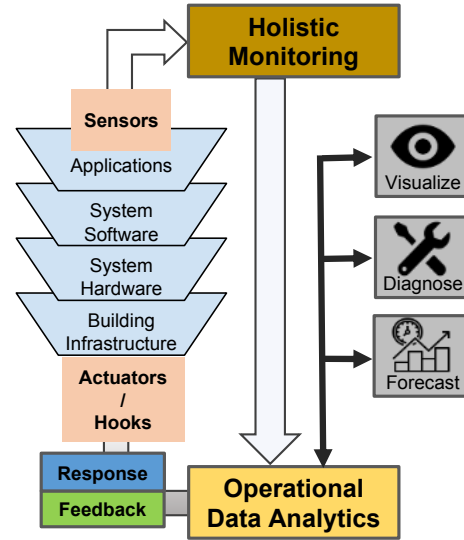


Fig. 1: Vision of holistic monitoring and operational data analytics.

about an element of interest, for example, an application. *Analyze* and *Plan* refer to analyzing the collected data and planning an appropriate response, for example, checkpointing. *Execute* refers to carrying out the planned response through the use of response hooks. *Knowledge* is pervasive in the components of the *MAPE-K* loop, as shown in the “classic” (leftmost) loop in Fig. 2. It can include, for example, progress rate of an application compared with that of a previous run, as well as knowledge gained from assessing the effectiveness of the *Plan* and *Execute* phases of previous loop iterations.

Arbitrarily complex autonomous actions can be supported by different decentralized architectural design patterns for *MAPE-K* loops [28], in which the *MAPE* functionalities can be realized by multiple components that coordinate with one another in different ways. Some established patterns are illustrated in Fig. 2, which shows, from left to right, the *master-worker*, the *coordinated control*, and the *hierarchical control* patterns<sup>1</sup>. The first decentralizes only *Monitor* and *Execute*; the centralized *Plan* can achieve global objectives and guarantees but suffers from limited scalability, especially when managing a complex system. The coordinated control pattern relies on fully decentralized *MAPE* loops that control different parts of the managed system and have the potential of good scalability and robustness, but decentralized *Plan* policies may suffer from instability and side-effects due to indirect interactions. In the hierarchical control pattern, decentralized *MAPE* loops are organized in a hierarchy, with separation of concerns and time scales and aiming to improve scalability without compromising stability; however, division of control is not

<sup>1</sup>We omit the *Knowledge* component and how it is used and shared by the *MAPE* components to focus on the interaction between *MAPE* loop components and across *MAPE* loops and decrease the complexity. How *Knowledge* can be stored and exchanged among *MAPE* components paves the way for additional patterns.

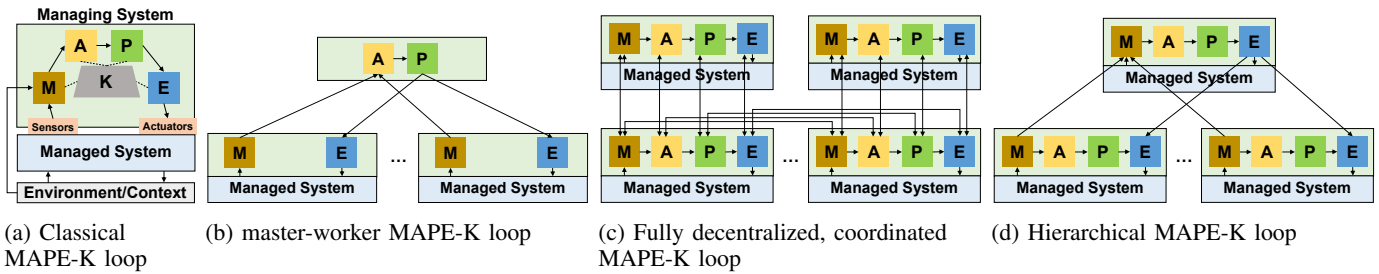


Fig. 2: Design Patterns for *MAPE-K* loops. Leveraging the *MAPE-K* formalism will facilitate application of the designs to MODA autonomy loops.

trivial.

Previous works have demonstrated the potential for autonomy loops to improve efficiency in HPC operations. Examples include throttling network or storage traffic in response to observed congestion or optimizing cache policies in response to observed memory access patterns. Progress has been limited, however, by available data, hooks, and opportunities for generality. To further such efforts, we seek to identify commonalities and conventions that could drive the development of widely reusable and interoperable infrastructure for MODA autonomy loops, leveraging the *MAPE-K* formalism and architectures. A unified and generalized approach to solutions would simplify integration of autonomy loop components and functionalities, supporting a number of complex subsystems, and would avoid an approach involving a confusing mix of disparate solutions. Considerations in determining designs for MODA cases would include access to monitored components and response options, response latency, desire for separability of functionality, need for coordination, and scalability and robustness.

Beyond design and development, establishing autonomy loops within HPC software stacks and infrastructure will also depend on satisfying concerns of security, trust, and validity. We therefore propose the following five key questions be considered in approaching autonomy loop use cases to help identify opportunities for broader impact:

- i) Can the autonomy loop be described in terms of high-level components with distinct responsibilities?
- ii) What interfaces or data formats would enable those components to be interchangeable?
- iii) What sort of open datasets would facilitate the use case?
- iv) How would validation be performed and user and system administrator trust be earned in order to enable autonomous actions?
- v) How would practitioners be incentivized to engage and provide additional support (such as providing functionality and hooks for data, feedback, and responses) and usage of the loops in production?

### III. DRIVING GENERALIZATION THROUGH INITIAL CASE

We posit that a collaborative community approach to the development of MODA autonomy loops will further development, adoption, and production deployment. To this end we have identified an initial set of specific use cases that could

benefit from use of autonomy loops and that target scenarios prevalent in production HPC. The diversity of cases is intended to enable exploration of commonalities that could drive wider interoperability. The cases are as follows:

- 1) *Maintenance*: Responses to system maintenance events to ensure continuity of running jobs.
- 2) *I/O QoS*: Refinement of a storage system whose users receive QoS allocations through the use of *MAPE-K* loops of decreasing size and increasing automation, from rough estimates over a research campaign to parametric alteration based on profiling. The goal would be to adapt QoS parameters based on the current application performance and system I/O load to decrease interference, reduce tail latency, and provide more consistent results for deadline dependent workflows.
- 3) *OST*: Response by an application, from continuous evaluation of storage back-end write performance, to close files using a poorly performing OST object storage target (OST), that is, a storage volume of a parallel filesystem such as Lustre). The application would then reopen them using different OSTs, or explicitly request to avoid that OST in a case where the filesystem would allow it.
- 4) *Misconfiguration*: Detection of misconfiguration of user jobs such as unintended mismatch of threads to cores, underutilization of CPUs or GPUs, or wrong library search paths. Depending on the type of misconfiguration, users could either be informed about their mistake along with suggestions for better configurations, or the misconfiguration could be corrected on the fly.
- 5) *Scheduler*: Modification of a job's allocated run time based on continuous evaluation of its projected time to completion. This would also be extended to enable the scheduler to signal an application to checkpoint based on the time needed to write a checkpoint and the time remaining in an allocation.

We will initially focus on the *Scheduler* case to define and develop our first set of common components. The *MAPE-K* autonomy loop is described here and illustrated in Fig. 3:

- *Monitor* progress of an application. This could be via markers that could be output by an application (e.g., simulation time-step) or via progress information based on function calls or any application-relevant convergence criterion.

- *Analyze* the progress relative to representative historical application run times, which would need to be collected and stored along with appropriate metadata. Given an application, a strategy is also required to map the application to a set of measurements of behavioral characteristics to enable comparison against past and future runs.
- *Plan* action to be taken. This should take into account prior *Knowledge* of running time and progress rate (which might have to be inferred from similar jobs with different input decks). This may also take into account system state and expected changes due to projected changes in workload and associated resource utilization.
- *Execute* the determined response. Although the determined response may be to inform an application that it needs to request a run time extension or even to make the request on behalf of an application, the scheduler may deny the request or provide a shorter extension than requested.
- Assess the *Knowledge* about the success of the *Plan* and refine the *Knowledge* through subsequent *Monitoring* of the job's progress, iterating the *MAPE-K* loop. Note that this needs awareness of whether or not the request was honored by the scheduler.

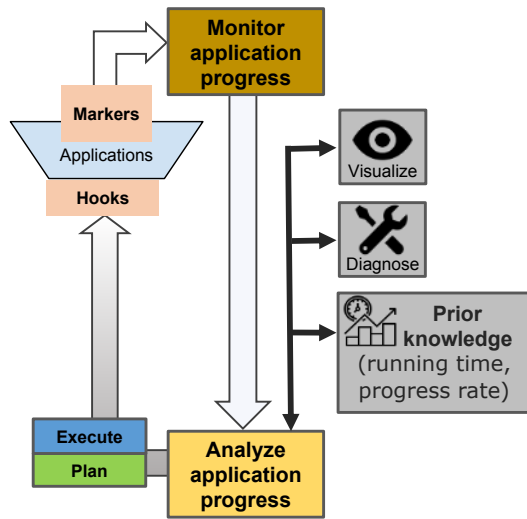


Fig. 3: Scheduler use case and its MAPE-K loop components.

The *Scheduler* case was chosen as the initial case for several reasons. First, we can gain insight into the variation of progress markers and run time through experimentation. This data should be straightforward to obtain and will be foundational to assessing the potential value of the work and developing the logic for the *Plan*.

Second, we can obtain reasonable initial functionality with a single “classical” autonomy loop per application with loosely coupled implementations needed for the interactions between *MAPE-K* components. Interactions necessary for the *Monitor* and *Analyze* phases could be simply done by having the application’s rank 0 drop time-steps periodically to a file or memory region to then be used in the progress assessment.

Thus the components performing the progress assessment and determining the action do not have to be in the same process space as the application. Also, the components performing the *Execute* phase for scheduler interaction only require the rights to invoke the scheduler to increase the wall time for the job. For typical HPC schedulers, such as SLURM [29], this is an existing command-line functionality.

Third, a few simple measurable quantities can be used to forecast time to completion which will be used, in conjunction with the remaining allocation time, to plan what action, if any, to take. These same quantities can also be used to assess the effectiveness of the *Plan* (e.g., over/under-estimation of a change to time allocation).

Fourth, a clear path exists to explore extensibility of the design and interactions from the simple prototype. Including an option for invoking asynchronous checkpointing would drive design of increasingly complex *Plan* and *Knowledge* components and application interactions.

Exploring extensibility will be foundational for developing our other cases as well. The following examples describe such relationships between the *Scheduler* case and our other four target cases: 1) the *Maintenance* case would use equivalent application interaction as invoking asynchronous checkpointing, 2) the *I/O QoS* case would utilize the same capability for storage/retrieval/comparison of behavioral attributes of an application (e.g., I/O bandwidth profile) along with application interaction with extension to provide guidance on appropriate times to perform I/O, 3) the *OST* case would again utilize the capability for storage/retrieval of behavioral attributes in order to have a reference for expected operation along with application interaction to inform another response, and 4) the *Misconfiguration* case would likewise require storage/retrieval of behavioral attributes and relationships to compare an application run with expectations along with application interaction and, potentially, scheduler interaction.

The development process for the *Scheduler* use case will follow the questions in Section II:

- The monitoring system will *Monitor* the application progress, potentially with help from the runtime system or instrumentation of the application. A yet-to-be-developed service will *Analyze* the progress and *Plan* the intervention. The scheduler will *Execute* the run-time extension.
- The actual interfaces will be determined during development, but tools and frameworks already exist that could be leveraged. The relationships, listed above, between the *Scheduler* and other cases will help determine possibilities for common interfaces and interoperability.
- We plan to release the exploratory datasets used to gain insight into the variation of progress markers and run-time variation as open datasets.
- Validation of the run-time extension will be clear through comparison of the time extension with the actual application run time. Trust by users and system administrators would require that other workloads and jobs were not adversely impacted by the extension mechanism. This could be done by additional controls, such as limits on the

number and overall time of extensions for a single application, and evaluations such as run time overestimations that would have resulted in untaken backfill opportunities.

- v) Adopting an autonomy loop that increases their jobs' execution success would incentivize users. Additional statistics, such as increase in completed and decrease in resubmitted jobs, would incentivize administrators to deploy it. Success in the *Scheduler* case could also motivate developers and users to implement hooks for a checkpointing response as well, since a job would not be extended indefinitely.

#### IV. DESIGN CHANGES FOR AUTONOMOUS MODA

Autonomy loops are the “killer case” for the MODA community. The ability to continuously make and enact data-driven decisions without requiring a human in the loop motivates the collection, analysis, and retention of holistic data at higher fidelity than ever before. This will cause both changes in and opportunities for design strategies for MODA autonomy loops.

Increases in core counts have long been seen as providing an opportunity to co-locate analytics closer to compute resources. However, this opportunity has not been widely realized. Our autonomy use cases target new and tighter interactions with applications and system software both for gathering information and enacting response. Ideally, an established standardized set of interfaces for each component type would make the components interchangeable and the loop(s) modular, however, different requirements and associated implementations (e.g., latency, sampling rates, cardinality, high availability for monitoring) may drive multiple interfaces and interactions. Therefore, *interoperability* and *interchangeability* are key design considerations. We will support these by ensuring well-defined and documented interfaces for interactions we develop.

Increasing possibilities for low-latency actions will provide more motivation for *in situ* analytics and decision-making, and hence storage in MODA designs. This will also drive more complex *MAPE-K* design patterns than first investigated in the *Scheduler* case. Distributed autonomy, where each component has some decision-making capability and decision-executing authority (as discussed in Section II), will be useful for robust and resilient operations. Agent-based models have shown this in the context of distributed systems [30]. Resilience is essential in HPC systems where operations must persist through component and subsystem failures.

Failure prediction and anomaly detection have long been MODA analysis goals. Our cases further analytics in continuous performance characterizations and comparisons. Our storage architecture decisions will then increasingly consider metadata representations for models, moving beyond traditional considerations of insert rates for raw time-series data.

Relatedly, our analyses will also be expanded to include determination of confidence in the models for decision-making and assessment of the effectiveness of the *Plan* and *Execute* phases. Confidence measures are required as we move beyond human-in-the-loop decision-making, particularly for safe operations of power and energy controls.

Note that autonomy loops in HPC operations do not have to replace the human-in-the-loop approach, and could complement it. A *human-on-the-loop* approach would have the loop continue without waiting for user and administrator input, but sending them notifications and explanation about decisions that allow for observing its effects when necessary [31]. The decision-making would then also include execution of contingency plans for when the humans are absent.

The HPC domain should look to AI tools, algorithms, and generated models to help drive the automated decision process. However, focus should be on careful selection of efficient models and modeling parameters that fit HPC data. For instance, the present outlook in the AI community is the use of large models with millions of parameters. However, such models may not be efficient when complex optimizations for real-time decisions must be made. In fact, the constantly evolving nature of the environment requires continual/lifelong AI that can evolve rapidly with small overhead [32]. Moreover, precision requirements must be built into these AI models, and the lack of *interpretability* and *explainability* must be addressed to get robust AI models. Furthermore, use of AI for autonomous loops may impact resources allocated to applications. In summary, simply applying the present AI tools and algorithms will not be sufficient and ample opportunities exist for further design and development.

#### V. CONCLUSION

In this position paper we have reported on the outcomes and paths forward from a recent Dagstuhl Seminar, seeking to carve a path for community progress on the development of autonomous feedback loops for MODA, based on the established formalism and architecture of *MAPE-K* loops in autonomous computing and self-adaptive systems.

We are presenting our position early in our work process in order to get input from the wider HPC community and to engage collaborators interested in determining interfaces for HPC autonomy loop components and developing interoperable and interchangeable components that utilize those interfaces.

Additionally, we seek to facilitate development by encouraging testbeds to be defined and made available to the community. The main obstacle to exploration of autonomy loops is the fear of potentially intrusive changes to the system behavior, often deemed unacceptable on production systems. Nevertheless, it may be possible to utilize stranded resources of HPC systems that are being decommissioned for a limited time, or include experiments during bring-up or extended maintenance. However, we believe that the best approach is to include MODA targets into the system definition itself. Such efforts will be propelled by well-defined modularization of *MAPE-K* components and their associated APIs, so that individual components can be replaced while preserving appropriate system boundaries, enabling appropriate auditing and trust levels. The OpenCUBE project [33] is aiming to provide such opportunities, by defining a process to submit *MAPE-K* loop experiments to be executed on their testbed system.

## ACKNOWLEDGMENT

The authors thank Schloss Dagstuhl for hosting Seminar 23171 “Driving HPC Operations With Holistic Monitoring and Operational Data Analytics.”

The work is jointly supported by the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 957407, DAPHNE).

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 (LLNL-CONF-851925).

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy’s National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

This work was supported by the U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research, under Contract DE-AC02-06CH11357.

This work was funded by the European Union under the Horizon Europe program’s OpenCUBE project, grant agreement 101092984.

## REFERENCES

- [1] A. Netti, W. Shin, M. Ott, T. Wilde, and N. Bates, “A conceptual framework for HPC operational data analytics,” in *2021 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, 2021, pp. 596–603.
- [2] A. Agelastos, B. Allan, J. Brandt, P. Cassella, J. Enos, J. Fullop, A. Gentile, S. Monk, N. Naksinehaboon, J. Ogden, M. Rajan, M. Showerman, J. Stevenson, N. Taerat, and T. Tucker, “The lightweight distributed metric service: A scalable infrastructure for continuous monitoring of large scale computing systems and applications,” in *SC ’14: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2014, pp. 154–165.
- [3] A. Netti, M. Müller, A. Auweter, C. Guillen, M. Ott, D. Tafani, and M. Schulz, “From facility to application sensor data: Modular, continuous and holistic monitoring with DCDB,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2019, pp. 1–27.
- [4] J. Li, G. Ali, N. Nguyen, J. Hass, A. Sill, T. Dang, and Y. Chen, “Monster: An out-of-the-box monitoring tool for high performance computing systems,” in *2020 IEEE International Conference on Cluster Computing*, ser. CLUSTER ’20, 2020, pp. 119–129.
- [5] G. K. Lockwood, W. Yoo, S. Byna, N. J. Wright, S. Snyder, K. Harms, Z. Nault, and P. Carns, “UMAMI: A recipe for generating meaningful metrics through holistic I/O performance analysis,” in *Proceedings of the 2nd Joint International Workshop on Parallel Data Storage & Data Intensive Scalable Computing Systems*, ser. PDSW-DISCS ’17. New York, NY, USA: ACM, 2017, pp. 55–60.
- [6] E. Arima, A. I. Comprés, and M. Schulz, “On the convergence of malleability and the HPC PowerStack: Exploiting dynamism in over-provisioned and power-constrained HPC systems,” in *High Performance Computing. ISC High Performance 2022 International Workshops*. Springer, 2022, pp. 206–217.
- [7] B. Yang, X. Ji, X. Ma, X. Wang, T. Zhang, X. Zhu, N. El-Sayed, H. Lan, Y. Yang, J. Zhai, W. Liu, and W. Xue, “End-to-end I/O monitoring on a leading supercomputer,” in *16th USENIX Symposium on Networked Systems Design and Implementation*, ser. NSDI ’19. Boston, MA: USENIX Association, 2019, pp. 379–394.
- [8] B. Behzad, S. Byna, Prabhat, and M. Snir, “Optimizing I/O performance of HPC applications with autotuning,” *ACM Trans. Parallel Comput.*, vol. 5, no. 4, 2019.
- [9] N. Bourassa, W. Johnson, J. Broughton, D. M. Carter, S. Joy, R. Vitti, and P. Brandt, “Operational data analytics: Optimizing the national energy research scientific computing center cooling systems,” in *Proceedings of the 48th International Conference on Parallel Processing: Workshops*, 2019, pp. 1–7.
- [10] J. Brandt, A. Gentile, J. Mayo, P. Pebay, D. Roe, D. Thompson, and M. Wong, “Resource monitoring and management with OVIS to enable HPC in cloud computing environments,” in *2009 IEEE International Symposium on Parallel & Distributed Processing*, 2009, pp. 1–8.
- [11] O. Aaziz, B. Allan, J. Brandt, J. Cook, K. Devine, J. Elliott, A. Gentile, S. Hammond, B. Kelley, L. Lopatina, S. Moore, S. Olivier, K. Pedretti, D. Poliakoff, R. Pawlowski, P. Regier, M. Schmitz, B. Schwaller, V. Surjadidjaja, M. S. Swan, N. Tucker, T. Tucker, C. Vaughan, and S. Walton, “Integrated system and application continuous performance monitoring and analysis capability,” *Sandia report*, 9 2021. [Online]. Available: <https://www.osti.gov/biblio/1819812>
- [12] R. Izadpanah, N. Naksinehaboon, J. Brandt, A. Gentile, and D. Dechev, “Integrating low-latency analysis into HPC system monitoring,” in *Proceedings of the 47th International Conference on Parallel Processing*, ser. ICPP ’18. New York, NY, USA: ACM, 2018.
- [13] J. Brandt, D. DeBonis, A. Gentile, J. Lujan, C. Martin, D. Martinez, S. Olivier, K. Pedretti, N. Taerat, and R. Velarde, “Enabling advanced operational analysis through multi-subsystem data integration on Trinity,” in *Proc. Cray Users Group*, 2015.
- [14] J. Eitzinger, T. Gruber, A. Afzal, T. Zeiser, and G. Wellein, “Cluster-Cockpit — a web application for job-specific performance monitoring,” in *2019 IEEE International Conference on Cluster Computing*, ser. CLUSTER ’19. IEEE, 2019, pp. 1–7.
- [15] M. Isakov, E. d. Rosario, S. Madireddy, P. Balaprakash, P. Carns, R. B. Ross, and M. A. Kinsy, “HPC I/O throughput bottleneck analysis with explainable local models,” in *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, 2020, pp. 1–13.
- [16] S. Herbein, D. H. Ahn, D. Lipari, T. R. Scogland, M. Stearman, M. Grondona, J. Garlick, B. Springmeyer, and M. Taufer, “Scalable I/O-aware job scheduling for burst buffer enabled HPC clusters,” in *Proceedings of the 25th ACM International Symposium on High-Performance Parallel and Distributed Computing*, ser. HPDC ’16. New York, NY, USA: ACM, 2016, pp. 69–80.
- [17] G. Ali, S. Bhalachandra, N. J. Wright, M. Side, and Y. Chen, “Optimal GPU frequency selection using multi-objective approaches for hpc systems,” in *2022 IEEE High Performance Extreme Computing Conference*, ser. HPEC ’22, 2022, pp. 1–7.
- [18] F. Chowdhury, F. Di Natale, A. Moody, K. Mohror, and W. Yu, “DFMan: A graph-based optimization of dataflow scheduling on high-performance computing systems,” in *2022 IEEE International Parallel and Distributed Processing Symposium*, ser. IPDPS ’22, 2022, pp. 368–378.
- [19] H. Shahzad, A. Sanaullah, S. Arora, R. Munafo, X. Yao, U. Drepper, and M. Herbordt, “Reinforcement learning strategies for compiler optimization in high level synthesis,” in *2022 IEEE/ACM 8th Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC)*, 2022, pp. 13–22.
- [20] J. Ansel, S. Kamil, K. Veeramachaneni, J. Ragan-Kelley, J. Bosboom, U.-M. O’Reilly, and S. Amarasinghe, “Opentuner: An extensible framework for program autotuning,” in *Proceedings of the 23rd International Conference on Parallel Architectures and Compilation*, 2014, pp. 303–316.
- [21] M. R. Wyatt, S. Herbein, K. Shoga, T. Gamblin, and M. Taufer, “CanariO: Sounding the alarm on IO-related performance degradation,” in *2020 IEEE International Parallel and Distributed Processing Symposium*, ser. IPDPS ’20, 2020, pp. 73–83.
- [22] Z. Qiao, Q. Liu, N. Podhorszki, S. Klasky, and J. Chen, “Taming I/O variation on QoS-less HPC storage: What can applications do?” in *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE Press, 2020.
- [23] A. Patke, S. Jha, H. Qiu, J. Brandt, A. Gentile, J. Greenseid, Z. Kalbarczyk, and R. K. Iyer, “Delay sensitivity-driven congestion mitigation for HPC systems,” in *Proceedings of the ACM International Conference on Supercomputing*, ser. ICS ’21. New York, NY, USA: ACM, 2021, pp. 342–353.

- [24] J. Brandt, K. Devine, A. Gentile, and K. Pedretti, "Demonstrating improved application performance using dynamic monitoring and task mapping," in *2014 IEEE International Conference on Cluster Computing*, ser. CLUSTER '14, 2014, pp. 408–415.
- [25] "Driving HPC operations with holistic monitoring and operational data analytics (Dagstuhl Seminar 231710)." [Online]. Available: <https://www.dagstuhl.de/23171>
- [26] D. Weyns, *An Introduction to Self-adaptive Systems: A Contemporary Software Engineering Perspective*. John Wiley & Sons, Inc., 2020.
- [27] J. O. Kephart and D. M. Chess, "The vision of autonomic computing," *Computer*, vol. 36, no. 1, pp. 41–50, 2003.
- [28] D. Weyns, B. Schmerl, V. Grassi, S. Malek, R. Mirandola, C. Prehofer, J. Wuttke, J. Andersson, H. Giese, and K. Göschka, "On patterns for decentralized control in self-adaptive systems," in *Software Engineering for Self-Adaptive Systems II*, ser. LNCS. Springer, 2013, vol. 7475, pp. 76–107.
- [29] SchedMD, "SLURM workload manager." [Online]. Available: <https://slurm.schedmd.com/overview.html>
- [30] D. Ackley and D. Cannon, "Pursue robust indefinite scalability," in *13th Workshop on Hot Topics in Operating Systems (HotOS XIII)*. USENIX, 2011.
- [31] N. Li, S. Adepu, E. Kang, and D. Garlan, "Explanations for human-on-the-loop: A probabilistic model checking approach," in *Proceedings of the IEEE/ACM 15th International Symposium on Software Engineering for Adaptive and Self-Managing Systems*, ser. SEAMS '20. New York, NY, USA: ACM, 2020, p. 181–187.
- [32] O. Gheibi and D. Weyns, "Lifelong self-adaptation: Self-adaptation meets lifelong machine learning," in *Proceedings of the 17th Symposium on Software Engineering for Adaptive and Self-Managing Systems*, ser. SEAMS '22. New York, NY, USA: ACM, 2022, p. 1–12.
- [33] OpenCUBE, "Open-source cloud-based services on EPI systems," <https://horizon-opencube.eu/>.