INTERSECT: The Open Federated Architecture for the Laboratory of the Future

Christian Engelmann, Swen Boehm, Michael Brim, Jack Lange, Thomas Naughton, Patrick Widener, Ben Mintz, and Rohit Srivastava*
engelmannc@ornl.gov
Oak Ridge National Laboratory
Oak Ridge, Tennessee, USA

ABSTRACT

The open Self-driven Experiments for Science / Interconnected Science Ecosystem (INTERSECT) architecture connects scientific instruments and robot-controlled laboratories with computing and data resources at the edge, the Cloud or the high-performance computing center to enable autonomous experiments, self-driving laboratories, smart manufacturing, and artificial intelligence driven design, discovery and evaluation. Its a novel approach consists of science use case design patterns, a system of systems architecture, and a microservice architecture.

ACM Reference Format:
Christian Engelmann, Swen Boehm, Michael Brim, Jack Lange, Thomas Naughton, Patrick Widener, Ben Mintz, and Rohit Srivastava. 2023. INTERSECT: The Open Federated Architecture for the Laboratory of the Future. In *ICPP '23: 52nd International Conference on Parallel Processing, August 07–10, 2023, Salt Lake City, Utah, USA.* ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

As outlined in the U.S. Department of Energy (DoE)'s Artificial intelligence (AI) for Science reports [2, 7] and the DoE's Computational Facilities Research Workshop report [3], science breakthroughs with autonomous experiments, self-driving laboratories, smart manufacturing, and AI-driven design, discovery and evaluation require intelligent systems, instruments, and facilities that enable automation and reduce human-in-the-loop needs.

Machine-in-the-loop intelligence for decision-making reduces such human-in-the-loop needs with an autonomous online control, collecting experiment data, analyzing it, and taking appropriate operational actions for experiment steering or design. A common federated hardware/software architecture is needed that connects instruments with edge and center computing resources and autonomously collects, transfers, stores, processes, curates, and archives scientific data.

The INTERSECT open architecture connects scientific instruments and robot-controlled laboratories with computing and data resources at the edge, the Cloud or the high-performance computing center to enable such autonomous experiments and self-driving laboratories. Oak Ridge National Laboratory (ORNL) is currently developing and deploying several prototypes using this architecture, such as in autonomous additive manufacturing, autonomous continuous flow reactor synthesis, autonomous electron microscopy, and an autonomous robotic chemistry laboratory.

This poster summarizes the current status of the INTERSECT open architecture and the progress made a year after its initial publication in summer/fall 2022 [1, 4–6] to update the community on recent significant improvements in the science use case design pattern, system of systems (SoS) architecture and microservice architecture specifications, and to seek feedback from the community.

2 THE INTERSECT OPEN ARCHITECTURE

The approach is inspired by the U.S. Department of Defense Architecture Framework (DoDAF) [9] and its different architectural viewpoints, such as (i) operational scenarios, (ii) composition, interconnectivity and context, (iii) services and their capabilities, (iv) policies, standards and guidance, and (v) capability. However, the INTERSECT open architecture [4] separates these viewpoints differently for clarity as: (1) science use case design patterns [5], (2) a SoS architecture with some of the DoDAF viewpoints [6], and (3) a microservice architecture [1].

Science use cases (for autonomous experiments and self-driving laboratories, smart manufacturing, and AI-driven design, discovery and evaluation) are described as design patterns to identify and abstract the involved components and their interactions in terms of control, work, and data flow. The SoS architecture clarifies used terms, architectural elements, their interactions, and compliance. The microservice architecture maps the science use case design patterns to the SoS architecture with loosely coupled microservices and standardized interfaces.

This approach separates different granularities of architectural design decisions. At coarse granularity (design patterns), the overall objective of a self-driving laboratory is considered and how it is achieved. At mid-level granularity (SoS architecture), the instruments, robots, networks, and computing systems that are part of a self-driving laboratory are considered and how they interact with each other. At fine granularity, the individual microservices that orchestrate experiment control, data transfer, and data analysis are considered.

2.1 Science Use Case Design Patterns

The basic template for a science use case design pattern is defined in a loop control problem paradigm [5]. There are two classes of

---

ICPP '23, August 07–10, 2023, Salt Lake City, Utah, USA

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

science use case design patterns: strategy patterns and architectural patterns. Strategy patterns define high-level solution methods using experiment control architecture features at a very coarse granularity. Architectural patterns define more specific solution methods using hardware and software architecture features at a finer granularity.

While the architectural patterns do inherit the features of certain parent strategy patterns, they also address additional problems that are not exposed at the high abstraction level of the strategy patterns. A specific solution may require the composition of patterns, such as the Experiment Steering and Design of Experiments strategy patterns and correspondingly the Local Experiment Steering and Distributed Design of Experiments architectural patterns.

2.2 System of Systems architecture

The SoS architecture decomposes the federated hardware/software ecosystem into smaller and less complex systems and components within these systems [6]. It permits the development of individual systems and components with clearly defined interfaces, data formats, and communication protocols. This not only separates concerns and functionality for reusability, but also promotes pluggability and extensibility with uniform protocols and system/component life cycles.

Instead of developing individual monolithic solutions for each science use case, the SoS architecture provides one solution that can be easily adapted to different use cases using different compositions of systems. It offers operational and managerial independence of systems and of components within systems, geographical distribution with a physically distributed and federated ecosystem, emergent behavior based on the interplay between systems and components, and evolutionary development through pluggability and extensibility.

2.3 Microservice Architecture

The microservice architecture provides a catalog of infrastructure and experiment-specific microservices [1]. The microservices are defined to facilitate composition within the federated SoS architecture. INTERSECT infrastructure microservices represent common service functionality and capabilities, such as data management, computing, messaging, and workflow orchestration that are likely to be generally useful across many science ecosystems without the need for customization. Experiment-specific microservices, on the other hand, represent services whose implementation may require detailed application knowledge, such as experiment planning or steering services that require knowledge of experiment-specific control parameters and their associated constraints.

3 CONCLUSION

The INTERSECT open architecture has been implemented in part by the INTERSECT Software Development Kit project [8] and integrated in part in a co-design effort with the INTERSECT domain science projects at ORNL, such as autonomous electron microscopy. Ongoing work focuses on finalizing the INTERSECT open architecture, releasing a version 1.0, obtaining feedback from the community, and improving the architecture documentation.

Future work seeks to clarify the relationships between the INTERSECT science use case design patterns and execution patterns, scientific workflow motifs, and workflow execution patterns, which are not design patterns but categorize behavioral commonalities of workflows. Further research investigates cybersecurity policies and practices that affect the SoS architecture, such as experienced by multi-site experiment workflows. Other planned efforts focus on the concepts, data types, and interfaces for error reporting and handling in the SoS architecture and the microservice architecture.

ACKNOWLEDGMENTS

Research sponsored by the Laboratory Directed Research and Development Program's INTERSECT Initiative of Oak Ridge National Laboratory. This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (http://energy.gov/downloads/doe-public-access-plan).

REFERENCES

[1] Michael Brim and Christian Engelmann. 2022. *INTERSECT Architecture Specification: Microservice Architecture (Version 0.5).* Technical Report ORNL/TM-2022/2715. Oak Ridge National Laboratory, Oak Ridge, TN, USA. https://doi.org/10.2172/1902805

[2] Jonathan Carter, John Feddema, Doug Kothe, Rob Neely, Jason Pruet, and Rick Stevens. 2023. Advanced Research Directions on AI for Science, Energy, and Security. Available at https://www.anl.gov/ai-for-science-report.

[3] DOE 2020. *DOE National Laboratories' Computational Facilities – Research Workshop Report.* Technical Report ANL/MCS-TM-388. Argonne National Laboratory, Lemont, IL, USA. https://publications.anl.gov/anlpubs/2020/02/158604.pdf

[4] Christian Engelmann, Olga Kuchar, Swen Boehm, Michael J. Brim, Thomas Naughton, Suhas Somnath, Scott Atchley, Jack Lange, Ben Mintz, and Elke Arenholz. 2022. The INTERSECT Open Federated Architecture for the Laboratory of the Future. In *Communications in Computer and Information Science (CCIS): Accelerating Science and Engineering Discoveries Through Integrated Research Infrastructure for Experiment, Big Data, Modeling and Simulation.* 18th *Smoky Mountains Computational Sciences & Engineering Conference (SMC) 2022*, Vol. 1690. Springer, Cham, 173–190. https://doi.org/10.1007/978-3-031-23606-8_11 Acceptance rate 32.4% (24/74).

[5] Christian Engelmann and Suhas Somnath. 2022. *INTERSECT Architecture Specification: Use Case Design Patterns (Version 0.5).* Technical Report ORNL/TM-2022/2681. Oak Ridge National Laboratory, Oak Ridge, TN, USA. https://doi.org/10.2172/1896984

[6] Olga A. Kuchar, Swen Boehm, Thomas Naughton, Suhas Somnath, Ben Mintz, Jack Lange, Scott Atchley, Rohit Srivastava, and Patrick Widener. 2022. *INTERSECT Architecture Specification: System-of-systems Architecture (Version 0.5).* Technical Report ORNL/TM-2022/2717. Oak Ridge National Laboratory, Oak Ridge, TN, USA. https://doi.org/10.2172/1968700

[7] Rick Stevens, Valerie Taylor, Jeff Nichols, Arthur Barney Maccabe, Katherine Yelick, and David Brown. 2020. AI for Science Report. Available at https://www.anl.gov/cels/reference/ai-for-science-report-2020.

[8] Addi Malviya Thakur, Seth Hitefield, Marshall McDonnell, Matthew Wolf, Richard Archibald, Lance Drane, Kevin Roccapriore, Maxim Ziatdinov, Jesse McGaha, Robert Smith, John Hetrick, Mark Abraham, Sergey Yakubov, Greg Watson, Ben Chance, Clara Nguyen, Matthew Baker, Robert Michael, Elke Arenholz, and Ben Mintz. 2022. Towards a Software Development Framework for Interconnected Science Ecosystems. In *Accelerating Science and Engineering Discoveries Through Integrated Research Infrastructure for Experiment, Big Data, Modeling and Simulation*, Kothe Doug, Geist Al, Swaroop Pophale, Hong Liu, and Suzanne Parete-Koon (Eds.). Springer Nature Switzerland, Cham, 206–224.

[9] U.S. Department of Defense. 2010. The DoDAF Architecture Framework Version 2.02. https://dodcio.defense.gov/Library/DoD-Architecture-Framework