LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY

Methods and Experiences for Developing Abstractions for Data-intensive, Scientific Applications

Andre Luckow[1] and Shantenu Jha[2,3]

[1]Ludwig-Maximilian University, Munich, Germany
[2]RADICAL, ECE, Rutgers University, Piscataway, NJ 08854, USA
[3]Brookhaven National Laboratory, Upton, NY, USA

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Motivation

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY

Large number of scientific applications require high performance compute and data capabilities. Many challenges related to infrastructure and management exist:

- **Heterogeneous infrastructure increases complexity:** instruments, edge, fog, HPC, cloud, serverless, accelerators need to be integrated.
- **Elasticity:** dynamic, externally induced changes to resource demands
- **Scheduling and provisioning of resources in a complex and dynamic environment:** right amount of resources at right time
- **Different application components require different programming models:** HPC (MPI, OpenMP, CUDA/GPUs), Big Data (MapReduce, Streaming), Serverless
- **Limited interoperability between infrastructures:** Resource management of tied to specific API of system (e.g. HPC schedulers or cloud API)
- **Distributed Compute and Data:** Data sources distributed across different environments (IoT, cloud, HPC, serverless)

Why are abstractions important?

Multiple distinct applications with complex characteristics and infrastructure requirements.

The complexity of infrastructure and applications prevents critical scalability and make scientific progress.

Objective: Abstractions that can support multiple applications across different scales and heterogeneity.

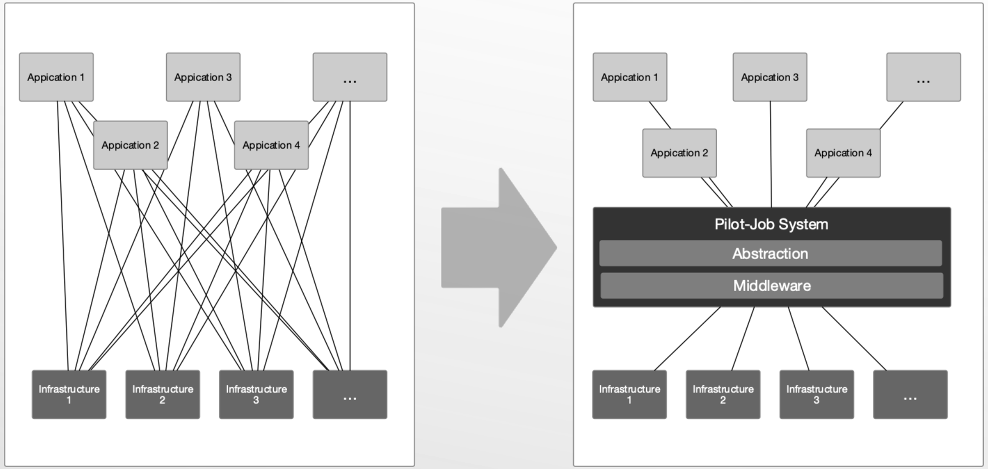

Challenge: Designing useful abstraction is challenging. Hiding complexity does not automatically lead to simple interfaces. Effective methods for developing and evaluation abstractions required!

Design Science Research for Abstraction Design

Relevance

Environment

Application
- Characteristics
- Application patterns
- Datasources: IoT, Instruments
- Dependencies
- Performance requirements
- …

Infrastructure
- HPC
- HTC
- Cloud
- …

Application Needs

Research

Design/Develop
Design **abstractions** for resource management on heterogeneous **infrastructure** for data-intensive **applications** enabling **performance** and **scale.**

Assess

Refine

Predict/Evaluate
- Case studies
- Conceptual modeling
- Performance measurement and modeling
- Mini-App Framework

Rigor

Knowledge Base

Applicable Knowledge

Foundations
- Parallel and distributed computing
- Frameworks
- …

Approaches
- Conceptual modeling
- Software engineering methods
- Design patterns
- Architecture models
- Performance measurements
- Performance models
- …

Application in appropriate environment

Addition to knowledge base

Problem Statement: Application Characteristics and Requirements

Work	Method	Scope	Description
Scientific Applications: Introducing distributed dynamic data-intensive (D3) science: Understanding applications and infrastructure [1].	Survey, Workshop, Literature Review	13 applications, 9 questions	Conceptual framework for analyzing and understanding distributed data-intensive applications scenarios
Big Data Ogres: Towards an understanding of facets and exemplars of big data applications [2, 3, 4]	Workshops, Literature Review	51 NIST Use Cases using 4 views and 40+ facets	Conceptual framework for characterizing commonalities and patterns in Big Data applications
Streaming computational science: Applications, technology and re-source management for HPC [5,6]	Workshops, Literature Review	12 applications, 4 applications categories	Characterization of streaming applications

[1] Shantenu Jha, Daniel S. Katz, Andre Luckow, Neil Chue Hong, Omer Rana, and Yogesh Simmhan. Introducing distributed dynamic data- intensive (d3) science: Understanding applications and infrastructure. Concurrency and Computation: Practice and Experience, 29(8), 2017.
[2] Shantenu Jha, Judy Qiu, Andre Luckow, Pradeep Kumar Mantha, and Geoffrey Charles Fox. A tale of two data-intensive paradigms: Applications, abstractions, and architectures. Proceedings of 3rd IEEE Internation Congress of Big Data, abs/1403.1528, 2014.
[3] Geoffrey C. Fox, Shantenu Jha,Judy Qiu, and Andre Luckow. Towards an understanding of facets and exemplars of big data applications. In Proceedings of Beowulf'14, Annapolis, MD, USA, 2014. ACM.
[4] Geoffrey C. Fox, Shantenu Jha, Judy Qiu, and Andre Luckow. Ogres: A Systematic Approach to Big Data Benchmarks. White paper at workshop on Big Data and Extreme-scale Computing (BDEC), 2015.
[5] Shantenu Jha, Lavanya Ramakrishnan, and Geoffrey Fox. Stream2016: Streaming requirements, experience, applications and middleware workshop final report. Office of Advanced Scientific Computing Research, DOE Office of Science, https://www.osti.gov/servlets/purl/ 1344785/, 2016.
[6] Geoffrey C. Fox, Devarshi Ghoshal, Shantenu Jha, Andre Luckow, and Lavanya Ramakrishnan. Streaming computational science: Applications, technology and resource management for hpc. http://dsc.soic. indiana.edu/publications/streaming-nysds-abstract.pdf, 2017.

LMU LUDWIG-MAXIMILIANS-UNIVERSITÄT MÜNCHEN

Problem Statement: Definition of 5 Application Scenarios Based on Analysis of >70 Applications

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY

	Task-Parallel	Data-Parallel	Dataflow	Iterative	Streaming
Description	Focus on functional de-composition into tasks and control flow	Decomposition of a problem into a diverse set of dependent and parallel tasks	Multiple processing stages modeled with a directed acyclic graph	Multiple generations of tasks with sharing of data between the generations	Processing of unbounded data feeds in near-realtime
Characteristics	Decomposition of a problem into a diverse set of de- pendent and parallel tasks	Embarrassingly parallel, loosely-coupled with minimal communication.	Multiple stages, loosely-coupled parallelism, global communication for shuffle operation	Loosely coupled par-allelism with global communication for up-dating machine learning model parameters	Data is processed in small batches often using data-parallel algorithms.
Examples	Molecular Dynamics [1], Ensemble-Kalman Filter [2], Scientific Gateways and Workflows [3]	Map-Only analytics [4], Molecular Data analysis Hausdorff Distance [5]	MapReduce for sequence alignment [6], Molecular Data analysis leaflet finder and RMSD [7]	Machine learning algorithms, K-Means [7]	Streaming for light source data [8]

[1] Andre Luckow, Shantenu Jha, Joohyun Kim, Andre Merzky, and Bettina Schnor. Adaptive Replica-Exchange Simulations. Royal Society Philosophical Transactions A, 2009.

[2] Yaakoub El-Khamra and Shantenu Jha. Developing autonomic dis- tributed scientific applications: A case study from history matching using ensemble kalman-filters. In Proceedings of the 6th International Con- ference Industry Session on Grids Meets Autonomic Computing, GMAC '09, pages 19–28, New York, NY, USA, 2009. ACM.

[3] Sharath Maddineni, Joohyun Kim, Yaakoub El-Khamra, and Shantenu Jha. Distributed application runtime environment (dare): A standards- based middleware framework for science-gateways. Journal of Grid Computing, 10(4):647–664, 2012.

[4] G. C. Fox, J. Qiu, S. Kamburugamuve, S. Jha, and A. Luckow. Hpc-abds high performance computing enhanced apache big data stack. In 2015 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, pages 1057–1066, May 2015.

[5] Ioannis Paraskevakos, Andre Luckow, Mahzad Khoshlessan, George Chantzialexiou, Thomas E. Cheatham, Oliver Beckstein, Geoffrey C. Fox, and Shantenu Jha. Task-parallel analysis of molecular dynamics trajectories. In Proceedings of the 47th International Conference on Parallel Processing, ICPP 2018, New York, NY, USA, 2018. ACM.

[6] Pradeep Kumar Mantha, Andre Luckow, and Shantenu Jha. Pilot-MapReduce: An Extensible and Flexible MapReduce Implementation for Distributed Data. In Proceedings of third international workshop on MapReduce and its Applications, MapReduce '12, pages 17–24, New York, NY, USA, 2012. ACM.

[7] Shantenu Jha, Judy Qiu, Andre Luckow, Pradeep Kumar Mantha, and Geoffrey Charles Fox. A tale of two data-intensive paradigms: Applications, abstractions, and architectures. Proceedings of 3rd IEEE Inter- nation Congress of Big Data, abs/1403.1528, 2014.

[8] Andre Luckow, George Chantzialexiou, and Shantenu Jha. Pilot-streaming: A stream processing framework for high-performance computing. IEEE eScience, 2018.

Pilot-Abstraction: a Common Resource and Task Management Abstraction for Data & Compute

Pilot-Job System: A system that generalizes a placeholder job to provide multi-level scheduling to allow application-level control over the system scheduler via a scheduling overlay.
The **Pilot-Abstraction** generalizes the concept of resource and task execution management across different applications (HPC, MapReduce, streaming) and infrastructures (edge, cloud, HPC, serverless).

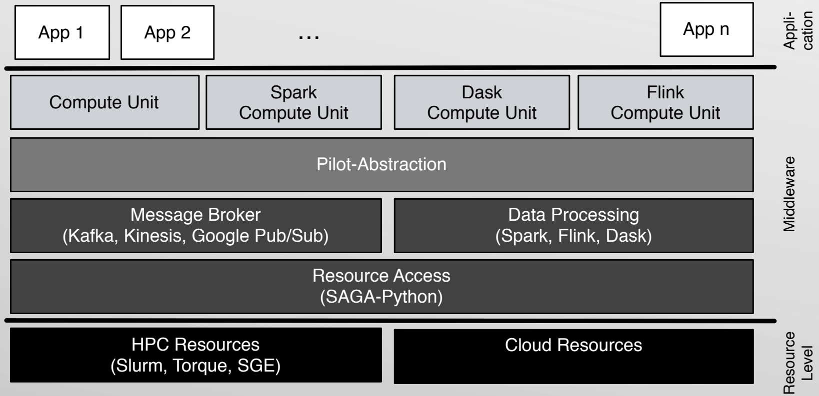

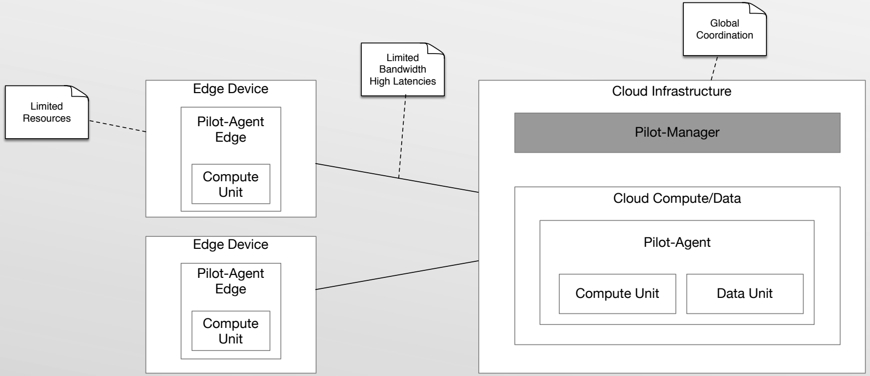

Design Science Research: Evaluation Process

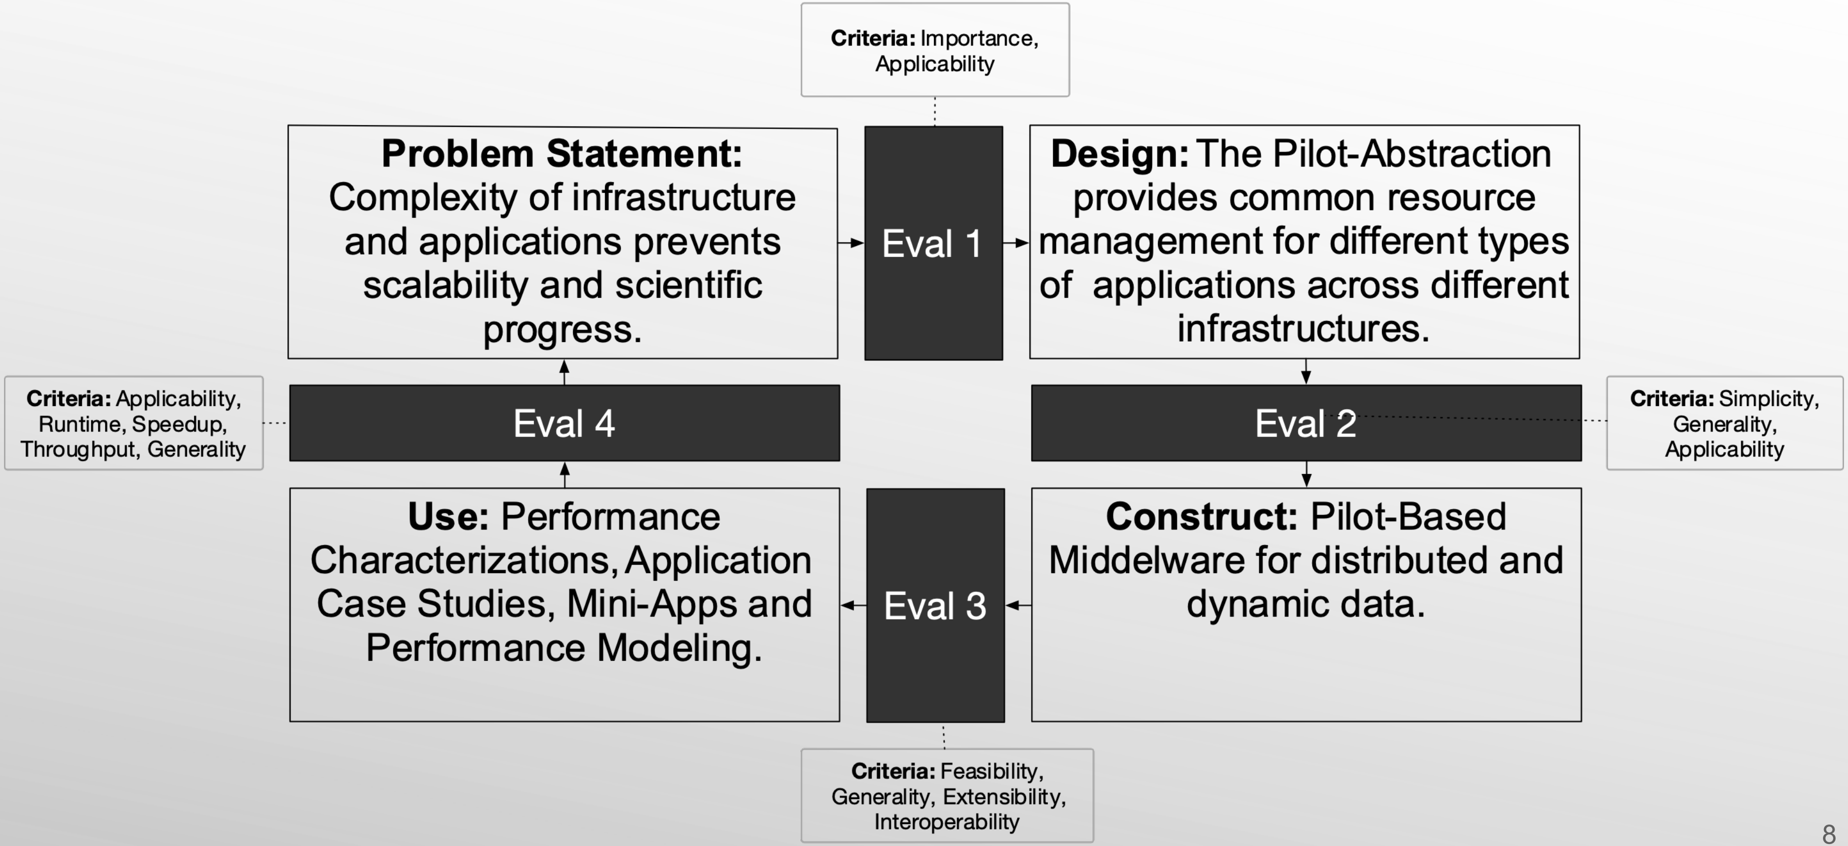

Summary of Obtained Results: Iterative Development and Refinement of Pilot-Abstraction

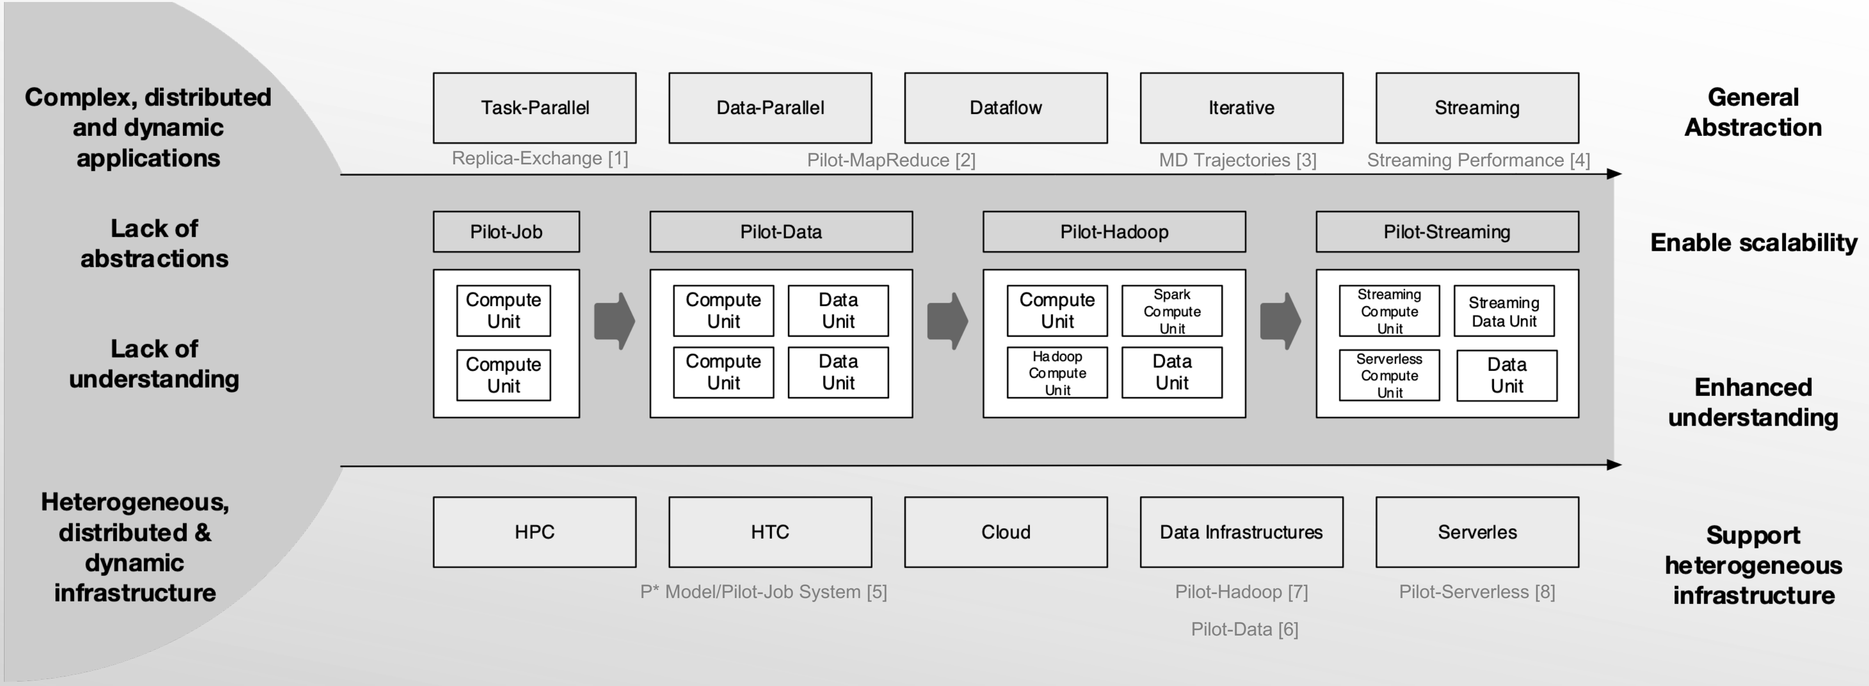

[1] Andre Luckow, Shantenu Jha, Joohyun Kim, Andre Merzky, and Bettina Schnor. Adaptive Replica-Exchange Simulations. Royal Society Philosophical Transactions A, 2009.
[2] Pradeep Kumar Mantha, Andre Luckow, and Shantenu Jha. Pilot- MapReduce: An Extensible and Flexible MapReduce Implementation for Distributed Data. In Proceedings of third international workshop on MapReduce and its Applications, MapReduce '12, pages 17–24, New York, NY, USA, 2012. ACM.
[3] Ioannis Paraskevakos, Andre Luckow, Mahzad Khoshlessan, George Chantzialexiou, Thomas E. Cheatham, Oliver Beckstein, Geoffrey C. Fox, and Shantenu Jha. Task-parallel analysis of molecular dynamics trajectories. In Proceedings of the 47th International Conference on Parallel Processing, ICPP 2018, New York, NY, USA, 2018. ACM.
[4] Andre Luckow and Shantenu Jha. Performance characterization and modeling of serverless and hpc streaming applications. In Proceedings of StreamML Workshop at IEEE International Conference on Big Data (IEEE BigData 2019), 2019.
[5] Andre Luckow, Mark Santcroos, Andre Merzky, Ole Weidner, Pradeep Mantha, and Shantenu Jha. P*: A model of pilot-abstractions. IEEE 8th International Conference on e-Science, pages 1–10, 2012. http://dx.doi.org/10.1109/eScience.2012.6404423.
[6] Andre Luckow, Mark Santcroos, Ashley Zebrowski, and Shantenu Jha. Pilot-data: An abstraction for distributed data. Journal of Parallel and Distributed Computing, 2014.
[7] Andre Luckow, Pradeep Kumar Mantha, and Shantenu Jha. Pilot- abstraction: A valid abstraction for data-intensive applications on hpc, hadoop and cloud infrastructures? CoRR, abs/1501.05041, 2015.
[8] Andre Luckow, George Chantzialexiou, and Shantenu Jha. Pilot-streaming: A stream processing framework for high-performance computing. IEEE eScience, 2018.