A Conceptual Framework for HPC Operational Data Analytics

3rd ISC HPC International Workshop on „Monitoring & Operational Data Analytics“ MODA22
Hamburg, 02/09/2026

Alessio Netti, Woong Shin, **Michael Ott**, Torsten Wilde, Natalie Bates

A Conceptual Framework for HPC Operational Data Analytics

Alessio Netti*, Woong Shin†, Michael Ott*, Torsten Wilde‡ , Natalie Bates§
*Leibniz Supercomputing Centre, {alessio.netti,michael.ott}@lrz.de
†Oak Ridge National Laboratory, shinw@ornl.gov
‡Hewlett Packard Enterprise, wilde@hpe.com
§Energy Efficient HPC Working Group, natalie.jean.bates@gmail.com

***Abstract*—This paper provides a broad framework for understanding trends in Operational Data Analytics (ODA) for High-Performance Computing (HPC) facilities. The goal of ODA is to allow for the continuous monitoring, archiving, and analysis of near real-time performance data, providing immediately actionable information for multiple operational uses. In this work, we combine two models to provide a comprehensive HPC ODA framework: one is an evolutionary model of analytics capabilities that consists of four types, which are descriptive, diagnostic, predictive and prescriptive, while the other is a four-pillar model for energy-efficient HPC operations that covers facility, system hardware, system software, and applications. This new fra[illegible] development and production deployments of ODA within leading-edge HPC facilities. Finally, we perform a comprehensive survey of ODA [illegible] in order to demonstrate its effectiveness.**

***Index Terms*—Exascale, Top500, HPC operations, Energy efficiency, Operational data**

and decisions in day-to-day operations [1]. Due to the sheer number and diversity of data streams that can be produced, ODA systems embrace large-scale data and leverage various analysis techniques to achieve this goal. In addition, ODA systems ultimately aim to provide systematic support for deriving optimal parameters that can improve *Key Performance Indicators* (KPIs). Yet, there are many difficulties in understanding, planning, designing, and implementing ODA systems, due to the degree of sophistication required to make sense out of large amounts of data. Coping with these challenges requires selecting from and experimenting with various techniques found in statistics, data mining, data science, machine learning and computer science. In many cases, HPC practitioners face difficulty in navigating through the abundance of techniques in developing use-cases and applications in their operations.

To address this challenge, we propose a conceptual frame-

A. Netti, W. Shin, M. Ott, T. Wilde and N. Bates:
A Conceptual Framework for HPC Operational Data Analytics
In Proceedings of the 2021 IEEE International Conference on Cluster Computing (CLUSTER), 2021
https://sites.google.com/view/ee-hpc-sop-2021/program

Mission

Many HPC sites are developing and deploying systems for Operational Data Analytics (ODA) to help them understand and optimize their HPC operations.

As a team, we want to learn and understand:

- What tools are the sites using for this?
- How are they using the collected data?
- What are the lessons learned?

We want to provide:

- A forum to discuss ODA deployments and use cases
- Guidance for the community to deploy similar systems
- Common terminology to foster discussion

Introduction

- *Operational Data Analytics* (ODA) uses monitoring to extract actionable **knowledge** on system **behavior**.
  - Can improve energy efficiency and reliability.
  - More and more data centers use ODA.
- However, ODA is a **broad** and **diverse** field:
  - Predictive tuning of CPU frequencies is ODA.
  - Diagnosing infrastructure failures with ML is ODA.
  - Simulating scheduling policies is ODA.
  - Computing a data center's PUE is ODA.

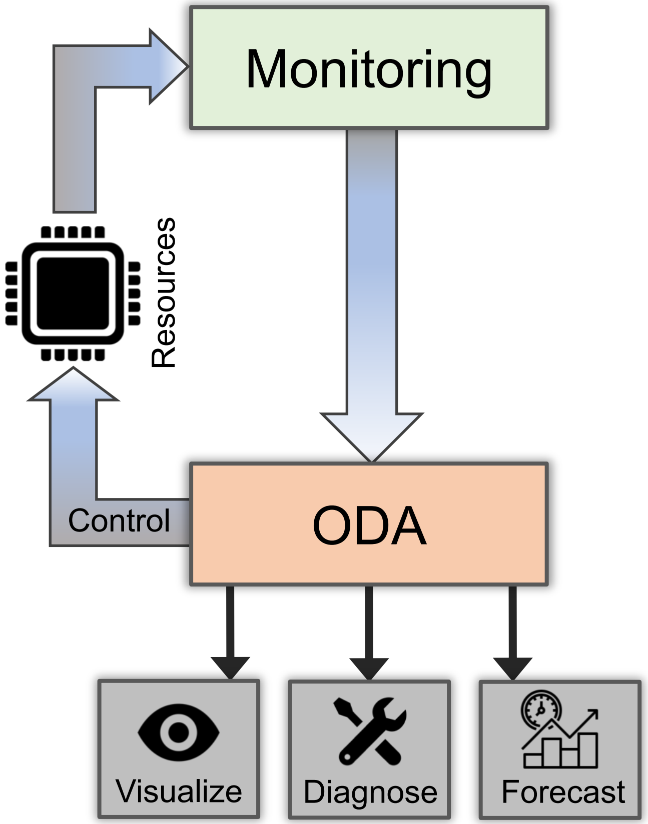

Motivation

- There is no **common language** to reason about ODA.
  - Research gaps and opportunities are difficult to identify.
  - System design and requirements are not standardized.
  - Adoption of ODA by data centers is cumbersome.
- Our **contributions** are the following:
  - A conceptual framework to help classify ODA use cases.
  - An extensive survey and categorization of ODA literature.
  - Demonstration of the framework on state-of-the-art use cases.

Designing a Framework for ODA

- Many possible **questions** about an ODA use case:
  - What is the functional **complexity** and data center scope?
  - How do we **decompose** it in simple, standard blocks?
  - Have other people already tackled a **similar** problem?
  - What are the deployment **requirements** and gains?
- We use two state-of-the-art frameworks as a **foundation**:
  - The "*4-Pillar Framework for Energy-Efficient HPC Data Centers*"*.*
  - The "*4 Types of Data Analytics Framework*".

A Conceptual Framework for HPC ODA

The 4-Pillar Framework

Pillar 1 Building Infrastructure

- Improve cooling efficiency
- Improve waste heat re-use
- Reduce power losses

Pillar 2 System Hardware

- Use newest processor technology
- Use energy-saving techniques
- Consider special accelerators for certain workloads

Pillar 3 System Software

- Optimize workload management
- Tune systems with respect to the applications' needs
- Shut down idle nodes

Pillar 4 HPC Applications

- Use the most efficient algorithms
- Tune libraries used by applications
- Use efficient programming paradigms

Facility ⟷ Users

[1] T. Wilde et al. "The 4-Pillar Framework for energy efficient HPC data centers".

The 4 Types of Data Analytics

- Model used by large consultancy firms to categorize **data analytics**.
- Consists of 4 *types*, which differ in the **functionality** they offer.
- Some types focus on **historical** events (*hindsight*), others on anticipating **future** ones (*foresight*).
- The types are not necessarily **ranked** by complexity.

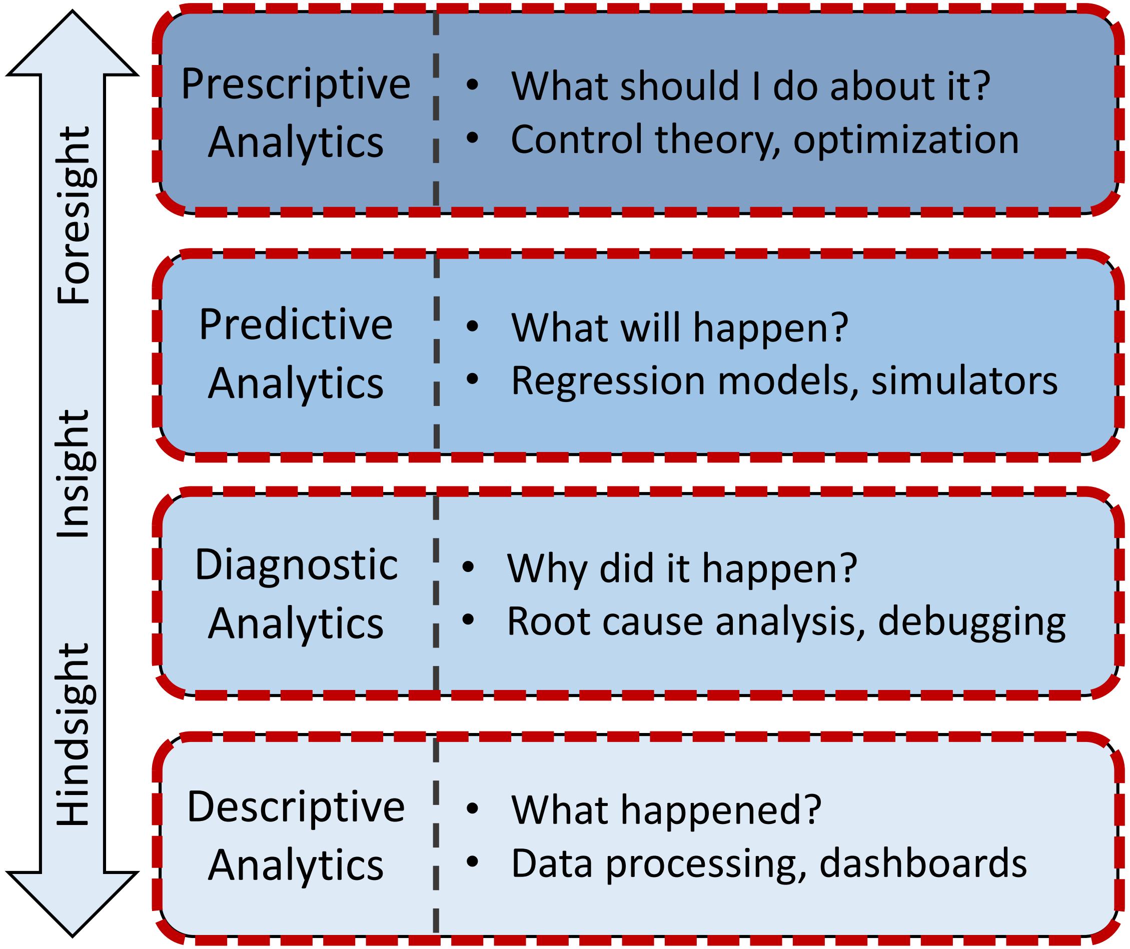

[2] K. Lepenioti et al. "Prescriptive analytics: Literature review and research challenges".

The 4x4 Conceptual ODA Framework

- We **combine** the 4-Pillar and the 4-Type models in a single framework.
- It consists of a 4x4 **matrix**:
  - The *pillars* in the horizontal axis describe the scope of ODA.
  - The *types* in the vertical axis describe ODA functionality.
- Any complex ODA system can be **decomposed** to fit the cells of the framework.

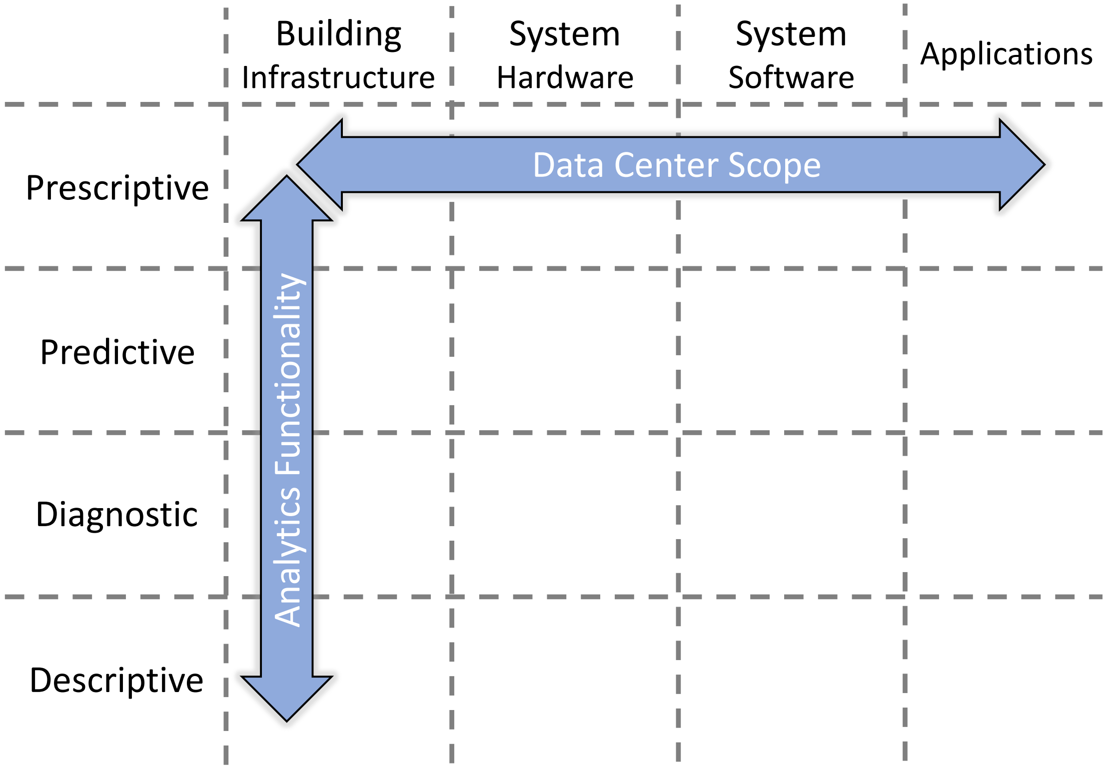

Applying the ODA Framework

- We conducted a **survey** of ODA research literature:
  - 70+ works analyzed and categorized.
  - ODA examples extracted for each category.
  - Provides an overall picture of the ODA field.
- We **applied** the framework to three state-of-the-art use cases:
  - Focus on complex ODA systems.
  - Discussion of the framework’s limitations.

Classifying ODA Research Literature

	Building Infrastructure	System Hardware	System Software	Applications
Prescriptive	• Switching between types of cooling • Tuning cooling knobs • Responding to anomalies	• Cooling optimization at the system level • CPU frequency tuning • Tuning hardware knobs	• Intelligent task placement • Plan-based scheduling • Power and KPI-aware scheduling	• Auto-tuning of HPC applications • Code improvement and recommendations
Predictive	• Predict data center KPIs • Predict cooling demand • Models for cooling performance	• Forecast sensors • Component failure prediction • Predict instruction mix	• Simulating HPC systems and schedulers • Predicting HPC workloads	• Predicting job durations • Predicting resource usage • Predicting performance profiles of code regions
Diagnostic	• Fingerprinting data center-level crises • Infrastructure anomaly detection • Stress testing	• Node anomaly detection • Root cause analysis at the system level • Diagnose network contention issues	• Detect data locality issues • Detect software anomalies • Diagnose OS noise	• Application fingerprinting • Identify application performance patterns • Diagnose code-level issues
Descriptive	• PUE calculation • Facility data processing • Facility-level dashboards	• ITUE calculation • System performance indicators • System-level dashboards	• Slowdown calculation • Scheduler-level dashboards	• Job performance models • Job data processing • Job-level dashboards

Analytics across Multiple Types

- Infrastructure anomaly detection (*diagnostic*) and cooling set-point tuning (*prescriptive*) at **ENI** [3].
- **Better** *prescriptive* decisions can be made with the help of *predictive* and *diagnostic* components.
- Higher technical **complexity**.
- Requires **fusion** of heterogeneous disciplines.

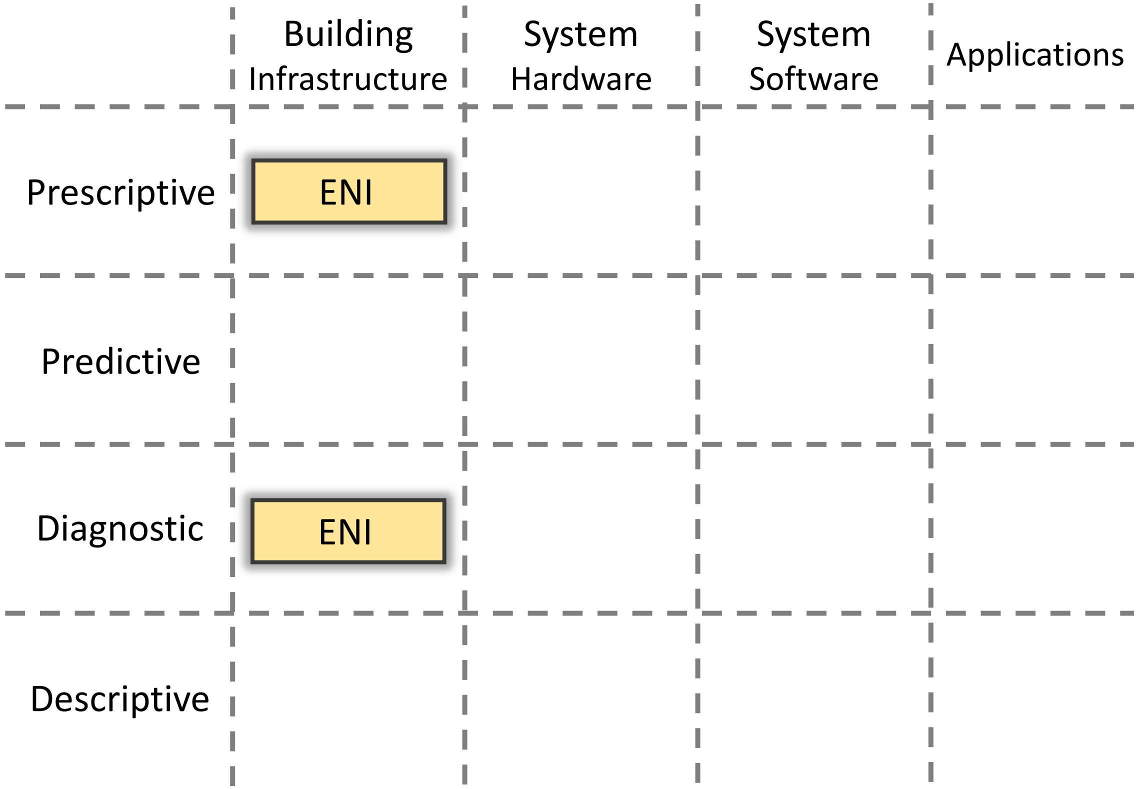

[3] L. Bortot et al. "Data centers are a software development challenge".

Implementing Multi-pillar ODA

- The **Powerstack** framework for power management (*prescriptive*) using data science (*predictive*) [4].
- Most ODA systems are **closed** and cover a single pillar (or *silo*).
- Multi-pillar designs must be **holistic** and integrate many levels of scope.
- Major operational **opportunities**.

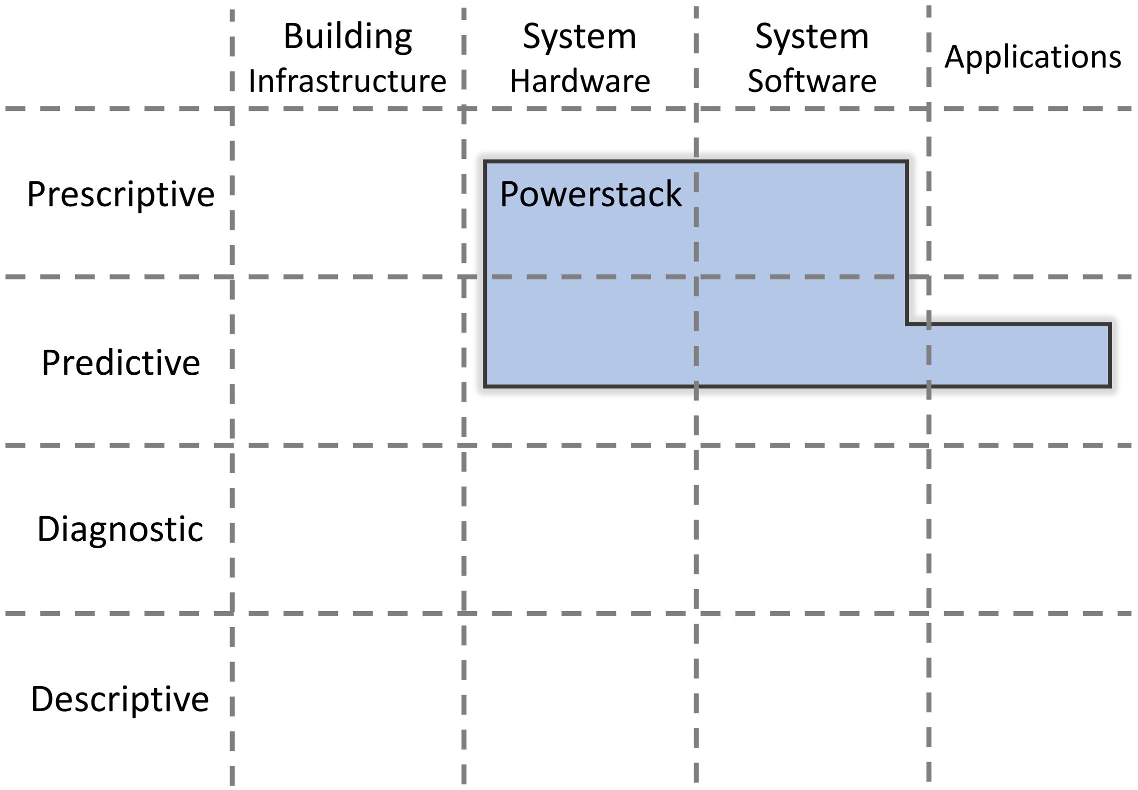

[4] X. Wu et al. "Toward an end-to- end auto-tuning framework in HPC powerstack".

ODA beyond Building Infrastructure

- Forecasting (*predictive*) and notifying (*prescriptive*) excessive power swings at **LLNL** [5].
- The electrical grid as an **extension** of the data center's infrastructure.
- Monitoring and control capabilities are **limited**.
- Practical implementation can be **challenging**.

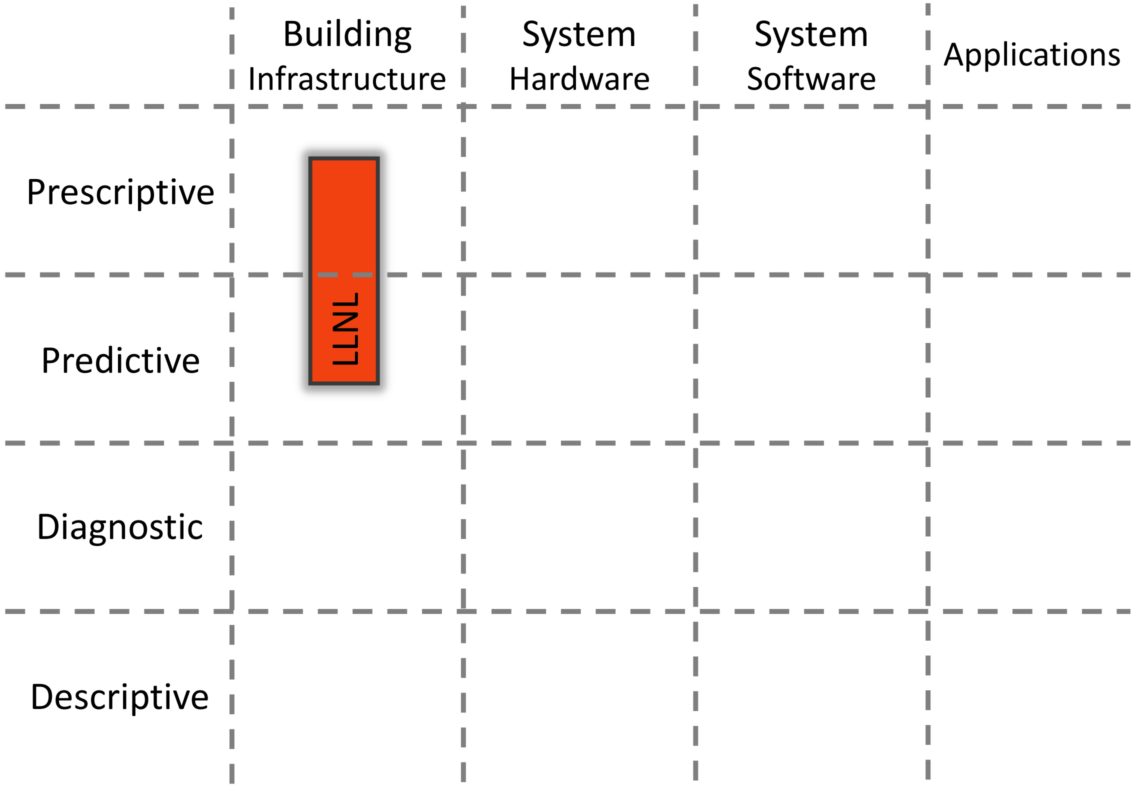

[5] G. Abdulla et al. "Forecasting extreme site power fluctuations using fast fourier transformation".

Conclusions

- Use of *Operational Data Analytics* (ODA) is becoming more and more **common** in HPC data centers.
- We propose a conceptual **framework** to classify ODA use cases according to their scope (*pillars*) and functionality (*types*).
- We aim to establish a common language to **simplify** discussion, analysis and adoption of ODA by the community at large.
- Thank you for your attention!

Acknowledgements: Alessio Netti was supported by the European Union's Horizon 2020/EuroHPC research and innovation programme under grant agreement No. 956560 (REGALE). This work was supported by, and used the resources of, the Oak Ridge Leadership Computing Facility, located in the National Center for Computational Sciences at ORNL, which is managed by UT Battelle, LLC for the U.S. DOE (under the contract No. DE-AC05-00OR22725).