SAND2019-7784D

Exceptional service in the national interest

A Machine Learning Approach to Understanding HPC Application Performance Variation

Burak Aksar, Ben Schwaller baksar|bschwal@sandia.gov

Background and Motivation

High-performance Computing (HPC) is essential for many engineering and scientific applications. HPC systems:

- Support multiple parallel operations
- Consist of thousands of compute nodes

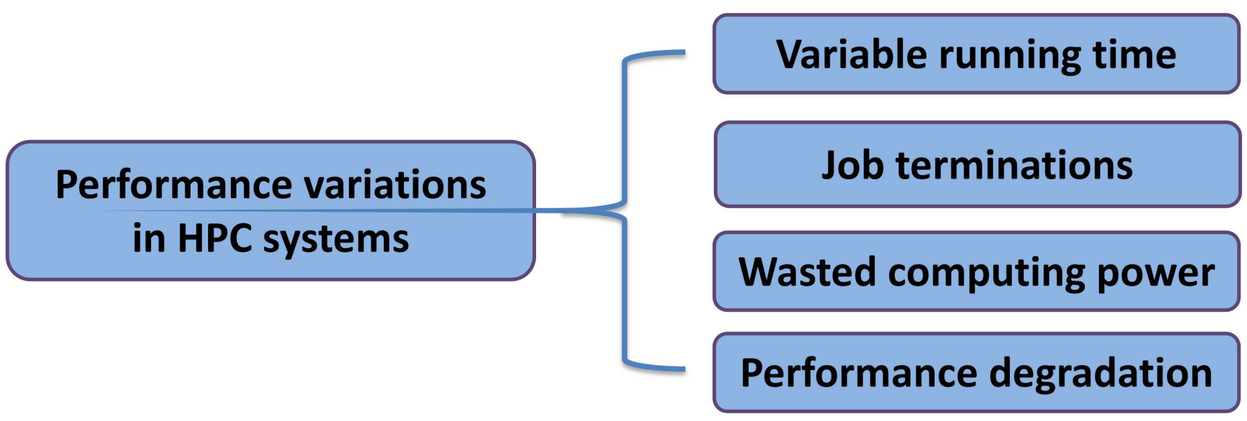

Figure 1: Performance variations can lead to several problems

Challenges

- "BIG" monitoring data makes near-real-time data analysis difficult
- A labeled dataset of HPC system anomalies does not exist in literature
- No established way to determine the ground truth health of a system

HPAS: an HPC Performance Anomaly Suite [1]

- Reproduces common root causes of performance variations in HPC systems
- Open-source: https://github.com/peaclab/HPAS/

Type	Name	Anomaly Behavior
Cache contention	*dial*	Arithmetic operations
CPU intensive process	*dcopy*	Cache read & write
I/O bandwidth contention	*iobandwidth*	File read & write
Network contention	*netoccupy*	Send messages between two nodes
Memory intensive process	*memeater*	Increasing memory allocation

Table 1: Some example synthetic anomalies

Anomaly Detection and Diagnosis [2]

Main Goals

- Investigate intra-node/application/system level performance anomalies
- Detect & diagnose performance anomalies in compute nodes
- Identify the types of the anomalies during the application run-time by using resource usage and performance metrics

Current and Future Directions

- Analyze performance variations in:
  - Real applications (LAMMPS, HACC, QMCPACK)
  - Real production machines (Cori, Astra, Blue Waters)
- Create an application-transparent and scalable anomaly detection framework

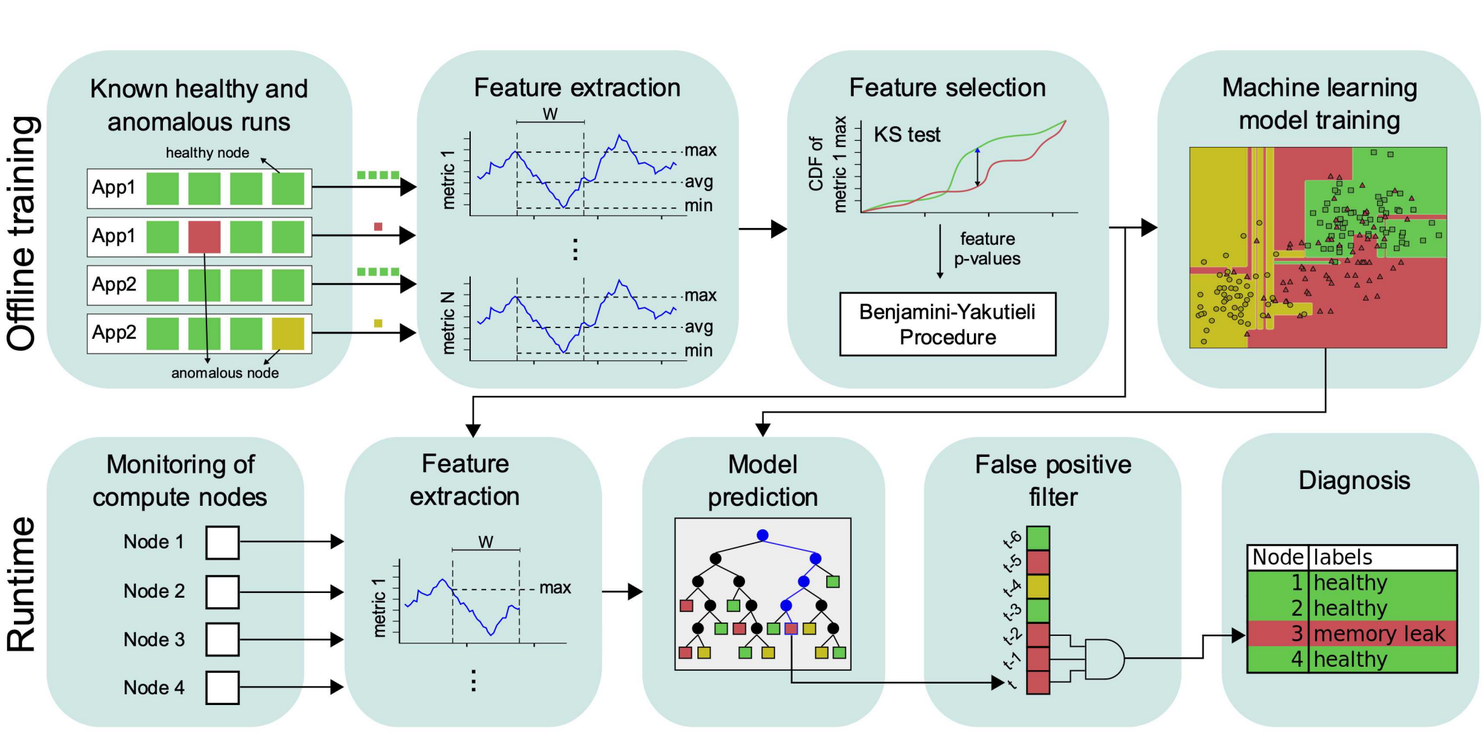

Figure 2: Overall system architecture

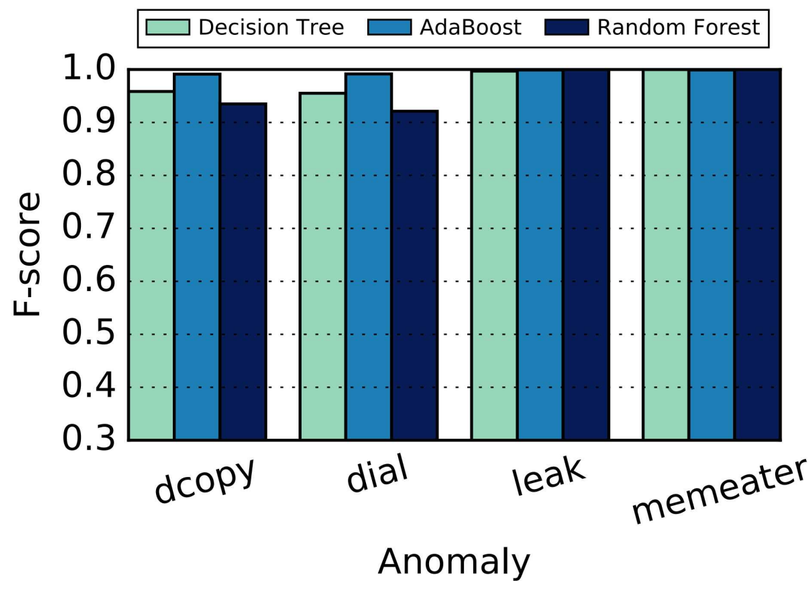

Figure 3: Anomaly detection F-score of various classifiers

References

[1] Ates, E., Zhang, Y., Aksar, B., Brandt, J., Leung, V. J., Egele, M., & Coskun, A. K. (2019). *HPAS: An HPC Performance Anomaly Suite for Reproducing Performance Variations*. 10. To appear in International Conference on Parallel Processing (ICPP), 2019.

[2] Tuncer, Ozan & Ates, Emre & Zhang, Yijia & Turk, Ata & Brandt, Jim & Leung, Vitus & Egele, Manuel & Coskun, Ayse. (2018). Online Diagnosis of Performance Variation in HPC Systems Using Machine Learning. IEEE Transactions on Parallel and Distributed Systems.

U.S. DEPARTMENT OF ENERGY

National Nuclear Security Administration

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and 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-NA-0003525. SAND #