Operational Data Analytics in HPC: Towards Anomaly Prediction

Martin Molan, Andrea Bartolini
DEI Department, University of Bologna, Italy

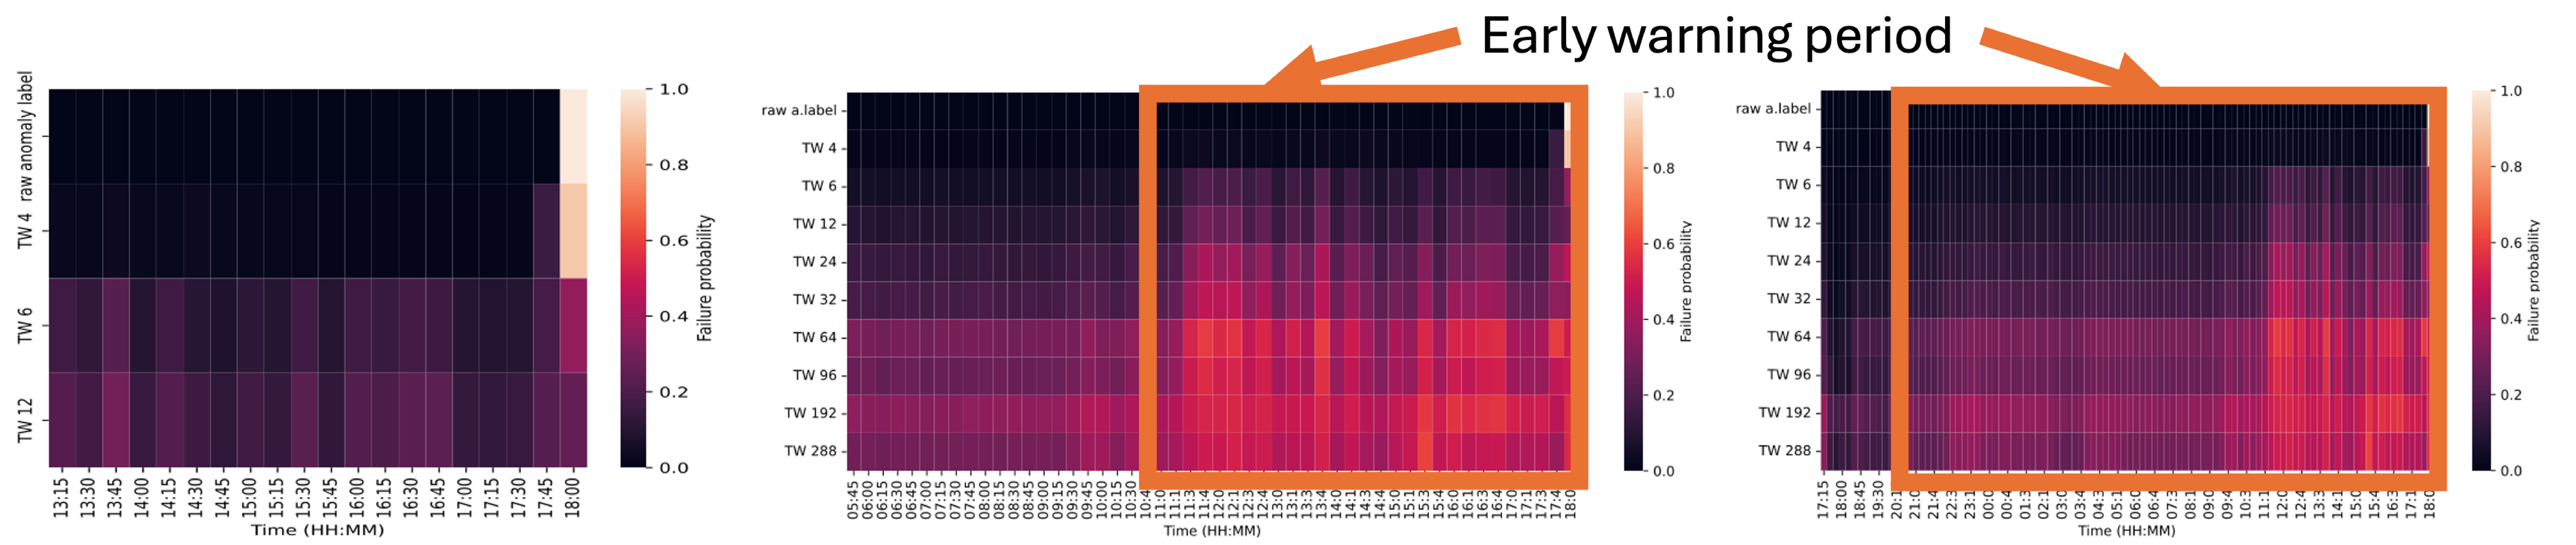

(a) Period of 20 timestamps or 5 hours. (b) Period of 50 timestamps or 12.5 hours. (c) Period of 100 timestamps or 25 hours.

Figure: The same rising-edge event, observed over multiple magnifications. The last timestamp is the occurrence of an anomaly; the rest are the preceding period. Prob fw is the probability for class 1 as generated by the GNN trained with the future window with the length of fw (the duration of the future window is *fw* $*$ 15 minutes).

Abstract

The transition to exascale in high-performance computing (HPC) systems motivates the introduction of machine learning methodologies for system administrators. This work examines self-supervised approaches for anomaly detection and the transition to supervised anomaly prediction with graph neural networks, presenting a model where compute nodes are represented as vertices in a line graph. Preliminary results show this model outperforms current models, predicting anomalies up to 24 hours in advance.

Introduction

The complexity of modern HPC systems requires machine learning-powered operational data analytics to support systems management, focusing on anomaly detection and prediction. This work focuses on node-level anomalies using telemetry data, addressing the challenge with a novel graph neural network approach for prediction.

Data Collection and Monitoring System

The Examon Holistic monitoring system collects real-time data on HPC systems' status and availability, enabling this study. The proposed graph-based models, depicting compute nodes' status, aim to replace older anomaly detection methods, integrating seamlessly into a digital twin dashboard for comprehensive monitoring.

Figure: A digital twin of the Marconi 100 supercomputer and the final goal for deploying a graph-based predictive model.

Data collection campaign

The output of the monitoring system is collected in a data collection campaign on the CINECA's Marconi100 Tier-0 supercomputer (Borghesi et al. 2023). It is a culmination of a multiple-year data collection and monitoring campaign by the University of Bologna. Anomaly detection model RUAD (Molan, Borghesi, Cesarini, et al. 2023) is demonstrated as a use case for the utility of the data, while the proposed anomaly detection approach uses its dataset for validation.

Methodology

The proposed anomaly detection approach encodes compute nodes' physical proximity in a line graph, with nodes connected to their immediate neighbors as presented in Figure 3. Node similarity has also been explored from the perspective of node segmentation (as proposed by Molan, Borghesi, Benini, et al. 2022) while the motivation for the use of the graph structure comes form the preliminary work of Molan, Khan, et al. 2023, showcasing the ability of graphs to capture the physical information about the compute nodes. This graph, representing nodes as vertices with monitoring data attributes, is fed into a graph convolutional neural network (GNN) for anomaly prediction, aiming to predict future anomalies within a defined window.

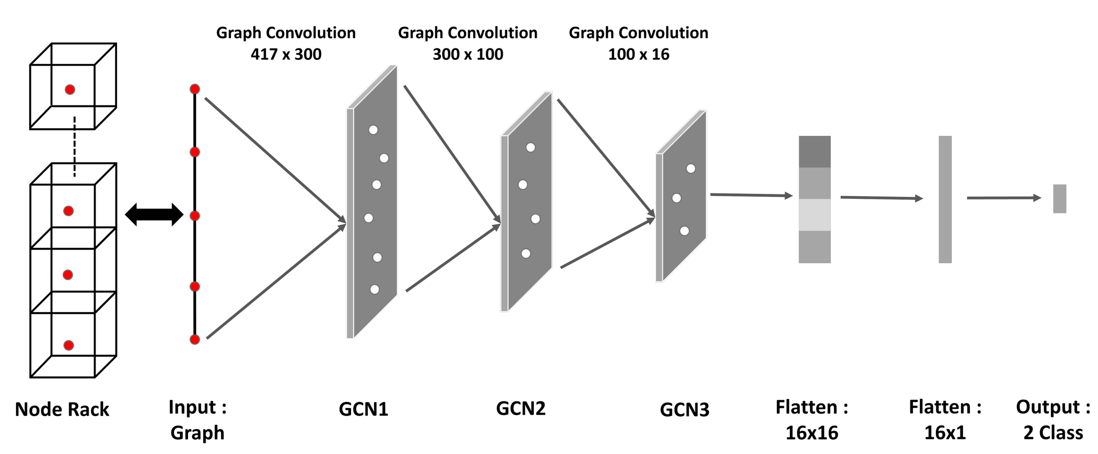

Figure: The structure of the GCN network exploits the organization of compute nodes in a rack.

Results

Our graph neural network model demonstrates superior performance on two compute racks of the Marconi 100 supercomputer, outperforming previous methods in anomaly prediction up to 24 hours in advance, with AUC scores significantly higher than the state-of-the-art anomaly detection model RUAD proposed by Molan, Borghesi, Cesarini, et al. 2023. The anomaly anticipation approach, as depicted in Figure 1, provides an early warning period for the upcoming anomaly.

(T+N)	Time-ahead	AUC of GNN
T+6	1.5hr ahead	0.8826
T+12	3hr ahead	0.8676
T+24	6hr ahead	0.8661
T+32	8hr ahead	0.8454
T+64	16hr ahead	0.7801
T+96	24hr ahead	0.7317
T+192	48hr ahead	0.6455
T+288	72hr ahead	0.6219

Table: The proposed GNN-based approach outperforms the previous anomaly detection approach, which achieved 0.86 AUC, up to 24 hours in advance.

Conclusions

Introducing graph structures into machine learning models for HPC operational data analytics enables the training of supervised models for anomaly prediction. This advancement opens possibilities for further modeling tasks, including thermal and job performance modeling, leveraging the physical proximity encoded in graph structures.

References

- Borghesi, Andrea et al. (May 2023). "M100 ExaData: a data collection campaign on the CINECA's Marconi100 Tier-0 supercomputer". In: *Scientific Data* 10.1, p. 288. ISSN: 2052-4463. DOI: 10.1038/s41597-023-02174-3. URL: https://doi.org/10.1038/s41597-023-02174-3.
- Molan, Martin, Andrea Borghesi, Luca Benini, et al. (2022). "Analysing Supercomputer Nodes Behaviour with the Latent Representation of Deep Learning Models". In: *Euro-Par 2022: Parallel Processing*. Ed. by José Cano and Phil Trinder. Cham: Springer International Publishing, pp. 171–185. ISBN: 978-3-031-12597-3.
- Molan, Martin, Andrea Borghesi, Daniele Cesarini, et al. (2023). "RUAD: Unsupervised anomaly detection in HPC systems". In: *Future Generation Computer Systems* 141, pp. 542–554. ISSN: 0167-739X. DOI: https://doi.org/10.1016/j.future.2022.12.001.
- Molan, Martin, Junaid Ahmed Khan, et al. (2023). "The Graph-Massivizer Approach Toward a European Sustainable Data Center Digital Twin". In: *2023 IEEE 47th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, pp. 1459–1464.

Acknowledgments

This research was partly supported by the EuroHPC EU projects Regale and EUPEX, the HE EU Graph-Massivizer project, and the PNRR's SPOKE 1: Future HPC & Big Data. Thanks to CINECA and collaborators Andrea Borghesi and Junaid Ahmed Khan for their support.

GRAPH MASSIVIZER

ALMA MATER STUDIORUM A.D. 1088