

# Learning anomalies from graph: predicting compute node failures on HPC clusters

Anonymous Full Paper  
Submission 51

## Abstract

Today, high-performance computing (HPC) systems play a crucial role in advancing artificial intelligence. Nevertheless, the estimated global data center electricity consumption in 2022 was around 1% of the final global electricity demand. Therefore, as HPC systems advance towards Exascale computing, research is required to ensure their growth is sustainable and environmentally friendly. Data from infrastructure monitoring can be leveraged to predict downtimes, ensure these are treated in time, and increase the overall system's utilization. In this paper, we compare four machine-learning approaches, three of them based on graph embeddings, to predict compute node downtimes. The experiments were performed with data from Marconi 100, a tier-0 production supercomputer at CINECA in Bologna, Italy. Our results show that the machine learning models can accurately predict downtime, surpassing current state-of-the-art models.

## 1 Introduction

The increasing research and deployment of artificial intelligence require massive hardware, which has strong implications for HPC and data center energy sustainability and makes efficient utilization of the resources even more critical [1]. Maintaining consistent availability of HPC resources is crucial to avoid negative impact on research and minimize the carbon footprint as well [2]. Correctly forecasting potential compute node downtimes is critical to this end. Predictive maintenance has been shown to enable the efficient planning of maintenance tasks to minimize operation downtimes and preserve the health of the entire system [3].

Machine learning has shown great promise in providing accurate downtime predictions in HPC and data center environments to realize predictive maintenance. Among the works describing this approach, we find Pelaez et al. [4], who described how clustering was applied to perform online failure prediction. Klinkenberg et al. [5] followed a different approach, leveraging a supervised machine learning model trained on monitoring data to predict lock events. More recently, Borghesi et al. [6] developed deep learning models to predict compute node downtimes in HPC systems.

The fact that network architectures are governed by the same organizing principles regardless of the science domain and can provide a unified representation of heterogeneous data is a compelling reason driving research at the intersection of network science and machine learning [7].

The approaches described above leverage sensor readings or logs to predict compute node downtimes. Nevertheless, such approaches miss much contextual information by not being able to include additional points of view, such as the kind of information being monitored or the sensor placement. Such information can be included through a graph representation. E.g., Molan et al. [8, 9] joined sensor readings for each point in time considering the sensors' topological location within a particular rack. Nevertheless, such a representation provides no information on the types of sensors being used. Furthermore, the graph representations provide no context regarding past sensor readings - something that could be relevant to understanding whether we are heading to a particular state. To mitigate these issues, we propose a different graph representation with nodes representing specific sensor types and associated sensors and subgraphs resulting from translating time series into networks following specific heuristics. Such graphs are then encoded into embeddings and used for downstream compute node downtime prediction.

In this research, we develop local and global machine-learning models that leverage graph embeddings summarizing domain knowledge regarding sensor types and readings to minimize system unavailability. We trained and evaluated the machine learning models on a subset of the publicly available data from the Marconi 100 supercomputer compiled by Borghesi et al. [10]. The models' performance was evaluated considering the AUC ROC scores. Our results show that the models trained with graph embeddings and sensor reading information performed best. Furthermore, global models exhibited a stronger performance than local ones. While sensor reading data and graph embeddings resulted in a stronger mean performance than using sensor reading data only, the difference remained consistently small across the forecasting horizons.

The rest of this paper is structured as follows: Section 2 presents related work, Section 3 introduces the dataset we worked on. Section 4 describes the methodology we followed and Section 5 the exper-

096 iments performed. Finally, Section 6 presents and  
097 discusses the results obtained and Section 7, we  
098 provides our conclusions.

## 099 2 Related work

100 The size, complexity, and heterogeneous architecture  
101 of contemporary HPC systems require introducing  
102 machine learning methodologies that support the  
103 work of the system administrators [11]. As char-  
104 acterized by Netti et al. [11], anomaly detection  
105 and prediction are among the applications with the  
106 most direct impact on the overall availability, and  
107 by extension, sustainability [12], of the HPC sys-  
108 tem. Consequently, much effort has been invested in  
109 building data-driven models for anomaly detection  
110 and, later, anomaly prediction and anticipation.

111 Based on the telemetry data, various node fail-  
112 ure detection approaches, such as RUAD [13] or  
113 PROCTOR [14], have been proposed. These ap-  
114 proaches, based on self-supervised learning, are only  
115 capable of recognizing the anomalies and failures  
116 that are taking place. They cannot provide fail-  
117 ure predictions, allowing the system administrators  
118 to manage the system proactively, including intel-  
119 ligent scheduling considering anticipated hardware  
120 failures and preventative maintenance. While there  
121 has been an effort to create an anomaly anticipation  
122 system for HPC compute node failures, such as the  
123 one proposed by Borghesi et al. [15], such systems  
124 can anticipate the failure but provide no estimation  
125 about the time frame for it. Failure prediction ap-  
126 proaches exist for components within the compute  
127 node, such as the work of Devesh Tiwari et al. [16]  
128 or Yu Liu et al. [17] focus on component failure  
129 prediction (disk failure specifically). Besides not  
130 being holistic and only covering a part of potential  
131 compute node failures, these approaches also have  
132 limited prediction windows [18].

133 The prediction model must include additional in-  
134 formation beyond the telemetry data to go beyond  
135 component failure prediction or node-level failure  
136 detection. The current state-of-the-art approach,  
137 GRAAFE [12], is based on the observation that the  
138 additional information about the physical layout of  
139 the compute nodes within a computing room aids  
140 in the ability to train the compute node anomaly  
141 prediction model. This information is encoded as a  
142 graph: each compute node is represented as a vertex  
143 connected to its nearest neighbors, and node teleme-  
144 try data is represented for each compute node as a  
145 vertex attribute. While different graph topologies  
146 have been tested, the optimal graph topology uncov-  
147 ered is a line graph representing a single compute  
148 rack in a compute room. Different approaches to  
149 anomaly detection with graphs have been applied to  
150 other domains and could be considered in HPC envi-  
151 ronments. Among them we find anomaly detection

with graphs approaches that aim to alleviate struc- 152  
tural distribution shifts and novel techniques for 153  
posing the problem as a temporal graph clustering 154  
problem [19, 20]. 155


## 3 Dataset 156

157 The source of our data is a collection of the Marconi  
158 100 supercomputer sensory data, which was gathered  
159 by Borghesi et al. and was made publicly accessible  
160 at <https://zenodo.org/records/7541722> [10]. More  
161 specifically, we focused on a subset of the original  
162 data stored in the 1.tar distribution file where in-  
163 formation about sixteen computing nodes in one  
164 rack of the system are available. The data is a two-  
165 dimensional table, where rows represent timestamps  
166 taken 15 minutes apart, ranging between March 9<sup>th</sup>  
167 2020 and September 28<sup>th</sup> 2022. This means roughly  
168 86 thousand rows for each compute node data. Ap-  
169 proximately 9 thousand rows have missing values.  
170 Columns, on the other hand, contain different sen-  
171 sor measurements of the system aggregated over the  
172 15 minute time interval, which include the power  
173 consumption of the fans and CPUs, the tempera-  
174 ture of the GPUs cores and memory, the voltage of  
175 the power supply and many more. To be more pre-  
176 cise, each sensor measurement is given in 4 columns,  
177 which store the minimum and maximum values, the  
178 average values, and the standard deviation of the  
179 measurement. Two additional columns are provided,  
180 one that stores the time information and the other  
181 one that annotates anomalies in the data. The last  
182 column has integer values, where zero means there  
183 was no anomaly, and a value greater than zero sig-  
184 nals there was an anomaly, which means the system  
185 was unavailable then. The timestamp column has  
186 date and time as values, and the remaining sensor  
187 columns have numerical values.

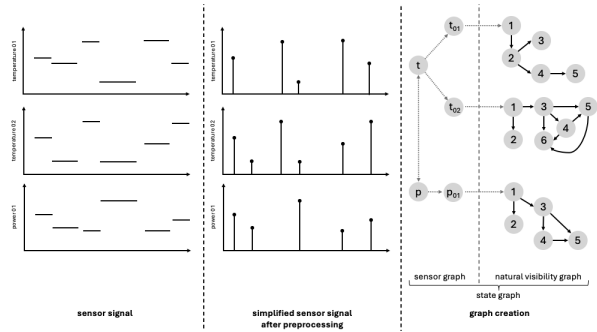
188 Accounting for all the sensors, there are 354  
189 columns, though we chose only a small subset of  
190 the sensors to focus on. Our point of interest is  
191 the columns with the average values of the mea-  
192 surements. Furthermore, we selected some columns  
193 from which to get our data. These columns include  
194 ambient\_avg, dimm0\_temp\_avg, fan\_disk\_power\_avg,  
195 gpu0\_core\_temp\_avg, gpu0\_mem\_temp\_avg,  
196 p0\_io\_power\_avg, p0\_mem\_power\_avg, p0\_power\_avg,  
197 ps0\_input\_power\_avg, ps0\_output\_curre\_avg,  
198 ps0\_output\_volta\_avg, fan0\_0\_avg and  
199 p0\_vdd\_temp\_avg.

## 4 Methodology 200

**Data preprocessing** To develop the features 201  
needed for training our model, we followed a se- 202  
ries of steps to preprocess the data and generate 203  
new representations of the information. 204



**Figure 1.** The Figure presents two signals: (i) an example of sensory data (black), and (ii) the post-processed signal once change detection and value replacement were performed (red).



**Figure 2.** The Figure presents three signals and processing stages. The sensor signal is preprocessed considering change level detection and clustering into bins. The resulting values are then mapped into natural visibility graphs. The natural visibility graphs can be further enriched with information regarding the sensor they belong to and its type.

205 First, we addressed any missing values in the  
 206 dataset by applying a forward-fill strategy, where  
 207 the last observed value was carried forward to re-  
 208 place any missing entries within each column. The  
 209 forward-fill strategy was chosen under the assump-  
 210 tion that if a sensor did not emit a value or if a  
 211 reading was lost, the most likely value to exist if  
 212 conditions did not change in between, it would be  
 213 the latest sensor reading.

214 Next, we utilized a change detection algorithm  
 215 based on a relative moving average and a predefined  
 216 threshold to identify significant changes in the data  
 217 for each column. Once the indices indicating changes  
 218 were identified, we replaced the values between two  
 219 consecutive change points with the mean of that  
 220 data segment. This process produced a simplified  
 221 dataset with distinct segments of different values for  
 222 each column. We exemplify such signal processing  
 223 in Fig. 1.

224 To handle segments where values varied by a non-  
 225 significant amount, we divided the range of values  
 226 in each column into five quantiles and then mapped  
 227 the values according to their respective quantiles.  
 228 This quantization made the data discrete, resulting  
 229 in each column containing only five possible val-  
 230 ues. The choice to perform discretization into five  
 231 possible values was empirical weighting the number  
 232 of states resulting from a particular number of se-  
 233 lected bins. Consequently, there is a finite number  
 234 of unique combinations of column values, which we  
 235 interpreted as distinct states in the system.

236 **Feature engineering** Four different kind of fea-  
 237 tures were created. First, given the compute node  
 238 states identified above, we performed a one-hot en-  
 239 coding representation of this data. Nevertheless,  
 240 given a one-hot encoded representation provides  
 241 only insight into that specific state but no context  
 242 on the preceding ones, we considered a more sophis-

243 ticated approach. Given a time series of length  $n$ ,  
 244 natural visibility heuristics can be applied applied  
 245 to create a graph  $G = (V, E)$ , where nodes  $V$  represent  
 246 time series data points and edges  $E$  represent visi-  
 247 bility relationships between those nodes. Then a deep  
 248 learning model can be used to transform the graph  $G$   
 249 into a vector representation  $\mathbf{H} = f_{\text{embed}}(G)$ . Choos-  
 250 ing  $k=5$  would encode information from each sensor  
 251 with the same number of columns as the one-hot  
 252 encoded approach while eventually conveying more  
 253 information. To this end, we constructed natural  
 254 visibility graphs for each sensor in a given state, con-  
 255 sidering the values of up to ten preceding states. In  
 256 Fig. 2 we show how a sample drawing on how sensor  
 257 signal was reduced into a sequence of unique values  
 258 and later to a visibility graph. The Graph2Vec ap-  
 259 proach was used to transform the graphs into vector  
 260 embedding representations.

261 Finally, we also created a slightly different graph  
 262 representation encoding information about the sen-  
 263 sors and joined it to the abovementioned visibility  
 264 graphs to the nodes representing the sensors from  
 265 which the sensor readings originate (see Fig. 2).  
 266 This way, the random walks would convey informa-  
 267 tion about current and past states and how these  
 268 relate to particular sensors and sensed aspects (e.g.,  
 269 temperature). These graphs were then turned into  
 270 an embedding using the Graph2Vec approach.

271 To create natural visibility graph embeddings, we  
 272 employed a Graph2Vec model trained over all of  
 273 the visibility graphs, regardless of the sensor. The  
 274 embedding creation process began with generating  
 275 tagged random walks over a graph, which served  
 276 as documents for training the Graph2Vec model.  
 277 Once the model was trained, new random walks  
 278 over the graphs were conducted, and these were  
 279 used to perform inference and obtain a vector from  
 280 the previously trained model. The final embedding  
 281 for each graph was obtained by averaging these vec-

tors, providing a robust representation of the graph’s structure. Since one state produced several natural visibility graphs (for each sensor, there was one natural visibility graph), the resulting embeddings were concatenated into one vector embedding for that state, following the same positional convention we used for the one-hot encoding schema. In addition, we also created the state graph embeddings following the same procedure as with the natural visibility graph embeddings, with the difference being that instead of using the natural visibility graphs, we used the state graphs. Because the state graphs were larger we decided to perform a random walk of length 20 per node.

**Model training** We trained two kinds of CatBoost classifiers: (i) models trained on data from each compute node separately (local models) and (ii) models trained on data from all of the nodes (global model). In the case of local models, there were, on average, 8700 available unique state instances, which we used to train the model, considering a 75/5/20 split for train, validation, and testing. The splits were made such that the information about the time of each instance was preserved and that the training data was recorded before the validation data, and the validation data was recorded before the testing data. The global model’s process was similar to that of the local models. First, we took the 75% training data from the local models and joined them into one training data for the global model. The same was done with the validation set. When we tested the global model, we did so for each compute node separately to compare the results between the local models and the global model node by node.

The CatBoost classifier was trained running for 700 iterations with a cross-entropy loss, a learning rate of 0.03 and an L2 regularization factor of 0.3. The local and global models were trained for three forecast horizons determined by state changes. Considering the states changed every 165 minutes on average, the compute node downtimes could be predicted to up to 495 minutes ahead.

**Model evaluation** We measured the classifiers’ discriminative performance with the AUC ROC metric. The score was computed at a compute node level and then summarized to report the mean, minimum (min), and maximum (max) values obtained across nodes for each experiment.

## 5 Experiments

We performed four experiments to understand how different features and graph representations affected the outcomes when predicting compute node downtimes. Our focus on graph representations is rooted in the fact that while much sensor data could be encoded by providing features that capture  $t_{n-1}$  values, graph representation allows the conveying of

information regarding the sensor time series and associating it to specific domain knowledge, such as sensor ID and sensor types. Three forecasting horizons were considered, predicting one to three states ahead. We detail them below:

**Experiment A** The purpose of the first experiment was to establish a baseline, considering the one hot encoded representation of the compute node states. The feature vector size was 65.

**Experiment B** The purpose of the second experiment was to determine whether a natural visibility graph representation of past sensor values provides additional information that could enhance the models performance. The natural visibility graph representation was used to encode sensor data and a Graph2Vec representation to turn those graphs into embeddings. A single Graph2Vec model was used to learn visibility graphs from all the sensors. The decision was made under the assumption that the embedding model would learn best from a higher number and variety of samples than would have learned if exposed only to visibility graphs of a single sensor. To make the results comparable against *Experiment A*, we created embeddings for each sensor considering a vector size 5. Therefore, the final feature vector size matched the original one with a size of 65.

**Experiment C** The purpose of the third experiment was to determine whether the information encoded in experiments A and B was complementary and, therefore, joining it would lead to better results. The intuition behind the experiment was that (i) the one-hot encoded representation only had information about the current state but missed the sensor values observed in the near past, and (ii) the visibility graphs only encode the topology of a time series but do not encode information regarding the actual time series values. The feature vector was created by concatenating the hot encoded features with the natural visibility graph embeddings. The feature vectors had a total length of 120 values.

**Experiment D** The purpose of the fourth experiment was to determine whether the representation from Experiment C could be enriched with an additional embedding representation that not only considered the visibility graphs but also encode information about the aspects that they measured and the particular sensor that provided the sensor values. To that end, we created an additional graph representing the sensors, their sensing domain, and the associated natural visibility graphs. This experiment used the one-hot encoded data, the natural visibility graph embeddings, and the state graph

391 embeddings. The state graph embeddings were gen-  
392 erated with the Graph2Vec model and given a size  
393 of 15, resulting in an overall feature vector size of  
394 135.

## 395 6 Results and discussion

396 Table 1 shows the results obtained across the com-  
397 pute nodes and breaks them down by experiment  
398 and whether a global or local model was trained.  
399 The results show that local models leveraging the  
400 one-hot encoded state features obtained the best  
401 mean results. Using natural visibility graph embed-  
402 dings alone resulted in poor performance. Neverthe-  
403 less, the Graph2Vec embeddings for natural visibility  
404 graphs proved useful in enhancing the worst-case  
405 performance across the local models and n+3 global  
406 model. The addition of the state graph resulted  
407 in a marginal improvement in some cases for local  
408 and global models (e.g., improvement for worst case  
409 in global models at n+3). Nevertheless, further ex-  
410 perimentation is required to draw conclusions in  
411 which cases do such embeddings help the model  
412 learn better and when they negatively affect the  
413 overall performance.

414 We have no explanation to why the graph embed-  
415 dings alone displayed poor performance and further  
416 research is required to reach a better understand-  
417 ing. Among possible reasons we consider the fact  
418 that the natural visibility graphs capture only the  
419 topological aspects of the time series and any in-  
420 formation related to their particular values is lost.  
421 This could be amended by adding complementary  
422 graph representations such as quantile graphs or  
423 ordinal partition graphs. Another reason could be  
424 the quality of the Graph2Vec embeddings which  
425 may depend on how the graph was sampled. A more  
426 thorough evaluation of hyperparameters is required  
427 to understand how the number of paths and their  
428 length affects the outcomes.

## 429 7 Conclusion

430 Enabling predictive maintenance at HPC and data-  
431 centers is critical to minimize downtimes and max-  
432 imize their utilization. Machine learning models  
433 have shown great promise accurately forecasting  
434 such events ahead of time. In this research we pro-  
435 pose a graph-based approach to predicting compute  
436 nodes downtime and evaluate it on public data from  
437 Marconi 100 - a tier-0 production supercomputer  
438 from CINECA. To that end, we first process the  
439 sensor reading data by treating missing values and  
440 discretizing them. We used natural visibility graphs  
441 to represent time series and include them into graph  
442 representations describing compute nodes and their  
443 sensors. Such graphs were converted into embed-

**Table 1.** Results for the models in different experiments. The 'time' column indicates how many steps ahead from the current state the model tries to predict:  $n + 1$  means the model tries to predict the target variable of the next state,  $n + 2$  means the model tries to predict the target variable of two states ahead. The 'mean', 'min', and 'max' columns provide the aggregated results computed over the results obtained from individual computational nodes. The results are divided into 'local' and 'global' according to the data used to train the model.

Exp.	Time	Local			Global		
		Mean	Min	Max	Mean	Min	Max
A	$n + 1$	0.8059	0.6786	0.9162	0.8131	0.7325	0.8984
	$n + 2$	0.7997	0.7023	0.9271	0.7812	0.7215	0.8888
	$n + 3$	0.7653	0.6561	0.8531	0.7510	0.6861	0.8623
B	$n + 1$	0.5567	0.5132	0.6371	0.5532	0.5174	0.6150
	$n + 2$	0.5603	0.5210	0.5987	0.5657	0.5183	0.5912
	$n + 3$	0.5719	0.5255	0.6399	0.5775	0.5393	0.6208
C	$n + 1$	0.7951	0.6938	0.8916	0.8043	0.7207	0.8896
	$n + 2$	0.7864	0.7341	0.9233	0.7882	0.7064	0.8911
	$n + 3$	0.7480	0.6618	0.8384	0.7564	0.6898	0.8492
D	$n + 1$	0.7974	0.6805	0.8945	0.8150	0.7197	0.8974
	$n + 2$	0.7670	0.6075	0.9112	0.7973	0.7077	0.8978
	$n + 3$	0.7398	0.5894	0.8239	0.7540	0.6891	0.8488

444 dings through Graph2Vec models and used down-  
445 stream to train a classifier and predict compute node  
446 downtimes. The results suggest that using graph  
447 embeddings could enhance the classifier performance  
448 in some cases. Nevertheless, further research is re-  
449 quired to understand when such representations are  
450 beneficial or detrimental to the overall models' per-  
451 formance.

## 452 References

- 453 [1] D. Zhao, S. Samsi, J. McDonald, B. Li, D.  
454 Bestor, M. Jones, D. Tiwari, and V. Gade-  
455 pally. "Sustainable supercomputing for AI:  
456 GPU power capping at HPC scale". In: *Pro-  
457 ceedings of the 2023 ACM Symposium on  
458 Cloud Computing*. 2023, pp. 588–596.
- 459 [2] B. Everman. "Improving carbon, cost, and  
460 energy efficiency of large scale systems via  
461 workload analysis". In: (2022).
- 462 [3] A. L. d. C. D. Lima, V. M. Aranha, C. J. d. L.  
463 Carvalho, and E. G. S. Nascimento. "Smart  
464 predictive maintenance for high-performance  
465 computing systems: a literature review". In:  
466 *The Journal of Supercomputing* 77.11 (2021),  
467 pp. 13494–13513.
- 468 [4] A. Pelaez, A. Quiroz, J. C. Browne, E. Chuah,  
469 and M. Parashar. "Online failure prediction for  
470 hpc resources using decentralized clustering".  
471 In: *2014 21st International Conference on  
472 High Performance Computing (HiPC)*. IEEE.  
473 2014, pp. 1–9.
- 474 [5] J. Klinkenberg, C. Terboven, S. Lankes, and  
475 M. S. Müller. "Data mining-based analysis of  
476 HPC center operations". In: *2017 IEEE In-*

- 477 *ternational Conference on Cluster Computing*  
478 (*CLUSTER*). IEEE. 2017, pp. 766–773.
- 479 [6] A. Borghesi, M. Molan, M. Milano, and A.  
480 Bartolini. “Anomaly detection and anticipa-  
481 tion in high performance computing systems”.  
482 In: *IEEE Transactions on Parallel and Dis-*  
483 *tributed Systems* 33.4 (2021), pp. 739–750.
- 484 [7] M. Pósfai and A.-L. Barabási. *Network science*.  
485 Citeseer, 2016.
- 486 [8] M. Molan, J. Ahmed Khan, A. Borghesi,  
487 and A. Bartolini. “Graph neural networks for  
488 anomaly anticipation in HPC systems”. In:  
489 *Companion of the 2023 ACM/SPEC Interna-*  
490 *tional Conference on Performance Engineer-*  
491 *ing*. 2023, pp. 239–244.
- 492 [9] M. Molan, M. S. Ardebili, J. A. Khan, F. Ben-  
493 eventi, D. Cesarini, A. Borghesi, and A. Bar-  
494 tolini. “GRAAFE: GRaph anomaly anticipa-  
495 tion framework for exascale HPC systems”. In:  
496 *Future Generation Computer Systems* (2024).
- 497 [10] A. Borghesi, C. D. Santi, M. Molan, M. S.  
498 Ardebili, A. Mauri, M. Guarrasi, D. Galetti,  
499 M. Cestari, F. Barchi, L. Benini, F. Ben-  
500 eventi, and A. Bartolini. *M100 dataset: time-*  
501 *aggregated data for anomaly detection*. Ver-  
502 sion 1.0.0. Zenodo, Jan. 2023. DOI: [10.5281/](https://doi.org/10.5281/zenodo.7541722)  
503 [zenodo.7541722](https://doi.org/10.5281/zenodo.7541722). URL: [https://doi.org/](https://doi.org/10.5281/zenodo.7541722)  
504 [10.5281/zenodo.7541722](https://doi.org/10.5281/zenodo.7541722).
- 505 [11] A. Netti, W. Shin, M. Ott, T. Wilde, and N.  
506 Bates. “A Conceptual Framework for HPC  
507 Operational Data Analytics”. In: *2021 IEEE*  
508 *International Conference on Cluster Comput-*  
509 *ing (CLUSTER)*. 2021, pp. 596–603. DOI: [10.](https://doi.org/10.1109/Cluster48925.2021.00086)  
510 [1109/Cluster48925.2021.00086](https://doi.org/10.1109/Cluster48925.2021.00086).
- 511 [12] M. Molan, M. S. Ardebili, J. A. Khan, F.  
512 Beneventi, D. Cesarini, A. Borghesi, and A.  
513 Bartolini. “GRAAFE: GRaph Anomaly Ant-  
514icipation Framework for Exascale HPC sys-  
515tems”. In: *Future Generation Computer Sys-*  
516 *tems* 160 (2024), pp. 644–653. ISSN: 0167-  
517 739X. DOI: [https://doi.org/10.1016/j.](https://doi.org/10.1016/j.future.2024.06.032)  
518 [future.2024.06.032](https://doi.org/10.1016/j.future.2024.06.032). URL: [https://www.](https://www.sciencedirect.com/science/article/pii/S0167739X24003327)  
519 [sciencedirect.com/science/article/](https://www.sciencedirect.com/science/article/pii/S0167739X24003327)  
520 [pii/S0167739X24003327](https://www.sciencedirect.com/science/article/pii/S0167739X24003327).
- 521 [13] M. Molan, A. Borghesi, D. Cesarini, L. Benini,  
522 and A. Bartolini. “RUAD: Unsupervised  
523 anomaly detection in HPC systems”. In: *Fu-*  
524 *ture Generation Computer Systems* 141 (2023),  
525 pp. 542–554. ISSN: 0167-739X. DOI: [https://](https://doi.org/10.1016/j.future.2022.12.001)  
526 [doi.org/10.1016/j.future.2022.12.001](https://doi.org/10.1016/j.future.2022.12.001).
- 527 [14] B. Aksar, Y. Zhang, and E. e. a. Ates. “Proc-  
528tor: A semi-supervised performance anomaly  
529 diagnosis framework for production hpc sys-  
530 tems”. In: *High Performance Computing: 36th*  
531 *International Conference, ISC High Perfor-*  
532 *mance 2021, Virtual Event, June 24–July 2,*  
533 *2021, Proceedings 36*. Springer. 2021, pp. 195–  
534 214.
- 535 [15] A. Borghesi, M. Molan, M. Milano, and A. Bar-  
536 tolini. “Anomaly Detection and Anticipation  
537 in High Performance Computing Systems”.  
538 In: *IEEE Transactions on Parallel and Dis-*  
539 *tributed Systems* 33.4 (2022), pp. 739–750. DOI:  
540 [10.1109/TPDS.2021.3082802](https://doi.org/10.1109/TPDS.2021.3082802).
- 541 [16] S. Lu, B. Luo, T. Patel, Y. Yao, D. Tiwari,  
542 and W. Shi. “Making Disk Failure Predic-  
543 tions SMARTer!” In: *Proceedings of the 18th*  
544 *USENIX Conference on File and Storage Tech-*  
545 *nologies*. FAST’20. Santa Clara, CA, USA:  
546 USENIX Association, 2020, pp. 151–168. ISBN:  
547 9781939133120.
- 548 [17] Y. Liu, Y. Guan, T. Jiang, K. Zhou, H. Wang,  
549 G. Hu, J. Zhang, W. Fang, Z. Cheng, and P.  
550 Huang. “SPAЕ: Lifelong disk failure prediction  
551 via end-to-end GAN-based anomaly detection  
552 with ensemble update”. In: *Future Genera-*  
553 *tion Computer Systems* 148 (2023), pp. 460–  
554 471. ISSN: 0167-739X. DOI: [https://doi.](https://doi.org/10.1016/j.future.2023.05.020)  
555 [org/10.1016/j.future.2023.05.020](https://doi.org/10.1016/j.future.2023.05.020).  
556 URL: [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S0167739X23002030)  
557 [science/article/pii/S0167739X23002030](https://www.sciencedirect.com/science/article/pii/S0167739X23002030).
- 558 [18] D. Jauk, D. Yang, and M. Schulz. “Predict-  
559ing Faults in High Performance Computing  
560 Systems: An in-Depth Survey of the State-of-  
561 the-Practice”. In: *Proceedings of the Interna-*  
562 *tional Conference for High Performance Com-*  
563 *puting, Networking, Storage and Analysis*. SC  
564 ’19. Denver, Colorado: Association for Com-  
565 puting Machinery, 2019. ISBN: 9781450362290.  
566 DOI: [10.1145/3295500.3356185](https://doi.org/10.1145/3295500.3356185). URL: [https://](https://doi.org/10.1145/3295500.3356185)  
567 [doi.org/10.1145/3295500.3356185](https://doi.org/10.1145/3295500.3356185).
- 568 [19] Y. Gao, X. Wang, X. He, Z. Liu, H. Feng, and  
569 Y. Zhang. “Alleviating structural distribution  
570 shift in graph anomaly detection”. In: *Proce-*  
571 *edings of the sixteenth ACM international con-*  
572 *ference on web search and data mining*. 2023,  
573 pp. 357–365.
- 574 [20] M. Liu, Y. Liu, K. Liang, W. Tu, S. Wang, S.  
575 Zhou, and X. Liu. “Deep temporal graph clus-  
576 tering”. In: *arXiv preprint arXiv:2305.10738*  
577 (2023).