

DyTwin: Federated Adaptive Digital Twins for Data Centers – Visualization and Anomaly Detection

Ebad Taheri
Hewlett Packard Labs
Milpitas, CA
ebad.taheri@hpe.com

Pedro Bruel
Hewlett Packard Labs
Milpitas, CA
bruel@hpe.com

Pavana Prakash
Hewlett Packard Labs
Milpitas, CA
prakash@hpe.com

Gourav Rattihalli
Hewlett Packard Labs
Milpitas, CA
gourav.rattihalli@hpe.com

Ninad Hogade
Hewlett Packard Labs
Fort Collins, CO
ninad.hogade@hpe.com

Aditya Dhakal
Hewlett Packard Labs
Santa Clarita, CA
aditya.dhakal@hpe.com

Rolando P. Hong Enriquez
Hewlett Packard Labs
Oxford, UK
rhong@hpe.com

Torsten Wilde
Hewlett Packard Enterprise
Berlin, Germany
wilde@hpe.com

Leo Popokh
Hewlett Packard Enterprise
Carrollton, TX
leonid.i.popokh@hpe.com

Dejan Milojicic
Hewlett Packard Labs
Milpitas, CA
dejan.milojicic@hpe.com

Cullen Bash
Hewlett Packard Labs
Milpitas, CA
cullen.bash@hpe.com

Abstract—Reliable and uninterrupted operation is crucial in supercomputers, especially during failures or inconsistencies i.e., anomalies. In this paper, we present a federated adaptive Digital Twin (DT) framework, with a focus on enhancing anomaly detection – a critical aspect of modern data center management. Our DT continuously monitors key metrics, detects anomalies powered by AI, and dynamically adjusts its monitoring parameters to ensure optimal performance. Using a dashboard, our system provides real-time alarms and detailed visualizations of detected anomalies, along with real-time visualization and forecast for selected metrics. Through a series of experiments, we validate the effectiveness of our approach in maintaining operational reliability and promptly identifying potential anomalies within the data center.

Index Terms—Digital Twin, Real-time Visualization, Data Center Management, Anomaly Detection, Federated Adaptive Systems

I. INTRODUCTION

A naive definition of Digital Twins (DTs) would involve for instance, creating a digital version of a physical asset and maybe include some level of communications between them. Although NASA’s Apollo program in the 1970’s is often credited for the implementation of the first DT, what they built was mostly a “physical twin”. That is, they made an identical physical copy of the actual spacecraft. Training, troubleshooting, and what-if scenarios during the mission were therefore performed using physical simulators and certainly not digital ones [1]. Nevertheless, the core idea was already there, and the concept of DTs evolved, adopting multiple definitions with their associated implementations. The unfolding of these concepts into an assortment of definitions have been reviewed elsewhere [2]–[4]. With an ongoing proliferation of DT definitions, the pragmatic way to proceed is probably avoiding the temptation of adding ever more technical DT definitions to the list. Although partially unsatisfactory, it might be better to simply design DTs to solve actual problems and worry about the relevance of a particular DT definition

later, if at all. Nevertheless, for the present contribution we operationally use the notion of data center digital twin as recently proposed by Athavale et al. [5]: “we define a data center digital twin (DCDT) as a virtual model that replicates the structure, context, and behavior of a data center. Continuously updated with data from its physical counterpart, a DCDT possesses predictive capabilities, and aids in informed decision-making to optimize operations, extend operating life, and realize value. A key aspect of a DCDT is the bidirectional interaction between the virtual model and the physical data center.”

In the context of modern data center management, DTs can play a critical role in not only simulating and optimizing operations but also in enhancing the detection and visualization of anomalies. As data centers scale and become more complex, the ability to monitor and detect irregularities in real-time becomes increasingly crucial. DTs, integrated with our anomaly detection mechanisms, allow for continuous monitoring and rapid identification of deviations from expected behavior, which is essential for maintaining operational reliability.

Furthermore, visualization is vital for interpreting the vast amount of data generated and for quickly assessing the health of the system. Anomaly detection, coupled with effective visualization, empowers operators to respond swiftly to potential issues, minimizing downtime and ensuring the stability of the infrastructure. This paper explores how DTs, anomaly detection, and visualization are interwoven to enhance the management and resilience of modern data centers.

Although it is indeed possible to do visualization and anomaly detection without DTs, implementing these features through a DT offers several advantages. A DT provides a comprehensive, integrated view of the entire system, allowing for more context-aware anomaly detection, considering various factors such as location, time, user activity and surrounding

conditions. DTs can correlate data from various sources, making it possible to detect complex anomalies that might be missed when monitoring individual components in isolation. Moreover, a DT can scale more easily than a traditional monitoring system since it can adapt to changes in the physical infrastructure without requiring extensive reconfiguration.

II. ADAPTIVE FEDERATED DIGITAL TWINS

The present work provides a general framework to model the behaviour of multiple data centers simultaneously by using an adaptive federated DT schema. The architecture for this framework is displayed in Figure 1. In the proposed architecture, time series data from the *Physical Twin* flows initially to the Prometheus monitoring system and database [6] and subsequently it can go directly to Grafana [7] for real time visualization or to the communication module of the *Digital Twins*. For each physical data center we create a corresponding digital twin that incidentally, is functionally divided into a digital twin at the node level (DT-N), a digital twin at the rack level (DT-R) and a last one at the data center level (DT-DC). Our architecture also features a global AI system that controls the monitoring and makes predictions at different levels and in different combinations. For every digital version of a data center, we also included a management module that communicates and synchronizes the predictions between all the other digital data centers in the federation. Additionally, the predictions from the AI module are handed over through the management module to a database connector to InfluxDB, which also keeps data from Prometheus.

In this work, by purposely adopting a federated architecture for DTs we implicitly recognize the need to train complex models that must comply with a variety of constraints. This is particularly important for our use case, which involves a group of individual data centers with their own administration, data, and security requirements. The decentralized nature of federated learning applied here allows for the creation of complex multi-level models for the group of data centers as a whole by using selected streams of data that are both, generated and stored in their originating data centers. Therefore, this approach structurally minimizes data transfer and increases overall data security while providing monitoring and predictive services. Moreover, federation enables additional insights into prediction, for example, if a specific application or memory consistently fails across data centers, it can be concluded more convincingly than in the case of a single system alone.

A. Potential applications and use cases of the proposed architecture

To illustrate the wider scope behind the construction of a DT-based system in our proposed architecture, we present here a non-exhaustive list of potential use cases.

Data center development and prototyping - Evaluate the effect of introducing changes in data centers or create baseline assumptions about data centers designed from scratch.

Workforce training - Staff responsible for various aspects of data center operations can be trained to maintain and trou-

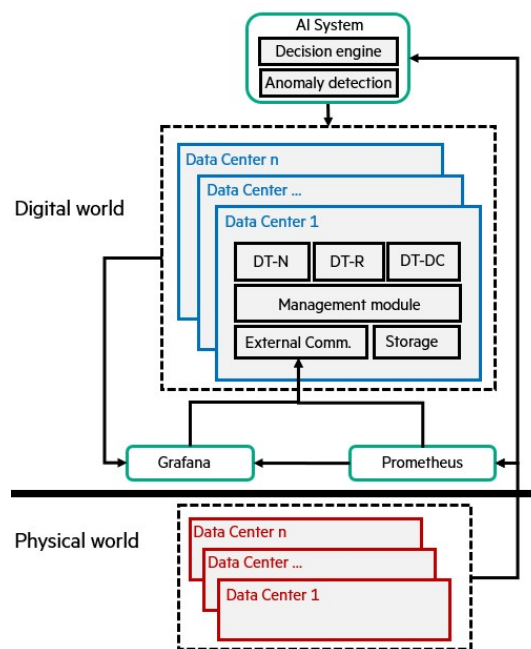


Fig. 1. Federated Digital Twin architecture. Time series data from the Physical Twins (data centers) are ingested by the Prometheus [6] and is also used by a global AI system that features a decision engine and anomaly detection models. For real-time visualization, Prometheus shares data with Grafana [7]. For further processing, Prometheus also shares data with the DT versions of the physical assets through their data communication modules. To account for the multi-level complexity of the physical assets, digital subsystems are created at the node level (DT-N), rack level (DT-R) and data center level (DT-DC). Data processed by the federated DTs are fed back to Grafana in form of models, predictions, and alerts.

bleshoot different problems using engaging DT experiences. Knowledge transfer to new personal can also be facilitated by DTs with the consequent increase of workplace safety.

Predictive maintenance - The models developed for the DTs can be used to predict failures of data center components. With this information, precise and dynamic maintenance schedules and/or component replacements can be optimized to increase financial savings and minimizing operational disruptions.

Insights from simulations - A series of simulations, stress testing and what-if scenarios can be used to understand the likely behaviours of the data centers when faced with specific conditions. The data produced in those tests and simulations, if validated, can be used in several strategies for data augmentation to improve the DT predictive capabilities.

Monitoring - Flexible and multi-level realtime monitoring of operational parameters with enhanced visualizations enable direct intervention by data center staff.

Detection of anomalies and cyber-physical attacks - Models trained at several levels can be used as generative models for anomaly detection. Historical data collected on cyber-physical attacks can eventually be used to train tailored security models.

B. Case Study: Anomaly Detection

In this section, we discuss and analyze anomaly detection as a use case of our DT framework. Anomalies in a data center can arise from various sources and must be promptly detected

to ensure optimal performance and avoid potential disruptions. Relevant anomalies might include unexpected temperature spikes which could indicate cooling system failures, or sudden drops in power supply, possibly due to faulty power distribution units. Other anomalies might include abnormal network traffic patterns, potentially signaling a security breach, or unusual latency in data processing, which could be a sign of hardware degradation or software malfunctions. By detecting such anomalies early, the DT enables proactive maintenance and response strategies, improving security, minimizing downtime and ensuring the continuous and efficient operation of the data center.

In this paper, we focus on anomalies that affect system performance and set alarms when such anomalies are detected. We consider a decision engine that can trigger various actions within our data center. One such action, which we discuss further in the next section, is altering the monitoring policy itself. For instance, when the DT detects possible anomalies, it increases the frequency of updates and monitors the data center more closely for any potential issues.

Generally, our DT detects anomalies by comparing the behavior of the physical twin against the DT. The motivation behind this approach is that the DT is trained and adapted to mimic the physical twin's behavior. When there is a divergence between the two, it suggests the presence of an anomaly, indicating that the current behavior of the system deviates from its past behavior.

In this case study, our focus is not on designing an optimal DT that must accurately represent its physical counterpart. Instead, as a proof of concept, our aim is to highlight an important use case of a DT. However, we believe that our initial model is a strong candidate for DT design, and we are interested in exploring its efficiency compared to other mathematical and machine learning models in the future. In this study, we employ Gaussian Process Regression with a Spectral Mixture Kernel [8] to model the behavior of the physical twin. The DT is equipped with an anomaly detection module that continuously compares the actual measurements from the physical twin with the predicted values from the DT counterpart. The difference between these measurements, termed the error, serves as the basis for our anomaly detection. Let y_t represent the actual measurement from the physical twin at time t , and \hat{y}_t represent the inference value from the DT. The error e_t at time t is defined as:

$$e_t = y_t - \hat{y}_t \quad (1)$$

Additionally, we define the accepted error threshold ϵ as:

$$\epsilon = \sigma \times \text{confidence} \quad (2)$$

where σ represents the standard deviation of the Gaussian Process predictions, and confidence is the desired confidence level.

If the absolute value of the error $|e_t|$ exceeds this threshold, the behavior of the physical twin is flagged as anomalous:

$$\text{Anomaly} \iff |e_t| > \epsilon \quad (3)$$

This approach leverages the probabilistic nature of Gaussian Process Regression to provide not only point estimates but also confidence intervals for predictions, allowing for more robust anomaly detection. By continuously monitoring the error margins, our system can dynamically adapt to varying operational conditions, enhancing the reliability and accuracy of the DT.

In our DT, we have designed a specialized module called the Decision Engine, which makes critical decisions not only for triggering actions but also for adapting the detection frequency. This adaptive mechanism is central to the scalability and efficiency of our approach.

When an anomaly is detected, the DT automatically switches to a high-frequency detection mode. This heightened monitoring allows for closer reliability and scrutiny of the system, enabling the DT to rapidly identify any subsequent anomalies or patterns that may emerge. By dynamically adjusting the detection frequency based on the likelihood of an anomaly, the system ensures that it remains vigilant when necessary, without expending unnecessary resources during normal operation.

This dynamic design is key to the scalability of our approach. Under normal conditions, there is no need for frequent updates to the models within the DT, which conserves computational resources and reduces overhead. However, when the system detects that an anomaly is more likely, the DT adapts by increasing the detection frequency, ensuring that potential issues are caught and addressed with greater precision.

This balance between efficiency and responsiveness allows our DT to maintain a high level of performance and reliability, making it an effective tool for monitoring complex systems and preventing failures in real-time.

III. PERFORMANCE EVALUATION

We defined two alarms for our DT: a yellow alarm that triggers when the error remains high for a short time (we consider 10 minutes for our experiment), and a red alarm that triggers when the error remains high for a long time (30 minutes in this experiment). We injected different rates of anomalies using stress-ng [9], affecting CPU utilization and ranging from 5% to 60%. The results are shown in Figure 2. We evaluated the true positive rate, which indicates when an alarm is correctly triggered while the data center is experiencing an anomaly, and the false positive rate, which indicates when the DT mistakenly sets on an alarm. False positives may occur either due to the inefficiency of the DT, which fails to generalize and mimic the behavior of its physical counterpart, or because of changes in the resources and data of the physical twin. We injected 10 random anomalies for each test scenario (e.g., 10 anomalies for the 5% scenario) at different times of the day. Each injected anomaly lasted for 1 hour.

Figure 2 shows the detection capability of our DT under the various anomaly scenarios. In Figure 2(a), we present an example of anomaly detection, where an anomaly was deliberately injected at time 20:30. The figure depicts the resulting




Fig. 2. Digital twin performance under various anomalies: (a) a detection showcase, (b) Rate of True Positives (TP) among various random anomalies injected during different times of the day, and (c) Rate of False Positives (FP)

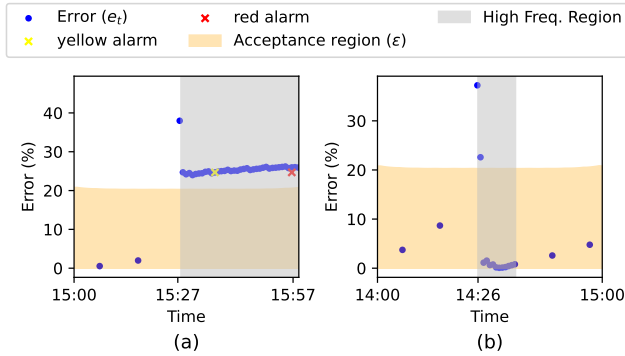


Fig. 3. Adaptive anomaly detection by the digital twin. (a) Anomaly injected at 15:25 leads to increased detection frequency and subsequent alarms. (b) high error at 14:26 prompts a temporary frequency increase, reverting after error absence.

Error e_t and Acceptance region ϵ as defined in equations 1 and 2. As expected, there is a notable increase in the error once the anomaly manifests. This occurs because the DT begins to observe a metric—in this case, CPU usage—that deviates significantly from the expected behavior of the physical twin. Such a deviation is critical, as it triggers alarms in the DT when the anomaly persists for a predefined duration, leading to the activation of both yellow and red alarms.

In Figure 2(b), we analyze the rate of true positives for 10 randomly injected anomalies occurring at different times of the day. The results are particularly promising: after injecting an anomaly that affects 40% of the CPU usage, the DT successfully detected 100% of the anomalies. Furthermore, the false positive rate remains impressively low, with both yellow and red alarms showing false positive rates below 1%. Depending on the use case the general acceptable false positive rate could be either data driven or coming from some human insights. In the case lower false positive is required, the error acceptance can be increased. The high true positive and low false positives underscore the robustness of our DT in accurately identifying anomalies while minimizing false alarms, thereby enhancing the reliability of the system.

Figure 3 illustrates our analysis of DT adaptation. In Figure 3(a), the anomaly detection process is shown. Before 15:27, the error rate is low because the DT expects the physical twin’s behavior to be normal. In this scenario, the detection frequency is set to 10 minutes. We injected an anomaly at 15:25, and

based on the detection frequency, the first error was detected at 15:27. Upon detecting the high error rate, the DT switched to a high-frequency detection mode, with a detection frequency of 1 minute. After 10 minutes, at 15:37, a yellow alarm was triggered, and subsequently, a red alarm was set at 15:57 as the error persisted.

In contrast, Figure 3(b) shows a scenario where no anomaly was injected. At 14:26, a high error was detected, prompting the DT to switch to high-frequency detection. However, since the error did not recur after two samples, the DT reverted to the normal detection frequency after 10 minutes.

A. Visualization

To effectively monitor and manage the large number of machines in our data center, we employ Grafana as our primary visualization tool, integrated with our DT. As shown in Figure 4, Grafana enables us to track and analyze CPU and memory utilization, as well as anomaly status, across the entire infrastructure, providing real-time insights into system performance and health. Please note that in this figure, we have selectively reduced the number of machines and metrics displayed to fit within the paper’s space constraints.

Our visualization setup features a color-coded status system, as shown in right side of Figure 4, where each machine is represented by a state indicator: green for normal operation, yellow for potential anomalies, and red for confirmed anomalies. Additionally, Grafana Alerts are configured to automatically notify operators when predetermined thresholds are exceeded, allowing for proactive responses to possible issues before they escalate. This system allows for a quick and intuitive assessment of the overall health of the data center, enabling rapid responses to any irregularities. The use of Grafana enhances our ability to maintain a high level of operational reliability by ensuring that anomalies are promptly identified and addressed.

IV. RELATED WORK

In [10], a DT is used for the intelligent management of data center networks. This DT abstracts the service model, particularly during device and network upgrades, and supports model management, configuration automation, and network element validation. While these applications are significant, we believe there are additional uses for DTs that warrant exploration. In [11], offline-trained machine learning models

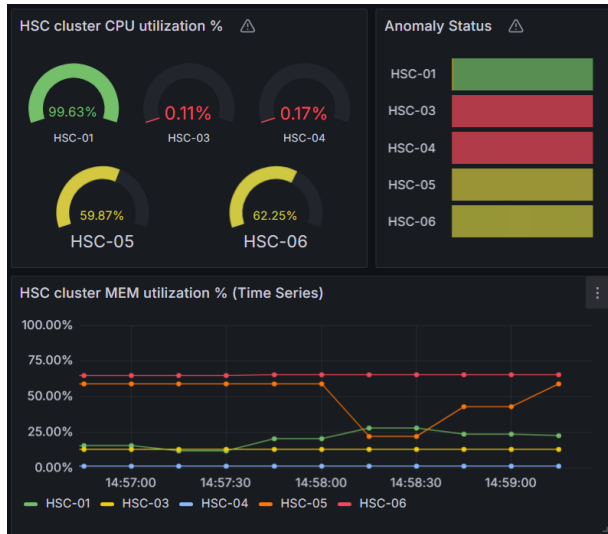


Fig. 4. Grafana dashboard integrated with our digital twin, to monitor CPU, memory utilization, and anomaly status across the data center. The color-coded indicators on the right side—green for normal, yellow for potential anomalies, and red for confirmed anomalies—enable quick, real-time assessments.

are integrated with DTs to enhance decision-making. Although this approach leverages valuable features of DTs, it does not address several key use cases, such as anomaly detection and visualization.

There have been successful traditional approaches to anomaly detection that do not use DTs. For instance Zhao et al. [12] uses anomaly detection models to predict hard drive failures in data centers, Pinciroli et al. [13] uses these models to predict industrial machine failures using multivariate time series, Lee et al. [14] focus on thermal anomaly detection using heat-generation modeling and thermal cameras, and Todd et al. [15] also utilizes patterns of real-time data to flag anomalies in HPC data center operation. The relevance of anomaly detection extends to network security, as highlighted by Ahmed et al. [16], who reviewed state-of-the-art network anomaly detection techniques. Classification-based methods, such as support vector machines (SVM) used by Balabine et al. [17], are popular for detecting network anomalies. However, maintaining an up-to-date normal profile is increasingly challenging in dynamic environments [16]. Yet none of them incorporate DT technology and therefore cannot benefit from the additional advantages of this technology.

DTs enhance these approaches by providing a comprehensive and adaptive monitoring environment. They enable real-time tracking of system performance and the ability to perform dynamic what-if analyses. This capability supports more informed decision-making and enhances overall system management. Our work includes an adaptive DT framework that adjusts monitoring and detection frequencies to balance scalability with detection accuracy.

V. DISCUSSION AND CONCLUSION

We presented a federated adaptive DT framework for anomaly detection in data centers. Our system effectively

monitors data center metrics, detects anomalies with high accuracy, and integrates with Grafana for real-time alerting. The experiments demonstrate the system’s capability to enhance operational reliability with minimal false positives. DTs can also significantly contribute to performance prediction, forecasting, and optimization within data centers.

One of the key advantages of using DTs in a data center environment is their ability to create a dynamic and continuously updated digital replica of the physical infrastructure. This real-time synchronization allows operators to monitor and analyze the performance of individual machines and systems with a high degree of accuracy. The inclusion of anomaly detection algorithms within the DT framework further strengthens this capability, enabling the system to identify deviations from normal operating conditions as soon as they occur. Anomaly detection is crucial in preventing system failures and optimizing maintenance schedules.

Visualization also plays a pivotal role in making these complex processes accessible and actionable for operators. Tools like Grafana, when integrated with DTs, offer intuitive dashboards that display real-time data in a clear and comprehensible manner. The color-coded status indicators, as discussed earlier, allow operators to quickly assess the health of the entire data center at a glance. This visual approach not only improves situational awareness but also enhances decision-making by providing a quick reference for understanding the current state of the system. Moreover, the ability to customize and filter the data displayed in these visualizations ensures that the most relevant information is always at the forefront, particularly in large-scale data centers.

Despite these advantages, the implementation of DTs, anomaly detection, and visualization is not without challenges. The accuracy of the DT relies heavily on the quality and timeliness of the data it receives. Any discrepancies or delays in data transmission can lead to a misalignment between the digital and physical twins, potentially compromising the effectiveness of anomaly detection and the reliability of visualizations. Additionally, the development and integration of anomaly detection algorithms require careful tuning to balance sensitivity and specificity, ensuring that true positives are captured without an excessive number of false positives.

Looking forward, the continued evolution of DTs, coupled with advances in machine learning and AI, promises to further enhance anomaly detection capabilities. As these technologies mature, we can expect more sophisticated models that not only detect anomalies but also predict potential failures before they occur. This predictive maintenance approach could revolutionize the way data centers are managed, shifting from reactive to proactive strategies.

Last but not least, we are in the process of integrating our monitoring capabilities with augmented/virtual reality (AR/VR) technologies. This integration aims to enhance the visualization experience by providing an immersive environment where administrators can interact with and manage the data center in a more intuitive and spatially-aware manner.

REFERENCES

- [1] R. Rosen, G. von Wichert, G. Lo, and K. D. Bettenhausen, "About the importance of autonomy and digital twins for the future of manufacturing," *IFAC-PapersOnLine*, vol. 48, no. 3, pp. 567–572, 2015, 15th IFAC Symposium on Information Control Problems in Manufacturing. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2405896315003808>
- [2] M. Sjarov, T. Lechler, J. Fuchs, M. Brossog, A. Selmaier, F. Faltus, T. Donhauser, and J. Franke, "The digital twin concept in industry – a review and systematization," in *2020 25th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*, vol. 1, 2020, pp. 1789–1796.
- [3] C. Semeraro, M. Lezoche, H. Panetto, and M. Dassisti, "Digital twin paradigm: A systematic literature review," *Computers in Industry*, vol. 130, p. 103469, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0166361521000762>
- [4] J.-F. Yao, Y. Yang, X.-C. Wang, and X.-P. Zhang, "Systematic review of digital twin technology and applications," *Visual Computing for Industry, Biomedicine, and Art*, vol. 6, no. 1, p. 10, May 2023. [Online]. Available: <https://doi.org/10.1186/s42492-023-00137-4>
- [5] J. Athavale, C. Bash, W. Brewer, M. Maiterth, D. Milojicic, H. Petty, and S. Sarkar, "Digital twins for data centers," *IEEE Computer*, p. to appear.
- [6] "Prometheus monitoring system," 2012, <https://prometheus.io/> [Accessed: 30.07.2024].
- [7] "Grafana: The open-source platform for monitoring and observability." 2024, <https://grafana.com/> [Accessed: 30.07.2024].
- [8] C. Williams and C. Rasmussen, "Gaussian processes for regression," *Advances in neural information processing systems*, vol. 8, 1995.
- [9] "stress-ng." [Online]. Available: <https://wiki.ubuntu.com/Kernel/Reference/stress-ng>
- [10] H. Hong, Q. Wu, F. Dong, W. Song, R. Sun, T. Han, C. Zhou, and H. Yang, "Netgraph: An intelligent operated digital twin platform for data center networks," in *Proceedings of the ACM SIGCOMM 2021 workshop on network-application integration*, 2021, pp. 26–32.
- [11] T. Goodwin, J. Xu, N. Celik, and C.-H. Chen, "Real-time digital twin-based optimization with predictive simulation learning," *Journal of Simulation*, vol. 18, no. 1, pp. 47–64, 2024.
- [12] M. Zhao, R. Furuhashi, M. Agung, H. Takizawa, and T. Soma, "Failure prediction in datacenters using unsupervised multimodal anomaly detection," in *2020 IEEE International Conference on Big Data (Big Data)*. IEEE, 2020, pp. 3545–3549.
- [13] N. O. Pinciroli Vago, F. Forbicini, and P. Fraternali, "Predicting machine failures from multivariate time series: An industrial case study," *Machines*, vol. 12, no. 6, p. 357, 2024.
- [14] E. K. Lee, H. Viswanathan, and D. Pompili, "Model-based thermal anomaly detection in cloud datacenters using thermal imaging," *IEEE Transactions on Cloud Computing*, vol. 6, no. 2, pp. 330–343, 2015.
- [15] A. Todd, A. Purkayastha, H. Egan, D. Sickinger, M. Eash, S. Serebryakov, J. Hanson, M. Slaby, N. Wunder, N. Guba *et al.*, "Artificial intelligence for data center operations (ai ops)," National Renewable Energy Lab.(NREL), Golden, CO (United States), Tech. Rep., 2021.
- [16] M. Ahmed, A. N. Mahmood, and J. Hu, "A survey of network anomaly detection techniques," *Journal of Network and Computer Applications*, vol. 60, pp. 19–31, 2016.
- [17] I. Balabine and A. Veleznitsky, "Method and system for confident anomaly detection in computer network traffic," Dec. 12 2017, uS Patent 9,843,488.