2023 IEEE 47th Annual Computers, Software, and Applications Conference (COMPSAC) | 979-8-3503-2697-0/23/$31.00 ©2023 IEEE | DOI: 10.1109/COMPSAC57700.2023.00224

The Graph-Massivizer Approach Toward a European Sustainable Data Center Digital Twin

Martin Molan, Junaid Ahmed Khan, Andrea Bartolini
Alma Mater Studiorum – Università di Bologna
Bologna, Italy
martin.molan2@unibo.it, junaid.ahmedkhan@studio.unibo.it, a.bartolini@unibo.it

Roberta Turra, Giorgio Pedrazzi
Cineca Consorzio Interuniversitario
Bologna, Italy
r.turra@cineca.it, g.pedrazzi@cineca.it

Michael Cochez, Alexandru Iosup
Vrje Universiteit Amsterdam
Amsterdam, Netherlands
m.cochez@vu.nl, a.iosup@vu.nl

Dumitru Roman
SINTEF AS
Oslo, Norway
dumitru.roman@sintef.no

Jože Rožanec
Institut “Jožef Stefan”
Ljubljana, Slovenia
joze.rozanec@ijs.si

Ana Lucia Vărbănescu
Universiteit Twente
Amsterdam, Netherlands
a.l.varbanescu@utwente.nl

Radu Prodan
Alpen Adria – Universität Klagenfurt
Klagenfurt am Wörthersee, Austria
radu.prodan@aau.at

***Abstract*— Modeling and understanding an expensive next-generation data center operating at a sustainable exascale performance remains a challenge yet to solve. The paper presents the approach taken by the Graph-Massivizer project, funded by the European Union, towards a sustainable data center, targeting a massive graph representation and analysis of its digital twin. We introduce five interoperable open-source tools that support this undertaking, creating an automated, sustainable loop of graph creation, analytics, optimization, sustainable resource management, and operation, emphasizing state-of-the-art progress. We plan to employ the tools for designing a massive data center graph, representing a digital twin describing spatial, semantic, and temporal relationships between the monitoring metrics, hardware nodes, cooling equipment, and jobs. The project aims to strengthen Bologna Technopole as a leading European supercomputing and big data hub offering sustainable green computing for improved societally relevant science throughput.**

Keywords—Data center, digital twin, green computing, sustainability, massive graph, graph processing, graph neural network, anomaly detection.

I. Introduction

Supercomputers are the backbone of high-performance computing (HPC), supporting computational science discoveries and massive engineering analyses. They maximize societal and economic impact through high computation and scientific throughput per investment. As the community focuses on peak performance in its race towards exascale machines, two critical factors limit HPC sustainability. Firstly, energy consumption is a vital factor in data centers' total cost of ownership (TCO) and a de-facto barrier to their peak performance. While data-driven heat dissipation models exist in digital devices, they do not capture complex spatiotemporal dependencies between the cooling equipment, computing nodes, and computational workloads. Secondly, system utilization is critical and directly impacts a supercomputer's productivity, quantified by the science, research, and innovation throughput. However, system utilization is hard to maximize while preserving fairness and fulfilling the requirements of jobs and workloads.

Modeling and understanding an expensive next-generation data center operating at a sustainable exascale performance before constructing it requires a comprehensive *massive graph (MG)* representation and analysis of its digital twin, describing spatial, semantic, and temporal relationships between the monitoring metrics, hardware nodes, cooling equipment, and jobs. Such MGs are universal abstractions that capture, combine, model, analyze, and process knowledge about real and digital worlds into actionable insights through digital twin representations. For societally relevant problems, such as exascale supercomputers, graphs are extreme data that require further technological innovations to meet the needs of the European data economy.

The *Graph-Massivizer project* [2], funded by the Horizon Europe research and innovation program of the European Union, targets the development of a high-performance and sustainable platform for information processing and reasoning based on the MG representation of extreme data in the form of general graphs, knowledge graphs, and property graphs. Graph-Massivizer addresses the *any-volume* graph challenge by supporting up to billions of vertices and trillions of edges and the *velocity* graph challenge of dynamically changing topologies. Finally, it proposes a novel *viridescence* graph challenge for sustainable processing at exascale speed.

Graph-Massivizer targets *sustainable science throughput* through scalable energy-aware, exascale operation and total traceable cost of ownership understanding, including sustainability indicators and their environmental effects (e.g., greenhouse gas emissions). The Graph-Massivizer tools will enable the creation of a novel, graph-based digital twin of a data center to further support the construction of sustainable exascale computing operational models for scientific discovery in the next decade.

This paper reuses parts of the original Horizon Europe project proposal Graph-Massivizer. Some paragraphs overlap articles in IEEE Cloud Summit 2022 [2] and GraphSys '23 workshop [3], aiming to connect with different audiences.

The paper has six sections. Section II covers the ambition of the Graph-Massivizer toolkit covering the sustainable lifecycle of processing extreme data as MG using a comprehensive state-of-the-art analysis. Section III presents the use case of employing the Graph-Massivizer tools for designing a data center digital twin for sustainable exascale computing. Section III.B presents an anomaly detection use

Authorized licensed use limited to: Vrije Universiteit Amsterdam. Downloaded on December 17,2024 at 16:31:18 UTC from IEEE Xplore. Restrictions apply.

case on the Marconi 100 (M100) digital twin. Section V elaborates on the expected impact of the proposed digital twin solution on European science, industry, and society, and Section VI concludes the paper.

II. Graph-Massivizer: State-of-the-Art and Ambition

Current graph processing platforms do not support diverse workloads, models, languages, and algebraic frameworks [1]. Existing specialized platforms are difficult to use by non-experts and suffer from limited portability and interoperability, leading to redundant efforts and inefficient resource and energy consumption due to vendor and even platform lock-in. While synthetic data emerged as an invaluable resource overshadowing actual data for developing robust machine learning (ML) analytics, graph generation remains a challenge due to extreme dimensionality and complexity. On the European scale, this state of practice is unsustainable and, thus, threatens the possibility of creating a climate-neutral and sustainable economy based on graph data. Making graph processing sustainable is essential but needs credible evidence [2].

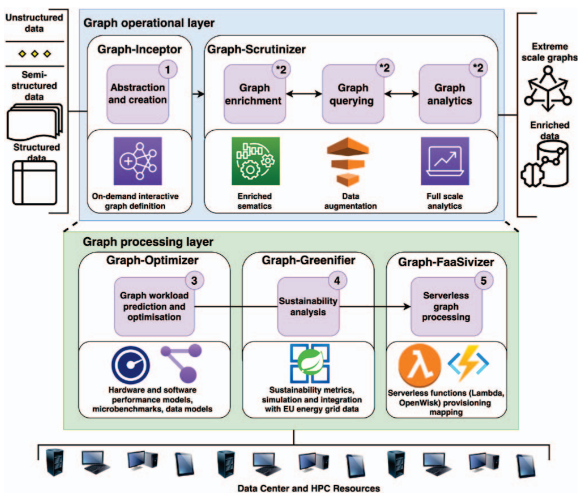

Figure 1: Graph-Massivizer architecture.

Graph-Massivizer proposes to research and develop a toolkit consisting of five interoperable open-source tools to support this need, creating an automated, sustainable loop of graph creation, analytics, optimization, sustainable resource management, and operation with unprecedented capabilities and quality (see Figure 1). This section introduces the tools but summarizes the state-of-the-art in the field and then concludes with each tool's novelty.

A. *Graph-Inceptor: Extreme Data Ingestion, Massive Graph Creation and Storage*

a) State-of-the-art: Many ETL (extract, transform, and load) products, such as Airbyte, Stitch, Fivetran, Matillion, Pipelinewise, Airflow, Singer, Meltano, and Hevo Data, provide multiple generic connectors to load and persist data. Nevertheless, they target only tabular data and do not support extreme-scale data. Furthermore, their data persistence connectors do not natively support graph storage, and possible integrations require generic interfaces. The dbt framework captures the best data practices, requiring versioning, testing, and deploying data transformations following a continuous delivery pipeline. However, it uses SQL for data transformations. At the same time, there is a need for a generic framework to support scripting languages and native connectors for graph storage, leveraging ML to enhance graph creation. Existing database solutions need extensions to help graph storage, such as the Titan [6] distributed graph database supporting labeled property graphs. In contrast, BlazeGraph and Amazon Neptune comply with the RDF/SPARQL paradigm and the Apache TinkerPop [7] graph stack but lack distributed storage and processing capabilities. None leverages ML to optimize distributed storage, processing, and querying.

b) Ambition: Graph-Massivizer develops a smart ETL component, learning from data content and load to drive further optimizations. It also proposes a graph-first ETL framework, providing connectors to specific graph storages and enhancing them with ML models to improve persistence and query execution. Finally, it researches streaming ML models and approximate online reasoning methods to avoid expensive batch model training over MG.

B. *Graph-Scrutinizer: Massive Graph Analytics and Reasoning*

a) State-of-the-art: Distributed frameworks, like Pregel [8] and Apache GraphX [9], provide limited-scale graph analytics; stand-alone sampling algorithms are only available as academic prototypes, while commercially distributed graph sampling implementations (e.g., AliGraph [10]) are not open source. Existing ML-based graph reasoning libraries are proprietary or work within a single node, lacking distributed processing. Relevant libraries such as GraphSAGE [11] create graph embeddings, while PyTorch Geometric [12] implements graph neural networks. Most graph storage solutions leverage the Gremlin language for provider-agnostic querying capabilities based on the Apache TinkerPop [13] graph stack.GQE [14] foundational work on multi-hop reasoning supports a limited set of operators, while its Query2Box [15] and BetaE extensions allow disjunctions and negations. EmQL improves on Query2Box with general entailment and count-min sketches. MPQE [16] enables graph representation of queries for reasoning, BiQE [17] solves multiple simultaneous reasoning tasks, and CQD [18] answers complex queries without explicit representations.

b) Ambition: Graph-Massivizer researches a distributed library combining multiple graph sampling strategies and graph storage connectors. It also investigates streaming ML and online probabilistic reasoning methods that provide advanced analytics and reasoning while avoiding regular model retraining over the MG. Graph-Scrutinizer decomposes reasoning tasks into sequences of smaller primitives trained without end-to-end learning.

C. *Graph-Optimizer: Graph Processing Workload Modelling with Performance and Energy Guarantees*

a) State-of-the-art: Graph and hardware-specific optimizations lead to orders of magnitude improvements in performance, energy, and cost over conventional graph processing methods [19][20][21]. Typical big data platforms, such as Apache MapReduce and Apache Spark, rely on generic primitives, exhibiting poor performance and high financial and environmental costs [19][22]. Even optimized basic graph operations (BGO) lack tools to combine them

Authorized licensed use limited to: Vrije Universiteit Amsterdam. Downloaded on December 17,2024 at 16:31:18 UTC from IEEE Xplore. Restrictions apply.

with real-world applications. Furthermore, graph topology and dynamics (i.e., changing the number and content of vertices and edges) lead to high variability in computational needs [23]. Primitive predictive models demonstrate they can enable algorithm selection and advanced auto-scaling techniques to ensure better performance [24], but no such models exist for energy.

b) Ambition: Graph-Massivizer develops a BGO repository optimized for different heterogeneous platforms (i.e., CPU, accelerators) and defines a methodology to design and accurately model BGO-based workloads. It further develops Graph-Optimizer, the first tool to provide accurate predictions and guarantees of performance and energy efficiency for graph processing workloads on heterogeneous systems. Its forecasts enable informed decisions for high-performance, sustainable graph processing.

D. *Graph-Greenifier: Sustainable and Energy-Aware Massive Graph Processing*

a) State-of-the-art: Energy consumption is a primary component of a data center's TCO, while power consumption and thermal dissipation limit the achievable peak performance [25]. The cooling system is central to reducing energy consumption costs and increasing peak performance but suffers from limited thermal design power. Hot water and free and chillerless cooling reduce the energy cost at the expense of increased complexity of the room thermal control [26]. Existing works predict and optimize the cooling efficiency based on operational and environmental parameters [27] but fail to capture spatiotemporal dependencies related to conduction and convection heat dissipation. Finally, they optimize energy based on simple metrics, such as resource utilization or dissipated heat. Similarly, computational benchmarks consider Watt consumption (e.g., in Green500) but not electricity production. Sustainability models do not capture complex prevalent operational phenomena, such as performance interference and failures, leading to limited predictive power and decision-making evidence [2].

b) Ambition: Graph-Massivizer defines accurate metrics, benchmarks, and green labeling techniques for sustainability and performance analysis of MG processing workloads. Graph-Massivizer enables energy-aware graph execution on large-scale infrastructure and codesigned hardware and produces evidence [2] for green graph processing by providing facts and model-based performance and sustainability analysis.

E. *Graph-Choreographer: Scalable Serverless Graph Analytics over a Codesigned Continuum Infrastructure*

a) State-of-the-art: Current serverless platforms, such as Apache OpenWhisk [28], rely on deployment frameworks like Kubernetes and Docker Swarm, which utilize greedy decision-making techniques, such as filtering nodes, and are incapable of hosting a given function and establishing a rating among the rest. Graphless [29] allows graph processing function deployment using push or pull operations on predefined worker resources using static and super-step schedulers. Related works employ resource partitioning [30] and centralized cluster-level heuristics to schedule latency and throughput-sensitive serverless applications but do not consider graph processing operations explicitly.

b) Ambition: Graph-Massivizer enables scalable energy-aware serverless execution of MG processing over extreme data, combining sustainability analysis and performance metrics, such as GHG emissions and throughput. It allows transparent provisioning over vast heterogeneous resources across the computing continuum, comprising HPC, Cloud, Edge, and specialized hardware.

III. Green Data Center Digital Twin

The complexity of modern pre-exascale supercomputers necessitates ML methodologies that support the work of the system administrators, built upon the large quantities of data collected by the sensors and monitoring systems of the HPC systems. Some ML-supported operational data analytics areas in HPC, such as anomaly detection, have matured and are operating in production systems. The transition to more complex tasks on the operational data analytics roadmap, however, necessitates the introduction of more sophisticated methodologies – specifically the processing of MGs.

A. *Dara Center Massive Graph*

The Graph-Massivizer tools leverage the holistic monitoring data of CINECA to produce a *data center MG (DC-MG)*, representing a digital twin describing spatial, semantic, and temporal relationships between the monitoring metrics, hardware nodes, cooling equipment, and jobs (see Figure 2). The DC-MG supports the deployment of performance prediction and what-if analysis using ML methods to change its configuration to maximize utilization or sustainability requirements and observe the effects in simulation. For this purpose, the data center digital twin relies on a mathematical and visually connected DC-MG model to locate undesired effects like highly demanded racks with prohibitive energy consumption. Graph-Massivizer will use the CINECA HPC facilities to investigate different clustered partitions of the HPC facility tuned for incoming workloads with accurate processing and queuing time estimates to improve its operational model based on these DC-MG analyses. The data center digital twin representation targets a sustainable computing operation at exascale by optimizing two parameters. Firstly, it improves the data center's power usage effectiveness (PUE) and GHG emissions by creating and training DC-MG to capture the spatiotemporal-ontological dependencies among computation, computing nodes, and cooling equipment and predict the impact of the spatial power distribution on cooling efficiency and cost. Secondly, it improves global resource utilization based on predictive workload, resource consumption, and job queuing models, maximizing the science throughput.

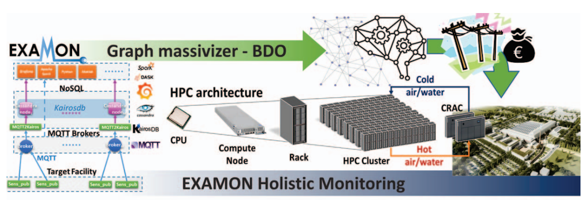

Figure 2. Data center digital twin.

Graph-Massivizer uses M100 scaled to the EuroHPC Leonardo pre-exascale supercomputer to design and validate the data center digital twin. The holistic monitored data includes approximately one million sensors producing 21,000 metrics per second on M100. The Leonardo supercomputer

Authorized licensed use limited to: Vrije Universiteit Amsterdam. Downloaded on December 17,2024 at 16:31:18 UTC from IEEE Xplore. Restrictions apply.

overpasses the complexity of the graph by two orders of magnitude compared to the current deployment (see Table 1).

Table 1. Extreme data characteristics of a data center digital twin at exascale.

Data characteristic	*Big HPC data state-of-the-art*	*Extreme HPC data dimension*
Volume	10 TB of monitoring data	1 PB of monitoring data
Velocity	21,000 metrics per second	2 million metrics per second
Variety	1 million metrics	10 million metrics
Viridescence	Unsustainable resource-intensive analytics	Sustainable energy-accountable graph analytics

B. *Examon Monitoring Framework*

The proper functioning of HPC systems relies heavily on their monitoring and data collection infrastructure. This infrastructure gathers information from various sources, including hardware sensors, software logs, and performance metrics. CINECA's M100 supercomputer uses the Examon monitoring framework developed in collaboration with the University of Bologna [31].

Examon [32] acquires many data inputs, including hardware sensors, which provide information on CPU load across all supercomputer node cores, metrics such as CPU clock speed, instructions executed per second, memory access frequency, and power consumption. Other sensor readings include fan speeds and temperature measurements for IT components and room conditions. Additionally, workload-related data, such as job submissions and their specific characteristics, are collected. System administrators utilize warning messages or alarms generated by diagnostic software in performing their duties to determine the availability status of compute nodes. Previous research demonstrated that combining these diverse metric sources into comprehensive dashboards provides a complete overview of the supercomputer's operation.

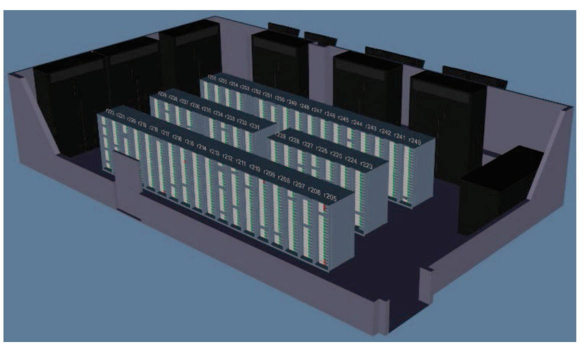

Figure 3: M100 digital twin representation.

Furthermore, integrating fundamental metrics and sophisticated data derived from ML models acquire demonstrated enhanced insights into the system [33]. One notable example involves generating a virtual replica of the M100 supercomputer as a three-dimensional interface that displays comprehensive information about all computational nodes, utilizing various hues to indicate their statuses. In Figure 4, for example, green nodes indicate a low failure probability, and red nodes indicate a high failure probability.

IV. Anomaly Prediction

Anomaly prediction is more advantageous in HPC than anomaly detection for proactive system management. Instead of solely relying on anomaly detection after their occurrence, implementing ML algorithms can enable a predictive analysis for potential anomalies before they occur, allowing system administrators to take preventive measures and effectively manage any negative impacts, ultimately leading to enhanced performance and reliability of the HPC systems.

A. *State-of-the-Art*

The state-of-the-art methodologies for anomaly detection train unsupervised per-node models [34] or use a semi-supervised autoencoder approach [33], which is unsuitable for supervised anomaly prediction learning. Supervised training for a per-node model is inappropriate for anomaly prediction as the anomalies are rare. Supervised training of a large model combining the data of all nodes in a supercomputer allows leveraging the collective knowledge from all nodes. Still, it fails to capture the specifics of an individual node and produces suboptimal results [33].

B. *Graph Neural Networks*

In Graph-Massivizer, the University of Bologna proposed a *graph neural network (GNN)* model for anomaly prediction [35]. GNNs are a deep learning algorithm well-suited for processing graph-structured data, common to HPC systems. Specifically, the method leverages the spatial proximity of compute nodes in an HPC room, uses GNNs to forecast anomalies, and conducts an experimental evaluation on a dataset collected from two M100 racks.

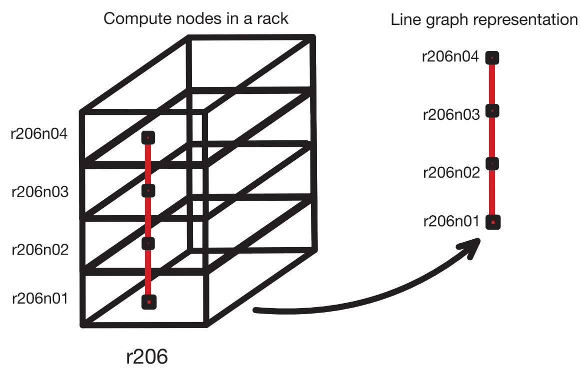

Figure 4: Line graph representation of compute racks.

Convolutional graph neural networks (GCN) refer to a distinct GNN for handling graph-structured information. GCNs perform convolution operations on the input graph data to gather knowledge from every node's nearby neighborhood and incorporate this into updating its representation. This iterative process continually updates these representations, eventually capturing the entire graph's global structure through several iterations. With this understanding, GCN algorithms apply to diverse applications such as node classification, trying to predict the probability of the anomaly (a label) for each compute node in a rack. Each rack is a line graph having nodes as vertices connected to those above and below the rack (see Figure 4). A node has the label 0 if no anomalies occur in a future window; otherwise, the label is anomalous. The graph representation of compute nodes takes advantage of the nodes' physical layout and creates an individual model for each computing rack. As explored in

Authorized licensed use limited to: Vrije Universiteit Amsterdam. Downloaded on December 17,2024 at 16:31:18 UTC from IEEE Xplore. Restrictions apply.

other approaches, the modular approach allows scalability with the future large exascale systems.

C. Evaluation

The anomaly prediction evaluated on M100 used a GNN trained and implemented using PyTorch on a single computing node comprising 32 IBM POWER9 cores, 256GB of memory, and four NVIDIA V100 GPUs with 16 GB of memory. The open dataset used for evaluation contains observations from two compute racks of M100 for 31 months, with 80% of data used for training and the remaining 20% reserved for testing. The GNN architecture underwent manual tuning by relying on background knowledge. It resulted in the following graph structure: graph convolution layer of shape (416,300), graph convolution layer of shape (300,100), graph convolution layer of shape (100,16), dense layer of shape (16,16), and dense layer of shape (16,1).

a) Anomaly detection: The anomaly detection results indicate an area-under-the-curve (AUC) of 0.86 in pure anomaly detection compared to 0.77 AUC for Recurrent Unsupervised Anomaly Detection (RUAD) [33] while predicting the anomalies by several hours. The prediction accuracy maintains good results of around 0.84 AUC with eight hours lookahead windows, slightly decreasing for more extended time windows. GNN achieves comparable results to RUAD up to 24 hours ahead (i.e., 0.73 AUC).

b) Anomaly prediction: Regarding prediction strength, GNN achieves comparable results with the last value anomaly predictor (in the absence of other more complex methods in the HPC literature), explained by the typical HPC anomaly duration of at least a few hours. The difference between the last value and GNN predictors becomes more significant for longer time intervals. In other words, anomaly prediction becomes more complex, and the last value predictor can explain anomalies. For longer prediction periods up to eight and 16 hours, the difference of AUC of the last value and GNN is around 0.1.

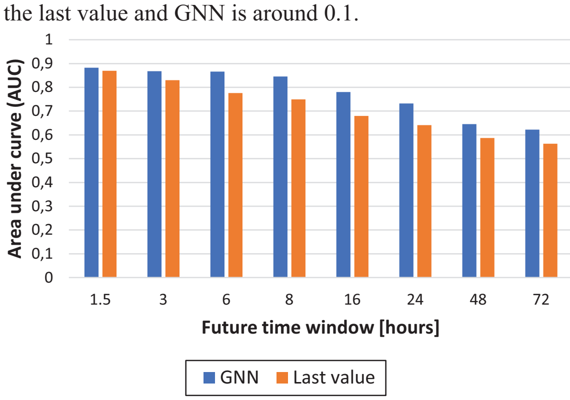

Figure 5: Anomaly prediction results.

V. Impact

The HPC market is central to European industrial development, with an average of €867 of increased revenues and €69 profit for each invested Euro. The European Commission has recognized the importance of HPC in the European data economy and expanded more than €8 billion in the forthcoming Multiannual Financial Framework (2021–2027). The sustainable use of this investment becomes, therefore, a priority.

A. Leading European Supercomputing and Big Data Hub

CINECA is a founding member of the *Bologna Technopole*, aggregating the next-generation supercomputers of CINECA, INFN (Istituto Nazionale di Fisica Nucleare), and ECMWF (European Centre for Medium-Range Weather Forecasts) powered at 60 MW and hosting over 4,000 scientific researchers. The green data center digital twin, researched by Graph-Massivizer, and demonstrated on the pre-exascale Leonardo system operating at approximately 6 MW (10%), provides essential information for HPC facilities' future engineering and sustainable operation. Critical to evidence-based TCO and green practices, capturing graph data of complicated topographical and temporal linkages between cooling equipment, compute nodes, and MG workloads (DC-MG). CINECA and the University of Bologna intend to apply their product in the codesign of the European pilot for exascale [36].

B. Sustainable Green Computing for Improved Societally Relevant "Science Throughput"

CINECA's primary goal is to maximize the societal and economic impact through high computation and scientific throughput per investment. Despite operating over 1,700 million core yearly processor hours, CINECA strives to increase the "science throughput" towards socially relevant fields. Last year, CINECA partnered with national and international researchers to solve urgent societal problems such as COVID-19 drug discovery and energy production from fusion reactors. CINECA uses Graph-Massivizer tools to pursue this goal by modeling the new Leonardo supercomputer as a digital twin and gaining unique insights into its sustainable operation and TCO through 100% capturing of spatial interaction, 10% PUE reduction, 10% lower power use, and 20% improved utilization.

C. Improved data center sustainability.

The Graph-Greenifier tool based on the open-source OpenDC simulator [37] proposes an accurate sustainability-performance analysis of extreme graph processing workloads on large-scale infrastructure, including a sustainability benchmark and green labeling of BGO. The tool aims to demonstrate the possibility of a two-fold improvement in data center energy efficiency and over 25% lower GHG emissions, considering Tier 1 and national suppliers' and public energy data. The tool avoids 'greenwashing' and strengthens Europe's ability to pursue real sustainability in practice by providing quantitative evidence.

VI. Conclusions

This paper presented the effort of the Graph-Massivizer project funded by the Horizon Europe research and innovation program of the European Union, with emphasis on a use case targeting a sustainable data center digital twin operating at exascale performance. The use case receives support from the Graph-Massivizer tools targeting the development of a high-performance and sustainable platform for information processing and reasoning based on the MG representation of extreme data in the form of general graphs, knowledge graphs, and property graphs. Preliminary work on using GNN for prolonged anomaly prediction in HPC systems permits a holistic examination of spatial and temporal connections among system components.

Future research concentrates on expanding the size of GNNs beyond a computing room and investigating diverse

Authorized licensed use limited to: Vrije Universiteit Amsterdam. Downloaded on December 17,2024 at 16:31:18 UTC from IEEE Xplore. Restrictions apply.

forms of interconnections within such graphs to enhance the accuracy, efficiency, and sustainability of anomaly prediction. The Graph-Massivizer toolset accommodates the augmented complexity and magnitude fusing GNN applications on MGs with the ingestion of extreme data and processing analytics developed by the project. The study marks an essential stride towards establishing reliable methods for anticipating anomalies in HPC systems while highlighting the immense potential of massive GNNs in this arena.

ACKNOWLEDGMENT

Graph-Massivizer receives funding from the Horizon Europe research and innovation program of the European Union. Its grant management number is 101093202.

REFERENCES

[1] S. Sakr, A. Bonifati, H. Voigt, A. Iosup, K. Ammar, R. Angles, W. Aref, M. Arenas, M. Besta, P. A. Boncz, K. Daudjee, E. Della Valle, S. Dumbrava, O. Hartig, B. Haslhofer, T. Hegeman, J. Hidders, K. Hose, A. Iamnitchi, V. Kalavri, H. Kapp, W. Martens, M. T. Özsu, E. Peukert, S. Plantikow, M. Ragab, M. R. Ripeanu, S. Salihoglu, C. Schulz, P. Selmer, J. F. Sequeda, J. Shinavier, G. Szárnyas, R. Tommasini, A. Tumeo, A. Uta, A. L. Varbanescu, H.-Y. Wu, N. Yakovets, D. Yan, E. Yoneki, "The future is big graphs: a community view on graph processing systems", Communications of the ACM, vol. 64, pp. 62–71, September 2021.

[2] R. Prodan, D. Kimovski, A. Bartolini, M. Cochez, A. Iosup, E. Kharlamov, Kože Rožanec, L. Vasiliu, A. L. Vărbănescu, "Towards Extreme and Sustainable Graph Processing for Urgent Societal Challenges in Europe", 2022 IEEE Cloud Summit, pp. 23–30, October 2022.

[3] GraphSys '23: The First Workshop on Serverless, Extreme-Scale, and Sustainable Graph Processing Systems, Coimbra, Portugal, April 2023, https://graph-massivizer.eu/events/graphsys23/.

[4] Screening of websites for 'greenwashing': half of green claims lack evidence, Press release, Brussels, January 2021, https://ec.europa.eu/commission/presscorner/detail/en/ip_21_269.

[5] dbt. https://www.getdbt.com/.

[6] Titan. Distributed Graph Database. https://titan.thinkaurelius.com/.

[7] Apache TinkerPop. https://tinkerpop.apache.org/.

[8] G. Malewicz, M. H. Austern, A. JC Bik, J. C. Dehnert, I. Horn, N. Leiser, G. Czajkowski, "Pregel: a system for large-scale graph processing", SIGMOD '10: Proceedings of the 2010 ACM SIGMOD International Conference on Management of data, pp. 135–146, June 2010.

[9] Apache Spark, https://spark.apache.org/graphx/.

[10] R. Zhu, K. Zhao, H. Yang, W. Lin, C. Zhou, B. Ai, Y. Li, J. Zhou, "AliGraph: a comprehensive graph neural network platform", Proceedings of the VLDB Endowment, vol. 12, issue 12, pp 2094–2105, August 2019.

[11] GraphSAGE: Inductive Representation Learning on Large Graphs. http://snap.stanford.edu/graphsage/.

[12] PyTorch Geometric. https://pytorch-geometric.readthedocs.io/.

[13] Apache TinkerPop. https://tinkerpop.apache.org/.

[14] W. L. Hamilton, P. Bajaj, M. Zitnik, D. Jurafsky, J. Leskovec. "Embedding logical queries on knowledge graphs", arXiv:1806.01445, June 2018.

[15] H. Ren, W. Hu, J. Leskovec. "Query2box: reasoning over knowledge graphs in vector space using box embeddings", arXiv:2002.05969, February 2020.

[16] D. Daza, M. Cochez. "Message passing query embedding", arXiv:2002.02406, June 2020.

[17] B. Kotnis, C. Lawrence, M. Niepert, "Answering complex queries in knowledge graphs with bidirectional sequence encoders", arXiv:2004.02596, February 2021.

[18] E. Arakelyan, D. Daza, P. Minervini, M. Cochez. "Complex query answering with neural link predictors", arXiv:2011.03459, March 2021.

[19] Y. Guo, M. Biczak, A. L. Varbanescu, A. Iosup, C. Martella, T. L. Willke, "How well do graph-processing platforms perform? An empirical performance evaluation and analysis", 2014 IEEE 28th International Parallel and Distributed Processing Symposium, pp. 395–404, May 2014.

[20] Y. Lu, J. Cheng, D. Yan, H. Wu, "Large-scale distributed graph computing systems: an experimental evaluation", Proceedings of the VLDB Endowment, vol. 8, issue 3, pp. 281–292, November 2014.

[21] A. Iosup, T. Hegeman, W. L. Ngai, S. Heldens, A. Prat-Pérez, T. Manhardto, H. Chafio, M. Capotă, N. Sundaram, M. Anderson, I. G. Tănase, Y. Xia, L. Nai, P. Boncz, "LDBC graphalytics: a benchmark for large-scale graph analysis on parallel and distributed platforms", Proceedings of the VLDB Endowment, vol. 9, issue 13, pp. 1317–1328, September 2016.

[22] F. McSherry, M. Isard, D. G. Murray, "Scalability! but at what cost?", HOTOS'15: Proceedings of the 15th USENIX conference on Hot Topics in Operating Systems, pp.14, May 2015.

[23] A. Uta, S. Au, A. Ilyushkin, A. Iosup, "Elasticity in graph analytics? A benchmarking framework for elastic graph processing", 2018 IEEE International Conference on Cluster Computing (CLUSTER), pp. 381–391, September 2018.

[24] Bellewitte. https://github.com/merijn/Belewitte.

[25] X. Wu, A. Marathe, S. Jana, O. Vysocky, J. John, A. Bartolini, L. Riha, M. Gerndt, V. Taylor, S. Bhalachandra, "Toward an End-to-End Auto-tuning Framework in HPC PowerStack", 2020 IEEE International Conference on Cluster Computing (CLUSTER), pp. 473–483, September 2020.

[26] H. Shoukourian, T. Wilde, H. Huber, A. Bode, "Analysis of the efficiency characteristics of the first high-temperature direct liquid cooled petascale supercomputer and its cooling infrastructure", Journal of Parallel and Distributed Computing, vol. 107, pp. 87–100, September 2017.

[27] J. Schleier-Smith, V. Sreekanti, A. Khandelwal, J. Carreira, N. J. Yadwadkar, R. Ada Popa, J. E. Gonzalez, I. Stoica, D. A. Patterson, "What serverless computing is and should become: the next phase of cloud computing", Communications of the ACM, vol. 64, issue 5, pp. 76–84, May 2021.

[28] Apache OpenWisk. https://openwhisk.apache.org/.

[29] L. Toader, A. Uta, A. Musaafir, A. Iosup, "Graphless: toward serverless graph processing", 2019 18th International Symposium on Parallel and Distributed Computing (ISPDC), pp. 66–73, June 2019.

[30] M. Szalay, P. Mátray, L. Toka, "Real-time task scheduling in a FaaS cloud", 2021 IEEE 14th International Conference on Cloud Computing (CLOUD), pp. 497–507, September 2021.

[31] A. Bartolini, F. Beneventi, A. Borghesi, D. Cesarini, A. Libri, L. Benini, C. Cavazzoni, "Paving the Way Toward Energy-Aware and Automated Datacentre", ICPP Workshops '19: Workshop Proceedings of the 48th International Conference on Parallel Processing, article 8, pp. 1–8, August 2019.

[32] A. Borghesi, A. Burrello and A. Bartolini, "ExaMon-X: A Predictive Maintenance Framework for Automatic Monitoring in Industrial IoT Systems," in IEEE Internet of Things Journal, vol. 10, no. 4, pp. 2995–3005, February 2023.

[33] A. Borghesi, M. Molan, M. Milano, A. Bartolini, "Anomaly detection and anticipation in high-performance computing systems", IEEE Transactions on Parallel and Distributed Systems, vol. 33, issue 4, pp. 739–750, April 2022.

[34] M. Molan, A. Borghesi, D. Cesarini, L. Benini, A. Bartolini, "RUAD: Unsupervised anomaly detection in HPC systems", Future Generation Computer Systems, vol. 141, pp. 542–554, April 2023.

[35] M. Molan, J. A. Khan, A. Borghesi, A. Bartolini, "Graph neural networks for anomaly anticipation in HPC systems", ICPE '23 Companion: Companion of the 2023 ACM/SPEC International Conference on Performance Engineering, pp. 239–244, April 2023.

[36] EUPEX, European pilot for exascale, https://eupex.eu/.

[37] F. Mastenbroek, G. Andreadis, S. Jounaid, W. Lai, J. Burley, J. Bosch, E. van Eyk, L. Versluis, V. van Beek, A. Iosup. OpenDC 2.0: Convenient Modeling and Simulation of Emerging Technologies in Cloud Datacenters. 2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid), pp. 226–236, May 2021.

Authorized licensed use limited to: Vrije Universiteit Amsterdam. Downloaded on December 17,2024 at 16:31:18 UTC from IEEE Xplore. Restrictions apply.