FALL 2022

Newsletter

The Global Computing Lab (GCLab) is located at the University of Tennessee and focuses on high performance computing (HPC) and its interdisciplinary functions across the sciences. As director of the GCLab, I work alongside a talented group of emerging researchers on a variety of projects that promotes the efficiency of computational algorithms and adaptive scheduling policies for graphics processing units (GPUs), cloud and volunteering computing.

Collaborating with scientists in the fields of chemistry, chemical engineering, pharmaceutical sciences, seismology and mathematics gives our team diverse and substantial opportunities to put our dynamic problem solving expertise to use. Our team consists of post-doctoral researchers who bring their skillful experience to the lab. Our students, made up of undergraduates, masters and PhDs, examine innovative approaches for applying solutions that challenge today's fast-paced environment of high performance computing.

HPC is a field that continues to make a significant impact across the sciences, as the demands placed on technology exponentially grows. GCLab strives to be at the forefront of these advances, leading the way with our visionary direction, and training the scientists of tomorrow to be prepared for what is to come next in computer science.

This special edition newsletter highlights one of our ongoing research projects and provides a listing of recent publications generated by the GCLab. It also gives readers an opportunity to meet our team members and learn about the lab's accomplishments, including awards and recognition for the work we do.

With this newsletter, I hope you get to know more about what we do here at GCLab and who we are as researchers.

Michela Taufer
Jack Dongarra Professor in High Performance Computing
The University of Tennessee
Electrical Engineering and Computer Science Dept.

FIND US @

@TauferLab

https://globalcomputing.group

Project Spotlight

National Science Data Fabric: A Platform Agnostic Testbed for Democratizing Data Delivery
Source of Support: NSF #2138811
Total Amount: $5,609,259 ($750,000 at University of Tennessee Knoxville), co-PI.
Project Period: 10/01/2021 - 09/30/2024
http://nsdf.sci.utah.edu

The National Science Data Fabric (NSDF) is a multi-institutional project with the objective to create a platform to manage the abundance of data-driven science. Funded by NSF and led by researchers at the University of Utah, the University of Tennessee, along with other academic institutions, NSDF has formed partnerships with a diverse pool of research institutions that output a high-volume of scientific information.

Investments from research agencies, including those from the National Science Foundation, have pathed the way to the creation of shared experimental and computing facilities, including those dedicated to synchrotron light sources, Lepton and Hadron colliders, and multispectral satellite imagery technology, to name a few. Past funding models and lack of cohesiveness between application domains have created a challenging scientific environment where data depositories and computational workflows are siloed into specific criterias, making cross-interdisciplinary integration an after-thought.

As this reliance on data-driven sciences continues to increase, this absence of integration imposes limitations on the scientists to effectively manage their output. Previous national reports have emphasized the need of adopting a manageable approach to solve this issue.

Identifying the growing increase of generated data coupled with the lack of integration among the scientific community, has created a need for a CI platform with equitable access to the science and its data. Researchers involved with NSDF intend to connect storage, compute and network elements using a software stack that provides scalable tools for end-users that are easy to use, integrate and scale. The additional layer of community-driven education and outreach will promote equitable access and engagement for an open network between universities, including minority-serving institutions.

With the creation of a National Science Data Fabric (NSDF), a trans-disciplinary approach to integrate data delivery with the ability for easy accessibility, is in development for a scientific environment where data is everything and everywhere.

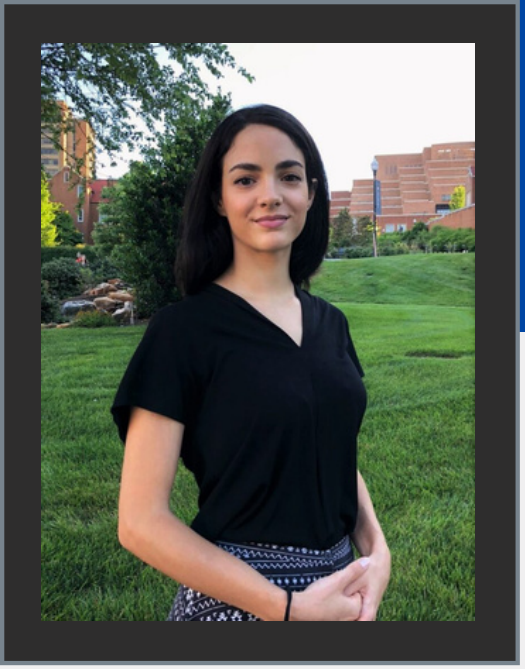

Researcher's Corner

with

Silvina Caino-Lores

Research Assistant Professor

After earning her Ph.D. in Computer Science and Technology at the Carlos III University of Madrid in Spain in 2019, Silvina joined the GC Lab as a Postdoctoral Research Associate. Most recently, Silvina was promoted to the position of Research Assistant Professor at the university.

We were pleased that she was able to share some insight as to her curious nature, and where she would like to venture in her future as a researcher.

What got you interested in computer science?

I was very curious about computers in general. I perceived them as magical machines, especially after playing around with LOGO (a drawing programming language) in school. Eventually I started tearing computers apart to see how they are built, and I got familiar with Linux and its infinite tinkering potential.

What is the main project you are focusing on currently?

Most of my work revolves around workflows and the technologies that enable them to tackle increasingly challenging problems in different scientific domains. My main project aims to accelerate knowledge discovery in molecular dynamics through modular in situ workflows. Our workflows allow domain scientists to select the methods that best fit a particular research problem and easily build a workflow that integrates these components. With this workflow, scientists can better understand the behavior of their simulations at runtime and make smart orchestration and simulation steering decisions.

What do you think are the most exciting areas in computer science today?

My first research work was in the area of convergence of high-performance computing and cloud computing from the perspective of the application. This was approximately eight years ago, and in those days the HPC community wasn't very supportive of these non-traditional approaches to scalability. The Big Data community -with much more influence from industry- was eager to find ways to accelerate their tasks, and eventually this influence transformed the landscape of HPC at all levels. I am very curious to observe the evolution of HPC in the context of converged computing. I see a lot of potential to disrupt all kinds of applications and I expect HPC to gain a lot of visibility and general public recognition as a result of this exposure.

On a more futuristic end, I have always been very interested in unconventional computing architectures. I suspect that, in the eternal quest for faster, larger and more power-efficient systems, we will see specialized accelerators and storage devices built with unconventional methods. I find biologically inspired architectures particularly exciting. In my opinion, neuromorphic electronics are already here to stay, and organic computing units are the natural next step towards self-organizing artificial intelligence.

Tell us something about yourself that is not computer science related.

I work with screens and ideas, so I like balancing the rest of my time getting in touch with the world and creating tangible things. I've got many hobbies, but cats are my weakness! I volunteer to look after cats in need as much as I can.

GCL

Awards & Recognitions

2022

Michela Taufer elected to Vice-Chair of SIGHPC

An international group within a major professional society, SIGHPC's mission is to help spread the use of HPC, raise the standards of the profession, and ensure a rich and rewarding career for people involved in the field. Michela will commence her term as Vice-Chair on July 1st.

IEEE Technical Community on Parallel Processing (TCPP) Outstanding Service and Contributions Award

Michela Taufer was recently recognized this past May by the IEEE, the world's largest technical professional organization for the advancement of technology, for her professional service contributions to the parallel, distributed, and high-performance computing community.

Tickle College of Engineering Research Achievement Award, University of Tennessee, Knoxville

Michela Taufer became the latest recipient of this distinguished faculty award, with a criteria focused on outstanding achievement and recognition in their respective research field on a national/international level.

2021

IBM Global University Program Academic Award

Congratulations to Michela Taufer, a 2021 recipient of the IBM Global University Program Academic Award. She received the Academic Award honor, in recognition for her collaborative research projects that promote scientific advancements with the technology community at-large.

Best Thematic Track Paper Award at the International Conference on Computational Science (ICCS)

Optimize Memory Usage in Vector Particle-In-Cell (VPIC) to Break the 10 Trillion Particle Barrier in Plasma Simulations (AIHPC4AS track) was selected as the best thematic track paper at the 2021 ICCS. Nigel Tan, Ph.D. Candidate and Michela Taufer were among the paper's authors that also included Robert Bird and Guangye Chen.

GCLab
Dr. Michela Taufer

Meet our Team

Silvina Caino-Lores, Ph.D., Research Assistant Professor

Jakob Luettgau, Ph.D., Postdoctoral Researcher
Jack Marquez, Ph.D., Postdoctoral Researcher
Ariel Rorabaugh, Ph.D., Postdoctoral Researcher
Naweiluo Zhou, Ph.D., Postdoctoral Researcher

Nick Bell, Research Scientist
Georgia Channing, Research Scientist

Ian Lumsden, Doctoral Student
Paula Olaya, Doctoral Student
Nigel Tan, Doctoral Student

Treece Burgess, Master's Student
Vanessa Lama, Master's Student
Kae Suarez, Master's Student

Cole Johnston, Undergraduate Student
Dominic Kennedy, Undergraduate Student
Jacob Leonard, Undergraduate Student
Ria Patel, Undergraduate Student
Lauren Proctor, Undergraduate Student
Brandan Roachell, Undergraduate Student

Barbara Fossum, Outreach Coordinator
Lauren Whitnah, Technical Writer
Grace Wisser, Project Coordinator

GCLab actively collaborates with the
Innovative Computing Laboratory (ICL) at UTK.

RECENT

Publications

2021-22 Journal Publications

Tu Mai Anh Do, Lo¨ic Pottier, Rafael Ferreira da Silva, Silvina Caino-Lores, Michela Taufer, and Ewa Deelman. Performance Assessment of Ensembles of In Situ Workflows under ResourceConstraints. Journal of Concurrency and Computation: Practice and Experience (CCPE), 2022. (Accepted).

Nigel Tan, Robert Bird, Guangye Chen, Scott V. Luedtke, Brian Albright, and Michela Taufer. Analysis of Vector Particle-In-Cell (VPIC) Memory Usage Optimizations on Cutting-Edge Computer Architectures. Journal of Computational Science, 2022. (Accepted).

Michael R. Wyatt II, Stephen Herbein, Todd Gamblin, and Michela Taufer. AI4IO: A Suite of AI-Based Tools for IO-Aware Scheduling. International Journal of High Performance Computing Applications (IJHPCA), 2022. (Accepted).

Ariel Keller Rorabaugh, Silvina Caino-Lores, Travis Johnston, and Michela Taufer. Building High-throughput Neural Architecture Search Workflows via a Decoupled Fitness Prediction Engine. IEEE Trans. Parallel Distributed Syst. (TPDS), 2022. (Accepted).

Stephen Herbein, Tapasya Patki, Dong H. Ahn, Sebastian Mobo, Clark Hathaway, Silvina Caino-Lores, James Corbett, David Domyancic, Thomas R. W. Scogland, Bronis R. de Supinski,and Michela Taufer. An Analytical Performance Model of Generalized Multi-Level Scheduling.International Journal of High Performance Computing Applications (IJHPCA), 2022. (Accepted).

Nicholas C. Mucci, Katarina A. Jones, Mengyi Cao, Michael R. Wyatt, Shane Foye, Sarah J. Kauffman, Gregory R. Richards, Michela Taufer, Yoshito Chikaraishi, Shawn A. Steffan, Shawn R. Campagna, Heidi Goodrich-Blair, and Christopher R. Anderton. Apex Predator Nematodes and Meso-Predator Bacteria Consume Their Basal Insect Prey through Discrete Stages of Chemical Transformations. mSystems, 0(0):e00312-22, 2022.

Ariel Keller Rorabaugh, Silvina Caino-Lores, Travis Johnston, and Michela Taufer. High Frequency Accuracy and Loss Data of Random Neural Networks Trained on Image Datasets. Datain Brief, Elsevier, 40:107780, 2022.

Robert F. Bird, Nigel Tan, Scott V. Luedtke, Stephen Lien Harrell, Michela Taufer, and Brian J. Albright. VPIC 2.0: Next generation particle-in-cell simulations. IEEE Trans. Parallel Distributed Syst., 33(4):952-963, 2022.

GCL

Publications

2021-22 Journal Publications

Patrick Bell, Kae Suarez, Dylan Chapp, Nigel Tan, Sanjukta Bhowmick, and Michela Taufer. ANACIN-X: A Software Framework for Studying Non-Determinism in MPI Applications. Software Impacts, Elsevier, 10:100151, 2021.

Duncan A. Brown, Karan Vahi, Michela Taufer, Von Welch, Ewa Deelman, Lorena A. Barba, and George K. Thiruvathukal. Reproducing GW150914: The First Observation of Gravitational Waves From a Binary Black Hole Merger. IEEE Comput. Sci. Eng. (CiSE), 23(2):73-82, 2021.

Tu Mai Anh Do, Loic Pottier, Silvina Caıno-Lores, Rafael Ferreira da Silva, Michel A. Cuendet, Harel Weinstein, Trilce Estrada, Michela Taufer, and Ewa Deelman. A Lightweight Method for Evaluating in situ Workflow Efficiency. J. Comput. Sci. Elsevier, 48:101259, 2021.

Hector Carrillo-Cabada, Jeremy Benson, Asghar M. Razavi, Brianna Mulligan, Michel A. Cuendet, Harel Weinstein, Michela Taufer, and Trilce Estrada. A Graphic Encoding Method for Quantitative Classification of Protein Structure and Representation of Conformational Changes. IEEE ACM Trans. Comput. Biol. Bioinform., 18(4):1336-1349, 2021.

Dylan Chapp, Nigel Tan, Sanjukta Bhowmick, and Michela Taufer. Identifying Degree and Sources of Non-Determinism in MPI Applications Via Graph Kernels. IEEE Trans. Parallel Distributed Syst. (TPDS), 32(12):2936-2952, 2021.

Mario Guevara, Michela Taufer, and Rodrigo Vargas. Gap-free Global Annual Soil Moisture: 15 km grids for 1991-2018. Earth System Science Data, 13(4):1711-1735, 2021.

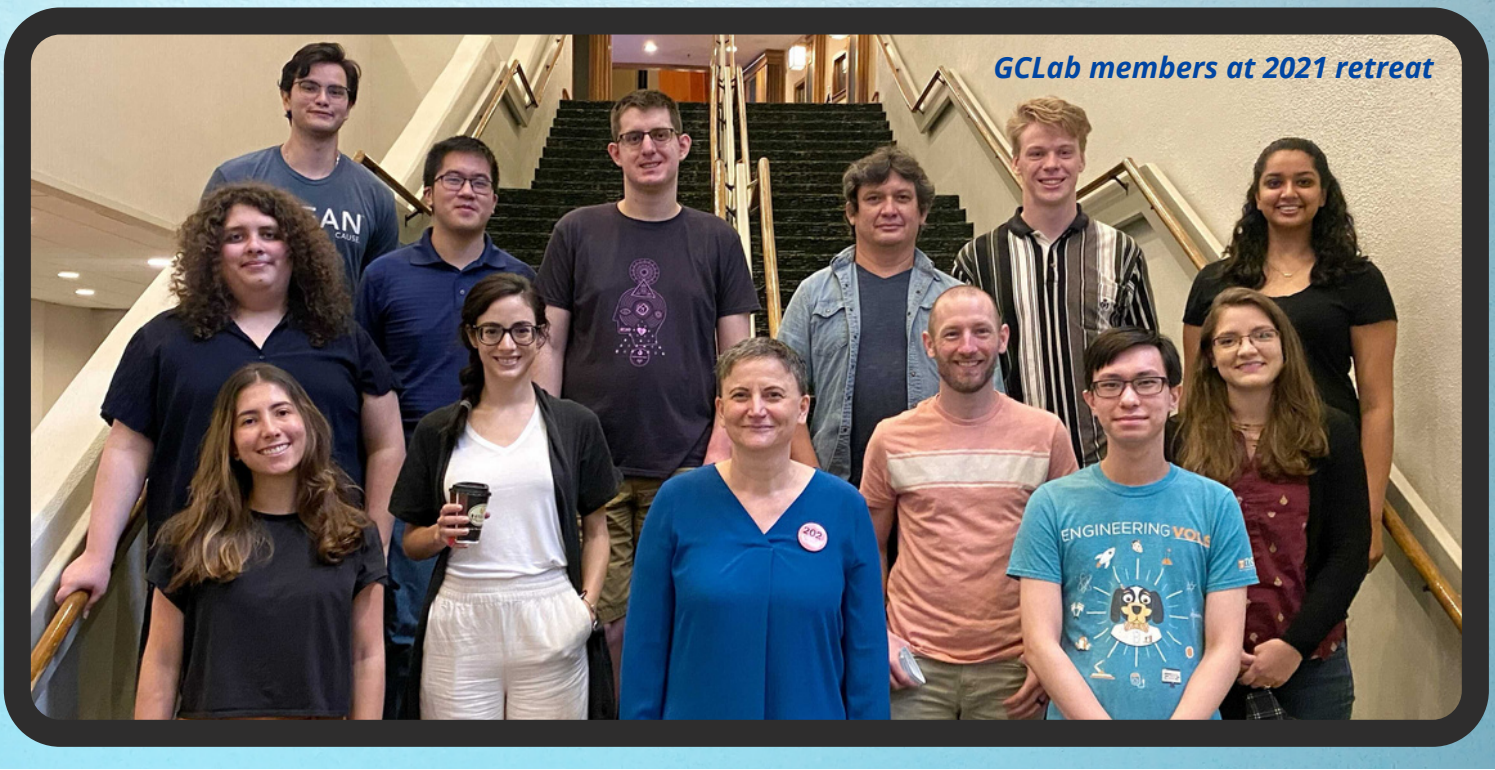

RECENT

GCLab
Dr. Michela Taufer

Publications

2021-22 Conference Publications

Paula Olaya, Jakob Luettgau, Naweiluo Zhou, Giorgio Scorzelli, Jay Lofstead, Valerio Pascucci, and Michela Taufer. NSDF-FUSE: A Testbed for Studying Object Storage via FUSE File Systems. In Proceedings of the 31st International ACM Symposium on High-Performance Parallel andDistributed Computing (HPDC), pages 1–2, Minneapolis, Minnesota, June 2022. IEEE Computer Society. (Short paper).

Jakob Luettgau, Paula Olaya, Naweiluo Zhou, Giorgio Scorzelli, Valerio Pascucci, and Michela Taufer. NSDF-Cloud: Enabling Ad-Hoc Compute Clusters Across Academic and Commercial Clouds. In Proceedings of the 31st International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC), pages 1–2, Minneapolis, Minnesota, June 2022. ACM. (Short paper).

Michela Taufer, Ewa Deelman, Rafael Ferreira da Silva, Trilce Estrada, Mary Hall, and Miron Livny. A Roadmap to Robust Science for High-throughput Applications: The Developers' Perspective. In Proceedings of the IEEE Cluster Conference (CLUSTER), pages 1–2, Portland, Oregon, September 2021. IEEE Computer Society. (Short paper).

Ross Ketron, Jacob Leonard, Brandan Roachell, Ria Patel, Rebecca White, Silvina Caino-Lores, Nigel Tan, Patrick Miles, Karan Vahi, Ewa Deelman, Duncan A. Brown, and Michela Taufer. A Case Study in Scientific Reproducibility from the Event Horizon Telescope (EHT). In Proceedings of the 20th IEEE International Conference on eScience, pages 1–2, Innsbruck, Austria, September 2021. IEEE Computer Society. (Short paper).

Michela Taufer, Ewa Deelman, Rafael Ferreira da Silva, Trilce Estrada, and Mary Hall. A Roadmap to Robust Science for High-throughput Applications: The Scientists' Perspective. In Proceedings of the 20th IEEE International Conference on eScience, pages 1–2, Innsbruck, Austria, September 2021. IEEE Computer Society. (Short paper).

Tu Mai Anh Do, Loic Pottier, Rafael Ferreira da Silva, Silvina Caino-Lores, Michela Taufer, and Ewa Deelman. Assessing Resource Provisioning and Allocation of Ensembles of In Situ Workflows. In Proceedings of the Fourteen International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), Chicago, Illinois, USA, August 2021. IEEE Computer Society.