Julian Bellavita

Computer Science PhD Student at Cornell University · DOE CSGF Fellow

jbellavita@cs.cornell.edu | https://github.com/Hooninator | https://www.linkedin.com/in/julian-bellavita-5876551ba/

Education

Cornell University *Ithaca, NY*

PhD, Computer Science *08/2023 - present*

- Advisor: Dr. Giulia Guidi
- Concentration: Scientific Computing
- GPA: 3.9 + $\mathcal{O}(\mu)$

University of California, Berkeley *Berkeley, CA*

BA, Computer Science *08/2019 - 05/2023*

- Major GPA: 3.97

Research Experience

Graduate Researcher, Cornell University *Ithaca, NY*

Advisor: Dr. Giulia Guidi *08/2023 - Present*

Large-scale multi-GPU clustering with sparse linear algebra. Parallelizing sparse linear algebra kernels. Randomized CP decomposition.

Graduate Researcher, Oak Ridge National Laboratory *Oak Ridge, TN*

Advisor: Dr. Ramakrishnan Kannan *05/2025 - 08/2025*

Mixed precision tensor processing kernels on GPUs.

Visiting PhD Student, University of Trento *Trento, Italy*

Advisor: Dr. Flavio Vella *01/2024 - 08/2024*

Kernel k-means on GPUs with sparse linear algebra.

Undergraduate Researcher, University of California, Berkeley *Berkeley, CA*

Advisor: Dr. James Demmel *05/2022 - 05/2023*

Generating BLAS 3 kernels for vector architectures via user-scheduled compilers.

Student Research Assistant, Lawrence Berkeley National Lab *Berkeley, CA*

Pagoda Project – CLaSS Group *2022-2023*

Task-based formulation of sparse Cholesky decomposition with GPU acceleration.

Student Research Assistant, Lawrence Berkeley National Lab *Berkeley, CA*

Scientific Data Management Group *2021-2023*

Studying access patterns of high-energy physics datasets on distributed storage systems.

Publications

Julian Bellavita, Thomas Pasquali, Laura Del Rio Martin, Flavio Vella, and Giulia Guidi. 2025. ”Popcorn: Accelerating Kernel Kmeans on GPUs through Sparse Linear Algebra”. In The 30th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming **(PPoPP 2025)**.

Thomas McFarland, **Julian Bellavita**, and Giulia Guidi. 2025. ”Parallel GPU-Enabled Algorithms for SpGEMM on Arbitrary Semirings with Hybrid Communication”. Short paper. In Proceedings of the 16th ACM/SPEC International Conference on Performance Engineering **(ICPE 2025)**.

Adrián Castelló, **Julian Bellavita**, Grace Dinh, Yuka Ikarashi, Hector Martínez. ”Tackling the Matrix Multiplication Micro-Kernel Generation with Exo.” In IEEE/ACM International Symposium on Code Generation and Optimization (**CGO 2024**).

Julian Bellavita, Mathias Jacquelin, Esmond G. Ng, Dan Bonachea, Johnny Corbino, and Paul H. Hargrove. ”symPACK: A GPU-Capable Fan-Out Sparse Cholesky Solver.” In Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis (**SC-W 2023**).

Julian Bellavita, Caitlin Sim, Kesheng Wu, Alex Sim, Shinjae Yoo, Hiro Ito, Vincent Garonne, Eric Lancon, ”Understanding Data Access Patterns for dCache System.” In 26th International Conference on Computing in High Energy & Nuclear Physics (**CHEP 2023**)

Julian Bellavita, Alex Sim, Kesheng Wu, Inder Monga, Chin Guok, Frank Würthwein, and Diego Davila. 2022. ”Studying Scientific Data Lifecycle in On-demand Distributed Storage Caches.” In Fifth International Workshop on Systems and Network Telemetry and Analytics (**SNTA 2022**).

Awards

Fellowships

2024 **DOE Computational Science Graduate Fellowship (DOE CSGF)**, United States Department of Energy

2023 **Cornell Fellowship**, Cornell University

Awards

2024 **Best Student Presentation**, The 2nd CINI Summer School on. High Performance Computing and Emerging Technologies

2022 **2nd Place**, ACM/IEE Student Research Competition, Undergraduate Division

Grants

2025 **NSF Student Travel Grant**, The 30th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming (PPoPP 2025)

Presentations

Talks

"Block Leverage Scores for Cache-Friendly Randomized CP Decomposition", RNLA Workshop at IPAM, 2025

"Algorithms for Computing a Tensor Times Matrix Chain in Mixed Precision", Oak Ridge National Laboratory Discrete Algorithms Group Seminar, 2025

"Multi-GPU Communication Schemes for Large-Scale Supercomputers", Cornell Systems Seminar, 2024

"Accelerating High-Dimensional K-Means Clustering on GPUs with Sparse Matrix Multiplication", University of Trento, 2024

"RDMA-Based Algorithms for Sparse Matrix Multiplication on GPUs", Cornell HPC Group, 2024

"Portable code generation and semi-automatic scheduling for BLIS microkernels with Exo", BLIS Retreat, 2022

Posters

"Forward Error Bounds and Efficient Algorithms for Computing a Tensor Times Matrix Chain in Low Precision on GPUs", ACM Student Research Competition at SC, 2025

"Mixed Precision Algorithms for Computing the Tucker Decomposition", CSGF Program Review, 2025

"Efficient Large-Scale Multi-GPU Clustering using Sparse Linear Algebra", IPDPS PhD Form, 2025

Mentoring

2025-Now **Pablo Raigoza**, MEng Student, Cornell University
2024-Now **Thomas McFarland**, Undergraduate, Cornell University
2025-2026 **Nolan Lizmi**, Undergraduate, Cornell University
2025 **Alexander Schatzberg**, Undergraduate, Cornell University
2024-2025 **Matthew Rubino**, MEng Student, Cornell University
2024-2025 **Nakul Iyer**, MEng Student, Cornell University
2024-2025 **Andrew Chang**, MEng Student, Cornell University
2024-2025 **Noam Benson-Tilsen**, Undergraduate, Cornell University

Services

2025 **SPAA 2025**, Junior Program Committee Member
2025 **CGO 2025**, Artifact Review Committee Member