



 EDITORIAL

In this issue:

Website	2
Social Media	2
Promotional Materials	2
Past Events	3
Communication, Dissemination and Outreach	4
Publications	4
Resources	5
Upcoming Events	6
More information	6

Dear readers,

Launched in April 2021, this 3-year project is funded by the European High Performance Computing Joint Undertaking (EuroHPC JU) under the 2019 call of Extreme Scale Computing and Data Driven Technologies for research and innovation actions.

The SparCity project aims at creating a supercomputing framework that is providing efficient algorithms and coherant tools specifically designed for maximizing the performance and

energy efficiency of sparse computations on emerging HPC systems, while also opening up new usage areas for sparse computations in data analytics and deep learning.

SparCity delivers a coherent collection of innovative algorithms and tools for enabling both high efficiency of sparse computations on emerging hardware platforms.

The SparCity Coordination,
Didem Unat



In-person Meetings | June 1-3, 2022 (Istanbul, Turkey) and June 15-16, 2023 (Lisbon, Portugal)

PAST EVENTS


Collaboration Meeting | June 6-7, 2023

SparCity was present at the EuroHPC Collaboration Meeting in Turin on June 6-7 to discuss the collaboration plans raised among projects (joint workshops, BoFs and publications). Prof. Didem Unat presented the SparCity project and its latest results.



In-person Meeting | June 15-16, 2023

On June 15 and 16, the SparCity members joined together in Lisbon for a two-day Meeting to share and discuss the progress made on each work package and plan the upcoming months.




Workshop on Tools for Data Locality, Power and Performance (TDLPP 2023) | August 29, 2023

This workshop was co-located with [Euro-Par 2023](#), in Limassol, Cyprus. The goal of the TDLPP workshop was

to provide a venue for developers and users of tools that address the important topic of memory access optimization. While hardware continues to evolve and high-bandwidth memory becomes available in accelerators and mainstream CPUs, the gap between compute capability (in terms of arithmetic operations per second) and the speed of memory (in terms of access latency or amount of bytes transferred) continues to widen. Tools are thus needed to help developers understand the behavior of their codes to support them with optimizing and modeling their applications. This is especially true in application areas that involve sparse matrices, tensors, or graphs.

The SparCity team (Fatih Taşyaran, Sabancı University; Osman Yasal, Koç University and José António Carvalho Freire Morgado, INESC-ID) conducted a Tool Demo Session (in-person and online) about SuperTwin: A Digital Twin for HPC Machines.



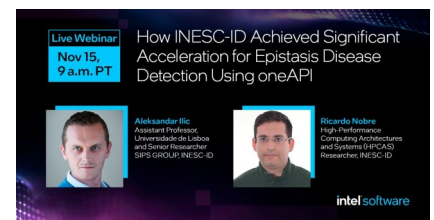
Code Together Podcast: Accelerating Epistasis Detection - How oneAPI Supports Genetics Researchers | October 16, 2023

On October 16, the INESC-ID team members Prof, Aleksandar Ilic and Dr. Ricardo Nobre, participated in the “Code Together” Podcast to talk about how their work INESC-

ID helps with Epistasis Disease detection and how oneAPI solutions speed up those solutions by 9x.

Intel Software Webinar | November 15, 2023

The SparCity team members Aleksandar Ilic and Ricardo Nobre from INESC-ID gave a webinar to explain how oneAPI on Intel Xeon CPU Max Series sped up INESC-ID’s bioinformatics application.



Supercomputing 2023 | November 12-17, 2023

SparCity was present in SC 2023, the International Conference for High Performance Computing, Networking, Storage, and Analysis, which was held in Denver, Colorado.



At the beginning of the week, Sergej Breiter from LMU Munich gave a presentation on Modeling Data Locality of Sparse Matrix-Vector Multiplication on the A64FX. James D Trotter (Simula Research Laboratory) had a talk on Bringing Order to Sparsity: A Sparse Matrix Reordering Study on Multicore CPU.

◆ Workshop on Future is Sparse

The SparCity Workshop on Future is Sparse, co-located with SC23, was held on November in Denver!

We would like to thank the speakers: Professors John Owens, Richard Vuduc, Hartwig Anzt and Tal Ben-Nun, for the wonderful talks.

We also want to thank the more than 65 participants who joined us during this half-day to discuss the latest research developments and outline major open problems in the field of HPC and AI.

COMMUNICATION, DISSEMINATION & OUTREACH

◆ #MeetTheTeam and #CodeVideos Series

New videos were published where you can learn more about a different team member and the work developed within SparCity.

 [#MeetTheTeam videos](#)
[#CodeVideo videos](#)

◆ [Women Scientists of Future](#) | June 12, 2023

Dr Didem Unat attended a panel about Women Scientists of Future, targeting female students at Turkish Universities.

◆ Annual EUGAIN Workshop | September 20, 2023



The SparCity PI Didem Unat gave a talk about “Breaking Barriers and Building Bridges: Advancing Gender Equality in Academia” at the Annual EUGAIN Workshop collocated with ACM’s womENCourage 2023 Conference in Norway!



◆ MAELSTROM Workshop | November 7, 2023

The workshop on MAELSTROM Dissemination was a fruitful session to know more about the work developments on the EuroHPC partner projects. Dr Tugba Torun presented the results on our project.

PUBLICATIONS

◆ Alexandre Rodrigues, Leonel Sousa and Aleksandar Ilic (2023). Performance Modelling-driven Optimization of RISC-V Hardware for Efficient SpMV. Proceedings of the International workshop on RISC-V for HPC. ISC High Performance Conference. DOI:[10.1007/978-3-031-40843-4_36](https://doi.org/10.1007/978-3-031-40843-4_36).

◆ Miguel Graça, Diogo Marques, Sergio Santander-Jiménez, Leonel Sousa, and Aleksandar Ilic (2023). Interpreting High Order Epistasis Using Sparse Transformers. In 2023 IEEE/ACM Conference on Connected Health: Applications, Systems and Engineering Technologies (CHASE), pp. 114-125. DOI: [10.1145/3580252.3586982](https://doi.org/10.1145/3580252.3586982).

◆ James D. Trotter, Johannes Langguth and Xing Cai (2023). Targeting performance and user-friendliness: GPU-accelerated finite element computation with automated code generation in FEniCS. Parallel Computing, vol. 118. DOI: [10.1016/j.parco.2023.103051](https://doi.org/10.1016/j.parco.2023.103051).

◆ James D. Trotter, Sinan Ekmeçibaşı, Johannes Langguth, Tugba Torun, Emre Düzakın, Aleksandar Ilic, and Didem Unat (2023). Bringing Order to Sparsity: A Sparse Matrix Reordering Study on Multicore CPUs. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '23). Association for

Computing Machinery, New York, NY, USA, Article 31, 1–13. DOI: [10.1145/3581784.3607046](https://doi.org/10.1145/3581784.3607046).

◆ Sergej Breiter, James D. Trotter, and Karl Furlinger (2023). Modelling Data Locality of Sparse Matrix-Vector Multiplication on the A64FX. In Proceedings of the SC '23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis (SCW '23). Association for Computing Machinery, New York, NY, USA, 1334–1342. DOI: [10.1145/3624062.3624198](https://doi.org/10.1145/3624062.3624198).


RESOURCES

SparseBase

SparseBase provides a comprehensive framework that can enhance any sparse data workflow. The library encompasses I/O, preprocessing, experimentation, and analysis into a single, optimized, and easy-to-use API. Its simple interface and its data type and representation agnosticism mean it can fit into any and every workflow involving sparse data with minimal coding overhead.

SparseBase provides reading and writing to the most common sparse data file formats, both from disk and from online repositories like SuiteSparse and FROSTT. It also contains many state-of-the-art algorithms for sparse data preprocessing operations including reordering, partitioning, and clustering. Additionally, it supports multiple architectural settings in addition to multi-core CPUs, e.g., GPUs. All the aforementioned functionalities are implemented optimally to accelerate sparse data workflows. Sparse data has a very unique property in that it can be stored in many physical representations, with each being more optimal for certain operations. SparseBase enables defining multiple implementations for each functionality, each optimized for a certain data representation.

Additionally, operations in SparseBase come bundled with a function-matching utility that takes users' input data and carries out any needed format conversions. This allows users to work with a sparse object as an abstraction, while still leveraging the many representations it can take under the hood.



The figure above illustrates the function-matching capability. An input object is represented in Coordinate (COO) format and is passed to the Reverse Cuthill-McKee (RCM) reordering module. The latter recognizes that the input format isn't supported for this reordering operation, and converts the input to a supported format, which, in this case, is the Compressed Sparse Row (CSR) format. This enables the users to access the full spectrum of sparse data functionalities in an optimal manner, and without needing to explicitly manage data formats.

SparseBase is bundled with an experimentation pipeline to examine how a certain workflow can benefit from different types of preprocessing techniques. Users specify data sources and preprocessing algorithms, and supply their workflows as kernels, and the library executes a complete experiment using all the provided modules.

The user provides multiple input objects which the library can read from disk and from online repositories. They also specify which reordering algorithms they wish to experiment with, and finally, they pass their own workflow kernels, which, in this case, are different sparse matrix-vector multiplication implementations. The experimentation pipeline executes the powerset of the passed data,

preprocesses, and kernels, and returns timings, analytical results, and auxiliary data back to the user. Moreover, to allow further exploration of the sparse data space, the library provides an engine for extracting features from sparse data. The engine includes fused kernels that extract multiple features concurrently for added optimality.



<https://github.com/sparcityeu/sparsebase>

Libmtx

Libmtx is a useful complement to the SparseBase framework. While SparseBase is written in C++ and particularly well suited for integration into C++-based application, Libmtx is written in plain C and can be easily incorporated in most software environments. While SparseBase uses the C++ abstractions to address most types of sparse data, Libmtx is specifically tailored for sparse vectors and matrices.



<https://github.com/simulahpc/libmtx>



UPCOMING EVENTS

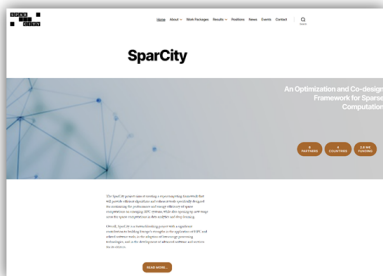
◆ HiPEAC'24
Munich, Germany
January 17-19, 2024

◆ EuroHPC Projects Shaping Europe's
HPC Landscape
Munich, Germany
January 18, 2024

◆ SparCity In-person Meeting,
Munich, Germany
January 22-24, 2024




WEBSITE



<https://sparcity.eu>

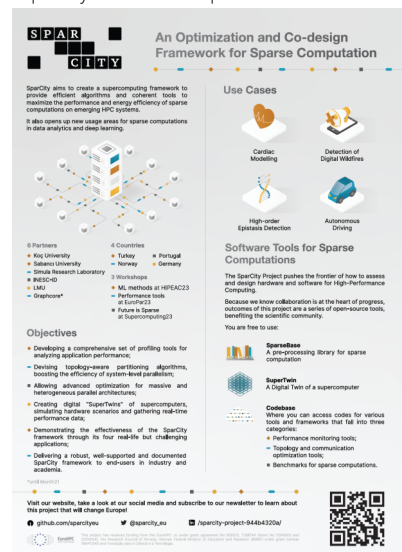


SOCIAL MEDIA




PROMOTIONAL MATERIALS

SparCity released new promotional material



REPOSITORY



<https://github.com/sparcityeu>

MORE INFORMATION

<https://sparcity.eu>

sparcity-project-group@ku.edu.tr



EuroHPC
Joint Undertaking

SPONSORED BY THE



Federal Ministry
of Education
and Research



The Research Council
of Norway

FCT Fundação
para a Ciência
e a Tecnologia

This project has received funding from the European High-Performance Computing Joint Undertaking under grant agreement No 956213.

Koç University and Sabancı University are supported by the Turkish Science and Technology Research Centre Grant No 120N003 and 220N254, respectively.

Simula and Graphcore are supported by the Research Council of Norway. Ludwig-Maximilians-Universität München (LMU) is supported by the German national Funding agency (BMBF). INESC-ID is supported by Fundação para a Ciência e a Tecnologia (FCT).