

Accelerators and Communications for HPC (AccelCom)



Our team drives leading-edge research and development around 3 main pillars: (1) Accelerators/coprocessors in HPC, (2) Programmability of heterogeneous memory systems, and (3) Inter-node communications. We collaborate closely with the major accelerator/coprocessor and network vendors for HPC: NVIDIA, Intel, and Mellanox. We run the BSC/UPC NVIDIA GPU Center of Excellence and collaborate with the Intel-BSC Exascale Lab. activities. We organize locally international events such as the [PUMPS+AI Summer School](#), PATC Courses on [CUDA/OpenACC](#) and [use of heterogeneous memory systems](#), and the [MareNostrum](#) hackathon, and collaborate in the organization of related international conferences and workshops such as SC and IEEE Cluster.

Objectives

Our broad mission is driving top-notch research and development around accelerators and communications technology for HPC. This can be framed in three main points:

- Helping vendors advance their hardware architecture and software stack.
- Designing and developing programming models, runtime systems, and libraries supporting accelerated computing and networking.
- Developing and porting coprocessor-accelerated and distributed-memory applications.

Group News

[Mar. 2024] Group Manager Antonio J. Peña, proposed by Spain, appointed as an Observer of the Research and Innovation Advisory Group (RIAG), part of the Industrial and Scientific Advisory Board at the EuroHPC Joint Undertaking, for a two-year term

[Mar. 2024] A very warm welcome to our new Sr. Research Engineer Lena Martens! She will take over the responsibilities of our inference engine on the HomE project

[Nov. 2023] Antonio J. Peña receives the [2023 Agustín de Betancourt y Molina by the Spanish Royal Academy of Engineering!](#)

[Nov. 2023] ODOS gets its way into the UCX (Unified Communication X) project of the UCF Consortium:

- <https://github.com/openucx/ODOS>

[Nov. 2023] AccelCom deeply involved in everyone's favorite SC23 conference!

- Muhammad Usman:
 - M. Usman, S. Iserte, R. Ferrer, and A. J. Peña, "DPU offloading programming with the OpenMP API", in [The Ninth Workshop on the LLVM Compiler Infrastructure in HPC \(LLVM-HPC\)](#)
- Sergio Iserte:
 - BoF Speaker: [Unified Communication X \(UCX\) Community](#)
 - BoF Speaker: [Enabling I/O and Computation Malleability in High-Performance Computing](#)
 - Martín, Castillo, Aliaga, Iserte, "Efficient data redistribution for malleable applications", in [ExaMPI23 Workshop](#)
 - Security committee member
 - Infrastructure Liaison for Tech Program
- Antonio Peña:
 - Tutorial presenter: [Leveraging SmartNICs for HPC Applications](#)
 - Workshop organizer: [Fourth Workshop on Heterogeneous Memory Systems \(HMEM\)](#)
 - Panelist at Panel "[Runtimes and Workflow Systems for Extreme Heterogeneity: Challenges and Opportunities](#)"
 - SC23 Finance Liaison for Technical Program
 - SC Steering Committee Meeting

[Nov. 2023] ODOS (OpenMP DPU Offloading Support) first released!

- ODOS Team: Muhammad Usman, Sergio Iserte, and Antonio J. Peña
- [Press Release](#)
- [Download](#)

[Nov. 2023] AccelCom welcomes Sergi Rovira as our HE expert!

[Dec. 2022] Antonio J. Peña obtains a BSC Distinction at the BSC Annual Meeting 2022.

[Dec. 2022] Antonio J. Peña and Álvaro García present "HomE. Enabling Private Deep Learning, an ERC Consolidator Grant" within the BSC Talks at the BSC Annual Meeting 2022.

[Dec. 2022] Paper "A symbolic emulator for shuffle synthesis on the NVIDIA PTX code" by Kazuaki Matsumura, Simon García de Gonzalo, and Antonio J. Peña accepted in The 32nd ACM SIGPLAN International Conference on Compiler Construction (CC), Montréal, Canada, Feb. 2023

[Dec. 2022] Antonio J. Peña is appointed as oneAPI Community Forum Steering Committee Member.

[Nov. 2022] Antonio J. Peña is selected as SC Steering Committee Member for the term 2023-2026. Antonio becomes the first Committee Member from a Spanish institution in the top conference in our field.

[Nov. 2022] AccelCom's SC22 participation:

- Paper "OmpSs-2 and OpenACC interoperation" in collaboration with our colleagues at INESC-ID from the EPEEC project successfully presented at the WACCPD Workshop by former AccelCom member Orestis Korakitis - well done!
- Antonio J. Peña participates as a panelist in the "Memory Heterogeneity in High Performance Computing" Panel

[Nov. 2022] Antonio J. Peña to talk about "ecoHMEM: Improving Object Placement Methodology for Hybrid Memory Systems in HPC", in Huawei Annual Compute Architecture Innovation Summit 2022.

[Nov. 2022] Antonio J. Peña to participate in the Defense Committee of PhD Candidate Aaron Call at UPC

[Oct. 2022] We welcome attendees of to the Antonio J. Peña and Marc Jordà to "Parallel Programming Workshop" of BSC's PRACE Advanced Training Centre (PATC). Lecturers from AccelCom are Antonio J. Peña (introductory lecture and MPI) and Marc Jordà (Debugging)

[Sep. 2022] We welcome Álvaro García, Sr. Researcher, to our group, the first incorporation to the ERC project - we get started!

[Aug. 2022] Paper on ecoHMEM accepted for publication at IEEE Cluster 2022 is a Best Paper Candidate

[July 2022] Big welcome to the latest incorporations to the group: Nandhana Sakthivel and Muhammad Usman, Jr. Research Engineers, Suraj Shirvankar, MS Student, and Kamil [Przybyszewsky, undergrad intern](#)

[July 2022] Paper on ecoHMEM accepted for publication at IEEE Cluster 2022

[June 2022] ERC Consolidator Grant HomeE confirmed to start September <https://www.hpcwire.com/off-the-wire/bsc-researcher-antonio-j-pena-awarded-erc-consolidator-grant-for-the-home-project>

[Mar. 2022] Proud to announce the talk at GTC'22 by our PhD Candidate [Kazuaki Matsumura](#): "[JACC: Automatically Retargeting OpenACC Kernels for Multi-GPUs](#)"

[Feb. 2022] Big welcome to the team [Dr. Hatem Elshazly](#), who has joined the team as a Postdoctoral Researcher

[Dec. 2021] Antonio J. Peña to give ExHET's Workshop *keynote* at PPOPP'22: <https://excl.ornl.gov/ppopp-exhet-2022/>

[Nov. 2021] Antonio J. Peña has been confirmed a [Ramón y Cajal Fellowship](#). This is the most prestigious and highest-level named fellowship by the Spanish government for researchers.

[Nov. 2021] Don't miss SC21's AccelCom participation:

- Tutorial [Efficient Distributed GPU Programming for Exascale](#), with Simon Garcia de Gonzalo
- A. J. Peña to be attending in-person as SC21's Communications Liaison for the Technical Program

[Sep. 2021] The international press is echoing our work on enabling large homomorphically encrypted DNN models: [AI Business](#), [Today.in-24](#), [datanami](#)

[Aug. 2021] Congrats to our Team Lead Antonio J. Peña for being [provisionally selected](#) to receive a [Ramón y Cajal fellowship](#), the highest-grade named fellowship by the Spanish government. He ranked 2nd

in the area "Information and Communication Technologies", scoring 99.5/100.

[July 2021] Congrats to our Team Lead Antonio J. Peña for being elevated to IEEE Sr. Member

[June 2021] Check out Guillermo Lloret's presentation on Homomorphically Encrypted DL inference at the [HMEM Workshop 2021](#)

[June 2021] We welcome undergrad summer intern Pablo Izquierdo from Universidad de Cantabria.

[June 2021] AccelCom just released [JACC](#) (Just In Time OpenACC framework). Big kudos to its developer Kazuaki Matsumura and his mentor Simon Garcia de Gonzalo!

[June 2021] Don't miss our participation at ISC'21:

- Our Team Lead Antonio J. Peña to present "[Breaking the DRAM size wall for DNN inference and homomorphic encryption](#)" in the Intel Vendor Use-Case
- Our Team Lead Antonio J. Peña, along with Pavan Balaji, Torsten Hoefler, and Yanfei Guo, to give the "Advanced MPI Programming Tutorial" [\[1\]](#) [\[2\]](#)
- Our PhD Student Kazuaki Matsumura to participate in the [PhD Forum](#)

[Apr. 2021] Check out our paper "[Enabling Homomorphically Encrypted Inference for Large DNN Models](#)" accepted for publication in IEEE Transactions on computers

[Apr. 2021] Congrats to our Team Lead Antonio J. Peña for being named [ACM Sr. Member](#)

[Apr. 2021] Don't miss our team's participation in the GPU Technology Conference 2021:

- [Accelerating Machine Learning Applications Using CUDA Graph and OpenACC](#), by Leonel Toledo, Pedro Valero-Lara, and Antonio J. Peña
- [A Tale of Two Programming-Models: Enhancing Heterogeneity](#), Productivity, and Performance through OmpSs-2 + OpenACC Inter-Operation, by Simon Garcia de Gonzalo

[Feb. 2021] Congrats to our Team Lead Antonio J. Peña for being named SC22 BoF Chair!

[Dec. 2020] AccelCom in the news: [Boosting memory capacity and performance while saving megawatts](#), by The Next Platform

[Nov. 2020] Congrats to our Postdoc Simon Garcia de Gonzalo for coauthoring [SC20's Best Paper](#)

[Oct. 2020] Congrats to our former Researcher Pedro Valero-Lara for receiving an [IEEE-CS TCHPC Early Career Researchers Award for Excellence in HPC](#)

[Jul. 2020] AccelCom in the news: [Building an ecosystem for heterogeneous memory supercomputing](#), by The Next Platform

[Jul. 2020] AccelCom in the news: [BSC research accelerates HPC workloads with less power-hungry DRAM](#), by HPCwire

[Jul. 2020] Don't miss [Antonio J. Peña's podcast on exascale computing by EPEEC](#)

[Apr. 2020] [Antonio J. Peña's interview on the EPEEC project](#), by insideHPC

[Nov. 2019] Welcome our new Postdoc Simon Garcia de Gonzalo joining on a STARS Fellowship

[Oct. 2018] The [EPEEC H2020 Project](#), led by our Team Lead Antonio J. Peña starts

Barcelona Supercomputing Center - Centro Nacional de Supercomputación

Source URL (retrieved on *16 abr 2024 - 08:31*): <https://www.bsc.es/ca/discover-bsc/organisation/scientific-structure/accelerators-and-communications-hpc>