Invited talk #4:
Title: Enabling Large Scale Simulations For Particle Accelerators
Speaker: Konstantinos Iliakis, CERN

Abstract:
The stagnation of Moore's law on transistor scaling and the implications of Amdahl's law on scalability indicates that nowadays more than ever, high-performance computing necessitates meticulous software optimization, fine-tuning and integration of hardware accelerators. This talk illuminates these concepts through the real-world example of code BLonD, the state-of-art beam dynamics simulator. BLonD, developed at CERN, transitioned from a single-node optimized, multi-threaded code to an MPI-over-OpenMP distributed architecture. Applying techniques such as dynamic-load balancing and approximate computing, simulation durations were reduced from weeks to hours. To meet the ever-growing demand for extensive simulation workloads, a distributed, GPU-accelerated version was also implemented, delivering remarkable improvements in latency and throughput. The dramatic reduction in execution time and user-friendly design of BLonD empowers scientists to conduct complex and accurate simulations, crucial for overcoming technological limitations, planning accelerator upgrades, and shaping the future of particle accelerators.

Brief CV:
Konstantinos Iliakis received his diploma in Electrical and Computer Engineering (ECE) from the National Technical University of Athens (NTUA), Greece. In his thesis implemented a relaxed synchronisation mechanism between the map and reduce phases of a state-of-art MapReduce runtime for shared-memory multi-/many-cores. He then conducted his PhD dissertation under the supervision of Prof. D. Soudris, Microlab, NTUA, and in collaboration with the European Organization for Nuclear Research (CERN), in the area of High Performance Computer Architecture and Data Complex Application Acceleration. In his PhD research, he focused on novel, Out-of-Order GPU micro-architectures, as well as on the parallelization and acceleration of scientific simulator software. His research interests include high performance computing, GP-GPU microarchitectures, and applications of deep learning for controls and diagnostics in experimental physics. In 2020 he received the best paper award in the 17th ACM International Conference on Computing Frontiers (CF'20).