Ph.D. in Information Technology
Thesis Defense

February 11th, 2026
at 14:00
Conference Room "Emilio Gatti" – building 20

Ian DI DIO LAVORE – XXXVIII Cycle

Workload-Agnostic Programming in HPC: A Tale of Challenges and Trade-offs

Supervisor: Prof. Marco Domenico Santambrogio

Abstract:

The increasing heterogeneity and complexity of modern High-Performance Computing (HPC) architectures have created fundamental challenges in developing programming models that balance productivity with performance. As a result of systems evolving toward distributed multi-GPU configurations with complex memory hierarchies and diverse interconnect topologies, the gap between high-level programming abstractions and efficient hardware utilization grows. This dissertation focuses on workload-agnostic programming approaches that enable developers to harness the potential of heterogeneous HPC systems without losing familiar programming paradigms or requiring extensive hardware expertise. We explore two complementary strategies to address these challenges. The first approach extends traditional C++ programming models by enhancing the SHAD library, introducing novel abstractions for distributed memory management, topology-aware execution, and transparent hardware acceleration while maintaining compatibility with Standard Template Library (STL) interfaces. The second approach democratizes access to GPU acceleration through polyglot runtime systems, developing frameworks that automatically manage multi-GPU execution and scale workloads across distributed systems without requiring specialized programming knowledge. The research contributes several key innovations to the field of HPC programming models. We develop topology-aware scheduling mechanisms that adapt to diverse hardware configurations, introduce flexible memory allocation abstractions that optimize data locality in distributed environments, and design runtime systems capable of transparently scaling GPU workloads beyond single-node limitations. Our extensive evaluation on production supercomputers demonstrates that carefully designed abstractions can effectively navigate the inherent trade-offs between programmability and performance. Through both practical implementations and theoretical analysis, we demonstrate how workload-agnostic systems can make HPC accessible to broader application domains while maintaining competitive performance.

Beatrice BRANCHINI – XXXVIII Cycle

On the Role of Reconfigurable Systems in Quantum Computing: the Case of Quantum Error Correction

Supervisor: Prof. Marco Domenico Santambrogio

Abstract:

Quantum computing is a novel paradigm that can potentially disrupt the Computer Science field. It promises to overcome the gap between the limitations of current classical processors and the performance demanded by modern applications. Nevertheless, significant challenges persist in constructing reliable quantum computers. Qubits within today's Quantum Processing Units (QPUs) remain highly unstable and susceptible to noise originating from various sources, ultimately preventing the reliable execution of many applications on quantum processors. To address these issues, the current generation of quantum devices integrates Quantum Error Correction (QEC) mechanisms to recover the computation, thus representing a critical component toward fault tolerance. However, the QEC process faces several constraints, with the latency of error decoding being the most critical limitation. In this scenario, Field Programmable Gate Arrays (FPGAs) have arisen as a promising platform for deploying QEC engines. Their deterministic latency and spatial computing capabilities enable them to meet strict timing requirements, maintaining low power consumption at the same time. Existing approaches leveraging FPGAs for QEC have delivered encouraging results, yet substantial room for improvement remains. In particular, the field lacks consensus on the most significant criteria to deploy and assess QEC engines on FPGA, leaving the research landscape fragmented. Furthermore, beyond the pressing needs of latency optimizations, future solutions must also remain adaptable to emerging QPU technologies to keep pace with this rapidly evolving domain. Lastly, abstraction frameworks and modeling tools are essential to equip quantum engineers with the means to effectively customize QEC engines, ensuring their practical usability.

In this context, this dissertation investigates how FPGA technology can advance the field of quantum computing through efficient deployment of QEC engines, enabling the full potential of this emerging paradigm. Specifically, it introduces a series of methodologies to develop high-performance and energy-efficient QEC decoders by leveraging FPGAs. The dissertation begins with a comprehensive review of the State of the Art and, building on this analysis, proposes a set of evaluation pillars to guide the benchmarking and evaluation of novel solutions. It then presents two distinct methodologies for designing FPGA-based QEC decoders. The first, QASBA, implements a gold-standard quantum error decoding algorithm entirely in hardware, ensuring high accuracy. To enhance its practical applicability, an automation framework is also introduced to enable quantum engineers to customize QASBA-based decoders according to the characteristics of the target quantum system. The second methodology forms the basis of QUEKUF, which combines the efficiency of a heuristic decoding algorithm with the computing power of FPGAs to achieve superior performance. Also in this case, to facilitate practical deployment, QUEKUF incorporates a latency-oriented resource optimization model that assists final users in selecting the most effective architectural configuration for their specific quantum environment. This dissertation then contextualizes these contributions within the State of the Art, evaluating their advantages and highlighting opportunities for further refinement. Lastly, it concludes by consolidating the

proposed contributions, outlining the limitations of the current work, and suggesting promising future research directions.

Mirko COGGI – XXXVIII Cycle

HIGH-PERFORMANCE ALGORITHMS AND FRAMEWORKS FOR GRAPH-BASED PANGENOMICS: METHODS, APPLICATIONS, AND TECHNOLOGY TRANSFER

Supervisor: Prof. Marco Domenico Santambrogio

Abstract:

The field of computational genomics is undergoing a profound transformation driven by the emergence of pangenome graphs as a more expressive alternative to linear references. These structures hold the potential to mitigate reference bias and improve the representation of complex genomic variation. However, their adoption remains hindered by significant computational challenges, lack of standardized evaluation frameworks, and limited interoperability with downstream analytics pipelines.

This dissertation contributes a comprehensive and modular framework for sequence-to-graph alignment, designed to adapt dynamically to the topological properties of genome graphs. The framework is articulated across three stages: a preprocessing phase that partitions and refines graph topology, introducing the first method for compact and accurate representation of copy number variations; a processing layer that integrates specialized aligners and includes a novel implementation of the Graph Wavefront Alignment algorithm with traceback support, as well as patented techniques for GPU-accelerated alignment on cyclic structures; and a postprocessing module that addresses the limitations of current variant representation formats through a GPU-accelerated library for transforming VCF data into structured, analysis-ready formats.

In parallel, this work introduces the first systematic benchmarking methodology for sequence-to-graph aligners, combining qualitative KPIs and quantitative assessments across real and synthetic datasets. This initiative lays the foundation for reproducible and extensible evaluation practices within the field.

Recognizing the translational potential of these contributions, the final part of the thesis proposes and validates a methodology for supporting deep tech technology transfer in academia. Developed in collaboration with NECSTLab (Politecnico di Milano) and tested through the GenoGra case study, this framework provides structured guidance for transforming scientific results into scalable, research-driven startups.

Through the integration of algorithmic innovation, high-performance computing strategies, and entrepreneurial methodology, this thesis advances both the technical foundations and the translational pathways for next-generation genomic analysis.

PhD Committee

Prof. Maurizio Magarini, Politecnico di Milano

Prof. Dionisios N. Pnevmatikatos, National Technical University of Athens

Prof. Juergen Becker, ITIV