PhD in Information Technology and Electrical Engineering
Università degli Studi di Napoli Federico II

PhD Student: Gianluca Sabella

Cycle: XXXVIII

Training and Research Activities Report

Year: First

Tutor: Prof. Carlo Sansone

Co-Tutor: Prof. Elvira Rossi

Date: December 09, 2023

1. Information:

- **PhD student: Gianluca Sabella**
- **DR number: 996970**
- **Date of birth: 28/06/1994**
- **Master Science degree: Computer Science** **University: University of Naples Federico II**
- **Doctoral Cycle: XXXVIII**
- **Scholarship type: *no funded scholarship***
- **Tutor: Prof. Carlo Sansone**
- **Co-tutor: Prof. Elvira Rossi**

2. Study and training activities:

Activity	Type[1]	Hours	Credits	Dates	Organizer	Certificate[2]
IoT Data Analysis	Course	12	4	9-13-16-20-23-27/01 2023	Dr. Raffaele Della Corte	Y
Spring School on Transferable Skills	Course	9,5	2	23-24/05 2023	Department of Pharmacy University of Naples Federico II	Y
First course about the porting on GPUs of code and algorithms	Course	15	1,5	19-20-21/06 2023	ICSC – Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing	Y
Ethics and AI	Course	12	2,4	4-11-18-25/10 8-15/11 2023	Italian Society for Ethics of AI (SIpEIA)	Y
Percorso per il rafforzamento delle competenze sulla progettazione europea	Course		3,4	14-28/09 12-26/10 09-23/11/2023	Ministero dell'Università e della Ricerca, Ateneo Federico II	Y
The Fifth International School on Open Science Cloud (SOSC 2023)	School	34	3,4	23-24-25-26-27/10/2023	ICSC – Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing	Y
Principi architetturali - TOGAF 1	Seminar	3	0,6	30/01/2023	5G Academy, DIETI, Capgemini	Y
Principi architetturali - TOGAF 2	Seminar	3	0,6	09/02/2023	5G Academy, DIETI, Capgemini	Y
Blockchain in business and 5G in business	Seminar	3	0,6	13/02/2023	5G Academy, DIETI, Capgemini	Y
How to Publish Under the CARE-CRUI Open Access Agreement with	Seminar	1,5	0,3	05/04/2023	CARE-CRUI and IEEE	Y

IEEE						
Nanoneuro: the power of nanoscience to explore the frontiers of neuroscience	**Seminar**	**1**	**0,2**	**03/05/2023**	**prof. C. Forestiere, DIETI, Unina**	**Y**
Digital & Green Create New Value Together	**Seminar**	**2,5**	**0,5**	**15/05/2023**	**Ing. Carmine Piccolo, UniNa**	**Y**
5G Tecnologia, Innovazione, Territorio e Diritto	**Seminar**	**2**	**0,4**	**27/05/2023**	**Prof. Nicola Pasquino, UniNa**	**Y**
Traffic Engineering	**Seminar**	**1**	**0,2**	**23/06/2023**	**Prof. V. Persico, DIETI, Unina**	**Y**
Ricerca e formazione nella società della transizione digitale	**Seminar**	**5**	**1**	**22/09/2023**	**Prof. Stefano Russo, DIETI**	**N**
La ricerca con Amazon Web Services Opportunità e Strumenti	**Seminar**	**2**	**0,4**	**29/09/2023**	**Prof. Carlo Sansone, UniNa**	**Y**
Nuove frontiere nel mondo dei Data Center Edge e HPC - UniNa long-term experience With RITTAL racks, liquid cooling systems, chillers	**Seminar**	**1**	**0,2**	**21/09/2023**	**Prof. Guido Russo, UniNa**	**Y**
MATLAB Tools in INFN for Open Science, Reproducible Workflows and Parallel Computing	**Seminar**	**2**	**0,4**	**30/11/2023**	**MathWorks**	**Y**

1) Courses, Seminar, Doctoral School, Research, Tutorship
2) Choose: Y or N

2.1. Study and training activities - credits earned

	Courses	Seminars	Research	Tutorship	Total
Bimonth 1	4	1,8	6	0	11,8
Bimonth 2	0	0,3	6	0	6,3
Bimonth 3	3,5	1,3	6	0	10,8
Bimonth 4	0	0	5	0	5
Bimonth 5	5,4	1,6	8	0	15
Bimonth 6	3,6	0,4	8	0	12
Total	16,5	5,4	39	0	60,9
Expected	30 - 70	10 - 30	80 - 140	0 – 4.8	

3. Research activity:

My research work focuses on two distinct aspects, the goal of which is their convergence. The primary focus of my research revolves around benchmarking for the evaluation of heterogeneous computing resources. High-performance computers continue to increase in speed and capacity. Alongside these advancements, architectures are progressively becoming more intricate, featuring multi-socket and multi-core central processing units (CPUs), multiple graphical processing unit (GPU) accelerators, and multiple network interfaces per node.

In the realm of scientific computing (SC), high-level benchmarks are designed to assess the system's overall performance, encompassing CPU, memory, final disk drives, and all available computing devices (e.g., GPUs). These tools are intended to evaluate how well the computing environment can deliver the high performance required by SC applications [1, 2, 3, 4]. Tests conducted by these tools are often used both to assess overall system performance and to compare the performance of different systems. For instance, the Linpack benchmark is a high-level benchmark used to evaluate computer system performance in terms of its ability to solve large-scale problems, as evidenced by its inclusion in the Top500 ranking [5]. During my first year of doctoral studies, we embarked on the initial stages of developing a new benchmark inspired by Linpack. This benchmark is based on reformulating the Schur Complement of a linear equation system solution.

Algorithm 1 The *"Schur complement"*-based algorithm for the linear system $Ax = y$ solution

procedure SchurComplementSolution(A, y, x)
    Input: A, y
    Output: x
    $L_{A_{11}}, U_{A_{11}} \leftarrow$ Compute LU factorization of A_{11} ▷ **Task n. 1.** $O\left(\frac{2}{3}n_1^3 - \frac{1}{2}n_1^2\right)$
    $z \leftarrow$ Solve $A_{11}z = y_1$ by mean of $L_{A_{11}}$ and $U_{A_{11}}$ ▷ **Task n. 2.** $O\left(2n_1^2\right)$. See equation (8)
    $E \leftarrow$ Solve $A_{11}E = A_{12}$ by mean of $L_{A_{11}}$ and $U_{A_{11}}$ ▷ **Task n. 3.** $O\left(2n_1^2 n_2\right)$. See equation (4)
    $w \leftarrow$ Compute the BLAS2 product $w = y_2 - A_{21}z$ ▷ **Task n. 4.** $O\left(n_2\left(n_1+1\right)\right)$. See equation (8)
    $S \leftarrow$ Compute the BLAS3 product $S = A_{22} - A_{21}E$ ▷ **Task n. 5.** $O\left(n_2^2\left(n_1+1\right)\right)$. See equation (4)
    $L_S, U_S \leftarrow$ Compute LU factorization of S ▷ **Task n. 6.** $O\left(\frac{2}{3}n_2^3 - \frac{1}{2}n_2^2\right)$
    $x_2 \leftarrow$ Solve $Sx_2 = w$ by mean of L_S and U_S ▷ **Task n. 7.** $O\left(2n_2^2\right)$. See equation (8)
    $u \leftarrow$ Compute the BLAS2 product $u = y_1 - A_{12}x_2$ ▷ **Task n. 8.** $O\left(n_1\left(n_2+1\right)\right)$. See equation (9)
    $x_1 \leftarrow$ Solve $A_{11}x_1 = u$ by mean of $L_{A_{11}}$ and $U_{A_{11}}$ ▷ **Task n. 9.** $O\left(2n_1^2\right)$. See equation (9)
end procedure

This tool is designed to assess the performance of heterogeneous computing resources. Its goal is not to replace the established HPL benchmark commonly used for evaluating HPC platforms, but rather to contribute to the ongoing discussion about integrating benchmarks better suited to actual system workloads. Initial evaluations, primarily focusing on performance scalability, were conducted in a computing environment based on multiple nodes of NVIDIA GP-GPUs interconnected via InfiniBand [6]. However, this approach could also adapt to other acceleration technologies such as ROCm for AMD GPUs or oneAPI for Intel accelerators. The tests conducted demonstrate that the performance of this benchmark is comparable to tools primarily leveraging GPU [7].

Enhancing task distribution among computational components and resolving implementation issues should enhance evaluation effectiveness in heterogeneous systems within the scientific computing domain.

Additionally, focusing on leveraging MPI CUDA-aware implementations or similar strategies is part of my forthcoming work to potentially improve performance.

Simultaneously, I've collaborated on the ATLAS experiment following my affiliation with the National Institute of Nuclear Physics, contributing as a member of WP2 and WP5 of the SPOKE 2 research program FUNDAMENTAL RESEARCH & SPACE ECONOMY within the ICSC (National Research Center in HPC, Big Data, and Quantum Computing) project. This initiative is funded by the European Union under NextGenerationEU, the Italian Ministry of University and Research, and the National Recovery and Resilience Plan (PNRR) 'ItaliaDomani.'

Specifically, my efforts involve crafting a flexible and scalable computing infrastructure, leveraging container technology and Kubernetes for orchestration. Unlike virtual machines that emulate complete hardware with separate operating systems for each application, containers share the host's operating system, isolating only the application's resources. This offers greater flexibility, lower overhead, and faster deployment compared to VMs.

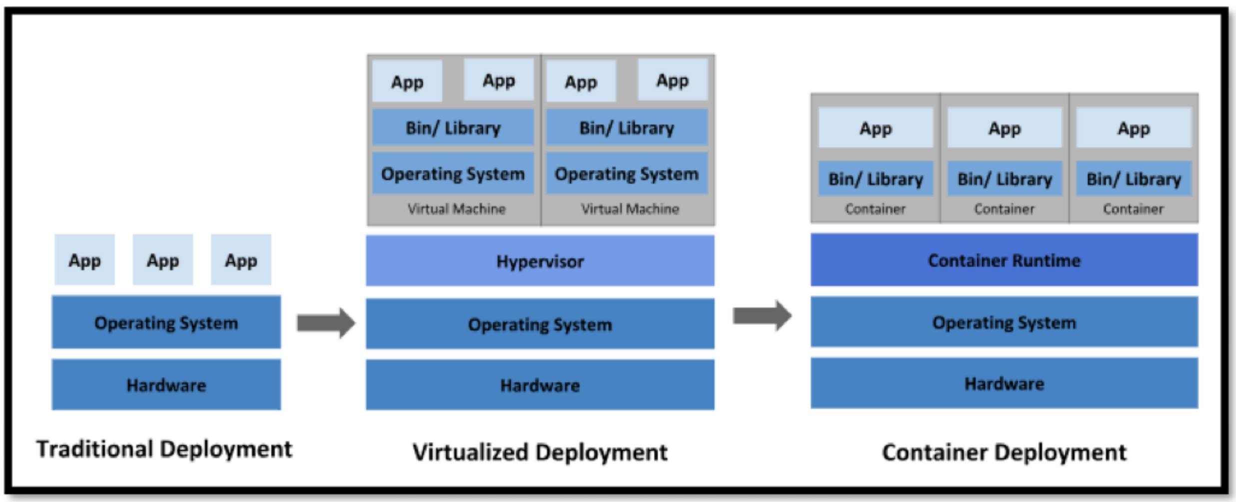

Even though the infrastructure is designed to accommodate a wide array of applications, an analysis facility system was chosen as a use case. It was considered a suitable example to test the infrastructure's functionality as it encompasses various services.

Among the components implemented for this use case are:

- MinIO with High-Performance Object Storage: An object storage infrastructure was instantiated to enable users to efficiently store and access data within the cluster. This service is designed to ensure secure and scalable management of user data.

- JupyterHub: An instance of JupyterHub was implemented to provide users with an interactive analysis environment. Access is managed through administrator-approved registration, ensuring the environment's security. Once registered, users can access a personalized instance of JupyterLab, providing them with a platform ready for data analysis.

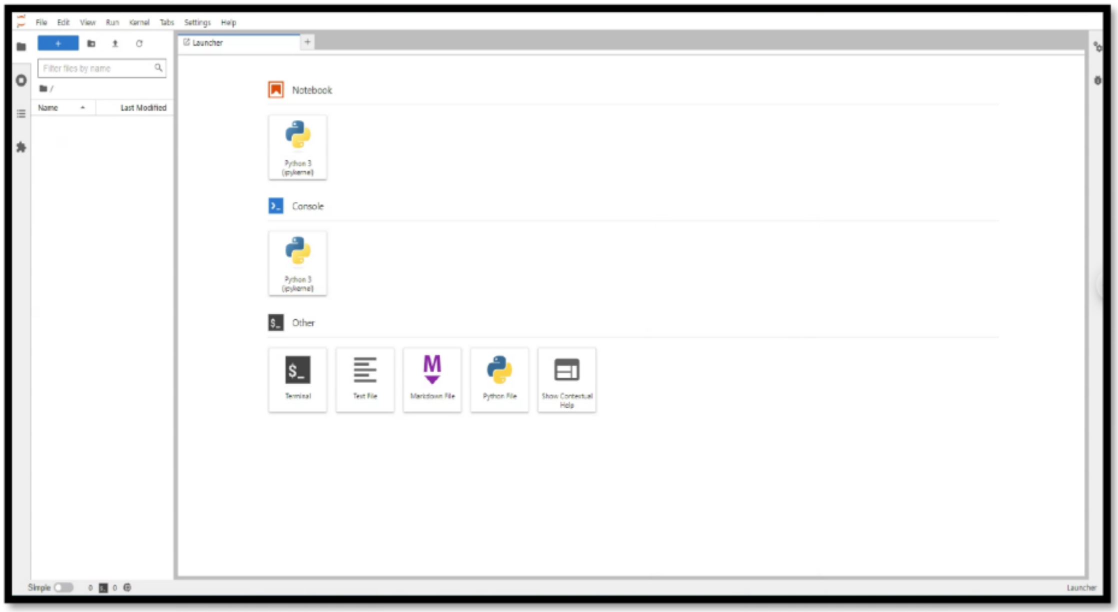

- JupyterLab and Dask: The JupyterLab environment was configured to enable users to perform advanced analysis using Python libraries for data science.

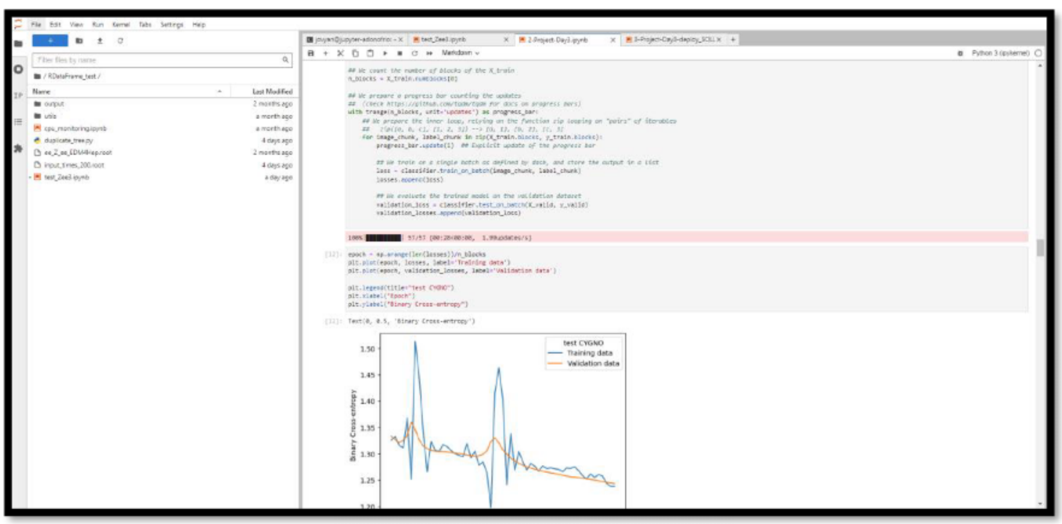

Scalability in processing operations is ensured by leveraging Dask, enabling the distribution and execution of operations across multiple nodes in the cluster. DASK can use the most popular batch systems (SLURM, HTCondor, etc. ...) or is provided with a native scheduler if necessary (our choice).

Cluster Configuration

The cluster comprises an architecture with 6 nodes orchestrated by the Kubernetes system. One node function as both the master and worker, while the other 5 nodes operate exclusively as workers.

- Master/Worker Node: It serves the role of coordinating and managing the cluster, alongside performing data processing operations. Its specifications include:
    - VCPU: 8 (1 physical with 8 cores)
    - Disk: 50 Gigabytes
    - RAM: 16 Gigabytes
    - OS: Alma Linux 9.2
- Worker Nodes: These nodes provide additional processing resources and support distributed and parallel operations within the analysis infrastructure. They have the following specifications:
    - VCPU: 32 (1 physical with 32 cores)
    - Disk: 100 Gigabytes
    - RAM: 64 Gigabytes
    - OS: Alma Linux 9.2

Current Scenario and Future Perspectives

My work aims to emphasize the need to expedite the development and implementation of applications in high-performance computing contexts. These procedures can often be lengthy and complex due to the intricacy involved in setting up appropriate infrastructures and specific development environment configurations. The demand for faster processes is crucial to adapt to the dynamic pace of the market and efficiently meet the needs of end-users. This demonstrative setup, utilizing containers and Kubernetes for orchestration, provides a tangible example of the cluster's capabilities, with the flexibility to expand to support additional use cases or specialized applications that require a flexible and scalable computing infrastructure.

The goal for the coming year will be to evaluate the benefits of offering a rapid and user-friendly application deployment in terms of its impact on performance. Additionally, studying the overhead of a high-performance system in the case of virtualized resources will be explored. To achieve this goal and comprehend ways to minimize this impact, the utilization of an appropriate benchmarking system will be fundamental.

Additionally, the plan is to further expand the developed methodologies, exploring new scenarios and application domains to assess the integration of GPUs and the application of deep learning techniques. Specific contexts such as large dataset analysis, predictive modeling in scientific or industrial domains, and tackling complex problems using advanced machine learning algorithms will be considered. The objective might be to demonstrate the effectiveness and versatility of the proposed solutions in real-world contexts and situations where performance and scalability are crucial

Ref.

[1] Laccetti G, Lapegna M, Mele V, Romano D, Murli A. A Double Adaptive Algorithm for Multidimensional Integration on Multicore Based HPC Systems. International Journal of Parallel Programming 2012; 40(4): 397-409. doi: 10.1007/s10766-011-0191-4

[2] Bertero M., Bonetto P., Carracciuolo L., Laccetti G et al. . MedIGrid: a medical imaging application for computational Grids. In: IEEE. ; 2003:1213457. doi: 10.1109/IPDPS.2003.1213457

[3] Carracciuolo, L. et. al . Implementation of a non-linear solver on heterogeneous architectures. Concurrency and Computation: Practice and Experience 2018; 30(24): e4903. doi: 10.1002/cpe.4903

[4] D'Amore L, Constantinescu E, Carracciuolo L. A Scalable Space-Time Domain Decomposition Approach for Solving Large Scale Nonlinear Regularized Inverse Ill Posed Problems in 4D Variational Data Assimilation. J. Sci. Comput. 2022; 91(2). doi: 10.1007/s10915-022-01826-7

[5] The Top 500 ListWebsite. https://www.top500.org/

[6] Carracciuolo L, Bottalico D, Michelino D, Sabella G, Spisso B. Benchmarking a High-Performance Computing Heterogeneous Cluster. In: Springer International Publishing; 2023: 101–114. doi: 10.1007/978-3-031-30445-3_9

[7] Carracciuolo, L, Mele, V, Sabella, G. Toward a new linpack-like benchmark for heterogeneous computing resources. Concurrency Computat Pract Exper. 2023;e7962. doi: 10.1002/cpe.7962

4. Research products:

Carracciuolo, L, Mele, V, Sabella, G. Toward a new linpack-like benchmark for heterogeneous computing resources. Concurrency Computat Pract Exper. 2023;e7962. doi: 10.1002/cpe.7962 Published

5. Conferences and seminars attended

ATLAS Italia Computing annual meeting, Genova, Italy, 27/11/2023 to 29/11/2023,
presenting author - Le attività del T2 di Napoli con Kubernetes: il caso d'uso delle analysis facilities.

SPOKE 2 annual meeting, Bologna, Italy, 18/12/2023 to 20/12/2023,
presenting author - Developing and testing of a flexible and scalable high-rate analysis platform.

6. Activity abroad

None

7. Tutorship

None