ICSC and Spoke 2 - Where Are We Now?

Tuesday, 10 December 2024 - Thursday, 12 December 2024

Physics Dept and INFN, Catania

Book of Abstracts

Contents

81

The AIDA project: galaxy formation in alternative dark matter models

Author: Giulia Despali[1]

[1] *University of Bologna*

Corresponding Author: giulia.despali@unibo.it

Despite years of experimental searches, understanding dark matter, a cornerstone of modern cosmology, remains elusive. For a variety of models, telescope observations often provide the only constraints on the nature of dark matter, and thus, new upcoming observations will be instrumental in studying dark matter through the comparison with numerical predictions. Only recently, simulations with alternative dark matter models have moved towards modelling dm variations and baryons at the same time, essential to derive realistic predictions.
I will present the first results from the AIDA-TNG simulations, a new set of cosmological simulations including CDM, WDM and SIDM, together with a complete treatment of galaxy formation. Thanks to the comparison between CDM and ADM, as well as hydro and dmo runs, we can constrain the effects of dark matter models on quantities such as halo statistics and matter power spectrum, galaxy shapes and dynamics, or the properties of their satellites. Moreover, cosmological boxes provide large statistics of galaxies, allowing us to create mock observations at multiple scales that can be compared to real data and used as training sets.
The simulations have been partly run on Leonardo and are the result of a Flagship project of WP3.

Giorno preferito:

82

Lattice QCD in the exascale computing era

Author: Marco Cè[1]

[1] *Università di Milano-Bicocca*

Corresponding Author: marco.ce@unimib.it

Lattice QCD is the leading framework for systematically studying the non-perturbative regime of Quantum Chromodynamics (QCD), the theory of strong interactions. Today, lattice methods are essential for high-precision calculations—both current and future—of fundamental quantities in the Standard Model of particle physics. These advancements are made possible by highly-parallelized HPC simulations, driven by ongoing research into new algorithms informed by physics. Modern implementations can utilize up to 100 000 CPU cores or 1000 GPUs in parallel, running on pre-exascale systems with near-ideal scalability.

Giorno preferito:

11 Dicembre Pomeriggio

83

Machine Learning for event reconstruction in Super-Kamiokande

Author: Nicola Fulvio Calabria[1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Author: nicola.calabria@ba.infn.it

In this preliminary study I consider and explore the application of Machine Learning algorithms for reconstruction in Super-Kamiokande, the largest Water Cherenkov detector in the world. I simulated event samples to train a custom ResNet-18 based model whose perfomance is presented in this talk. The goal is the development of a Machine Learning based tool to be employed in proton decay analysis along with the official reconstruction software (fiTQun), which is based on Likelihood Maximization, to enhance reconstruction of faint rings in multi-ring events, ultimately improving signal selection efficiency.

Giorno preferito:

Introduction / **84**

Saluti INFN-CT / UniCT

Corresponding Authors: romano@lns.infn.it, alessia.tricomi@ct.infn.it

Giorno preferito:

85

Stato finanziario e amministrazione

Introduction / **86**

Saluti Fondazione

Corresponding Author: antonio.zoccoli@bo.infn.it

Introduction / **87**

Spoke 2 Status

Introduction / **88**

Situazione Finanziaria

Corresponding Author: simona.petronici@pi.infn.it

Introduction / 89

The ICSC Status

Sessione Pubblica / 90

Esempi di attivita'

Sessione Pubblica / 91

Tavola rotonda

Giorno preferito:

Work Package reports / 92

WP1

Corresponding Authors: leonardo.giusti@mib.infn.it, leonardo.cosmai@ba.infn.it

Work Package reports / 93

WP6

Corresponding Authors: antonio.stamerra@inaf.it, alessia.tricomi@ct.infn.it

https://docs.google.com/file/d/1806GOYJ9z7sHeOPI-6ZZ7Ih1Bl7A7XMu/edit?usp=docslist_api&filetype=mspresentation

Giorno preferito:

Work Package reports / 94

WP2

Corresponding Authors: lenzip@fi.infn.it, alberto.annovi@pi.infn.it

Work Package reports / 95

WP3

Corresponding Authors: marco.landoni@brera.inaf.it, marco.landoni@gmail.com

Work Package reports / 96

WP4

Corresponding Authors: simone.gennai@mib.infn.it, simone.gennai@cern.ch

Work Package reports / 97

WP5

Corresponding Authors: elly@na.infn.it, daniele.spiga@pg.infn.it

Links with Industries and external entities / 98

ML-BLEN

Author: Alessandro Bombini[1]

[1] *FI*

Corresponding Authors: alessandro.bombini.fi@gmail.com, alessandro.bombini@fi.infn.it

Giorno preferito:

Links with Industries and external entities / 99

Leonardo IDL

Corresponding Authors: nicolo.magini@ge.infn.it, nicolo.magini@leonardo.com

Giorno preferito:

Links with Industries and external entities / 100

Predictive maintenance

Corresponding Author: robertocornali15@outlook.com

Links with Industries and external entities / 101

GAIAS2

Corresponding Author: cristiano.bozza@sa.infn.it

Giorno preferito:

Links with Industries and external entities / 102

OC2

Links with Industries and external entities / 103

OC3

104

Optimal use of timing measurement in vertex reconstruction at CMS

Author: Ksenia de Leo[1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Author: ksenia.de.leo@ts.infn.it

The upgrade of the CMS apparatus for the HL-LHC will provide unprecedented timing measurement capabilities, in particular for charged particles through the Mip Timing Detector (MTD). One of the main goals of this upgrade is to compensate the deterioration of primary vertex reconstruction induced by the increased pileup of proton-proton collisions by separating clusters of tracks not only in space but also in time.
This contribution discusses the ongoing algorithmic developments to optimally exploit such new information, going beyond the initial studies at the time of the detector proposal, both from the physics and computational performance point of view.

Giorno preferito:

105

Towards virtual painting recolouring using Vision Transformer on X-Ray Fluorescence datacubes

Author: Alessandro Bombini[1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Author: alessandro.bombini.fi@gmail.com

In this contribution, we define (and test) a pipeline to perform virtual painting recolouring using raw data of X-Ray Fluorescence (XRF) analysis on pictorial artworks. To circumvent the small dataset size, we generate a synthetic dataset, starting from a database of XRF spectra; furthermore, to ensure

a better generalisation capacity (and to tackle the issue of in-memory size and inference time), we define a Deep Variational Embedding network to embed the XRF spectra into a lower dimensional, K-Means friendly, metric space.
We thus train a set of models to assign coloured images to embedded XRF images. We report here the devised pipeline performances in terms of visual quality metrics, and we close on a discussion on the results.

Giorno preferito:

106

An API for training Physics-Informed Kolmogorov-Arnold Networks using Nvidia Modulus SYM

Author: Alessandro Bombini[1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Author: alessandro.bombini.fi@gmail.com

In this contribution we discuss the novel neural network architecture, based on the Kolmogorov-Arnold Representation Theorem, dubbed "Kolmogorov-Arnold Network" (KAN), its variants (like the Chebyshev-KAN, the Jacobi-KAN, the FastKAN), and its applications to numerical resolution of PDEs via the Physics-Informed Neural Network (PINN) framework.

We discuss our implementation of the API for using seamlessly these network architectures within a widely adopted open source package for PINN, Nvidia Modulus SYM.

Giorno preferito:

107

Heavy flavours jet substructure

Authors: Andrea Ghira[1]; Simone Marzani[2]; Prasanna Kumar Dhani[3]; Oleh Fedkevych[2]

[1] *Università di Genova and Istituto Nazionale di Fisica Nucleare (INFN) sezione di Genova*

[2] *Istituto Nazionale di Fisica Nucleare*

[3] *IFIC, CSIC-UV*

Corresponding Authors: oleh.fedkevych@ge.infn.it, andrea.ghira@ge.infn.it, simone.marzani@ge.infn.it, dhani@ific.uv.es

In this talk we will focus on hadronic jets initiated by heavy flavours. In particular, we will examine different observables to study the dead cone effect, i.e. the suppression of collinear QCD radiation around massive quarks, and to investigate the sensitivity of different observable definitions to the presence of quark masses.

Our results are based on all-order resummed predictions at next-to-leading logarithmic (NLL) accuracy which provide a more precise description of the transition near the dead-cone threshold. Furthermore, we will compare our analytical findings with Monte Carlo predictions.

Giorno preferito:

108

Implementation of low latency, fast inference neural networks on FPGA for trigger-like systems

Author: Salvatore Loffredo[1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Author: salvatore.loffredo@na.infn.it

In many research and industrial settings, achieving fast, low-latency algorithmic responses is essential. To meet the demands of the upgraded LHC and future High Energy Physics (HEP) detectors, quick and powerful triggers are necessary. In recent years, machine learning algorithms have been widely applied to such tasks, and more recently, solutions based on Field Programmable Gate Arrays (FPGAs) have proven to be effective, offering reduced latency and power consumption compared to GPUs not only in HEP field.
This work presents an overview of the development of fast neural networks on FPGAs for High-level and Level-0 trigger systems in the ATLAS experiment. Additionally, to investigate the performance and scalability of the algorithm on multi-FPGA systems, an AMD Alveo cluster is currently being constructed at the INFN-Naples site.
This work is within the use case "Ultra-fast algorithms running on FPGA"within the WP2 framework.

Giorno preferito:

109

Quasi interactive analysis of big data with high throughput - where are we now?

Authors: Francesco Giuseppe Gravili[1]; Tommaso Diotalevi[2]

[1] *Istituto Nazionale di Fisica Nucleare*

[2] *Università e INFN, Bologna*

Corresponding Authors: diotalevi@bo.infn.it, francesco.giuseppe.gravili@le.infn.it

Following the last lightning talk given at the 2023 Spoke2 Annual Meeting, about the flagship use-case UC2.2.2 "Quasi interactive analysis of big data with high throughput"of the Work Package 2 "Experimental High Energy Physics", this contribution will report the status of the ongoing efforts. Starting from the scientific production and going through the common cloud infrastructure deployment (in light of the resources provisioned by the RAC committee), the talk will provide an intermediate checkpoint towards the last year of the project.
A survey of the ongoing activities will also be provided, carried out mainly by Spoke 2 members, outlining the next steps of the flagship, in synergy with the architectural support provided by Work Package 5.

Giorno preferito:

11 Dicembre Pomeriggio

110

Building a Versatile HPC Data Center: a STILES lesson

Author: Fabio Ragosta[1]

[1] *Università degli studi di Napoli "Federico II"*

Corresponding Author: fabio.ragosta@unina.it

High-performance computing (HPC) has become indispensable for addressing the complex challenges of modern scientific research. From processing massive datasets to running simulations with millions of variables, HPC supports advancements across a range of disciplines. This presentation will provide an overview of the design and implementation of a new HPC data center, focusing on its ability to meet the computational and storage demands of diverse scientific workflows.
The data center is tailored to accommodate various user needs. For instance, astrophysical scientists analyze large-scale images from ground-based telescopes (e.g., ESO in Chile, GranTeCan at the Canary Islands), space telescopes (Hubble, Webb), and radio telescopes. These images, often exceeding millions of pixels, require advanced AI tools for object recognition and physical property extraction after calibration. Material physicists, on the other hand, use state-of-the-art mathematical models to simulate and optimize alternative jet fuels, generating thousands of tests that need efficient storage and rapid comparison. High-energy physicists process dozens of terabytes of data from experiments such as ATLAS (CERN) or Belle II (KEK), requiring sequential access, high-speed networking, and parallel file systems for efficient analysis.

This HPC infrastructure integrates GPUs, OpenMP-based parallelism, and optimized caching mechanisms to ensure high performance despite budgetary constraints. As a key application example, we will showcase its role in astrophysics, specifically for ELT and SKA, demonstrating how it enables cutting-edge research and fosters interdisciplinary collaboration.

Giorno preferito:

11 Dicembre Pomeriggio

111

Accelerating LHCb VELO U2 Design with Machine Learning: A Fast Simulation Approach for the TimeSPOT Sensor

Author: Zhihua Liang[1]

Co-authors: Adriano Lai [1]; Andrea Contu [1]; Andrea Lampis [1]; Angelo Loi [1]; Rudolf Oldeman [1]; Stefano Cabiddu [1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Authors: andrea.contu@ca.infn.it, angelo.loi@ca.infn.it, rudolf.oldeman@ca.infn.it, andrea.lampis@ca.infn.it, zhihua.liang@ca.infn.it, stefano.cabiddu@ca.infn.it, adriano.lai@ca.infn.it

The VELO Upgrade 2 (VELO U2) for the LHCb experiment relies on the TimeSPOT sensor for enhanced timing resolution and radiation hardness. Efficient simulation is crucial for optimizing the sensor design and predicting performance. Traditional Geant4/TCoDe simulations, while accurate, are computationally intensive, hindering large-scale design studies. This talk presents a novel approach using machine learning to develop a fast and accurate simulation for the TimeSPOT sensor.

We trained a Multi-Layer Perceptron (MLP) on a dataset of 400,000 simulated events from Geant4/TCoDe. The model accurately predicts the sensor response (charge collection and time of arrival) using particle properties as input. Importantly, the ML-based simulation achieves a remarkable speedup of 10^4 - 10^5 compared to traditional methods. This significant acceleration enables rapid exploration of sensor placement within the VELO U2 geometry, contributing to optimized design decisions.

We have successfully integrated this fast simulation model into the Gauss framework using TMVA SOFIE, paving the way for widespread adoption within the LHCb collaboration. This presentation

will outline the model development, performance results, integration process, and the potential impact on VELO U2 design and analysis. Future directions, including the exploration of more advanced ML models and comprehensive validation, will also be discussed.

Giorno preferito:

11 Dicembre Pomeriggio

112

Extending Rucio to support external metadata catalog

Author: Luca Pacioselli[1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Author: luca.pacioselli@pg.infn.it

In this talk, I will present the progress made in deploying and customizing Rucio as a scalable, metadata-integrated prototype Data Lake as a solution for Data Management (DM) of the Interoperable Data Lake (IDL) project. Rucio is an open-source DM software designed for large-scale scientific experiments such as those in high-energy physics (HEP). Initially developed by CERN for the ATLAS experiment, it is a robust solution for managing any type of data across geographically distributed storage systems.

The talk will focus on a custom did-metadata plugin developed to enable communication with an external database, AyraDB, to handle project-specific metadata with a predefined structure. This communication was achieved using an API client developed by CherryData, supporting both ingestion and query operations.

Finally, we developed a custom Rucio client to enable relevant features for the IDL project, not allowed by the default client, such as a combined file upload and metadata assignment.

An overview of the future steps will be presented, particularly focusing on the integration of a JupyterHub system into the Kubernetes cluster that seamlessly communicates with the Data Lake.

Giorno preferito:

12 Dicembre Mattina

113

High rate analysis benchmarks

Author: Tommaso Tedeschi[1]

[1] *Università e INFN Perugia*

Corresponding Author: tommaso.tedeschi@pg.infn.it

The current High Rate Analysis platform offers a general purpose environment where analyzers can scale up their computations. The platform has proven to be able to support diverse use cases: recently, a CMS Coffea-based benchmark analysis has successfully been tested, adding up to the already tested ROOT's RDataFrame benchmark workflows, allowing for an initial comparison between the two

approaches. This presentation will focus on such comparison results, achieved on an Italian grid site, seamlessly integrated in the platfrom via interLink[1].

[1] https://github.com/interTwin-eu/interLink/tree/main

Giorno preferito:

11 Dicembre Pomeriggio

114

Data traceability model implementation in the Rucio data lake

Author: Domingo Ranieri[1]

Co-authors: Alessandro Costantini [1]; Barbara Martelli [1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Authors: barbara.martelli@cnaf.infn.it, alessandro.costantini@cnaf.infn.it, domingo.ranieri@cnaf.infn.it

In recent years, blockchain has emerged as a promising technology for managing trusted information and facilitating the management of critical data by businesses and the public while maintaining high levels of security. Permissioned blockchains, unlike permissionless ones, restrict access to a select group of authorized entities, ensuring a controlled and secure environment. It is particularly valuable in sectors where data sensitivity and confidentiality are crucial.

As part of the Innovation Grant named "Interoperable Data Lake" (IDL), WP4 aims at leveraging blockchain to ensure the traceability and validity of the data and related metadata stored in the data lake. In such respect, a private blockchain has been defined and implemented in a way that stores the hash of the file's content, the hash of its associated metadata, and the Data IDentifier (DID) provided by Rucio (the CERN solution for distributed data management) for each file present in the data lake. During data retrieval, this information is read from the blockchain and used to verify the consistency of the data.

In this regard, Hyperledger Fabric was carefully evaluated during the definition process and subsequently implemented into the software platform to enable distributed ledger implementation. Hyperledger Fabric allows the creation of smart contracts in various programming languages (i.e. Go, JavaScript, and Java) and features a modular architecture to customize the components of the blockchain network.

As part of the integration process, a plugin to perform read and write operations on the blockchain (triggered by the Rucio client whenever predefined events occur) has been developed and integrated into Rucio.

The solutions implemented in the above-described software platform aimed at leveraging blockchain to ensure data traceability and validity, together with the related integration process, will be presented and described during the presentation.

Giorno preferito:

12 Dicembre Mattina

115

Graph Anomaly Detection with GNNs: A Case Study on Predictive Maintenance for ENI's Industrial Machinery

Authors: Elisabetta Ronchieri[1]; Giovanni Zurlo[1]; Roberto Cornali[1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Authors: robertocornali15@outlook.com, elisabetta.ronchieri@cnaf.infn.it, giovanni.zurlo@cnaf.infn.it

In mission-critical industries, continuous operation exposes equipment to wear that can escalate into costly failures and downtimes. Early anomaly detection is, therefore, essential to maintaining seamless operations. Graph Neural Networks (GNNs) are emerging as a powerful tool for predictive maintenance, offering unparalleled understanding of data from complex interconnected systems of components found in industrial sectors. This talk presents a case study focused on ENI S.p.A.'s rotating machinery, such as high-performance gas compressors and turbines.

Our approach combines Generative Adversarial Networks, a reconstruction-based framework, and GNNs for effective, interpretable anomaly detection. By representing multivariate time series data as graphs, we employ an adversarially regularized graph auto-encoder to capture normal operating conditions. This design enhances interpretability by allowing users to visualize network connections and pinpoint anomalous nodes, leading to actionable insights and improved maintenance strategies.

Giorno preferito:

11 Dicembre Pomeriggio

116

Declarative framework for analysis definition and implementation

Authors: Alberto Annovi[1]; Andrea Rizzi[1]; Paolo Mastrandrea[1]; Tommaso Boccali[1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Authors: andrea.rizzi@pi.infn.it, paolo.mastrandrea@cern.ch, tommaso.boccali@pi.infn.it, alberto.annovi@pi.infn.it

The software toolbox used for "big data" analysis in the last few years is rapidly changing. The adoption of software design approaches able to exploit the new hardware architectures and improve code expressiveness plays a pivotal role in boosting data processing speed, resources optimisation, analysis portability and analysis preservation.
The scientific collaborations in the field of High Energy Physics (e.g. the LHC experiments, the next-generation neutrino experiments, and many more) are devoting increasing resources to the development and implementation of bleeding-edge software technologies in order to cope effectively with always growing data samples, pushing the reach of the single experiment and of the whole HEP community.

The introduction of declarative paradigms in the analysis description and implementation is growing interest and support in the main collaborations. This approach can simplify and speed-up the analysis description phase, support the portability of the analyses among different datasets/experiments and strengthen the preservation and reproducibility of the results. Furthermore this approach, providing a deep decoupling between the analysis algorithm and back-end implementation, is a key element for present and future processing speed, potentially even with back-ends not existing today.

A framework characterised by a declarative paradigm for the analysis description and able to operate on datasets from different experiments is under development in the frame of the ICSC (Centro

Nazionale di Ricerca in HPC, Big Data and Quantum Computing, Italy). The Python-based demonstrator provides a declarative interface for the implementation of an analysis of HEP event data, with support for different input data formats and for the extension to longer chains of analysis tasks. Status and development plan of the demonstrator will be discussed.

Giorno preferito:

117

LoopIn: a code for automating scattering amplitudes calculations

Authors: Jonathan Ronca[1]; Manoj Kumar Mandal[1]; Marco Bigazzi[2]; Pierpaolo Mastrolia[1]; William J. Torres Bobadilla[3]

[1] *Istituto Nazionale di Fisica Nucleare*

[2] *rwth Aachen University - Department of physics*

[3] *Max-Planck-Institute for Physics*

Corresponding Authors: jonathan.ronca@na.infn.it, manojkumar.mandal@pd.infn.it, bigazzi@physik.rwth-aachen.de, pierpaolo.mastrolia@pd.infn.it, torres@mpp.mpg.de

We present "LoopIn", a framework aimed at automating calculations of multi-loop scattering amplitudes for pQFT. It has been designed to have only few user input (process, number of loop, ps-points), from which it can provide numerical values for interference terms, helicity amplitudes or form factors.
"LoopIn" implements a modular structure, for which public codes can be interfaced as modules, for performing individual calculation procedures. It also relies on additional built-in methods to provide the minimal calculation setup and trivial parallelization.

Giorno preferito:

11 Dicembre Pomeriggio

118

Super-Resolution Surrogate Model for Accelerated Geant4 Simulations

Author: Giuseppe Gallo[1]

Co-authors: Giuseppe A. P. Cirrone [2]; Serena Fattori [3]; Valentina Ientile [4]; Alberto Sciuto [3]; Alessia Rita Tricomi [5]

[1] *Università degli Studi di Catania, Catania, Italia*

[2] *Istituto Nazionale di Fisica Nucleare, Laboratori Nazionali del Sud; Centro Siciliano di Fisica Nucleare e di Struttura della Materia, Catania, Italia*

[3] *Istituto Nazionale di Fisica Nucleare, Laboratori Nazionali del Sud, Catania, Italia*

[4] *Istituto Nazionale di Fisica Nucleare, Sezione di Catania, Catania, Italia*

[5] *Università degli Studi di Catania; Istituto Nazionale di Fisica Nucleare, Sezione di Catania; Centro Siciliano di Fisica Nucleare e di Struttura della Materia, Catania, Italia*

Corresponding Authors: giuseppe.gallo@unict.it, sciuto.a@lns.infn.it, cirrone@lns.infn.it, alessia.tricomi@ct.infn.it, fattori@lns.infn.it, valentina.ientile@ct.infn.it

WP6-Spoke2 of ICSC - Italian Research Center on HPC, Big Data and Quantum Computing has a use case dedicated to the integration of machine learning models to enhance Geant4. This is a well-known simulation framework in medical physics that can reproduce particle interactions down to the micrometer scale and below. However, the resources required scale linearly with the complexity of the system being simulated, limiting its effectiveness. Hadrontherapy has been identified as a crucial initial application. It provides users with a valuable tool for calculating dose and linear energy transfer distributions in water or other materials. Our work begins with the generation of a high-fidelity data set using a dense voxel-eluted water phantom to provide a reference for primary and secondary particle interactions. A generative model is then developed and trained on this dataset to reproduce the inherent spatial correlations observed. The goal is to develop a super-resolution surrogate model that can significantly improve the resolution of prediction from lower-resolution input data, with a consistent reduction in the resources required.

Giorno preferito:

119

BoGEMMS-HPC: development of Geant4 simulations in High-Performance Computing environments

Authors: Alex Ciabattoni[1]; Valentina Fioretti[2]; Simone Lotti[3]

[1] *Università di Bologna, INAF OAS Bologna*

[2] *INAF OAS Bologna*

[3] *INAF/IAPS*

Corresponding Authors: simone.lotti@iaps.inaf.it, valentina.fioretti@inaf.it, alex.ciabattoni@inaf.it

Geant4 is a C++ toolkit library for the transport of particles through matter and it is commonly used for the simulation of high energy space missions, allowing for the evaluation of their performance and driving the instrument design. The increasing complexity of the new technology involved for the high-energy Universe observation requires the development of modern large-scale simulations and presents new challenges in managing increasingly large datasets, as well as higher demands in storage and computational resources. This evolving landscape drives the implementation of next-generation High-Performance Computing (HPC) techniques and efficient I/O transfers for the Geant4 simulations. We propose the development of an open-source multi-threading (MT) and multi-node Geant4-based simulation framework with ad-hoc I/O interfaces (e.g. run-time input configuration, databases) based on the Bologna Geant4 Multi-Mission Simulator (BoGEMMS), an astronomy-oriented Geant4-based application developed at INAF OAS. The BoGEMMS-HPC framework uses the Geant4 built-in MT library to evenly distribute the simulation events (i.e. primary particles) on different threads, which write the output to a common FITS file or SQLite database. The event-level multi-node parallelism is instead achieved using the (Open MPI) G4MPI library, where the tasks write to independent output files. The performance speed-up has been tested using the CINECA Leonardo resources. Exploiting the 112 cores of a single Leonardo node, we achieved a maximum speed-up (with respect to a serial execution) of about 70 using G4MPI and of about 12 using MT, whose main bottleneck derives from the simultaneous writing from the threads to the same file. We used as a test case the simulation of the Anti-Coincidence System of COSI, a NASA Small Explorer satellite mission with a planned launch in 2027. In this presentation, we show the results on the performance tests, the advantages/bottlenecks between different parallelization strategies and output data formats, and the development plans for the future.

Giorno preferito:

120

Hyperparameter Optimization for Deep Learning Models Using High-Performance Computing

Author: Muhammad Numan Anwar[1]

Co-authors: Domenico Diacono ; Francesco Grancagnolo ; Guang Zhao ; Linghui Wu ; Marcello Abbrescia ; Mingyi Dong ; Nicola De Filippis ; Shengsen Sun

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Authors: nicola.defilippis@ba.infn.it, muhammad.anwar@ba.infn.it

Clusters counting in a drift chamber represents a highly promising breakthrough in particle identification (PID) techniques for particle physics experiments. In this paper, we trained neural network models, including a Long Short-Term Memory (LSTM) model for the peak-finding algorithm and a Convolutional Neural Network (CNN) model for the clusterization algorithm, using various hyperparameters such as loss functions, activation functions, numbers of neurons, batch sizes, and different numbers of epochs. These models were trained utilizing high-performance computing (HPC) resources provided by the ReCas computing center. The best LSTM peak-finding model was selected based on the highest area under the curve (AUC) value, while the best CNN clusterization model was chosen based on the lowest mean square error (MSE) value among all configurations. The training was conducted on momentum ranges from 0.2 to 20 GeV.
The trained models (LSTM and CNN) were subsequently tested on samples with momenta of 4 GeV/c, 6 GeV/c, 8 GeV/c, and 10 GeV/c. The simulation parameters included 10% Helium (He) and 90% Isobutane (C4H10), a cell size of 0.8 cm, a sampling rate of 2 GHz, a time window of 800 ns, 5000 events, and a 45-degree angle between the muon particle track and the z-axis (sense wire) of the drift tube chamber. The testing aimed to evaluate the performance of the LSTM model for peak finding and the CNN model for clusterization.

Giorno preferito:

121

SAIFIN (Satellite data and Artificial Intelligence for FINtech)

Author: Alberto Garinei[1]

[1] *Università Guglielmo Marconi - Dipartimento di Scienze Ingegneristiche*

Corresponding Author: a.garinei@unimarconi.it

One of the most promising methods for evaluating trading strategies relies on the analysis of satellite data. The primary objective of SAIFIN (Satellite data and Artificial Intelligence for FINtech) is to develop an AI-based algorithmic trading system capable of identifying financial trading strategies by leveraging information retrieved from both web sources and satellite data.
The system will be organized into three main components: data acquisition and preprocessing, nowcasting and predictive analysis, and execution and management of trading strategies. The first component focuses on collecting and processing raw data from satellite imagery—such as optical images, radar, and multispectral sensors—to assess macroeconomic indicators like agricultural production, logistics, and energy consumption. It also gathers web-based data, including economic news, social media sentiment analysis, and financial market feeds. Advanced AI technologies like Convolutional Neural Networks (CNNs) for image processing and Transformer models for Natural Language Processing (NLP) are employed, alongside distributed high-performance computation (GPUs) to handle large volumes of data efficiently.
The second component involves nowcasting and predictive analysis to generate forecasts of economic variables and financial markets across various time scales. AI models such as Long Short-Term Memory (LSTM) networks and Gated Recurrent Units (GRUs) model temporal relationships in economic data. GPU-accelerated frameworks like PyTorch and TensorFlow are utilized for efficient computation, enabling high-frequency forecasts for intraday trading, short-to-medium-term

forecasts for swing trading, and long-term analyses for strategic investments.
The third component implements automated trading strategies based on signals generated by the predictive analysis module. It features a decision engine using Reinforcement Learning (RL) models like Deep Q-Networks (DQN) and Proximal Policy Optimization (PPO) to optimize buying and selling decisions. Simulation and backtesting are conducted to evaluate strategies, analyzing risks and performance. Automation is achieved through integration with financial broker APIs for real-time order execution. The system manages multiple time scales, from high-frequency trading requiring rapid forecasting supported by GPUs to medium-term trading that identifies seasonal or cyclical trends, as well as strategic investments employing ensemble approaches that combine AI models with traditional indicators.
This approach, which combines satellite data analysis with high-performance computing, enables nowcasting of critical variables to achieve better outcomes in high-frequency trading (HFT), short-term, and medium-to-long-term trading scenarios. The integration of standard commercial models with insights derived from the nowcasting trading system is expected to yield superior results, particularly during exogenous shocks, such as global pandemics or conflicts with significant geopolitical implications.
The need for effective trading strategies arises from the complex dynamics that drive financial markets, making it challenging to devise robust approaches. The execution of models involving large datasets and multiple variables requires substantial computational power, parallelization, and efficiency. To address this, the project will develop algorithms specifically designed for GPU acceleration. The system's architecture enables scalable and flexible operations, adaptable to the rapid and unpredictable dynamics of financial markets, ensuring that trading can be conducted across different time scales, from intraday to longer-term periods.

Giorno preferito:

11 Dicembre Pomeriggio

122

The Frequency-Hough project: algorithm optimization and HPC for present and future coninuous gravitational-wave searches

Author: Lorenzo Pierini[1]

Co-authors: Pia Astone [1]; Marco Serra [1]; Stefano Dal Pra [1]; Cristiano Palomba [1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Authors: stefano.dalpra@cnaf.infn.it, pia.astone@roma1.infn.it, marco.serra@roma1.infn.it, lorenzo.pierini@roma1.infn.it, cristiano.palomba@roma1.infn.it

The search for continuous and persistent gravitational waves emitted by isolated and rotating neutron stars is a top priority for current and future ground-based detectors. However, those searches are typically bounded in sensitivity by their high computational costs. In this talk, I will introduce the flagship use case devoted to the Frequency-Hough algorithm, which performs a blind search for unknown sources from any position in the sky. I will show the status of the project at Milestone 9 and the perspectives for the next months.

Giorno preferito:

123

GWTboost: boosting the efficiency of gravitational–wave transient detections

Authors: Agata Trovato[1]; Edoardo Milotti[1]; Giacomo Principe[1]; Giuseppe Troian[2]

[1] *Istituto Nazionale di Fisica Nucleare*

[2] *University of Trieste*

Corresponding Authors: agata.trovato@ts.infn.it, giacomo.principe@inaf.it, edoardo.milotti@ts.infn.it, giuseppe.troian@phd.units.i

Unmodeled methods are extensively employed to detect generic gravitational wave transients (GWTs), including signals that lack precise theoretical modeling. Interestingly, these methods also demonstrate competitive efficiency when applied to well-modeled gravitational wave signals, such as those from compact binary coalescences (CBCs), rivaling the widely used matched-filter approach. In this presentation, I will discuss the GWTboost non-flagship use-case, which seeks to enhance the detection and post-processing of gravitational wave transients. This project focuses on advancing core detection algorithms, optimizing them for high-performance computing environments, improving waveform reconstruction techniques, and refining CBC parameter estimation during post-processing. These developments aim to strengthen the robustness and accuracy of current gravitational wave analysis pipelines.

Giorno preferito:

124

FaDER - Assisting real-time track reconstruction for LHC experiments

Authors: Alessandro Zaio[1]; Andrea Coccaro[1]; Carlo Schiavi[1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Authors: andrea.coccaro@ge.infn.it, alessandro.zaio@ge.infn.it, carlo.schiavi@ge.infn.it

The real-time track reconstruction task for LHC experiments shows a processing time which increases significantly as a function of the average number of proton-proton collisions per bunch crossing. The future upgrade to the High-Luminosity LHC (HL-LHC), with way higher levels of simultaneous collisions, could thus lead to a considerable growth in computational cost for the current trigger algorithms. To face this issue, a machine-learning-based technique to assist tracking by filtering out background hits is presented and characterized, as part of the FaDER project. The algorithm is based on a Convolutional Neural Network architecture, to target final deployment on FPGA boards.

Giorno preferito:

12 Dicembre Mattina

Lightning Talks / 125

Lightning Talk: Heavy flavours jet substructure

Corresponding Author: andrea.ghira@ge.infn.it

Lightning Talks / 126

Lightning Talk: SAIFIN (Satellite data and Artificial Intelligence for FINtech

Lightning Talks / 127

Lightning Talk: FaDER - Assisting real-time track reconstruction for LHC experiments

Corresponding Authors: carlo.schiavi@ge.infn.it, andrea.coccaro@cern.ch

Giorno preferito:

Lightning Talks / 128

Lightning Talk: The AIDA project: galaxy formation in alternative dark matter models

Corresponding Author: giulia.despali@unibo.it

Lightning Talks / 129

Lightning Talk: Lattice QCD in the exascale computing era

Corresponding Author: marco.ce@unimib.it

Lightning Talks / 130

Lightning Talk: Machine Learning for event reconstruction in Super-Kamiokande

Corresponding Author: nicola.calabria@ba.infn.it

Lightning Talks / 131

Lightning Talk: Optimal use of timing measurement in vertex reconstruction at CMS

Corresponding Author: ksenia.de.leo@ts.infn.it

Lightning Talks / 132

Towards virtual painting recolouring using Vision Transformer on X-Ray Fluorescence datacubesLightning Talk:

Corresponding Author: alessandro.bombini.fi@gmail.com

Lightning Talks / 133

Lightning Talk: An API for training Physics-Informed Kolmogorov-Arnold Networks using Nvidia Modulus SYM

Corresponding Author: alessandro.bombini.fi@gmail.com

Lightning Talks / 134

Lightning Talk: Implementation of low latency, fast inference neural networks on FPGA for trigger-like systems

Corresponding Author: salvatore.loffredo@na.infn.it

Lightning Talks / 135

Lightning Talk: Quasi interactive analysis of big data with high throughput - where are we now?

Corresponding Author: francesco.giuseppe.gravili@le.infn.it

Lightning Talks / 136

Lightning Talk: Building a Versatile HPC Data Center: a STILES lesson

Corresponding Author: fabio.ragosta@unina.it

Lightning Talks / 137

Lightning Talk: Accelerating LHCb VELO U2 Design with Machine Learning: A Fast Simulation Approach for the TimeSPOT Sensor

Corresponding Author: zhihua.liang@ca.infn.it

Lightning Talks / 138

Lightning Talk: High rate analysis benchmarks

Corresponding Author: tommaso.tedeschi@pg.infn.it

Lightning Talks / 139

Lightning Talk: Graph Anomaly Detection with GNNs: A Case Study on Predictive Maintenance for ENI's Industrial Machinery

Corresponding Author: giovanni.zurlo@cnaf.infn.it

Lightning Talks / 140

Lighning Talk: Declarative framework for analysis definition and implementation

Corresponding Author: paolo.mastrandrea@cern.ch

Lightning Talks / 141

Lightning Talk: LoopIn: a code for automating scattering amplitudes calculations

Corresponding Author: jonathan.ronca@na.infn.it

Lightning Talks / 142

Lightning Talk: Super-Resolution Surrogate Model for Accelerated Geant4 Simulations

Corresponding Author: giuseppe.gallo@unict.it

Lightning Talks / 143

Lightning Talk: BoGEMMS-HPC: development of Geant4 simulations in High-Performance Computing environments

Corresponding Author: alex.ciabattoni@inaf.it

Lightning Talks / 144

Lightning Talk: Hyperparameter Optimization for Deep Learning Models Using High-Performance Computing

Corresponding Author: muhammad.anwar@ba.infn.it

Lightning Talks / 145

Lightning Talk: The Frequency-Hough project: algorithm optimization and HPC for present and future continuous gravitational-wave searches

Corresponding Author: lorenzo.pierini@roma1.infn.it

Giorno preferito:

Lightning Talks / 146

Lightning Talk: GWTboost: boosting the efficiency of gravitational–wave transient detections

Corresponding Author: giacomo.principe@inaf.it

Lightning Talks / 147

Lightning Talk: Extending Rucio to support external metadata catalog

Corresponding Author: luca.pacioselli@pg.infn.it

Lightning Talks / 148

Lightning Talk: Data traceability model implementation in the Rucio data lake

Corresponding Author: domingo.ranieri@cnaf.infn.it

149

Satellite Data Management for advanced environmental applications

Authors: Alessia Rita Tricomi[1]; Enrico Chiarelli[2]; Fabio Mantovani[2]; Gioacchino Alex Anastasi[3]; Giuseppe Piparo[4]; Irem Nedime Elek[2]; Kassandra Giulia Cristina Raptis[2]; Matteo Albéri[2]; Virginia Strati[2]

[1] *Università di Catania & INFN Catania & CSFNSM*

[2] *Department of Physics and Earth Science, University of Ferrara & INFN, Sezione di Ferrara*

[3] *Università di Catania & INFN Catania*

[4] *Istituto Nazionale di Fisica Nucleare*

Corresponding Authors: alessia.tricomi@ct.infn.it, gialex.anastasi@dfa.unict.it, strati@fe.infn.it, giuseppe.piparo@ct.infn.it

Efficient satellite data management is essential to improve environmental monitoring and support various space economy applications. As part of Working Project 6 (WP6) 'Cross-Domain Initiatives and Space Economy'of Spoke 2, we developed a custom library designed to simplify the entire workflow from downloading to pre-processing and analysis of satellite data.

The library facilitates the creation and management of datasets using multispectral images from the Sentinel-2 satellites of the Copernicus constellation. We demonstrate its utility through two case studies: the segmentation of wildfire-affected areas and the early detection of vineyard diseases. In the first case, we integrate Sentinel-2 imagery with wildfire information from the Copernicus Emergency Management Service. In the second case, we combine satellite data with geolocation and disease impact assessments from recent airborne surveys in the context of the regional PERBACCO project.

To further enhance our analytical capabilities, we are expanding our methods to incorporate additional data sources from other satellites within the Copernicus constellation. This improvement allows for more comprehensive datasets. By integrating diverse satellite and aerial data through our custom library, we aim to enhance the scalability and effectiveness of environmental monitoring applications, contributing significantly to disaster response and sustainable agricultural practices.

Giorno preferito:

150

Deep learning applications for the analysis of Sentinel-2 satellite imagery

Authors: Alessia Rita Tricomi[1]; Enrico Chiarelli[2]; Fabio Mantovani[2]; Gioacchino Alex Anastasi[3]; Giuseppe Piparo[4]; Irem Nedime Elek[2]; Kassandra Giulia Cristina Raptis[2]; Matteo Albéri[2]; Virginia Strati[2]

[1] *Università di Catania & INFN Catania & CSFNSM*

[2] *Department of Physics and Earth Science, University of Ferrara & INFN, Sezione di Ferrara*

[3] *Università di Catania & INFN Catania*

[4] *Istituto Nazionale di Fisica Nucleare*

Corresponding Authors: alessia.tricomi@ct.infn.it, giuseppe.piparo@ct.infn.it, gialex.anastasi@dfa.unict.it, strati@fe.infn.it

In recent times, the combination of remote sensing and machine learning applications has led to great perspectives in the context of space economy. A flagship use case "AI algorithms for (satellite) imaging reconstruction" has been therefore established within the Working Group 6 (WP6) "Cross-Domain Initiatives and Space Economy" under Spoke 2, to focus on the analysis of satellite and aerial images.
In this contribution we presents the work accomplished in the segmentation of wildfire-affected areas by means of deep learning techniques, in particular Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) Networks, tested separately and together in a UNet-like architecture with ConvLSTM layers. A custom library, specifically designed to optimize the entire workflow from download to pre-processing and analysis, has been employed to create a dataset of multispectral images from the Sentinel-2 satellites, in combination with the information about wildfires provided by the Copernicus Emergency Management Service. The results are more than encouraging, proving the power of such methodologies in environmental monitoring and disaster response.
Moreover, the analysis of a dataset of satellite images for the early detection of vineyard diseases is currently underway. Such a dataset has been built using the geolocation data and disease impact assessments from a recent airborne survey (8-19 July '24) covering approximately 2000 hectares across the Emilia-Romagna region, performed in the context of the PERBACCO project. This collaboration aims to integrate high-resolution, airborne observations with satellite data to develop more robust, multi-scale monitoring solutions.

Giorno preferito:

Lightning Talks / 151

Lightning Talk: Satellite Data Management for advanced environmental applications

Corresponding Author: giuseppe.piparo@ct.infn.it

Lightning Talks / 152

Lightning Talk: Deep learning applications for the analysis of Sentinel-2 satellite imagery

Corresponding Author: gialex.anastasi@dfa.unict.it

153

Vector Boson Scattering with Machine Learning in Boosted Topologies at CMS

Author: Luca Della Penna[1]

[1] *Istituto Nazionale di Fisica Nucleare*

Corresponding Author: luca.dellapenna@pg.infn.it

Vector boson scattering (VBS) processes serve as a powerful probe for detecting potential deviations from the Standard Model (SM) of particle physics. Recently, there has been increasing interest in studying VBS, with a particular focus on incorporating the polarization states of the gauge bosons. While the reconstruction of fully leptonic final states is clearer, hadronic final states require a precise characterization of the final state topology.

At the CMS experiment, cutting-edge machine learning tools, such as ParticleNet and ParT, are employed to extract detailed jet properties, enabling more precise measurements of these processes. In this presentation, we will review the current status of VBS measurements, discuss the tools and techniques used at CMS, and outline our plans for the ongoing BOOST project (funded by Spoke2 - ICSC HUB) project, with a particular emphasis on VBS processes involving final states with jets in boosted topologies

Giorno preferito:

11 Dicembre Pomeriggio

Lightning Talks / 154

Lightning Talk: Vector Boson Scattering with Machine Learning in Boosted Topologies at CMS

Corresponding Author: luca.dellapenna@pg.infn.it

Sessione Pubblica / 155

Cosa e' il Centro Nazionale di Supercalcolo, Big Data e Quantum Computing ?

Lightning Talks / 157

A numerical framework to study topologically induced flow patterns in active emulsions

Corresponding Author: giuseppe.negro@ba.infn.it