# Welcome to ESC24

## the XV INFN International School on

# “Architectures, tools ad methodologies for developing efficient large scale scientific computing applications”

October 14 - 24  2024
Bertinoro (FC) Italy

Alberto Garfagnini – INFN-PD and Padova University

# On behalf of

## The Italian Institute for Nuclear Physics (INFN):

- Bologna, CNAF, Padova, providing most of the support

## The University of Bologna, Department of Physics and Astronomy

## The ESC24 lecturers

- and their Institutions

## Attendance this year

- 27 participants were selected
- 10 from foreign Institutions (9 from CERN)
- 17 from Italian Institutions

# Bertinoro in the past

Bertinoro is a nice medieval village, famous for its **hospitality** and therefore quite well suited as a location for a center like CeUB

- the name most likely comes from **“Castrum Brittinori”** (XI century) and it is probably due to the frequent stopping of pilgrims coming from Britain, in their way to Rome, used to take in the quiet surroundings of Bertinoro
- A legend says that Galla Placidia, daughter of the Roman emperor Theodosius I, drank local wine in a humble clay chalice and said

  “Non di così rozzo calice sei degno, o vino, ma di *berti in oro*”

- One of the monument in Bertinoro known as **Colonna delle Anelle** ("Column of the Rings" or "Column of hospitality") is a column in white stone with 12 rings erected in 1300 by the noble families to express **their commitment to hospitality**.
- Each one of the rings corresponded to a family
- foreigners arriving in town, could select the family to be hosted, by **tying** the horse bridles to the correspondingring

# Bertinoro today

Bertinoro still hosted a "**Hospitality Festival**". Held at the end of August or first week of September (August 29 – September 1, 2024)
it includes an entire  night of music, dances and events, some historic commemorations and the final Hospitality Rite (https://emiliaromagnaturismo.it/it/eventi/bertinoro-festa-ospitalita)

In this ceremony visitors **can be hosted** for lunch by a family in the town simply **taking one of the envelopes** tied to the rings of the Hospitality column (which inside has the name of the hosting family).

You will not get this opportunity this week, but I'm sure you will at least enjoy the **good food and wine** that Bertinoro will offer to you while you are staying here

# CeUB: University Residential Center of Bertinoro


# The CeUB facilities

- The Center was brought back in use with renovation work that started in 1991 in the fortress and in the guest house, followed by integration of additional buildings

- Today the Center offers:
  - **15 lecturing/meeting rooms** inside the Bishop's Fortress, the Rivellino, St. Sylvester church and the Theatre;
  - **2 computer labs**, 20 and 50 seats;
  - **86 bedrooms** (single, twin, double);
  - a **canteen with 200 seats** for breakfast, lunch and gala dinners
  - **120+ events** per year
  - up to **30.000** presences per year

# Why hosting the ESC school

- High Energy Physics has been **heavily relying** on computing since long time
  - for many years the scale of resources needed by HEP experiments was such that the computing centers of the hosting lab **were coping well** with the core computing data processing needs
  - the UA1/2 experiments that discovered the Z and W bosons at CERN were good examples
  - the model started to break at the end of the last century when experiments, like the BaBar experiment at the SLAC B-factory, had to deal with a **huge amount of data** and the computing power had to be scaled up by more than one order of magnitude w.r.t. **the initial estimates**
  - the investment needed started to grow very significant
  - for the first time it was felt necessary to distribute the processing of the data stored on tape to an external center

# The end of the “free ride”

Transistor counts follow an exponential growth

The highest transistor count in any IC chip is a deep learning engine called the Wafer Scale Engine 2 by Cerebras, using a special design to route around any non-functional core on the device: it has **2.6 trillion** MOSFETs, manufactured using TSMC's 7 nm FinFET process

Additional transistors will provide more cores

To benefit from future processors parallel workloads are mandatory

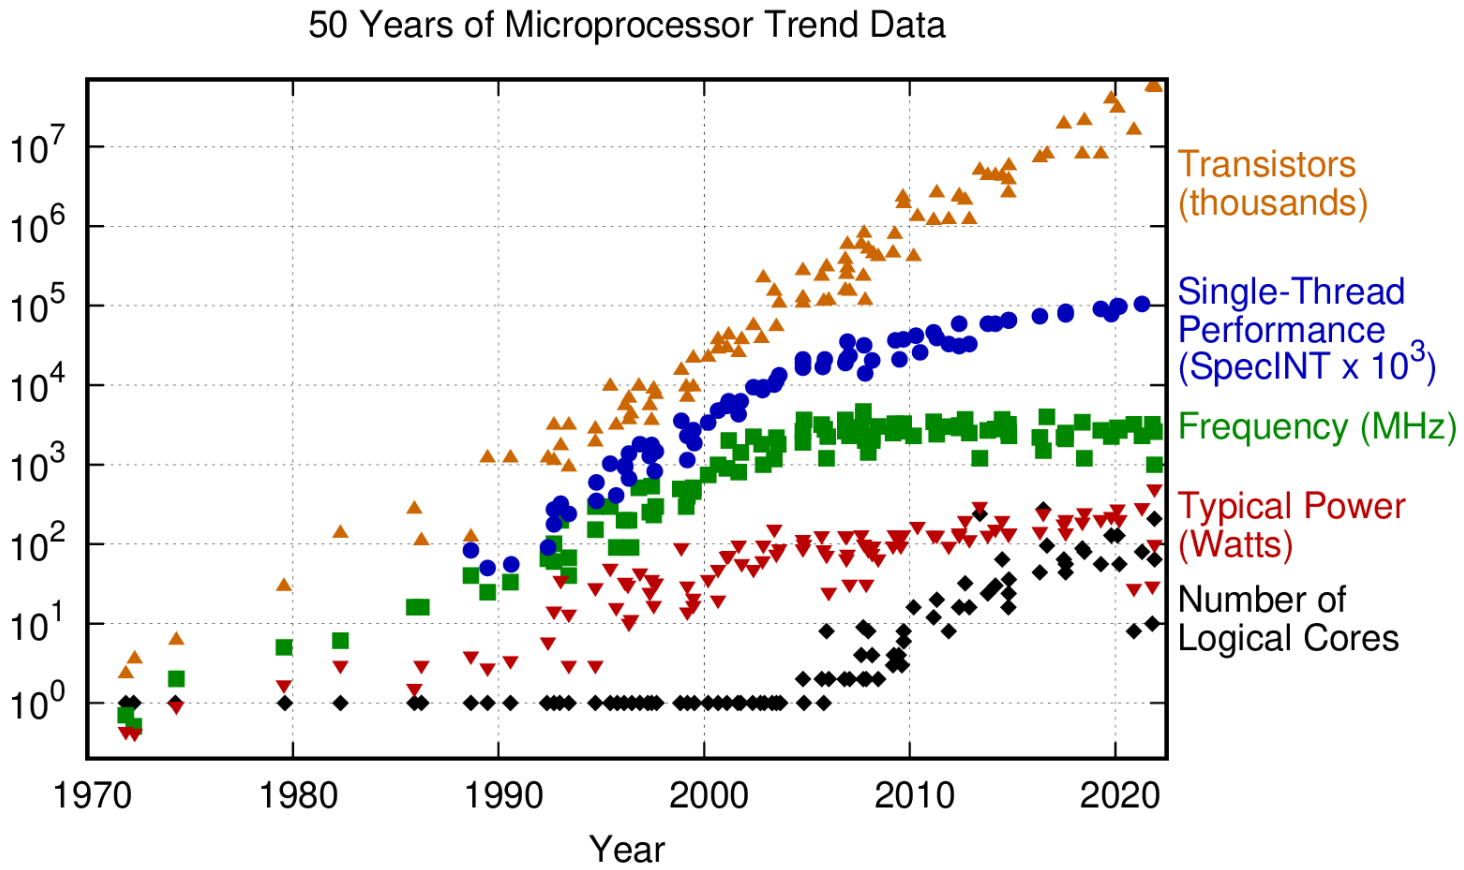


- Data and plot available at
- https://github.com/karlrupp/microprocessor-trend-data

# Cerebras WSE-3

*Fabrication process*
**5nm**

*Silicon area*
**46,225mm²**

*Transistors*
**4 Trillion**

*AI-optimized cores*
**900,000**

*Memory (on-chip)*
**44GB**

*Memory bandwidth*
**21PB/s**

*Fabric bandwidth*
**214Pb/s**

Cerebras WSE-3
4 Trillion Transistors
46,225 mm² Silicon

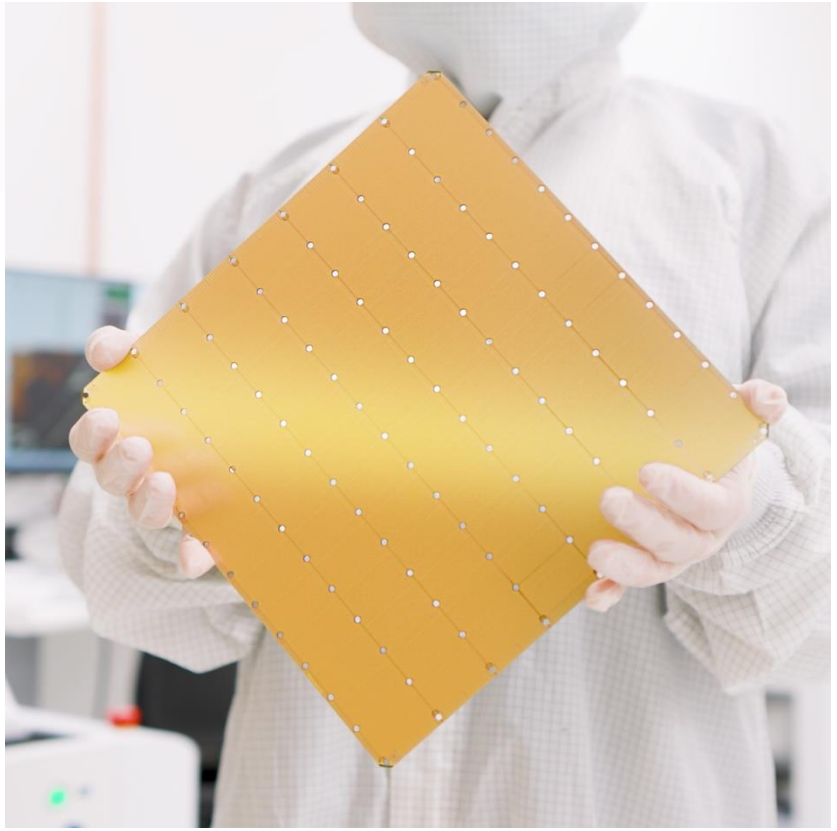

Largest GPU
80 BillionTransistors
814 mm² Silicon

The Cerebras Wafer Scale Engine (WSE) is a single, wafer-scale integrated processor that includes compute, memory and interconnect fabric.

# Why a computing school about efficiency

- The conception of this school was motivated by the awareness that efficient usage of computing resources in our field:
- had to be taken seriously, given the level of computing investment now required
- in the past was not always well understood and taken into proper consideration
- was becoming more and more challenging due in particular to the physical constraints in increasing scalar performance and the attempts to exploit anyway Moore's law with new processor architectures
  - many cores, co-processors,  GPU, vector units, new architectures (RISC-V), etc.
  - memory access getting more and more critical
  - ....

# Future will be more challenging

SKA Science Data Challenge: analyse a simulated datacube 1 TB in size, in order to find and characterise the neutral hydrogen content of galaxies across a sky area of 20 square degrees.

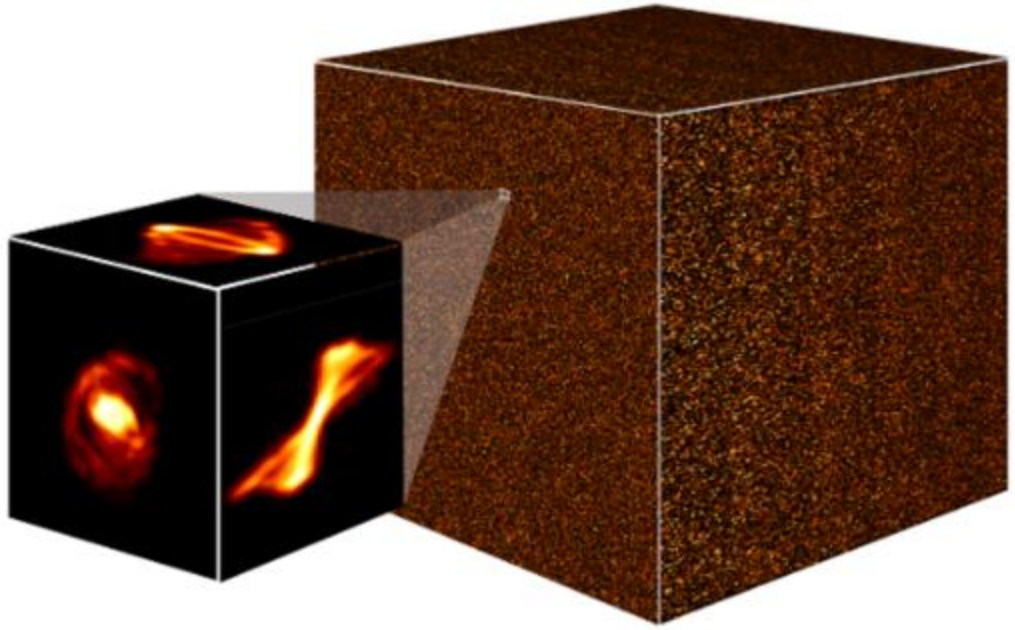

SKA 1TB data cube

https://www.skao.int/en

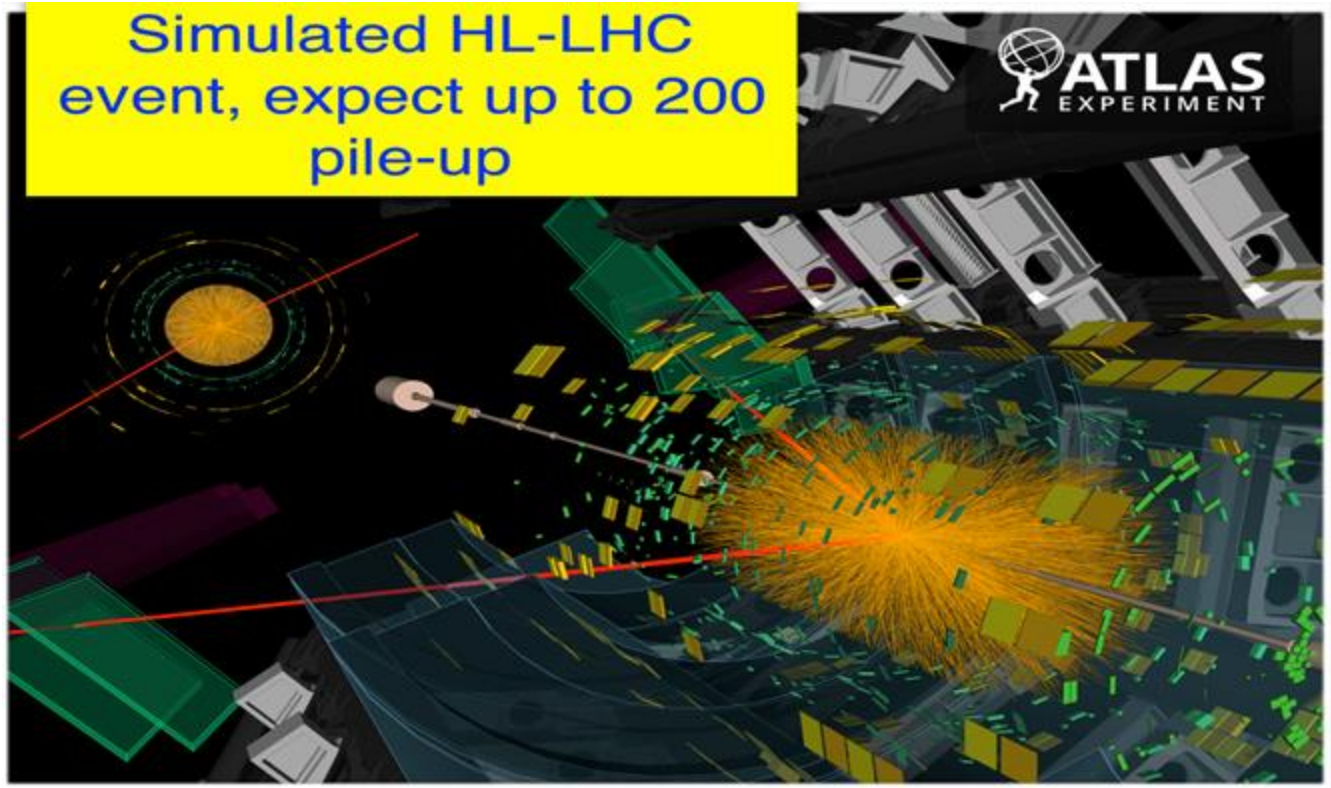


V. Kourlitis, "HL-LHC and Beyond Computing Challenges", https://cds.cern.ch/record/2861812/files/ATL-SOFT-SLIDE-2023-226.pdf

# TOP 500 HPC

- https://top500.org/lists/top500/2024/06/

| Rank | System | Cores | Rmax (PFlop/s) | Rpeak (PFlop/s) | Power (kW) |
|---|---|---|---|---|---|
| 1 | **Frontier** - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE<br>DOE/SC/Oak Ridge National Laboratory<br>United States | 8,699,904 | 1,206.00 | 1,714.81 | 22,786 |
| 2 | **Aurora** - HPE Cray EX - Intel Exascale Compute Blade, Xeon CPU Max 9470 52C 2.4GHz, Intel Data Center GPU Max, Slingshot-11, Intel<br>DOE/SC/Argonne National Laboratory<br>United States | 9,264,128 | 1,012.00 | 1,980.01 | 38,698 |
| 3 | **Eagle** - Microsoft NDv5, Xeon Platinum 8480C 48C 2GHz, NVIDIA H100, NVIDIA Infiniband NDR, Microsoft Azure<br>Microsoft Azure<br>United States | 2,073,600 | 561.20 | 846.84 | |
| 4 | **Supercomputer Fugaku** - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu<br>RIKEN Center for Computational Science<br>Japan | 7,630,848 | 442.01 | 537.21 | 29,899 |
| 5 | **LUMI** - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE<br>EuroHPC/CSC<br>Finland | 2,752,704 | 379.70 | 531.51 | 7,107 |

# CMS Patatrack R&D

Heterogeneous computing and performance portability a reality at LHC experiments

Hackathon focused on improving the performance portability of the CMS reconstruction software for charged particle trajectories.

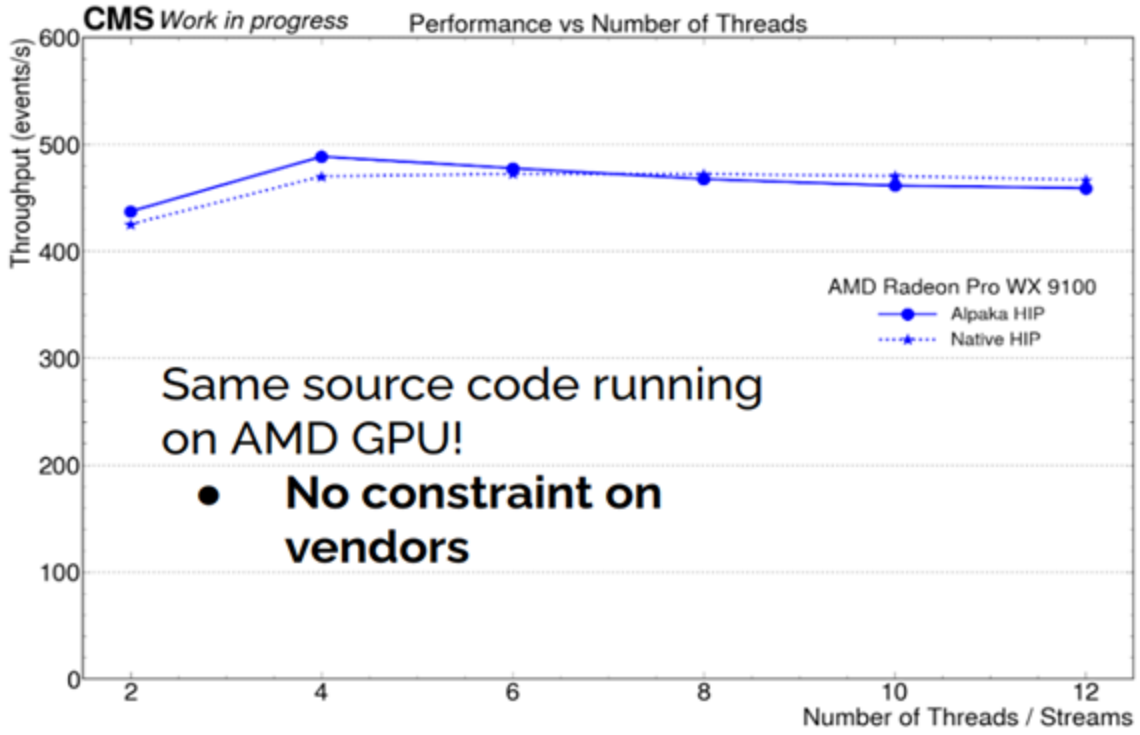


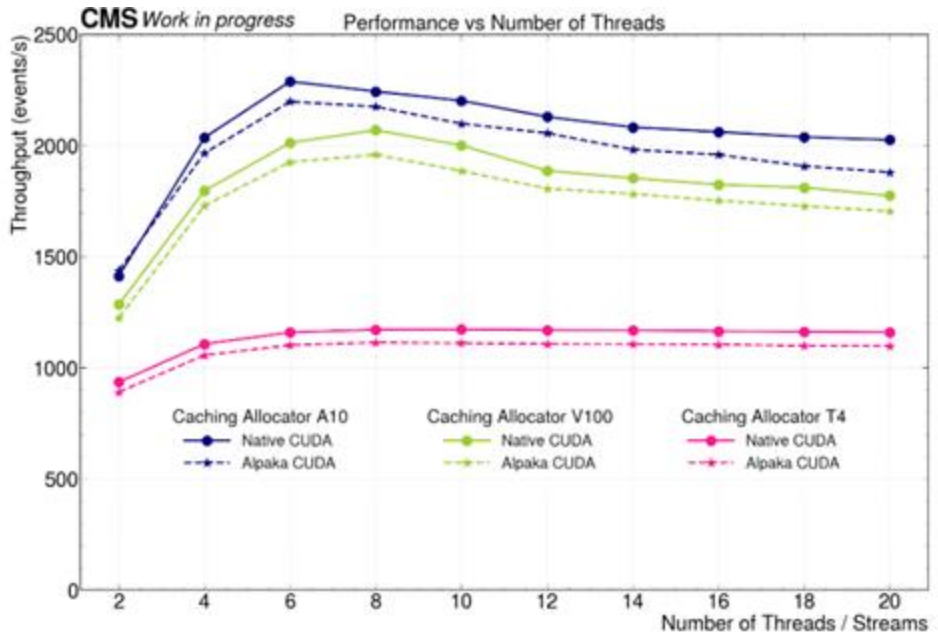


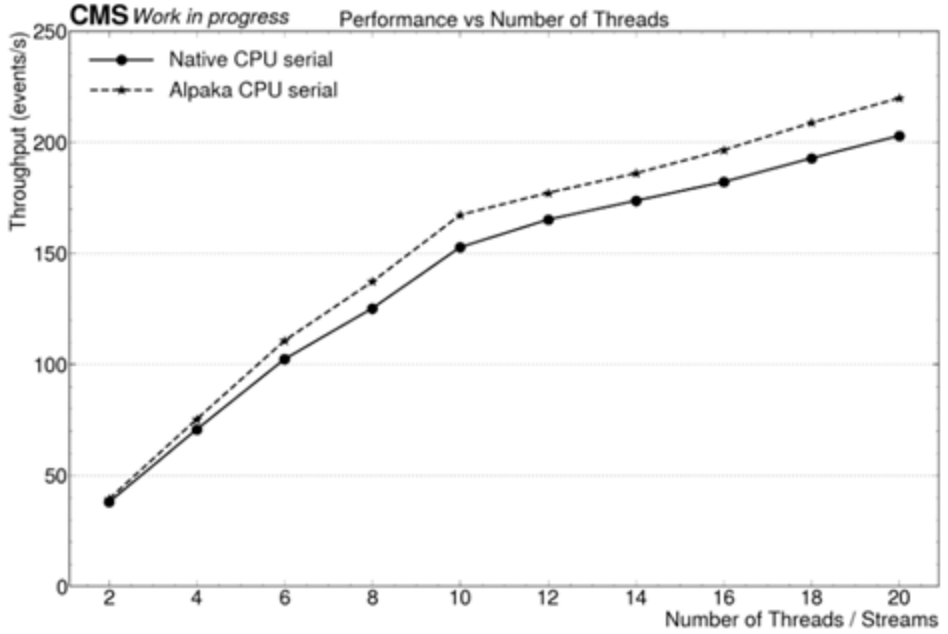

# A quick look at HPC

Exascale computing
- 1000 PFlops, 1000 Pbytes

However, **no supercomputer runs real applications faster than 10 percent of its maximum peak design speed**

The natural trend is towards **poorer and poorer efficiencies** as systems scale out to Exascale

“it is not **power** or **reliability** that are the **exascale challenges**: it’s **programmability of complex memory hierarchies**”
Bronis R. de Supinski Chief Technology Officer (CTO) for Livermore Computing

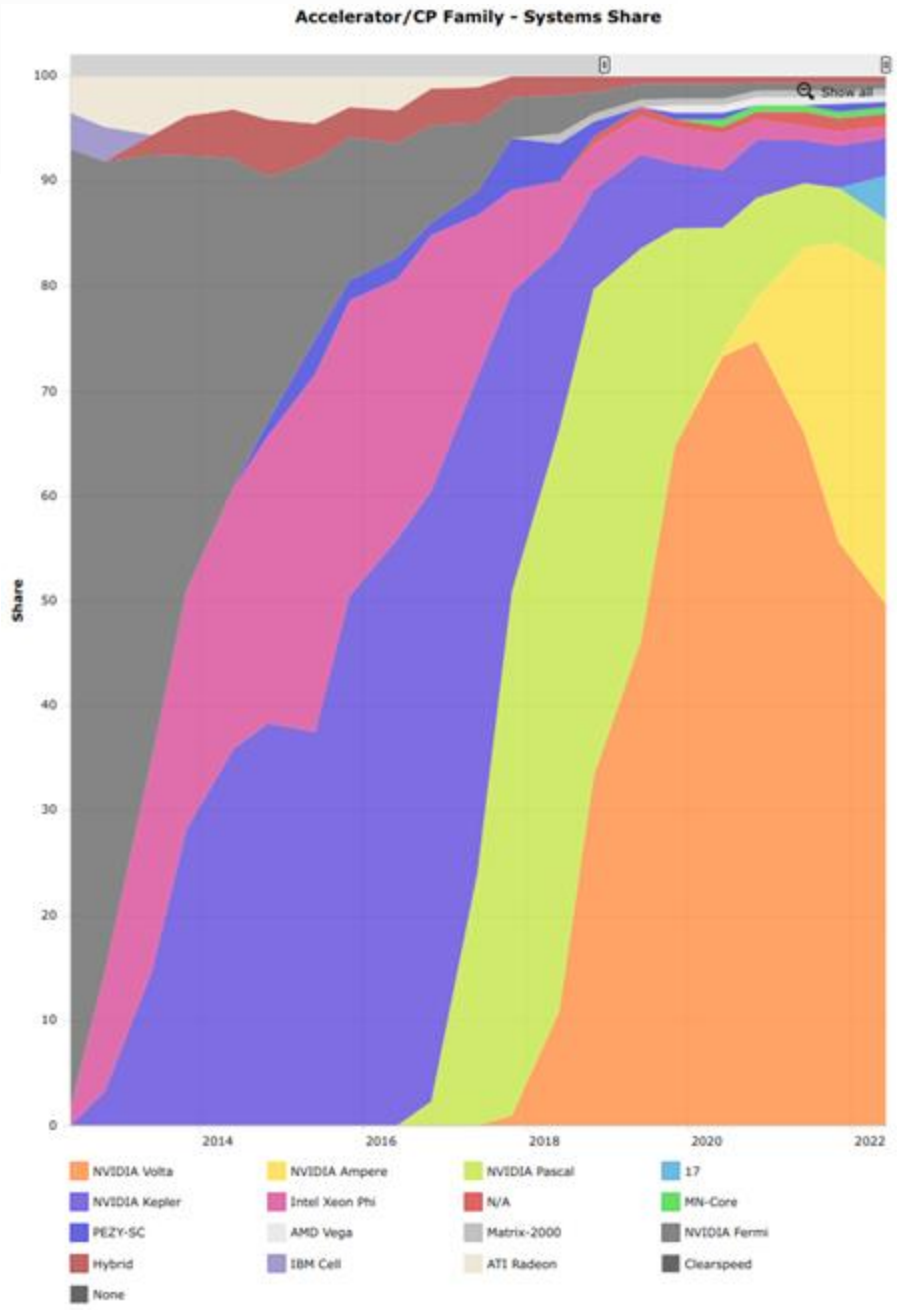

# The school evolution

- In the first editions of the School we tried to embrace **several key aspects** related to the efficient usage of computing resources in scientific applications, from **exploitation of modern CPUs** to **I/O related issues**
- however, we realized that the scope was too vast for a one-week-long school
- last year we experienced a longer (8 days) school to go deeper on few topics
- A longer school was welcomed by students, so this year you will have a longer time (11 days)
- a few years ago, we decided therefore to focus the School on the area where developments looked most disruptive and challenging: **parallel processing**

# ESC24 lecture plan in a nutshell

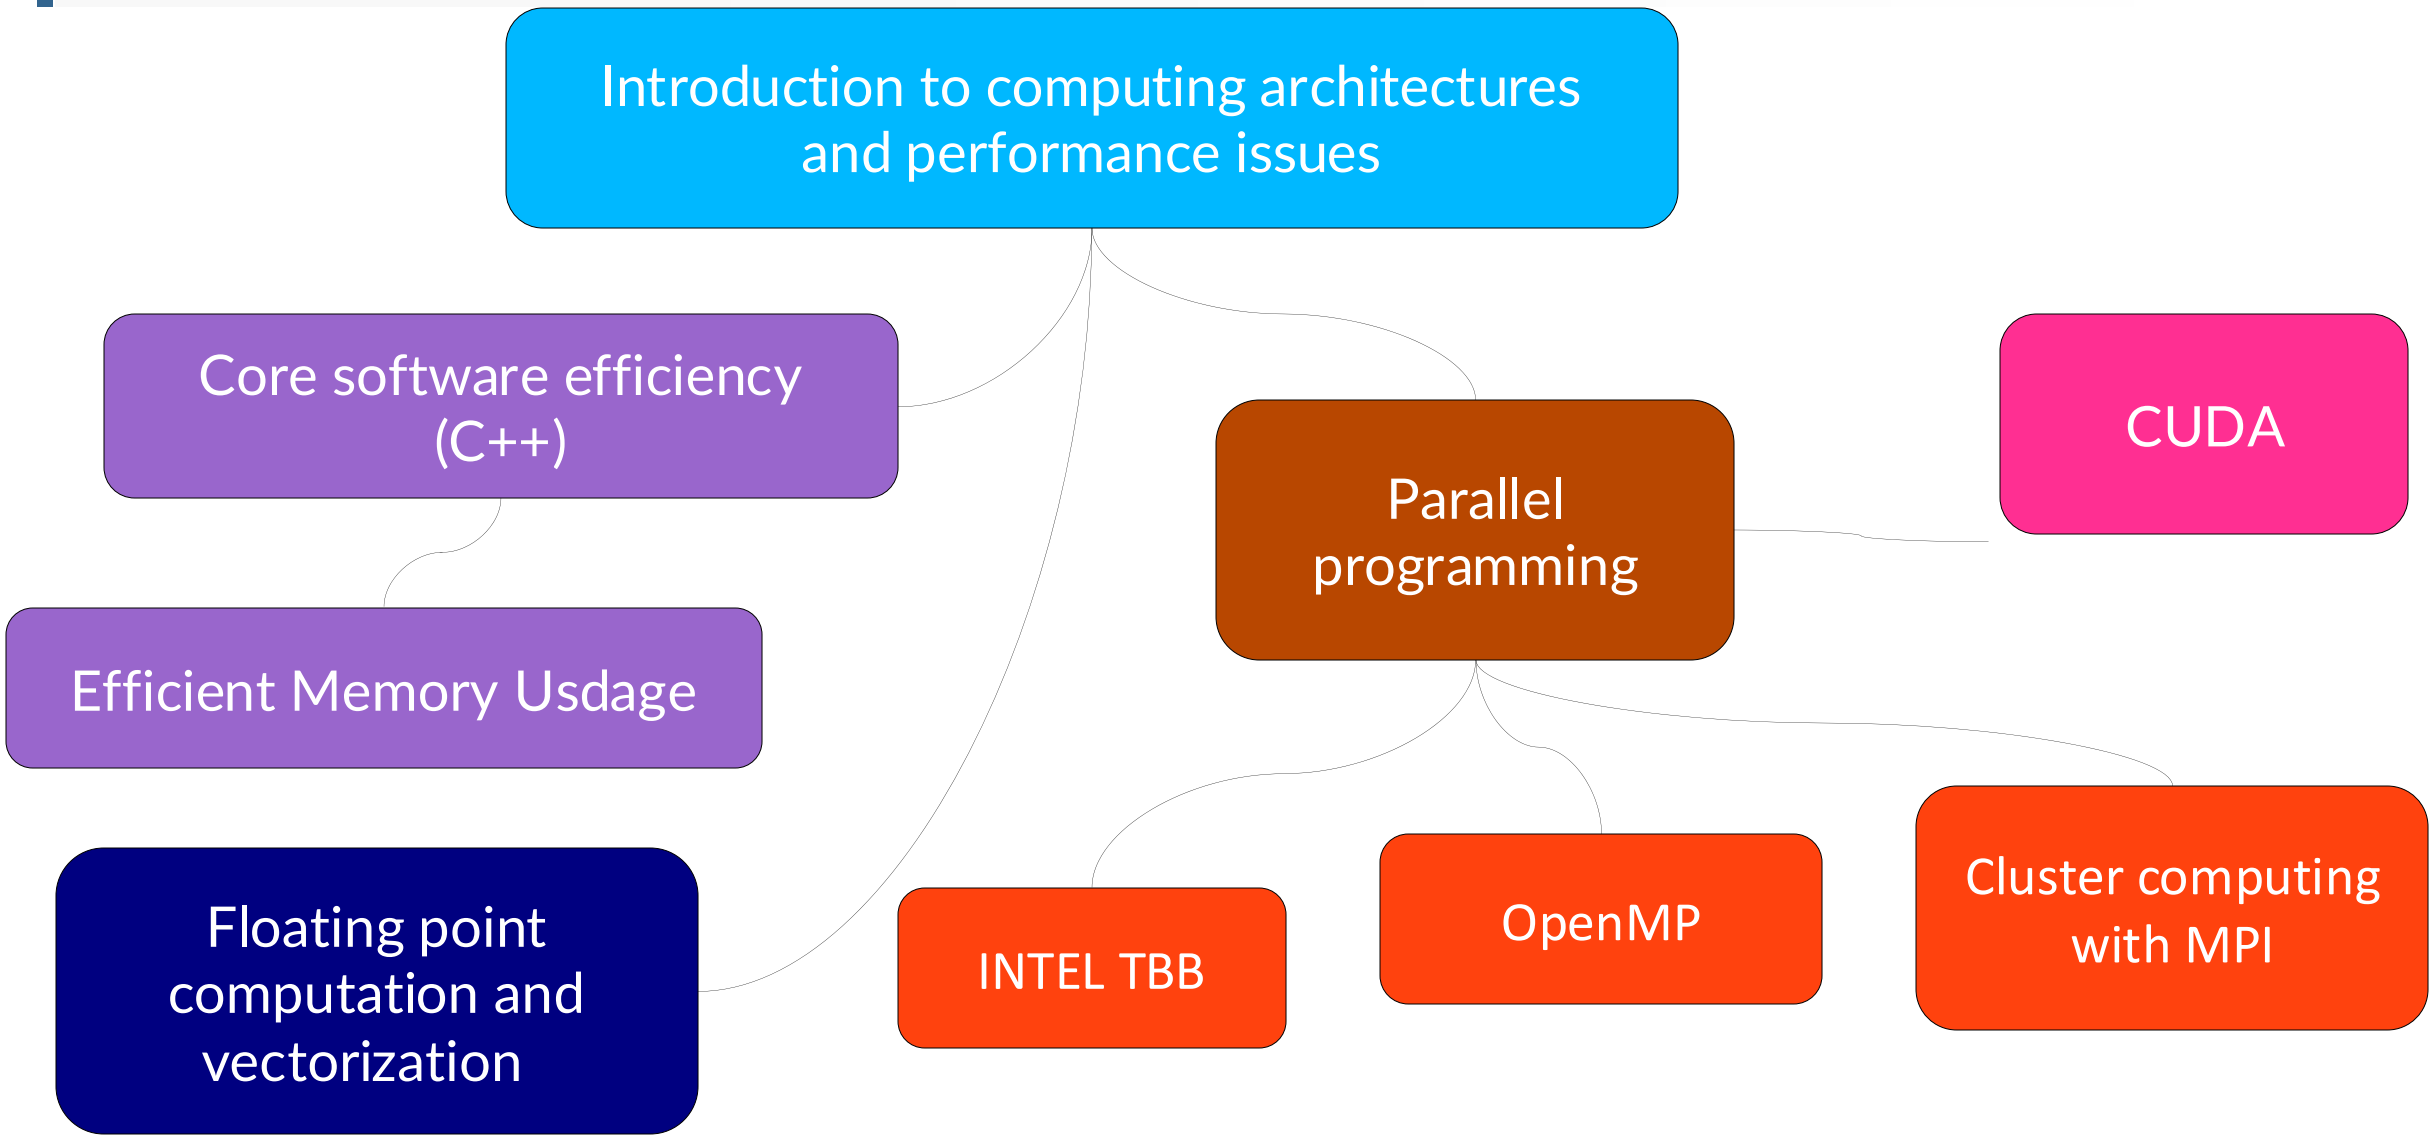

# Esc24 Social Event

**RAVENNA UNESCO SITES AND SOCIAL DINNER**

- 11.00 Departure from Bertinoro by bus
- 12.30 - 14.00 Free time for lunch
- 14.00 - 17.00 Ravenna guided visit in English
- 19.30 Social Dinner in Cervia at "Agriturismo Casa delle Aie"

# Evening Lectures

- lunedi' 21: "Opportunities and challenges of Artifical Intelligence", by Daniele Bonacorsi
- martedi' 22: "Use and abuse of random numbers" by Tim Mattson

+

- Dissemination lecture on underground physics, tbc

# Improving the School

- we have been striving to **improve the School** year after year
- the feedback we got from the students has always been **very useful**
- on Thursday, October 24 morning we will give you an opportunity to evaluate:
  - the **perceived quality** of various aspects of the School via the feedback questionnaire
  - the **competences you have acquired**
- through a final test
  - information collected is of course **very valuable to us too**

# Consolidation time

- to accommodate the need of students had to have some time for assimilating the material presented in the lectures there are some **consolidation time slots**

- **Almost every day, in the afternoon, you will have 2 hours to work on the exercises given by the lecturers**

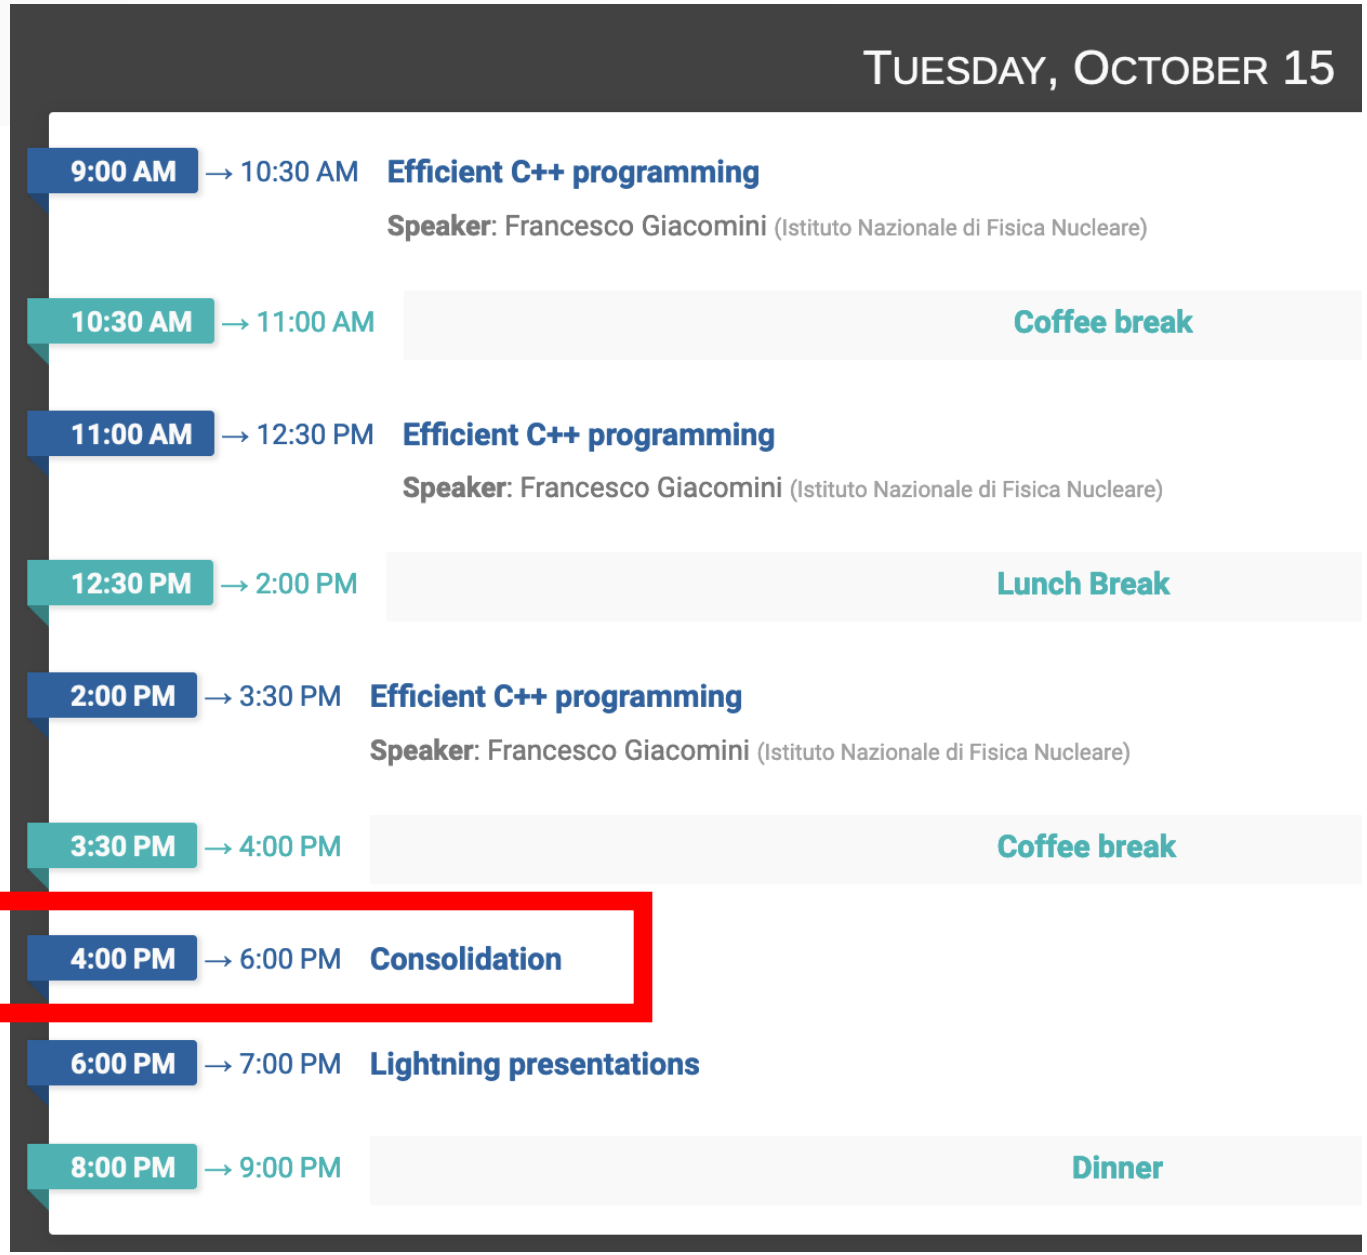

# School Computing Infrastructure

We believe that:

- **learning-by-doing** is an effective way of **learning**
- but it is also an effective way of **teaching**
- guiding students **to discover by themselves** new notions and concepts
- therefore a **computing infrastructure** suitable for supporting hands-on activities has been setup for you
- HPC cluster located at the CNAF INFN site

**This year:**

- each of you will get shared access, through a **login gateway** (bastion.cnaf.infn.it), to:
- **three Intel dual processors Linux (E5-2640v2 ivy bridge, 2x8 phys. cores 128 GB)** servers equipped with GPUs (2 x Tesla K20m/40m)
- one **Intel dual processor Linux (Gold 6148 sky lake, 2x20 phys. cores 192 GB)** server with 4xv100 GPU
- you should have received a mail message with the credentials to login into bastion; if not, check your spam folder, or let us know
- there is a dedicated web site with teaching material and exercises: **https://infn-esc.github.io/esc24/**

# Wireless and Social Networks

**Bertinoro wireless network** is available in all rooms including the guest house
[CeUB SSID and eduroam SSID]

## EMAIL

esc_contact@lists.infn.it

## TELEGRAM

Join the channel for logistic updates

ESC24 - https://t.me/+QBCNH8WvJnthYzc0

## INSTAGRAM

Follow ESC and tag the school

@infn_esc_school

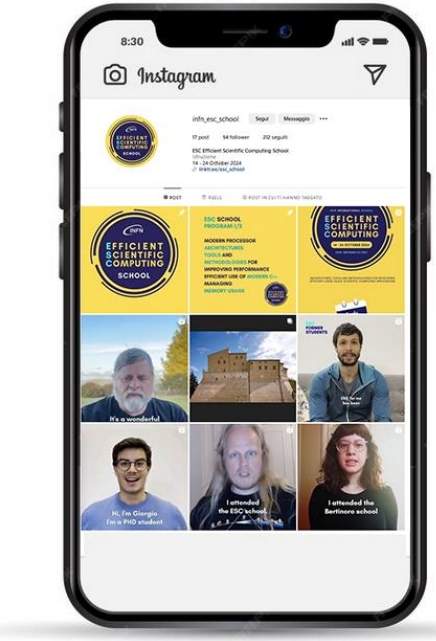

# We wish you a very pleasant and "efficient", great week in Bertinoro

# Backup