Joint APPEC/ECFA/NuPECC Computing Workshop

113th Plenary ECFA Meeting

Graeme A Stewart, CERN

Large Scale Scientific Computing in Europe

- JENA Symposium in Madrid (2022) discussions with funding agencies
  - Need to discuss strategy and implementation of European federated computing at future large-scale research facilities
  - Current situation is far from coherent
    - HEP: concept and implementation of HL-LHC computing is well established
      - Need to understand how the WLCG concept can be adapted to cope with the increased demands.
    - In nuclear physics, computing is currently organised an a facility basis
      - The community has limited access to the national computing centres
    - In astroparticle physics various, totally different, computing models for the distributed large-scale infrastructures exist
- Strong motivation for a workshop between the three communities to find *synergies* and *understand differences*
- **Scaling up coherently** for the next decades of science is the key point, thus:

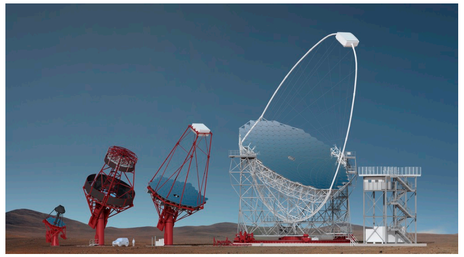

Joint APPEC/ECFA/NuPECC Computing Workshop in Bologna, June 2024

Key Topics to Cover

- What are the software and computing challenges in each field in the next decades?
- What European federated structures (will) exist?
  - How can we use them and work with them?
- How will the technology landscape evolve?
  - In what ways will this help or hinder our mission?
  - How can we train the next generation of people in the best tools and techniques?
    - And work with industry as well
- What federated computing models exist and how will they evolve?
  - How do software and distributed computing interact?
- What will be the impact of FAIR data and software policies; and an Open Science policy?

We have a European focus, with worldwide implications

Agenda...

14:30 **Challenges in (federated) Computing in Particle Physics**
Speaker: Concezio Bozzi (Istituto Nazionale di Fisica Nucleare)
20230612_JENA_H...

15:00 **Challenges in (federated) Computing in Astroparticle Physics**
Speaker: Stefan Schlenstedt (CTAO)
230612-JENAA-Sc...

15:30 **Challenges in (federated) Computing in Nuclear Physics**
Speaker: Johan Messchendorp (KVI-CART/University of Groningen)
JENA_Workshop_J...

European federated Computing

17:05 **The ESCAPE project for fundamental physics**
Speaker: Ian Bird (LAPP)
ESCAPE-JENA-Bol...

17:40 **Horizon Europe (EOSC...)**
remote presentation
Speaker: Giovanni Lamanna (LAPP)
JENA23_GLAMAN...

18:15 **Euro-HPC** ¶
remote presentation
Speaker: Evangelos Floros (EuroHPC JU)
EuroHPC - FLOROS....

18:50 **landscape of federated computing in the US**
Speaker: Frank K Wuerthwein (UCSD)
jena2023.pdf

(Federated) Computing Models (on national levels): (Federated) Computing Modles (on national levels)

11:00 **Introductionary statements by panel members**

11:30 **Round Table Disucssion**
Speakers: Andreas Petzold (KIT), David Britton (Glasgow), Fabio Hernandez (Lyon), Frank K Wuerthwein (UCSD), Mattias Wadenstein (Lund), Sabine Crépé-Renaudin (IN2P3)
20230612-Nordics.... BRITTON-230613_... JENA_Computing-... lanscape-germany-...

Panel discussions

Software developments (close to federated infrastructures)

14:00 **Introduction**

14:30 **Round Table Disucssion**
Speakers: Guenter Duckeck, Liliana Teodorescu (London), Mohammad Al-Turany (GSI), Tommaso Boccalli (IN
JENA sw round tab... small_experiments...

Computing Technologies: Computing Technologies

09:00 **Quantum Computing**
remote presentation
Speaker: Denis Lacroix (IN2P3)
Lacroix_JENABolo...

09:30 **Hardware Technologies and their evolution**
Speaker: Andrea Chierici (Istituto Nazionale di Fisica Nucleare
20230613 - Tech tr...

10:00 **High-performance RISC-V systems – recent development** ¶
Speaker: Andrea Bartolini (Bologna)
2023.Bartolini.JEN...

Computing and Society

16:00 **Education and Training in computing** ¶
Speaker: several speakers (few)
jena-computing-trai... JENA CSC.pptx Outreach_Open_Da...

16:30 **plenary discussion on Education and Training in Computing**

17:00 **Requirements on FAIR Data Management / Open Data Access / Outreach**
Speaker: Thomas Schoerner-Sadenius (PUNCH4NFDI / DESY)
20230612.JENA.up...

17:30 **open discussion on Resso**

12:00 **Sustainability in particle, nuclear and astroparticle physics - Spotlight on Computing**
Speaker: Kristin Lohwasser (Sheffield / CERN)
SustainComputing...

Challenges

- Particle Physics
  - High Luminosity LHC - enormous stress on software and computing from higher trigger rates and higher pileup
- Nuclear Physics
  - FAIR experiments will increase data rate and computing needs
  - Need to support many experiments, including small ones
  - Try to make the software and the computing as common a problem as possible - *with common solutions* (example: FAIR-ALICE collaboration)
- Astro(patricle) Physics
  - Observatory model - running and upcoming projects with data piped into existing computer centres
  - Need to generalise access and interfaces to give scientists access to data products and alerts
  - Vera Rubin, SKA - very high data rates

Current Federated Solutions

ESCAPE
European Science Cluster of Astronomy & Particle physics ESFRI research Infrastructures

- In terms of challenges, HEP has scaled up far more than other sciences have so far, due to the experiments at the LHC
  - Pioneering in terms of exploring new regimes of distributed computing
  - WLCG has proven to be a success
- ESCAPE project did good work to help generalise what HEP had developed and make it generally usable (FTS, Rucio, CVMFS, Storage Infrastructure)
  - Our tools need adapted to serve the specific needs of other sciences
  - But this can be done and it is a success for everyone
    - We become part of a larger, more sustainable, community
    - Other users benefit from mature, stress-tested tools
    - Sites do not have to support a zoo of different tools doing approximately the same thing
- ESCAPE now transitioned to a open collaboration of research infrastructures, including CERN

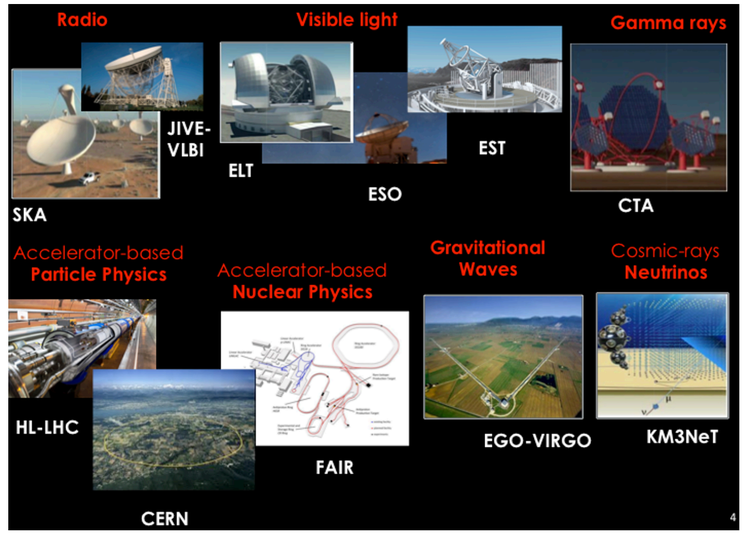

European Open Science Cloud and the FAIR Landscape

- Targeting the open sharing of scientific knowledge and the re-use of research outputs
  - Adopting FAIR principles to the whole lifecycle of data products, algorithms and software
- Some key aspects
  - Development of authentication and authorisation standards
  - Open access and data sharing across communities (from data lakes to Zenodo)
    - Also requires workflows to be sharable - cf. Virtual Research Environments and the JENAA EoIs in Dark Matter and Gravitational Waves
- To make FAIR a success it has to be possible, easy, normal practice, rewarded and, eventually, required
  - This is a significant technical and cultural challenge - it's certainly not 'free'
- Since the workshop EOSC OSCARS and EVERSE projects are both approved
  - CERN and other HEP institutes involved

High Performance Computing

- Very significant strategic investments are being put into HPC in Europe
  - EuroHPC has a budget of 3 Billion Euros to 2027
  - Aiming to build several exascale machines in the coming years
- There is a significant push for HEP and other sciences to make use of these machines
  - This can have a real impact on our available resources in HEP
    - e.g., Germany will no longer fund regional Tier-2 centres as traditional high-throughput facilities
    - Instead LHC computing moves to their new HPC centres
- However, the impedance mismatch can be severe…
  - HPCs traditionally run as silos, with little access in/out
    - Lack of ability to use CVMFS out of the box impedes software distribution
    - Job submission system can be exotic and not compatible with our distributed production engines
  - HPC hardware pushes towards significant use of GPU accelerators and away from CPUs
    - HEP is not in good shape to use this hardware today

Hardware Evolution

- Moore marches on…
  - Still getting increases in transistor density
  - With credible roadmaps for the next few years
    - But we are surely in the endgame…?
- AMD currently ruling in the CPU domain
  - New ARM data centre chips are competitive - particularly good in power efficiency
- NVidia dominate the GPU market
  - New datacenter architectures address the memory bandwidth issue
- Intel are a bit missing in action…
- Europe is also making strategic investments in RISC-V architecture and in chip making capacity

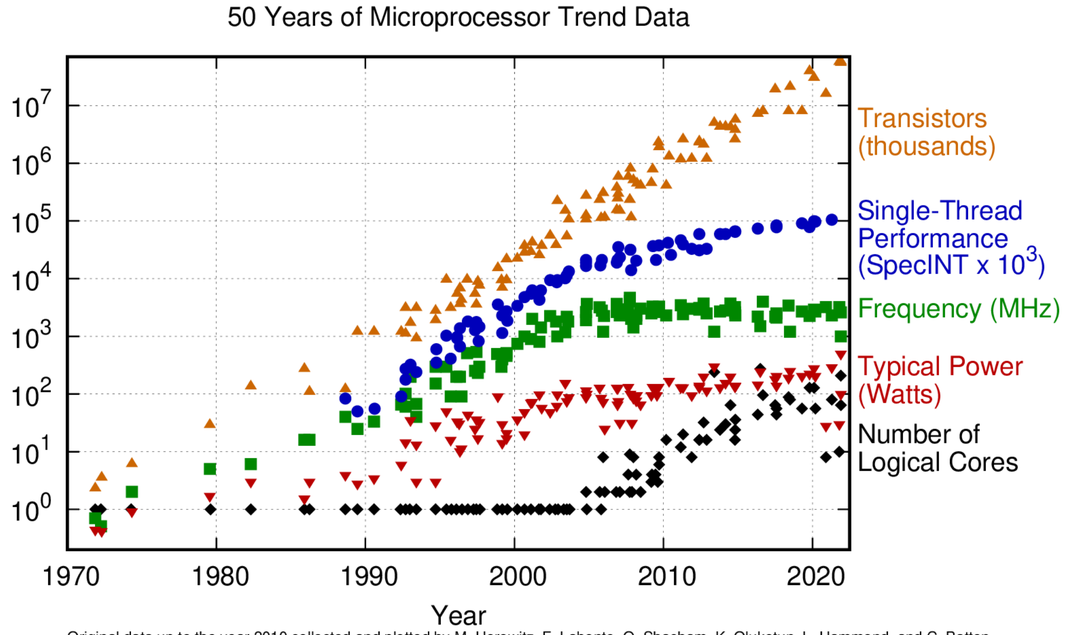

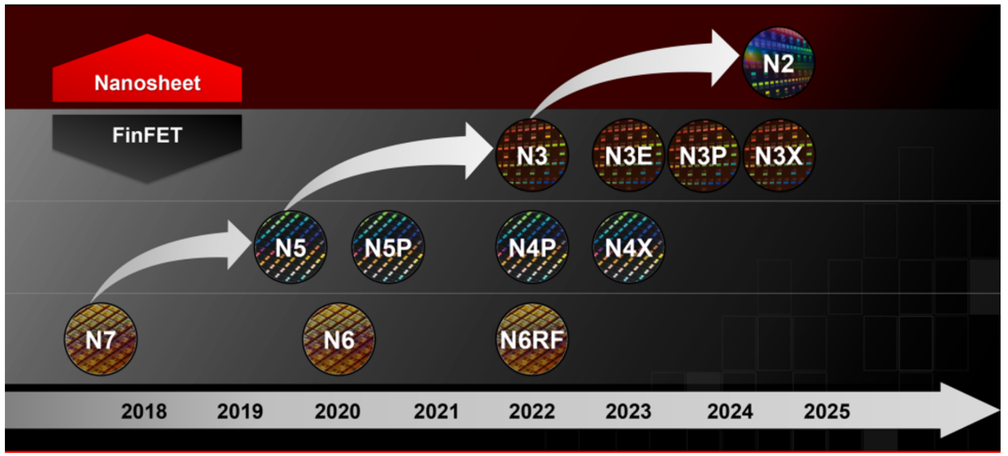

TSMC technology roadmap

Computing Discussions

- Federation and consolidation
  - Consolidation reduces costs, but forces alignment on particular solutions
  - Federation enables integration of diverse resources, risk sharing and sense of ownership
  - Federated model works better for WLCG and similar communities (Nordic, FR, DE)
    - But DE model moves away from university T2s (consolidated storage and HPCs)
- Incorporation of other resources will be needed
  - Substantial work needed to effectively use these resources:
    - Policy: understand how to pledge/plan these resources
    - Technical: development of common interfaces
  - N.B. This it will not make current resources disappear!
- Machine learning continues to grow in importance, but the resource implications are not yet that well understood

Software Discussions

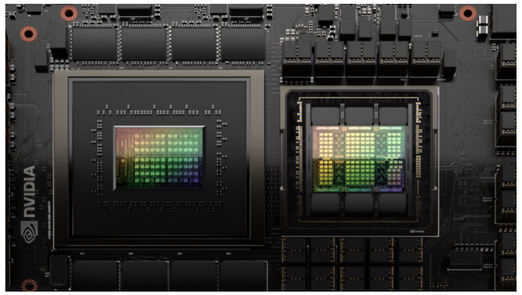

- Heterogeneity is inevitable
  - Constant battle to reduce CPU/unit of work → reduce *energy consumption* and *environmental costs*, financial costs
- Common solutions possible at the infrastructure level and in baseline codes (ROOT, Geant4, event generators, etc.)
- We require both algorithmic improvements and improved technical solutions
  - Radical changes seem to be needed (should we all program in CUDA? SYCL? Julia?)
- Some *techniques* can certainly be shared at the application level
  - Limited sharing in actual application code
- Small experiment support is a **big challenge**
  - Particularly true in nuclear physics
- Training is essential!
  - We have to equip people with the correct skills

Outcomes

- We started the process of asking the right questions
- To home in on the key points we will setup five working groups in the critical areas
  - *HPCs*
  - *Software and Heterogeneous Architectures*
  - *Federated Data Management, Virtual Research Environments and FAIR/ Open Data*
  - *Machine Learning and Artificial Intelligence*
  - *Training, Dissemination, Education*
- These should produce a white paper, from all three communities, that should be approved by the time of the next JENA Symposium, in spring 2025

HPCs

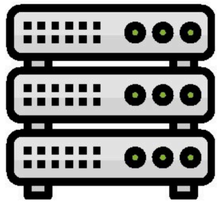

- Address the relationship with HPC centres and the integration of HPC resources with our computing infrastructures
- Engage at a higher level with *EuroHPC*
- Opportunity to *shape the evolution* and policies of HPC facilities towards the ENA sciences
  - Goal of both *augmenting the computing capacity* available for this community and facilitating the federation with existing data facilities
- In HEP the theory community also have significant needs and efficient codes are in development

ML and AI

- Rapid expansion in the last years in most fields of science, including the ENA domains
  - Follow the technologies in this fast evolving field; analyse the potential impact on the ENA *computing infrastructure needs*
- Focus will be to *quantify the resource needs* and to *define the interfaces* and services that are needed by physicists to run ML workloads (looking at both training and inference)
- Important workshop from EuCAIF planned for next year (European AI for Fundamental Physics Conference)

Federations and FAIR

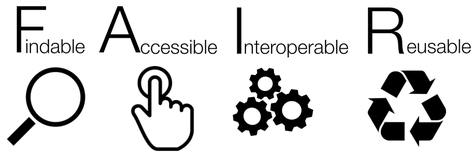

- The ESCAPE collaboration should be leveraged to *strengthen synergies* between the three sciences in Data management and federated identities
- Understand how distributed computing interacts with *FAIR principles*
- JENA chairs are in the ESCAPE advisory board and will recommend that *ESCAPE focuses on those areas* and the evolution of the tools and services for the next decade

Training, Dissemination, Education

- Training - *leverage the experience* in the HSF training initiative and find common ground with other sciences
  - Share/reuse material
- Dissemination - organise a *conference* on *scientific computing* similar to CHEP, but embracing more sciences (largely beyond ENA); work with the 5 science clusters
- Education - ECFA has an initiative on a European master program for detector physicists and engineers
  - Perhaps initiate a similar initiative on scientific software and computing

Software and Heterogeneous Architectures

- Heterogeneous architectures (as well as ARM, we have GPUs, perhaps FPGAs)
  - Effective use of these processors and increase the efficiency of our code by factors
- Need *Research Software Engineers* and *domain experts* that optimise the current code and also that engage in exploratory software R&D activities, *rethinking algorithms*.
  - There will be significant domain level differences in applications, but substantial overlap in skills and techniques
  - Convey the message to funding agencies that it is crucial to invest in *training*, *hiring* and *retaining* people with this profile
- This is one of the main opportunities to *address sustainability* (in all senses!)

Conclusions

- Successful meeting bringing together a lot of experts from all three areas
  - Probably need to broaden even more, e.g., astronomy will be a big player as well
- Many common points discovered where we can make a common case to funding agencies
  - This is a very useful opportunity for the software and computing communities
  - Federated infrastructures are a good model and generalise well
  - But not all resources should be the same, e.g., some sciences made have different resource needs and this should be acknowledged and properly funded
- Topical groups will form soon
  - *Volunteers welcome, of course!*
- There will be a lot of work to do next year to produce the white paper, with a high level of community engagement, in time for the next JENA Symposium