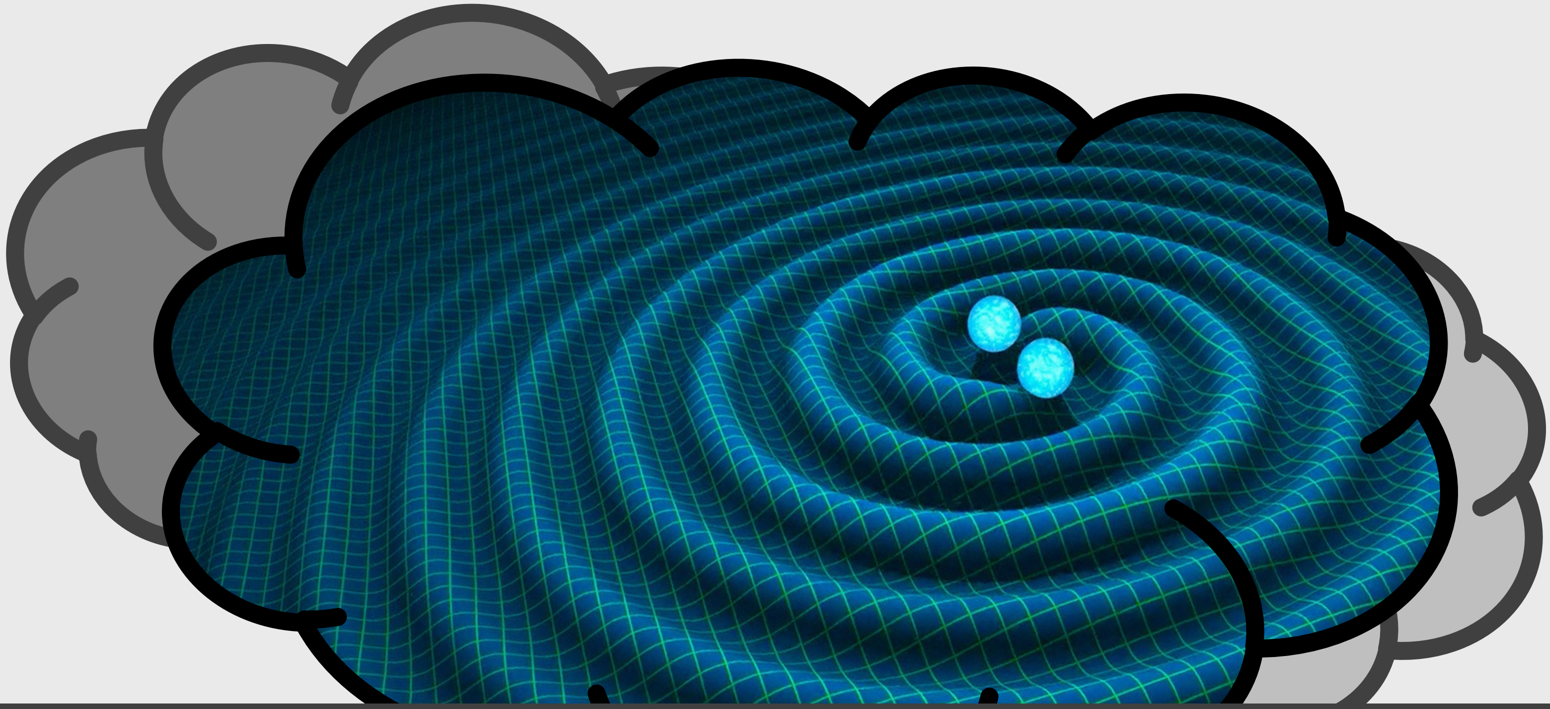

e-Infrastructure Board

Stefano Bagnasco, INFN

for the ET-EIB

ET Symposium
Budapest

Jun 7, 2022

Data and more data

- Raw interferometer data don't grow much with increasing instrument sensitivity
  - Current detectors write o(2PB)/year of raw data per detector
  - Pre-processed data for final user analysis is more than 1 order of magnitude smaller
  - In ET we expect about few tens of PB of raw data per year (baseline 6-interferometer design, more control channels,...)
  - No big deal today, piece of cake by 2035
- However, the amount of useful scientific information encoded in the data does grow a lot
  - And the computing power needed to wring it out (mostly from CBC Parameter Estimation)
  - Larger template banks, longer templates to fit in memory, overlapping events,...
  - Accurately estimating the computing power needs is itself a difficult task

1/10th of an LHC experiment

- As of 2022, the computing needs of the entire GW network are roughly o(10%) of an LHC experiment of today

- In ET the event rate will be 10^3 -10^4 times the current one
  - Analysis of the "golden" events (EM counterparts, high SNR or "special" events) would already be within reach using current technologies
    - O(500) events per year = 12.5MHS06-y per year, the same order of magnitude of a LHC experiment in Run 4
    - Target: 1/10th of an LHC experiment in Run 4

But: low-latency!

EINSTEIN TELESCOPE

Three computing domains

On-site infrastructure

Online

- Data acquisition and pre-processing
- Instrument control
- Environmental monitoring
- …

Plain old HTC (and some HPC)

Offline

- Deep searches
- Offline parameter estimation
- Template bank generation, NR
- …

Here's the fun

Low-latency

- Candidate search
- Sky localization
- LL parameter estimation
- Alert generation and distribution

The mandatory slide with boxes and arrows

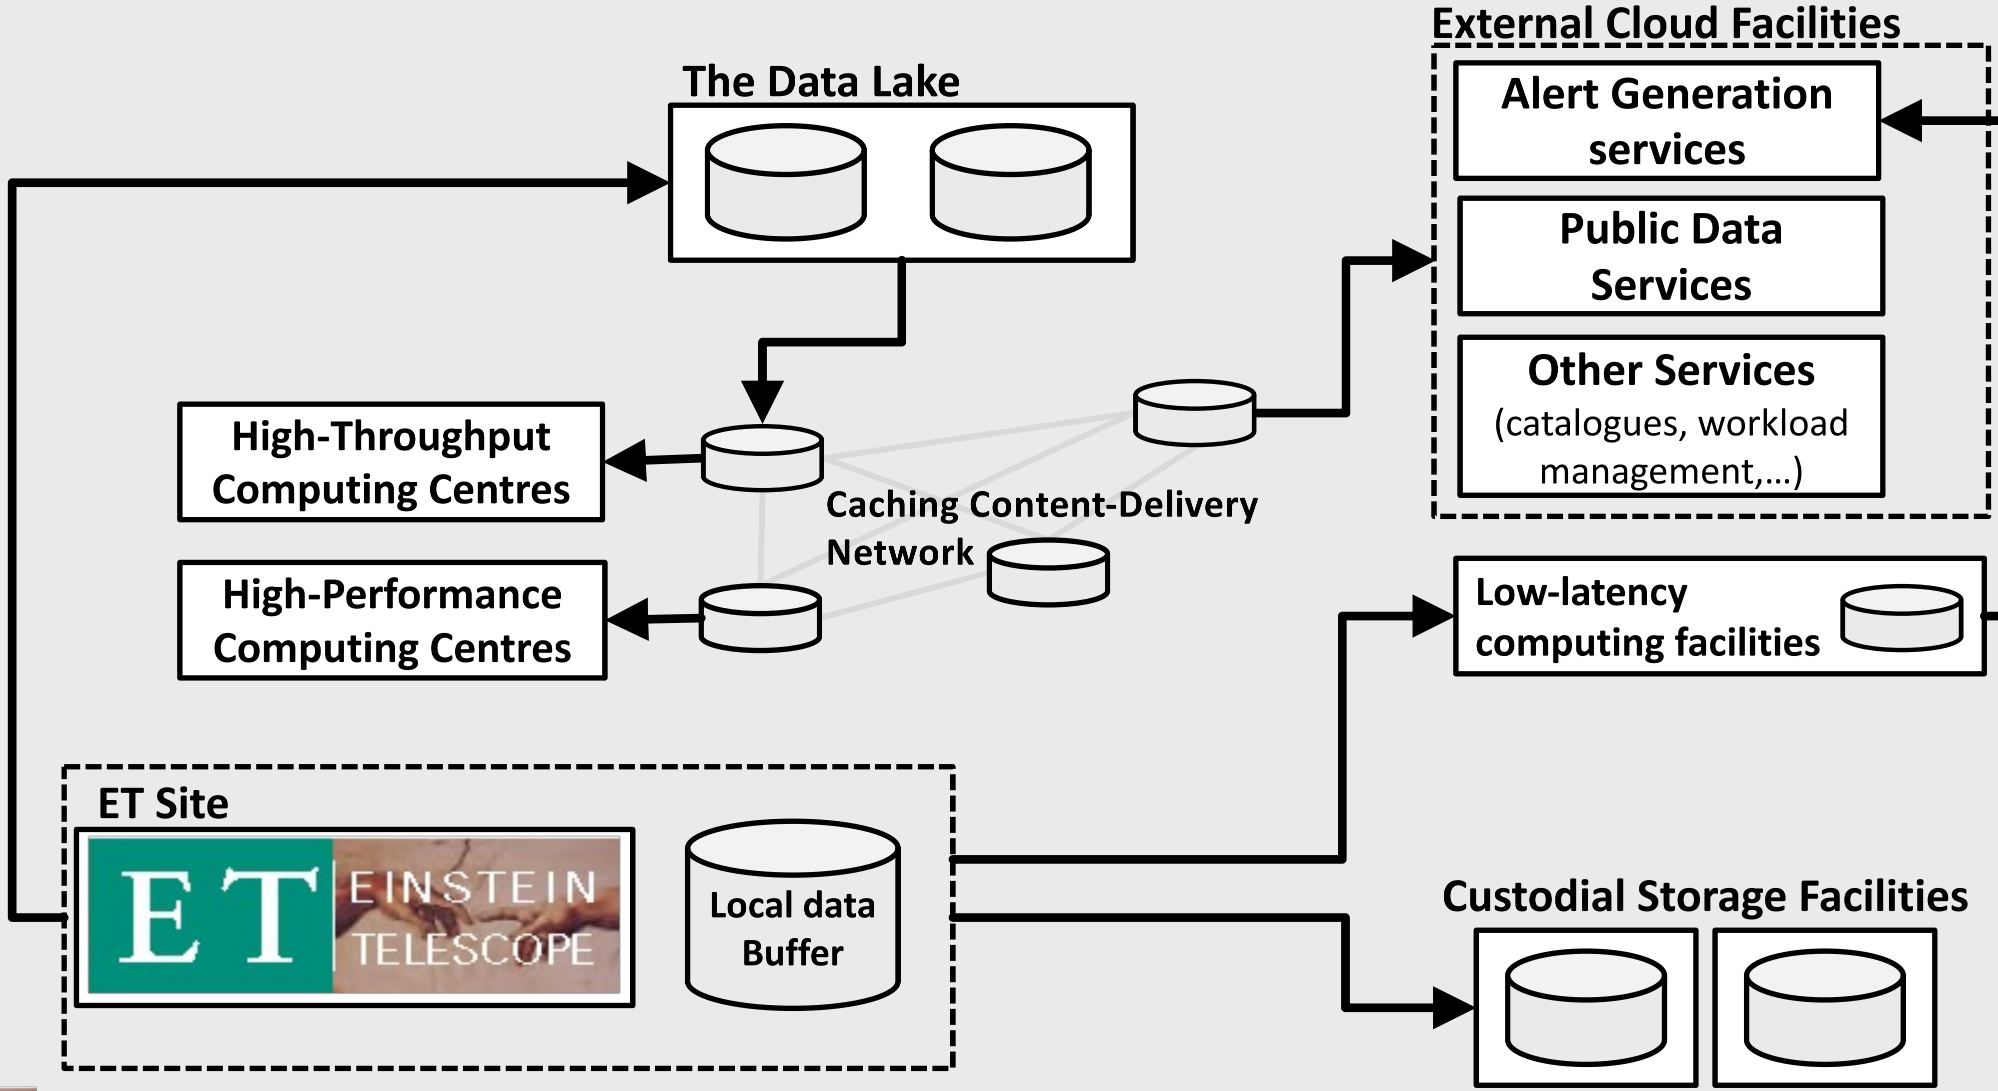

Figure from the ET ESFRI proposal

The mandatory slide with boxes and arrows

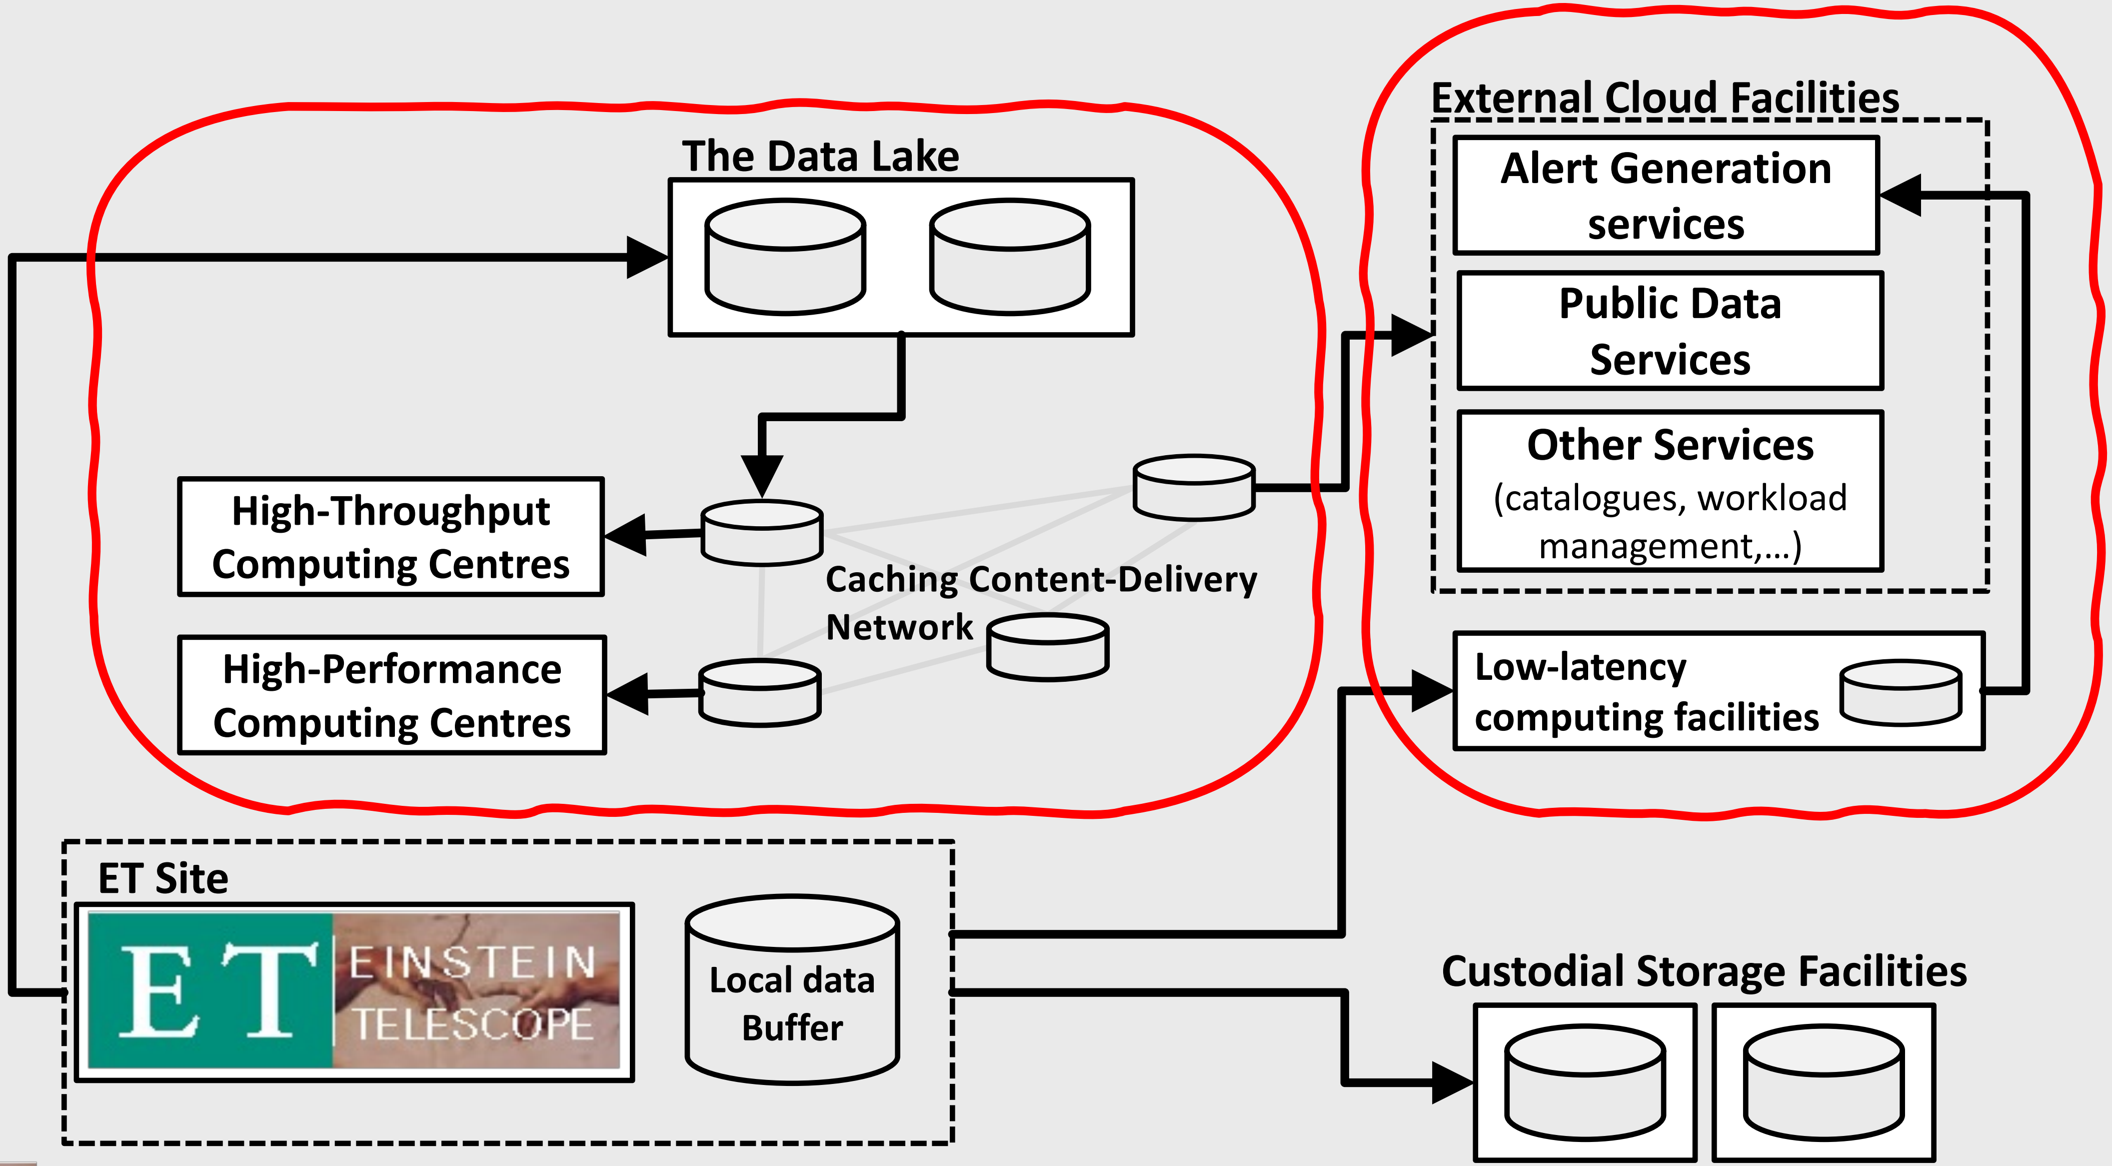

Figure from the ET ESFRI proposal

What will ML look like 10 years from now?

nature astronomy

ARTICLES

https://doi.org/10.1038/s41550-021-01405-0

Check for updates

Accelerated, scalable and reproducible AI-driven gravitational wave detection

E. A. Huerta[1,2], Asad Khan[3], Xiaobo Huang[3], Minyang Tian[3], Maksim Levental[2], Ryan Chard[1], Wei Wei[3], Maeve Heflin[3], Daniel S. Katz[3], Volodymyr Kindratenko[3], Dawei Mu[3], Ben Blaiszik[1,2] and Ian Foster[1,2]

The development of reusable artificial intelligence (AI) models for wider use and rigorous validation by the community promises to unlock new opportunities in multi-messenger astrophysics. Here we develop a workflow that connects the Data and Learning Hub for Science, a repository for publishing AI models, with the Hardware-Accelerated Learning (HAL) cluster, using funcX as a universal distributed computing service. Using this workflow, an ensemble of four openly available AI models can be run on HAL to process an entire month's worth (August 2017) of advanced Laser Interferometer Gravitational-Wave Observatory data in just seven minutes, identifying all four binary black hole mergers previously identified in this dataset and reporting no misclassifications. This approach combines advances in AI, distributed computing and scientific data infrastructure to open new pathways to conduct reproducible, accelerated, data-driven discovery.

Gravitational waves were added to the growing set of detectable cosmic messengers in the fall of 2015 when the advanced Laser Interferometer Gravitational-Wave Observatory (LIGO) detectors reported the observation of gravitational waves consistent with the collision of two massive, stellar-mass black holes[1]. Over the last five years, the advanced LIGO and advanced Virgo detectors have completed three observing runs, reporting over 50 gravitational wave sources[2,3]. As advanced LIGO and advanced Virgo continue to enhance their detection capabilities and other detectors join the international array of gravitational wave detectors, it is expected that gravitational wave sources will be observed at a rate of several per day[4].

An ever-increasing catalogue of gravitational waves will enable systematic studies to advance our understanding of stellar evolution, cosmology, alternative theories of gravity, the nature of supranuclear matter in neutron stars, and the formation and evolution of black holes and neutron stars, among other phenomena[5–11]. Although these science goals are feasible in principle given the proven detection capabilities of astronomical observatories, it is equally true that established algorithms for the observation of multi-messenger sources, such as template-matching and nearest-neighbour algorithms, are compute-intensive and poorly scalable[12–14]. Furthermore, available computational resources will remain oversubscribed, and planned enhancements will be outstripped rapidly with the advent of next-generation detectors within the next couple of years[15]. Thus, an urgent rethink is critical if we are to realize the multi-messenger astrophysics program in the big-data era[16].

To contend with these challenges, a number of researchers have been exploring the application of deep learning and of computing accelerated by graphics processing units (GPUs). Co-authors of this article pioneered the use of deep learning and high-performance computing to accelerate the detection of gravitational waves[17,18]. The first generation of these algorithms targeted a shallow signal manifold (the masses of the binary components) and required only tens of thousands of modelled waveforms for training, but these models served the purpose of demonstrating that an alternative method for gravitational wave detection is as sensitive as template matching and significantly faster, at a fraction of the computational cost.

Research and development in deep learning is moving at an incredible pace[19–37] (see also ref. [38] for a review of machine-learning applications in gravitational wave astrophysics). Specific milestones in the development of artificial intelligence (AI) tools for gravitational wave astrophysics include the construction of neural networks that describe the four-dimensional (4D) signal manifold of established gravitational wave detection pipelines, that is, the masses of the binary components and the z component of the three-dimensional spin vector in (m_1, m_2, s_1^z, s_2^z). This requires the combination of distributed training algorithms and extreme-scale computing to train these AI models with millions of modelled waveforms in a reasonable amount of time[30]. Another milestone concerns the creation of AI models that enable gravitational wave searches over hour-long datasets, keeping the number of misclassifications at a minimum[39].

In this article, we introduce an AI ensemble, designed to cover the 4D signal manifold (m_1, m_2, s_1^z, s_2^z), to search for and find binary black hole mergers over the entire month of August 2017 in advanced LIGO data[40]. Our findings indicate that this approach clearly identifies all black hole mergers contained in that data batch with no misclassifications. To conduct this analysis we used the Hardware-Accelerated Learning (HAL) cluster deployed and operated by the Innovative Systems Laboratory at the National Center for Supercomputing Applications. This cluster consists of 16 IBM SC922 POWER9 nodes, with four NVIDIA V100 GPUs per node[41]. The nodes are interconnected with an EDR InfiniBand network, and the storage system is made of two DataDirect Networks all-flash arrays with SpectrumScale file system, providing 250 TB of usable space. Job scheduling and resource allocation are managed by the SLURM (Simple Linux Utility for Resource Management) system. As we show below, we can process data from the entire month of

[1]Data Science and Learning Division, Argonne National Laboratory, Lemont, IL, USA. [2]University of Chicago, Chicago, IL, USA. [3]University of Illinois at Urbana-Champaign, Urbana, IL, USA. e-mail: elihu@anl.gov

ML is not yet a mainstream "tool of the trade" in GW, but a huge lot of R&D is already ongoing

- Efficiency & speed
  - Signal Classification
  - Parameter estimation
  - Noise glitch hunting
  - (Template bank generation)
- Technology exploitation
  - Use advanced hardware (GPU, TPU…)
  - FPGAs / custom hardware
- Automatization
  - Automatize standard procedure for Data Quality
  - Automated de-noising with synthetic noise from GANs
- Role of HPC (and ML-specific facilities) is expected to grow (Numerical Relativity, template bank production) and with impact of ML

The ecosystem (and some questions)

- First and foremost, other 3G facilities (and the 2G network)
  - CE (and LISA!)
  - Will there be an equivalent of the IGWN common computing infrastructure?
- Several EM and astroparticle initiatives coming of age in the same time frame
  - CTA, SKA, KM3Net, ELT in the ESFRI roadmap, and many more
  - They will all have stringent low-latency alert requirements, as producers or consumers (or both)
  - High rates will imply extreme automation in the generation and selection of triggers, and sophisticated scheduling algorithms
  - Will there be a MM-specific (possibly virtual) shared infrastructure like the WLCG?
  - The architecture of the next LL alert distribution system is being defined now!
- The EU is building the European Opens Science Cloud
  - Scientific Computing in the Digital Continuum
  - How concrete will it be in 2035?

e-Infrastructure Board mandate

«…to design, create and operate an evolving, efficient and functional e-infrastructure environment at a reasonable cost for the collaboration. Initially the focus will be the development of a Computing Model for the ET».

- Prepare a plan of the studies and activities that need to be undertaken for the development of the ET computing.
- Propose a computing model and its updates to the collaboration.

Out of scope: actual science code, physics and engineering tools

Shopping list

Data transfer and storage: safely and efficiently transfer all data to custodial storage and processing centres, including low-latency transfers;

Software packaging and distribution: manage software lifecycle, and make packages available ubiquitously;

Computing power: provide and manage computing resources (HTC and HPC) for the processing of data, in all computing domains;

Data distribution: make data available to worker nodes in computing centres anywhere, and possibly also to single workstations, including support to public releases of data;

High-availability service management: provide a platform for running the collaboration's services (e.g. alert generation services, event databases,...)

Data cataloguing and bookkeeping: organise all data and metadata and provide querying and discovering capabilities;

Job lifecycle management: provide a uniform job submission and runtime environment to research groups;

High-level workload management: keep a database of all jobs and allow the enforcement of priorities and scheduling strategies; provide support for organized large-scale data processing campaigns;

Monitoring and accounting: monitor local and distributed computing, checking performance and looking for issues, and provide reliable accounting both at the user/job and site level;

Authentication, Authorisation and Identity management: provide consistent AAI across all domains and activities.

Collaboration services: provide tools for efficient collaboration management, coordination, and outreach (e.g. document repositories, collaborative tools, administrative databases, communications,...)

The Computing Model

- The overall architecture of the e-Infrastructure, either as a single integrated system or as a few separate systems (e.g. instrument control and DAQ, low-latency, and offline)
- A documented way of evaluating the required computing power and storage space from the evolving scientific program of the collaboration
- Estimates of the involved costs and growth timelines
- A description of the data flows, with estimates for the needed network performances
- A description of the User Experience and workflows for relevant activities
- A description of the tools chosen or to be developed to provide all the required functionalities (foundation libraries, frameworks, middleware,...)
- Separate "Work Breakdown Structure" and "Implementation Plan" documents

The ET e-Infrastructure Board

Division 1: Software, frameworks, and data challenge support

Division 2: Services and Collaboration Support

Division 3: Computing and data model, Resource Estimation

Division 4: Multimessenger alerts infrastructure

TTG: Technology Tracking working Group

Division 1: Software, frameworks, and data challenge support

Define the software frameworks for ET computing workflows, the middleware for infrastructure, workload and data management. Develop software quality best practices and support their adoption with training and enforcement policies. Support code development in all computing domains. Provide computing support for mock-data challenges.

- Collaborate with OSB and ISB to define the data formats (both internal and for public release) and organized data processing workflows
- Support the development of the tools for the operation of the telescope
- Coordinate the development of common infrastructural tools and frameworks for the data-analysis, both offline and in low-latency
- Support the operation of large-scale computing campaigns
- Develop policies and best practices to ensure software quality, and encourage/enforce their adoption
- Organize a continuous training programme for both developers and users

Division 2: Services and Collaboration Support

Define and provide all the IT services needed for the administrative management of the Collaboration. Define and provide all the IT services needed for communication and collaboration within the Collaboration and outside.

- Provide collaborative tools for communication within the collaboration and to the outside
- Coordinate the operation of the collaborative and administrative tools for the management of the collaboration
- Define and provide a future-proof federated AAI infrastructure for the collaboration

Division 3: Computing and data model, Resource Estimation

Develop the Einstein Telescope Computing Model. Provide a running estimate of the computing resources needed for all computing domains.

- Develop a Work Breakdown Structure for the early stages of the preparation of the Computing Model and Cost Estimates
- Collaborate with OSB to define the initial activities to evaluate actual computing needs
- Liaise with the Numerical Relativity community

Division 4: Multimessenger alerts infrastructure

Design and develop the infrastructure needed for multi-messenger triggers management and distribution. Follow the development of software tools for low-latency computing.

- Coordinate the development of the tools for the low-latency alert generation and management system
- Participate in the technical development of the alert distribution infrastructure, by liaising with the wider astrophysical community

TTG: Technology Tracking working Group

Follow the evolution of hardware and software computing technologies. Organize regular occasions for inter-division updates.

- A transversal working group, coordinating TT activities across all four divisions.
  - Artificial Intelligence and Machine Learning
  - GPUs and HPC, FPGA and fancier architectures such as TPUs
  - Middleware tools and technologies
  - Quantum computing!

ET-PP (INFRA-Dev) WP8

Task 8.1: "TO data center" Conceptual design of the center in close collaboration with the instrument science board. Definition of the services provided by the center, delimitation against services realized with distributed computing.

Task 8.2: "Computing and Data Model" Development of the computing and data model in close cooperation with the instrument science board and observational science board of ET. Definition of the workflow from the instrument to the publication.

Task 8.3: "Resources" Estimate of the computing resources (computing power and data storage), the personnel, and the operational cost required for all aspects of ET computing. The potential for mitigation must be addressed.

Task 8.4: "Data Access Implementation" Guidelines for the data policy compliance, relevant to the data storage, access, process and distribution, on all relevant time scales, respecting the EU policies on open data.

ET-PP (INFRA-Dev) WP8 Deliverables

D8.1: Computing and Data Requirements (M18): Documentation of the inputs on the computing and data requirements received during the process;

D8.2: Computing and Data Model (M42): Final version of 8.1;

D8.3: Data Access Implementation Guidelines (M48): A document describing how to implement the policy for the storage and the access to the ET data, according to the data model.

WP8 Vs. EIB

	T8.1 T0 Data center	T8.2 Computing and Data Model	T8.3 Resources	T8.4 Data Access Implementation
D1 **Software, frameworks, data challenges**		Computing frameworks computing domains data formats	Resources for frameworks execution and data storage availability	Data availability Data releases format
D2 **Services and Collaboration Support**			Resources and infrastructure for services and collaborative tools	Tools for monitoring, AAI (IAM), data access
D3 **Computing and data model, Resource Estimation**	T0 storage and computing resources estimation	Computing model data model	Resources estimation	
D4 **Multimessenger alerts and low-latency infrastructure**	ET site infrastructure for alerts and low-latency	Low latency computing model	Resources for low-latency infrastructure	Tools for low latency results sharing

TTC
technology tracking committee

nadia.tonello@bsc.es

WP8 Vs. EIB

EIB Chairs: Patrice Verdier (IN2P3), SB (INFN)

WP8 leaders: Achim Stahl (U. Aachen), Sergi Girona (BSC) + Nadia Tonello (BSC)

- Joint WP8+EIB weekly call for coordination
- BSC is opening a position for a "system engineer"-like profile to follow the preparation of the computing model
- EIB is defining the division chairs (nearly there...) and populating the organigram
  - More about this later

M2Tech (INFRA-Tech) WP6

- M2Tech proposal submitted together with CTA, KM3NeT, Virgo, focused on multi-messenger astronomy
- WP6 “Efficient computing and algorithms”

Task 6.1 – Sustainable fast data reduction and processing

AI/ML-based processing, GPU and hardware accelerators,…

Task 6.2 – Operational Intelligence for instrument and data flows management

- ML-based tools and methods

Task 6.3 – Efficient tools for multimessenger alerts

- Common tools, standards, formats for alert management. Dynamic alerts database.

- Also: “Digital twin” concept in WP4
- See Edwige’s talk tomorrow

Collaboration services

- ET Member's Database
  - Developed by Gary Hemmings at EGO
  - Derived from the Virgo VMD
- IAM services
  - Identity and Access Management, SSO etc.
  - Evolving standards (e.g., x509 to tokens), not a simple task
  - Working group formed, headed by Cyfronet
- More services provided by EGO
  - GitLab, web, wiki,...
- Negotiating for extended support
  - And sending a request to funding agencies for 2 FTE with computing technician profiles

ET EINSTEIN TELESCOPE

Gathering requirements for computing resources from OSB, ISB, SPB,...

- Storage and access layer for site characterization data
  - Do we already need a Data Catalogue?
- CPU for Mock Data Challenge
  - Do we already need a global Workload Management System?
- Definition of an ongoing procedure for request collection
  - Important to centralize such requests and avoid direct negotiation with CCs
  - Also, need to establish support chain between users and CCs (exploit recently created Virgo network)

Personpower

- “Tentative” declared effort: ~6 FTE from about 25 people
  - Aachen University, University of Geneva, Budapest University of Technology and Economics, PIC, Wigner Research Centre, BSC, Cyfronet, ICCUB, KIT, Copernicus Astronomical Center, INFN (TO, NA, BO), CC-IN2P3, Institut des 2 Infinis de Lyon
  - As expected in this phase, small effort contributions from large number of people
- 3-5 further FTEs expected from prospective PhDs or project funding (M2Tech, ETIC,…)

A word of warning

- We need special professional profiles
  - Something between physical science and computer engineering
  - Not exactly "pipeline developers", not exacly "system architects"
- Such personpower is difficult to find
  - Skilled personpower for computing activities is scarce
  - Hard to train and keep, hard to hire: we are not competitive with industry!
- This is not a problem for the GW community only
  - And neither limited to the EU
- For example, HSF has some recommendations for that
  - Training, career incentives,...
  - We should plan also for that

Summary

- ET computing will not be huge by 2030s standards
  - but large enough to require careful planning
  - bad habits are quick to pick up and difficult to eradicate
- GW computing will be a major player in the 2030's computing landscape
  - Even if not one of the largest
  - We need to be represented in discussions that are already starting
- Many possible synergies with Virgo
  - Computing technology essentially a continuum O5 → Virgo nETX → ET
- Computing activities are starting (relatively) slowly
  - Coordination and communication with other boards (e.g. with OSB Division 10) is crucial
  - A decade seems a long time, but we need to start early...
  - ...and learn from old mistakes!