CALCOLO PER EINSTEIN TELESCOPE

Stefano Bagnasco, INFN Torino

Workshop sul Calcolo per l'INFN | Palau (SS) | May 20, 2024

LOOKING FOR...

Transient sources:

- CBC: Compact Binary Coalescence

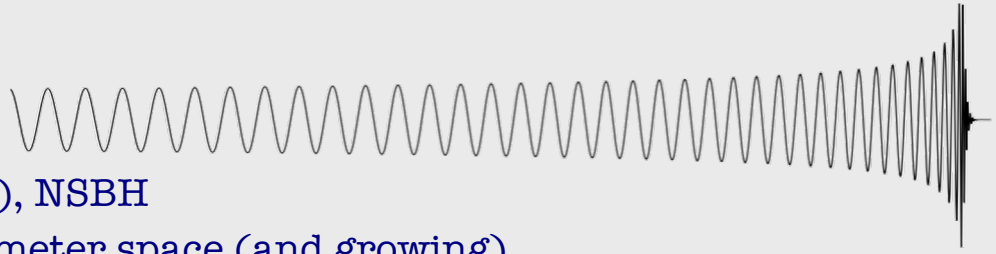

  - Binary Black Holes (BBH), Binary Neutron Stars (BNS), NSBH
  - Strongest emitters, well modelled for much of the parameter space (and growing)
  - Matched filtering very effective

- Burst: Unmodeled transient bursts

  - E.g., Core Collapse Supernovae (CCS, and anything else)
  - Weaker and no (or poor) model, so coherence methods more effective

Continuous sources:

- CW: Continuous waves

  - E.g., Asymmetric spinning neutron stars
  - Usually well-modelled
  - All-sky and targeted searches

- SGWB: Continuous stochastic background

  - Astrophysical & cosmological

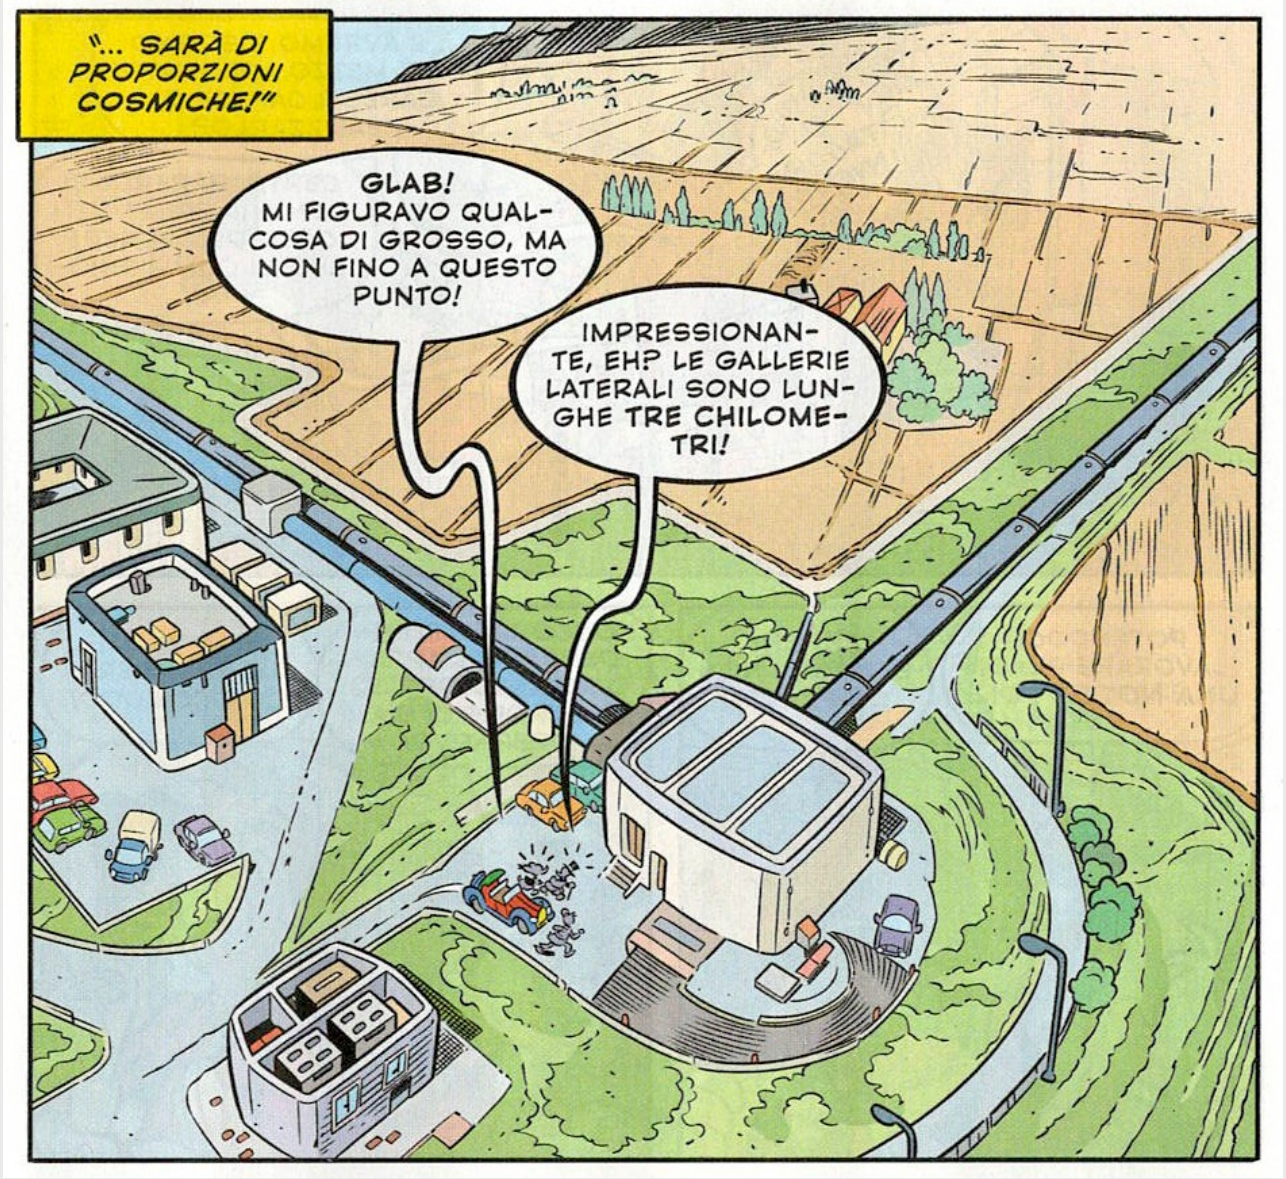

M. Bosco, G. Soldati, "Sogni d'oro zio Paperone"
Topolino **3538**:45-70 (2023)

FROM DISCOVERY TO OBSERVATION

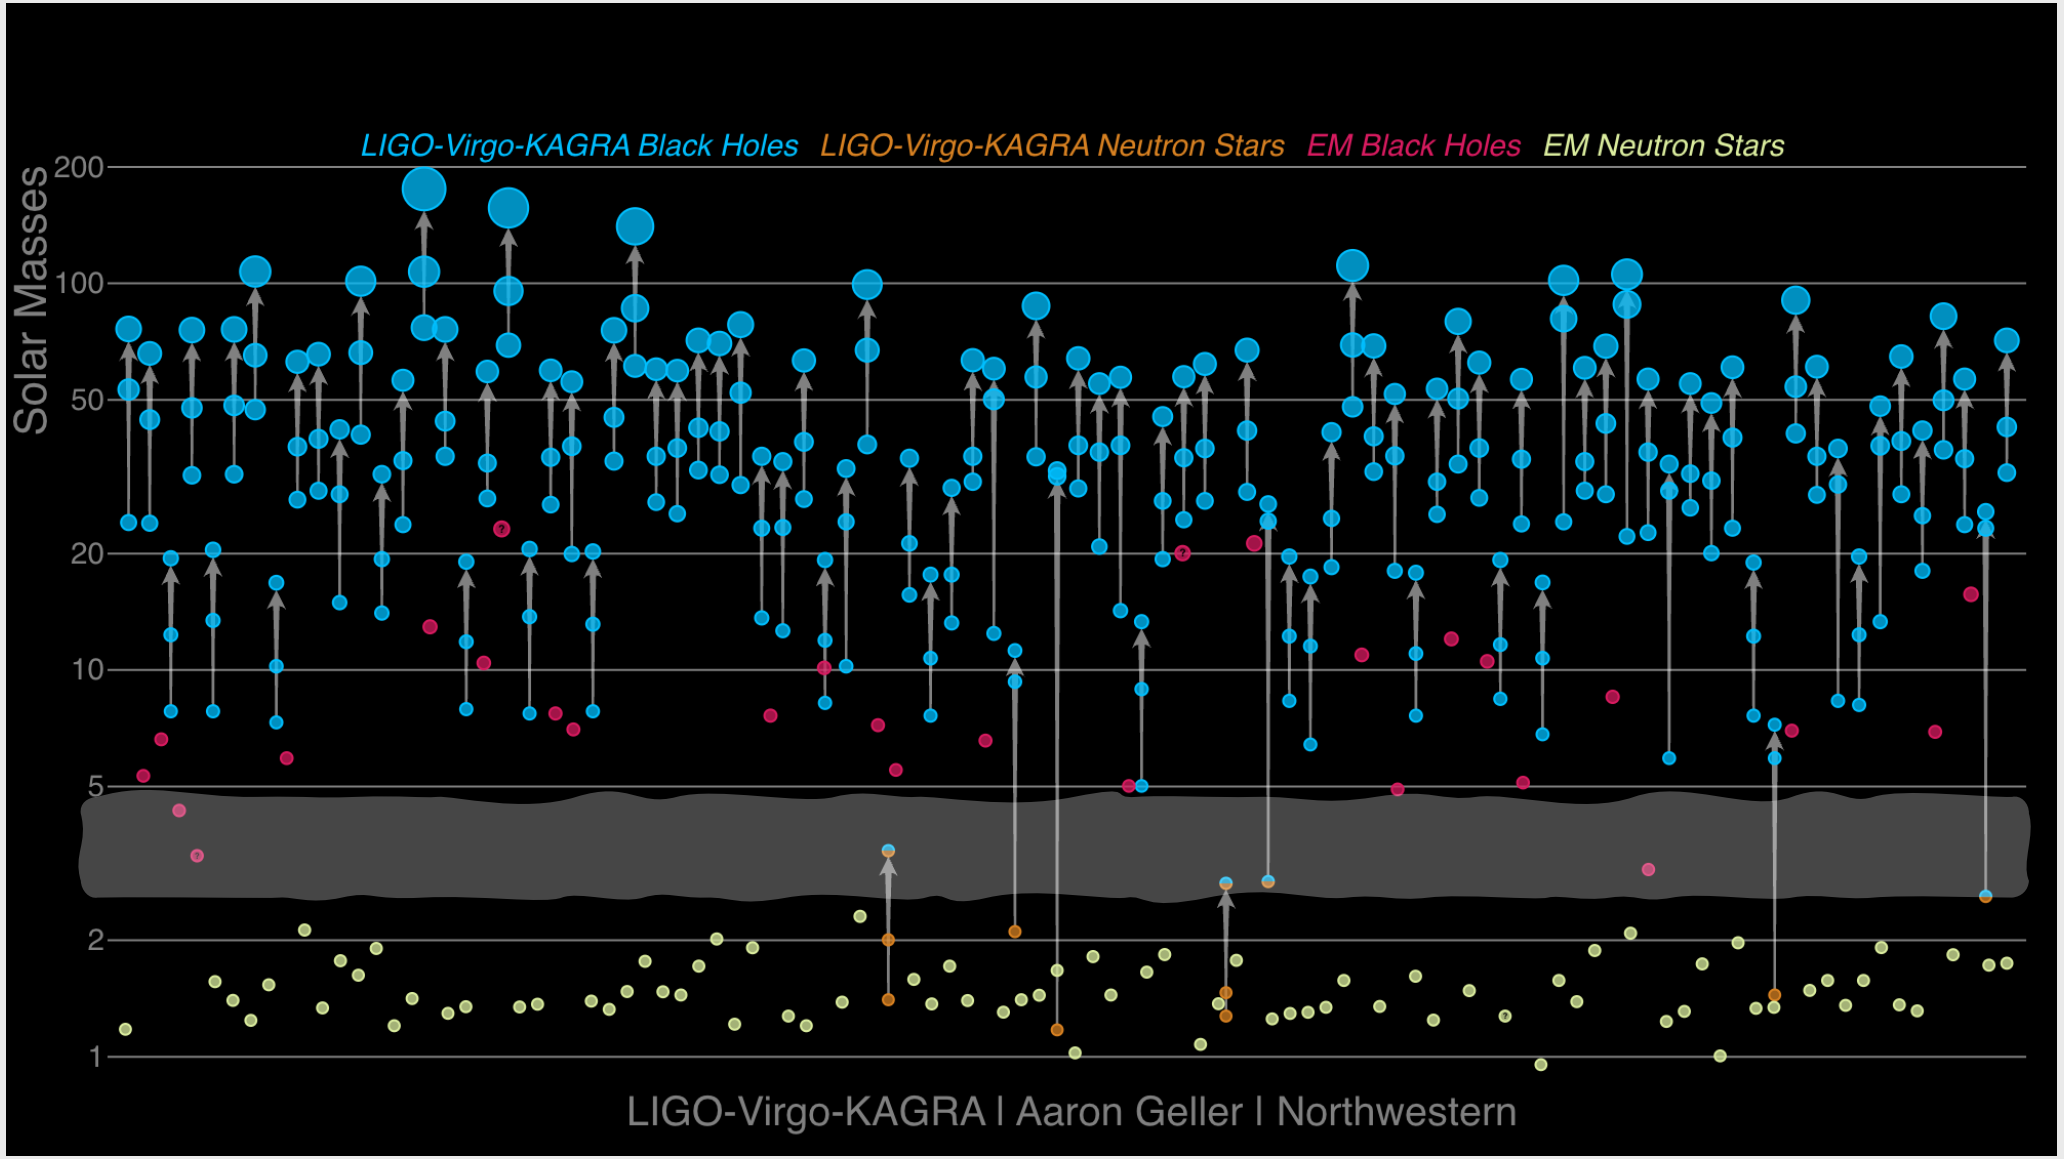

https://media.ligo.northwestern.edu/gallery/mass-plot

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
- Detector Characterization (DetChar)
- (Template bank generation)

Here's the fun

Low-latency

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

THE EINSTEIN TELESCOPE

- ET is the project aiming to realise the **European 3rd Generation Gravitational Wave observatory**

- ET has been a pioneer idea that defined the concept of 3rd generation GW observatory:
  - A sensitivity at least 10 times better than the (nominal) advanced detectors on a large fraction of the detection frequency band
  - Wideband (possibly wider than the current detectors) accessing the frequency band below 10Hz
  - High reliability and improved observation capability

- ET has a long and important history that formed first the ET community and now the ET project

- ET is also a formal scientific collaboration since June 2022

ET: WHERE

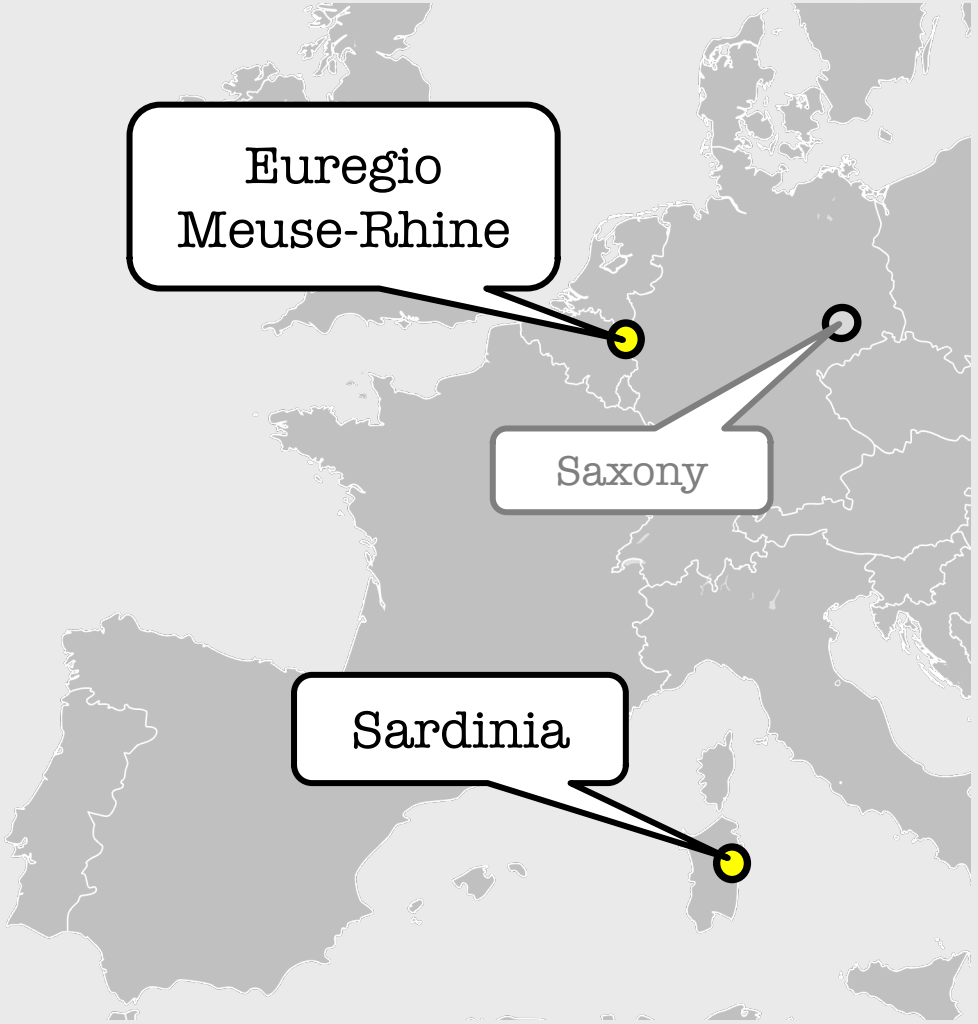

- Currently there are two candidate sites being characterized to host ET:
  - The Sardinia site, close to the Sos Enattos mine
  - The Euregio Meuse-Rhine site, close to the NL-B-D border
  - A third option in Saxony (Germany) was also proposed and is under discussion

ET: WHEN

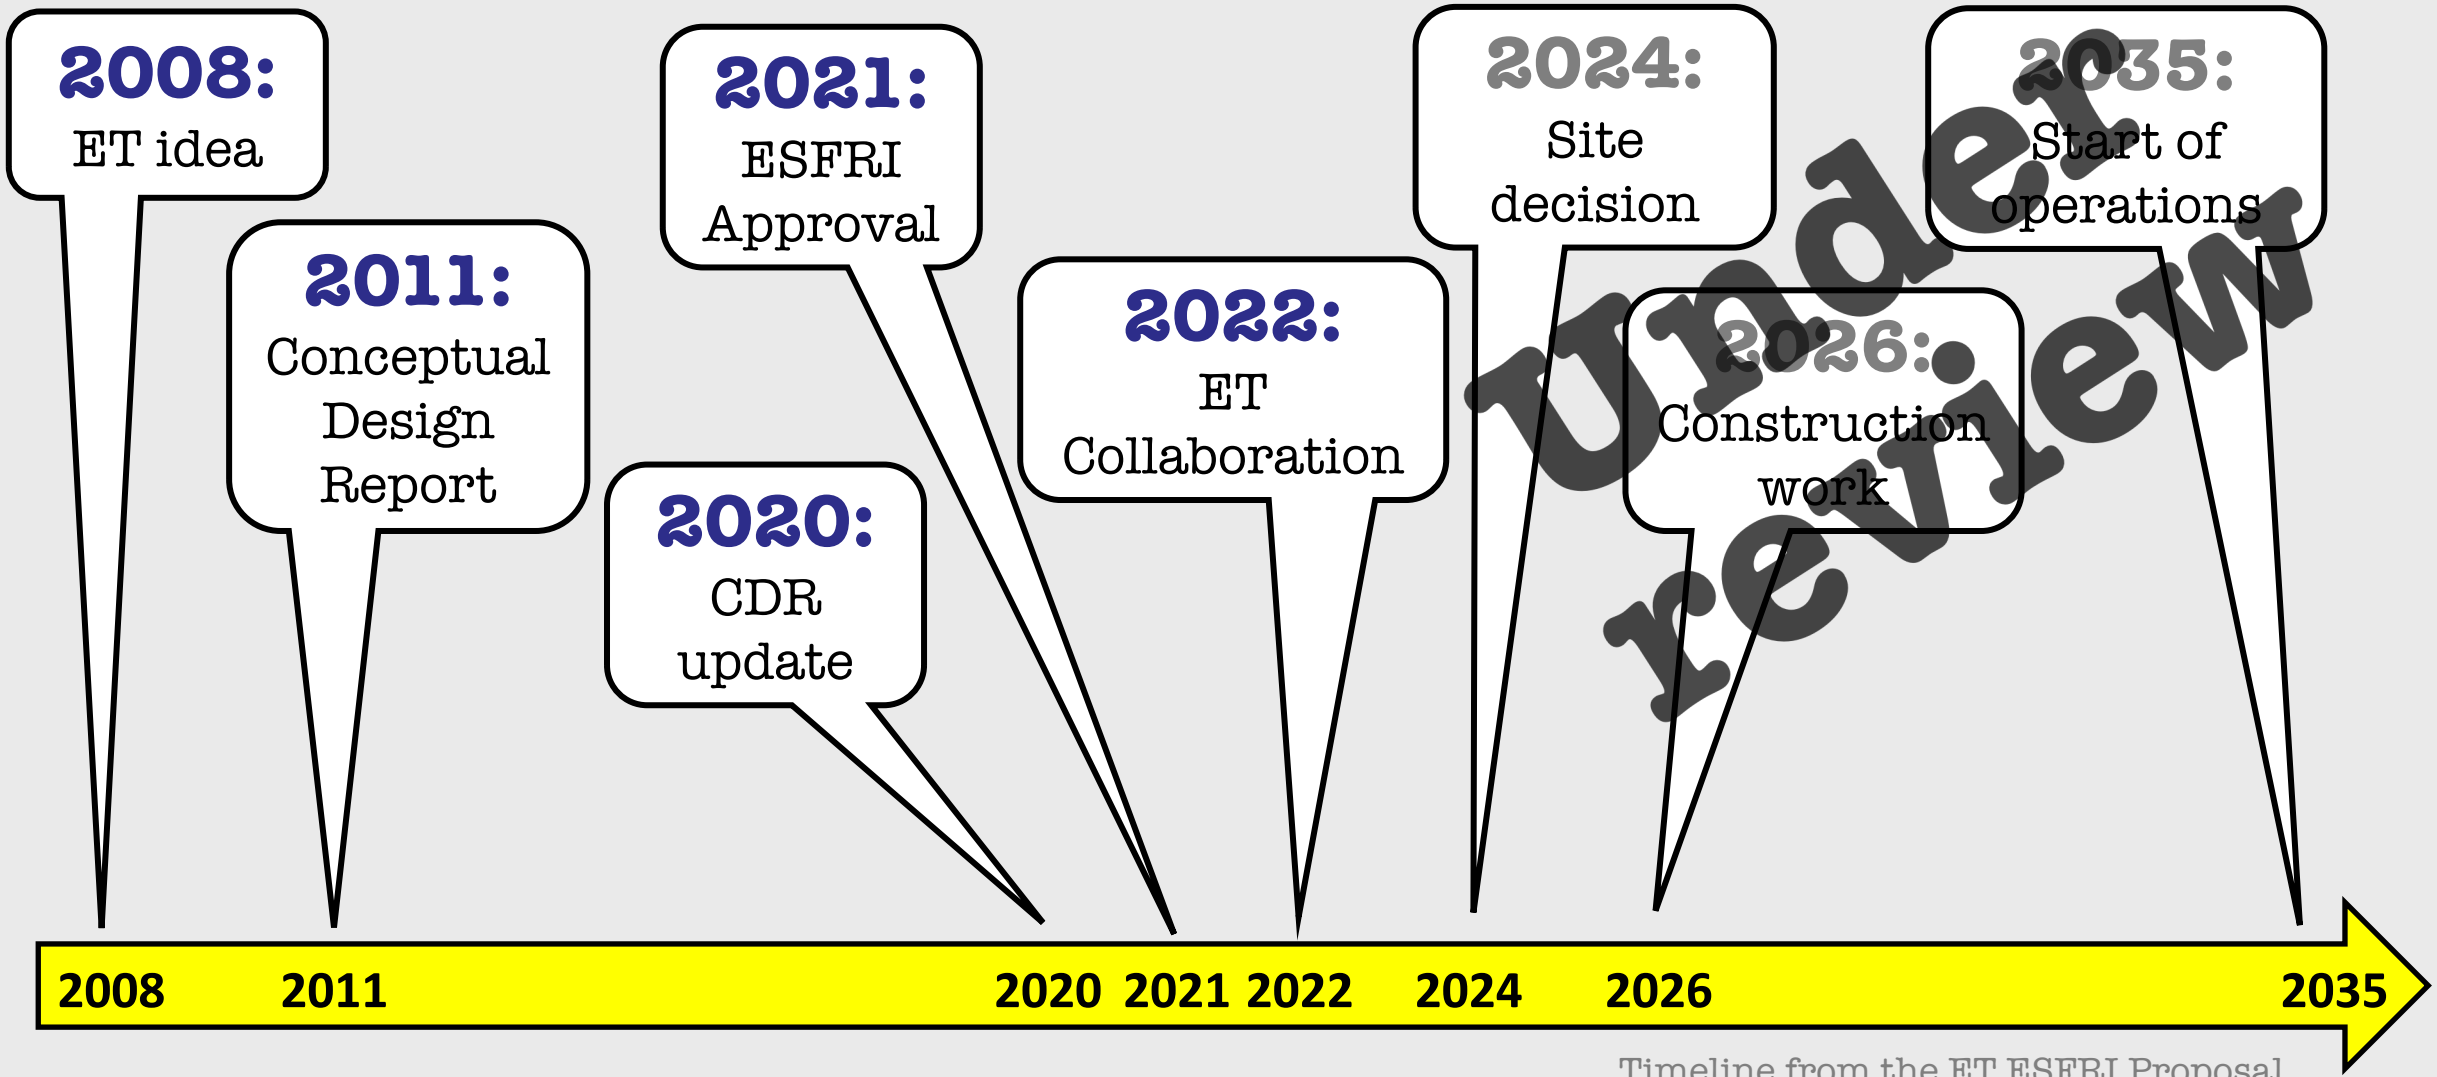

Timeline from the ET ESFRI Proposal

ET: HOW

- Three detectors in a **triangular** structure (baseline design)
  - Closed geometry allows the use of the null data stream
  - Extra detector adds redundancy and makes up for the non-right 60° angle
- Each detector ("red", "green" and "blue") consists of **two** Michelson interferometers
  - High-frequency and Low-frequency
- Alternative 2L (15km, misaligned, non-colocated) configuration under evaluation
  - Pros, cons and risks

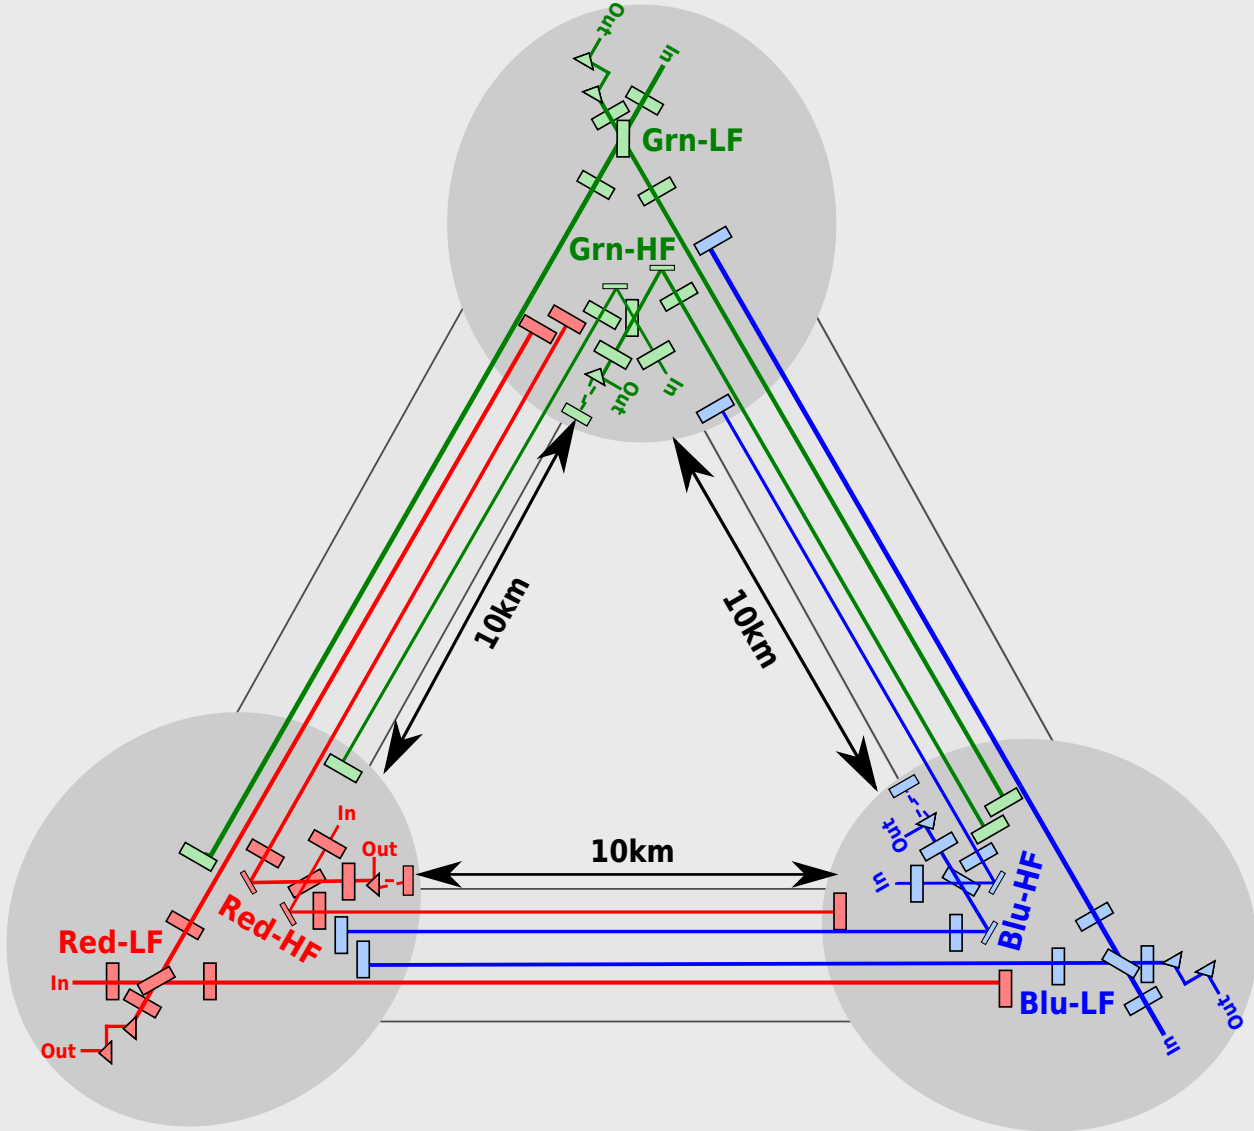

ENHANCED SENSITIVITY

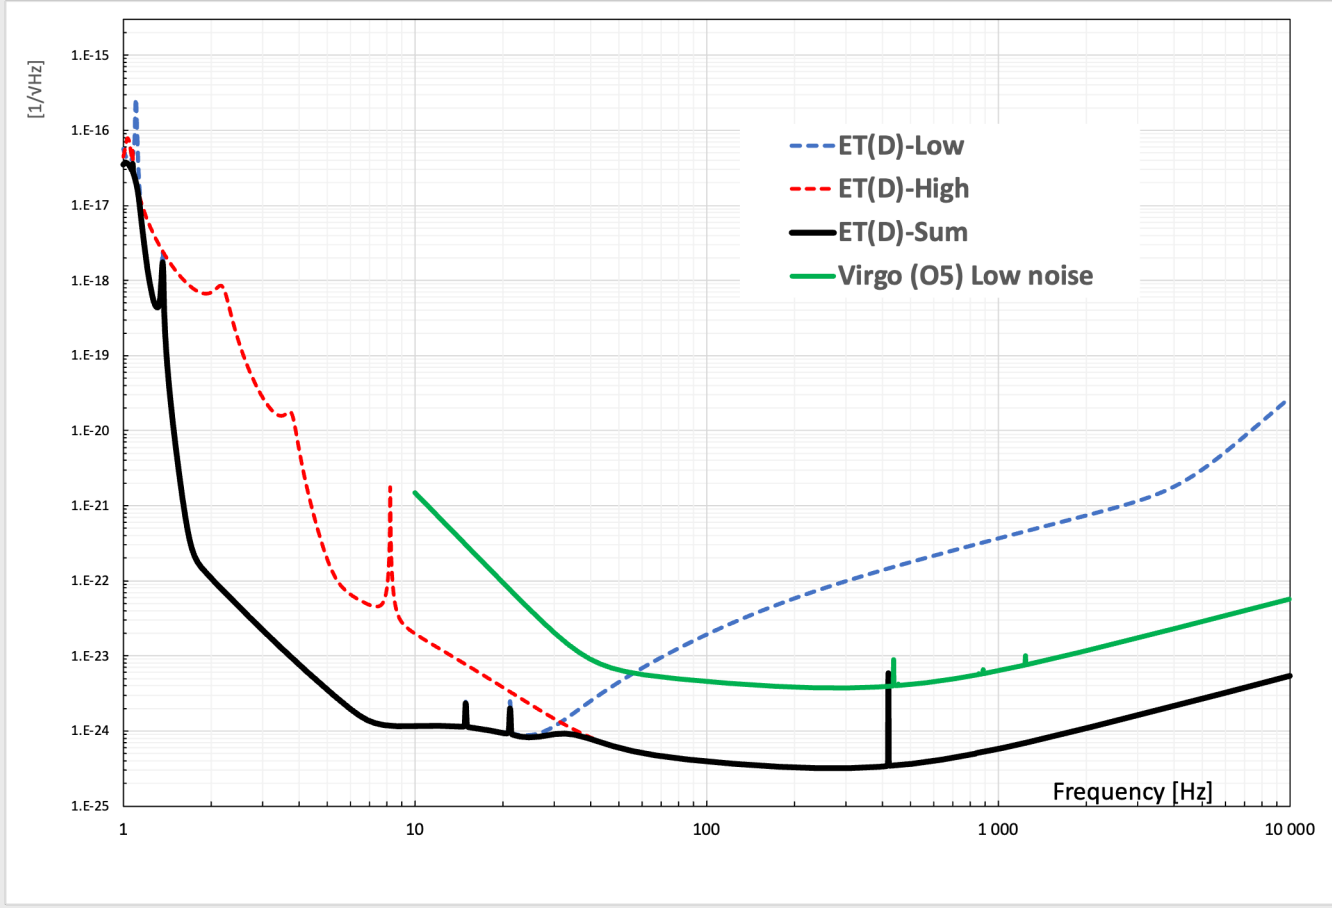

Binary Neutron Star "range"

[Mpc ≈ 3×10^{13} km ≈ 3×10^{6} ly]

	LIGO	Virgo	ET
O1	80		
O2	100	30	≈ 60
O3	100-140	40-50	
O4	160-190	(60-100)	
O5	230-325	150-260	
			>2000

ENHANCED SENSITIVITY

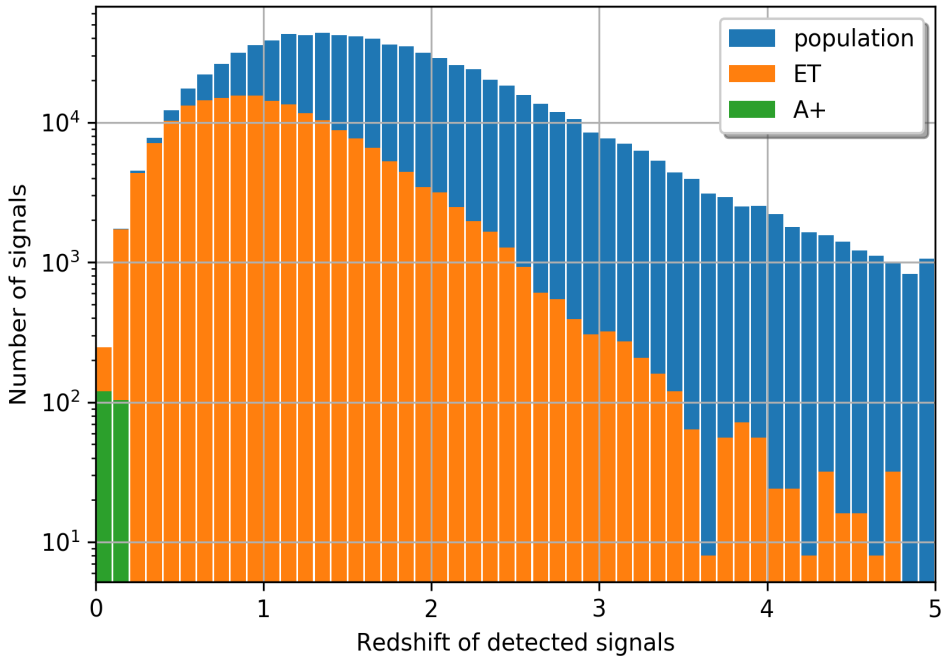

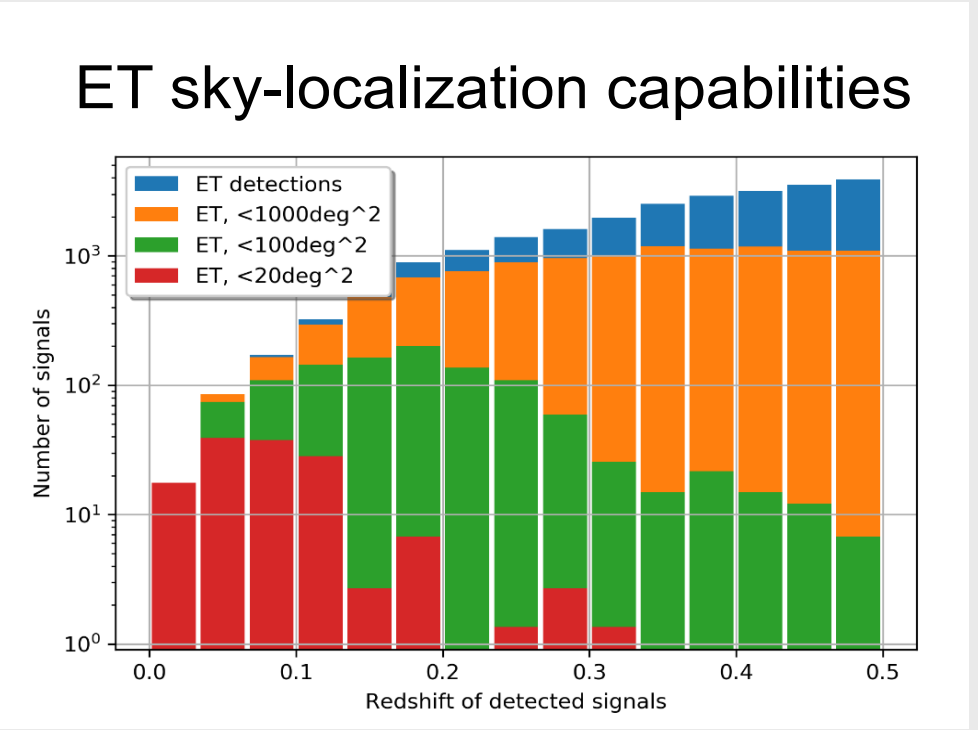

- 10^5 BBH detections per year
- 10^5 BNS detections per year

- o(100) detections per year with <20 deg^2
- Early warning by minutes (hours)

marica.branchesi@gssi.it

EARLY WARNING

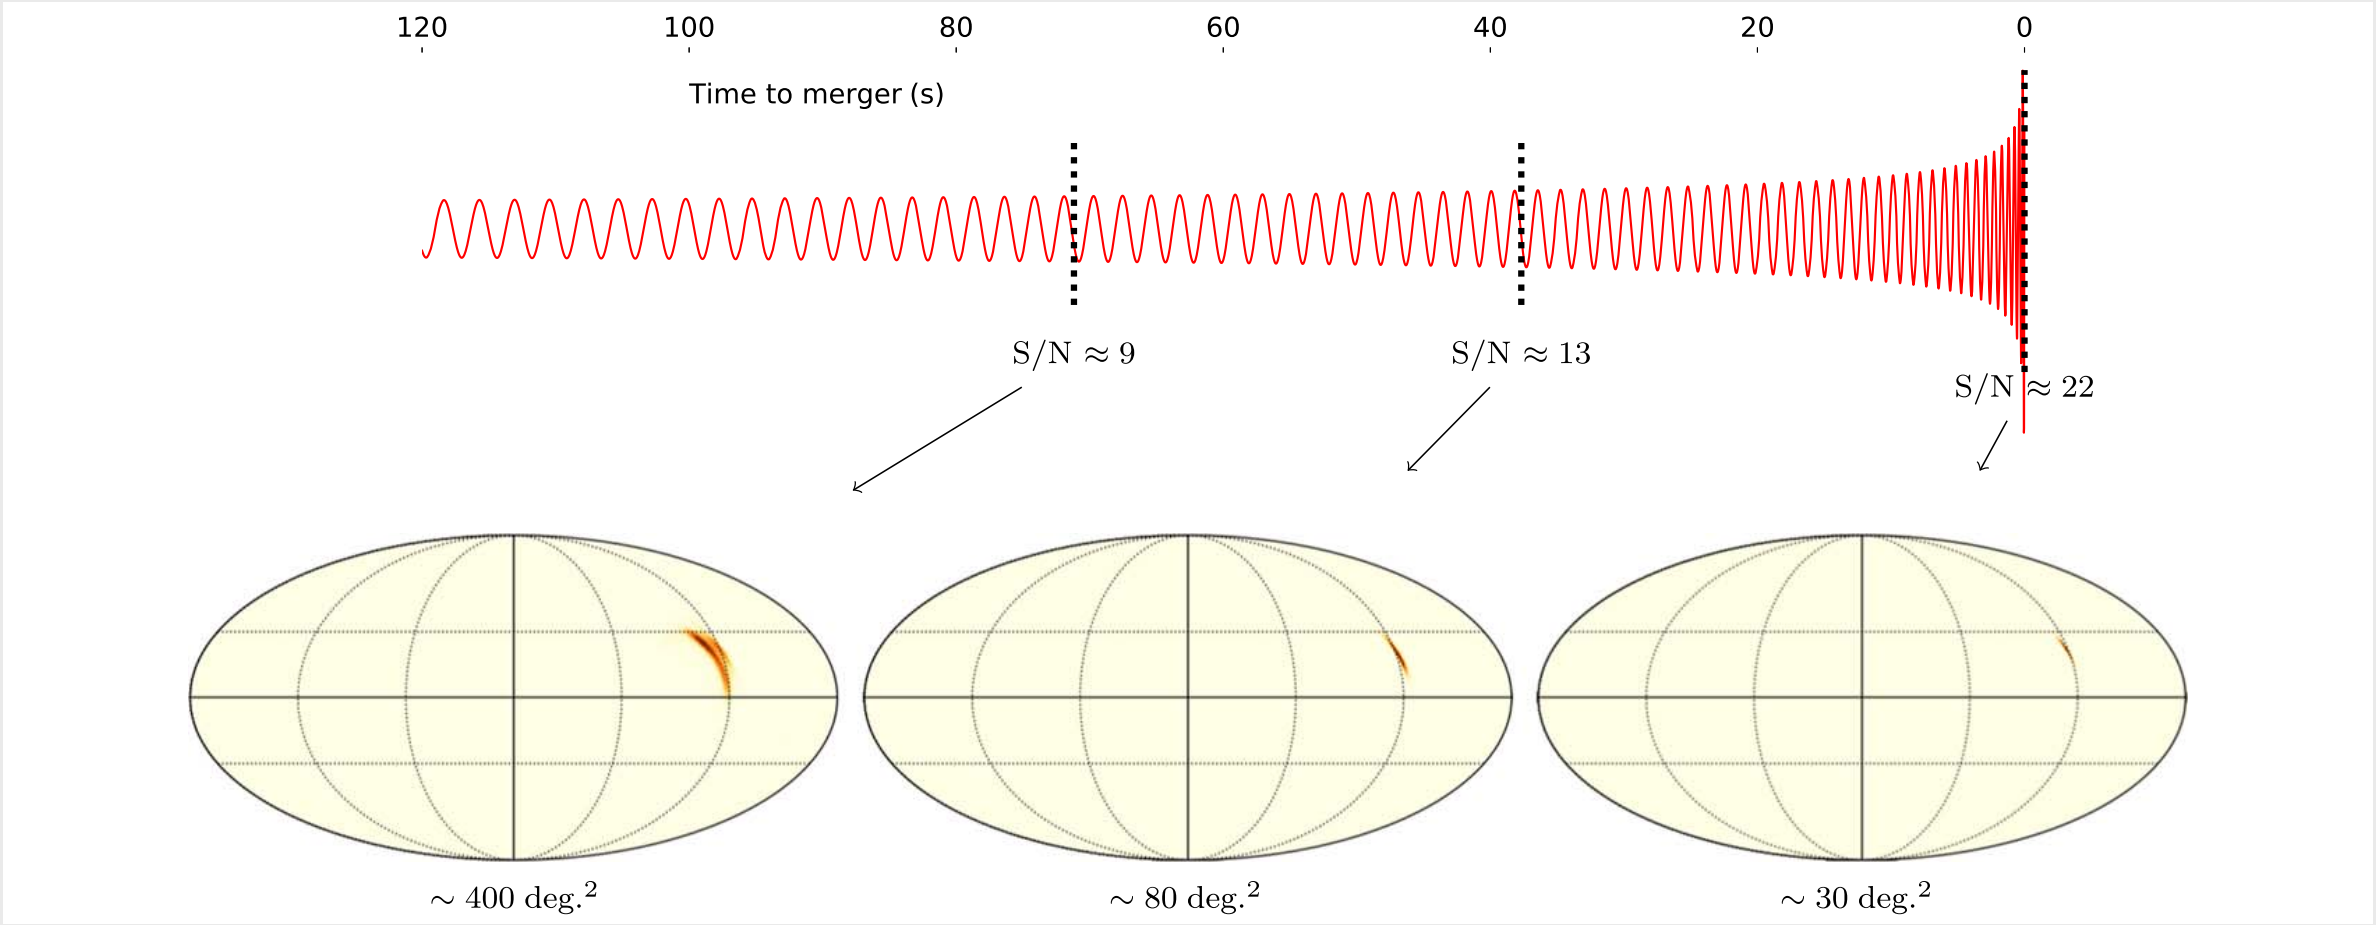

SETTING THE SCALE

- Raw interferometer data don't grow much with increasing instrument sensitivity
  - In ET we expect about few tens of PB of raw data per year (compare with o(2PB/year) today)
  - No big deal today, piece of cake by 2030's
- 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, correlated noise,...: accurately estimating the computing power needs is itself a difficult task
- The 10% of an LHC experiment mantra
  - "GW are 10% of ATLAS or CMS today, will be 10% of ATLAS or CMS in Run5"
  - This is more an ad claim than an actual target but helps having an idea of the scale

THE ECOSYSTEM

- First and foremost, other 3G facilities (and the 2G network)
  - CE (and LISA!)
  - Will the network work as it does today (i.e., exchanging LL data for collaborative analysis)?
- Multimessenger astronomy is the name of the game
  - CTA, SKA, KM3Net, ELT in the ESFRI roadmap, and many many more (e.g., Vera Rubin)
  - 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
- Also: “natively multimodal” analyses
  - Even in low-latency?
- Also: different underlying infrastructure
  - Likely dominated by large national HPC centres
- Also: very different role of AI technologies
  - Will be happening fast!

MANDATORY SLIDE WITH BOXES AND ARROWS

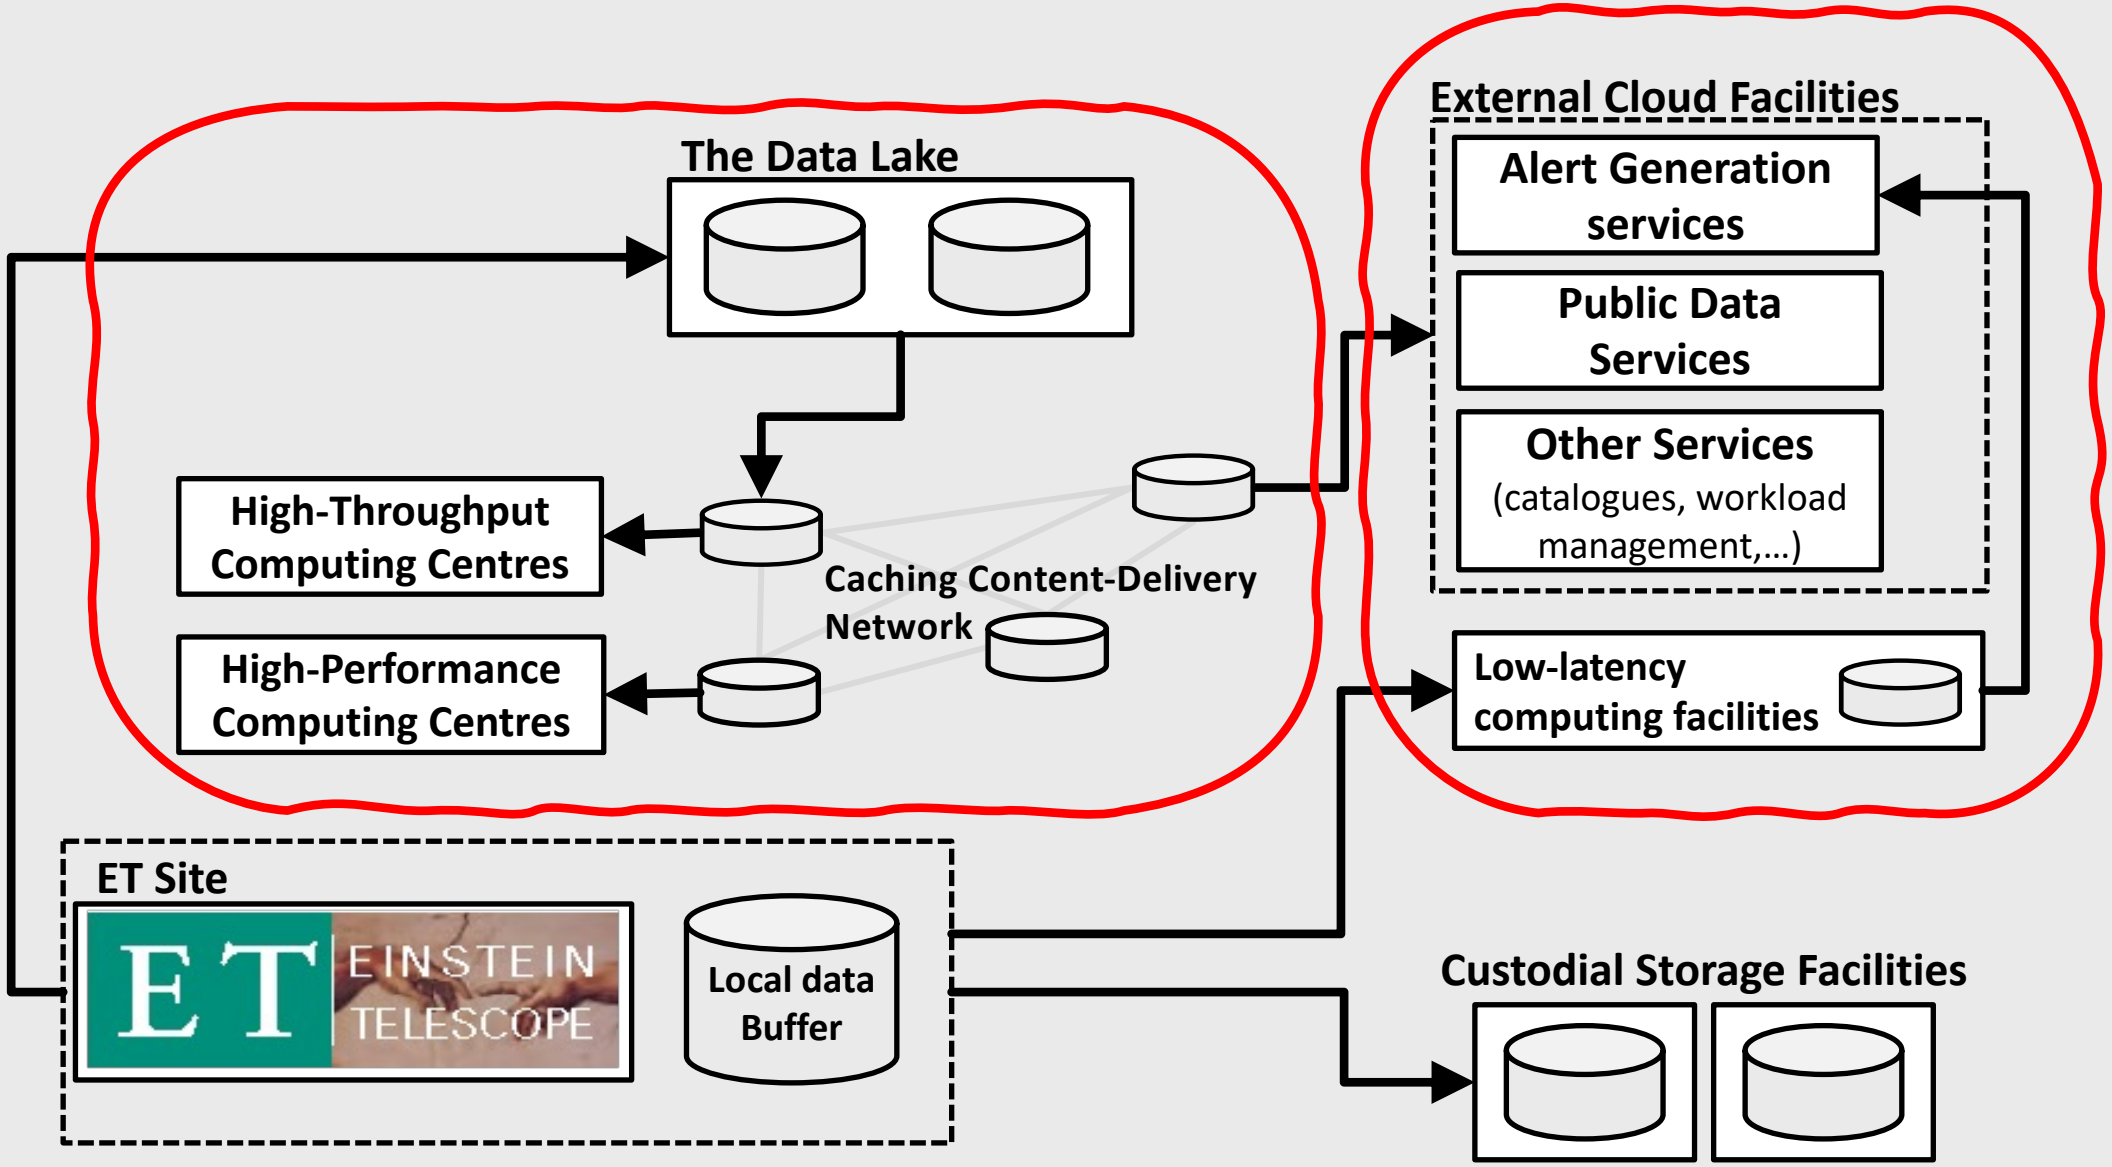

Figure from the ET ESFRI proposal

Synergies with Virgo: computing needs grow more or less as a continuum from Virgo O4 to O5 to Virgo-nEXT to ET, and technologies keep evolving.

- Distributed computing infrastructure
  - CPU power needs grow continuously with sensitivity (CBC PE)
  - ET already needs a working and evolving computing infrastructure (for MDCs, simulations,...)
- Low-latency alert distribution network
  - High rates imply high automation, long signals imply new features (e.g., continuous alert updates)
  - In the coming years the developments may be driven by running experiments, the GW community already needs to be present
- Sustainable computing
  - And, in general, technology tracking: heterogeneous computing, efficient algorithms, ML,...
  - Same message: development is a continuum

STRATEGY 1.0

- Use IGWN (=LVK) infrastructure as baseline
  - IGWN uses the European computing centres as an extension of the OSG (which is suboptimal…)
  - However, the functionality is there (OSDF + HTCondor)
- Use ESCAPE as the first toolbox
  - First the “Data Lake” (DIOS), then the Virtual Research Environment
  - Also, Virtual Observatory, streaming data,…
- Develop a common initial R&D program
  - Data Lake (Rucio) for data distribution
  - VRE/REANA for data access and job management
  - Using ET MDCs as testbeds

MOCK DATA CHALLENGES

- MDC as multipurpose tools
  - Develop and exercise analysis code and strategies
  - Build the data analysis community and bootstrap new groups
  - Educate the community in the use of common distributed computing tools and best practices
  - Iteratively test the distributed computing infrastructure

- Mock Data Challenge support
  - **MDC1:** provide data distribution layer (OSDF: CVMFS + cache) and survey the activities
  - **MDC2:** provide (possibly a set of) prototype tools for workload management etc.
  - **MDC3...n:** iterate

WORKFLOW EVALUATION KITS

- Independent packaged parts of the final architecture
  - Providing limited functionalities, possibly some as mere demonstrators
  - But actually to be released to users (i.e., they MUST be functional)
  - Different implementations may exist, with different tools/technologies used to provide same functionality
  - Integration of existing tools, with little bespoke developments, to map “kits” onto small(ish) projects
- Examples:
  - ESCAPE Datalake + RucioFS for data distribution
  - IAM-based AAI
  - ESCAPE Datalake + VRE interactive data analysis
  - OSDF + INFNCloud interactive data analysis
  - “Packaged” and quality-tested MDC data generation tool
  - HSF rich metadata tool

PROJECTS & MORE PROJECTS

- ET-PP
  - WP8 – very good collaboration with eIB
- ICSC_S2
  - See next slides
- M2Tech v2
  - WP6 – see next slides
- ESCAPE/OSCARS proposals
  - MADDEN (See Lia's poster)
  - ETAP (Université de Genève)
  - Streaming data for Low-latency?
- ETIC
  - See Luca's poster on CTLab/TechZoo

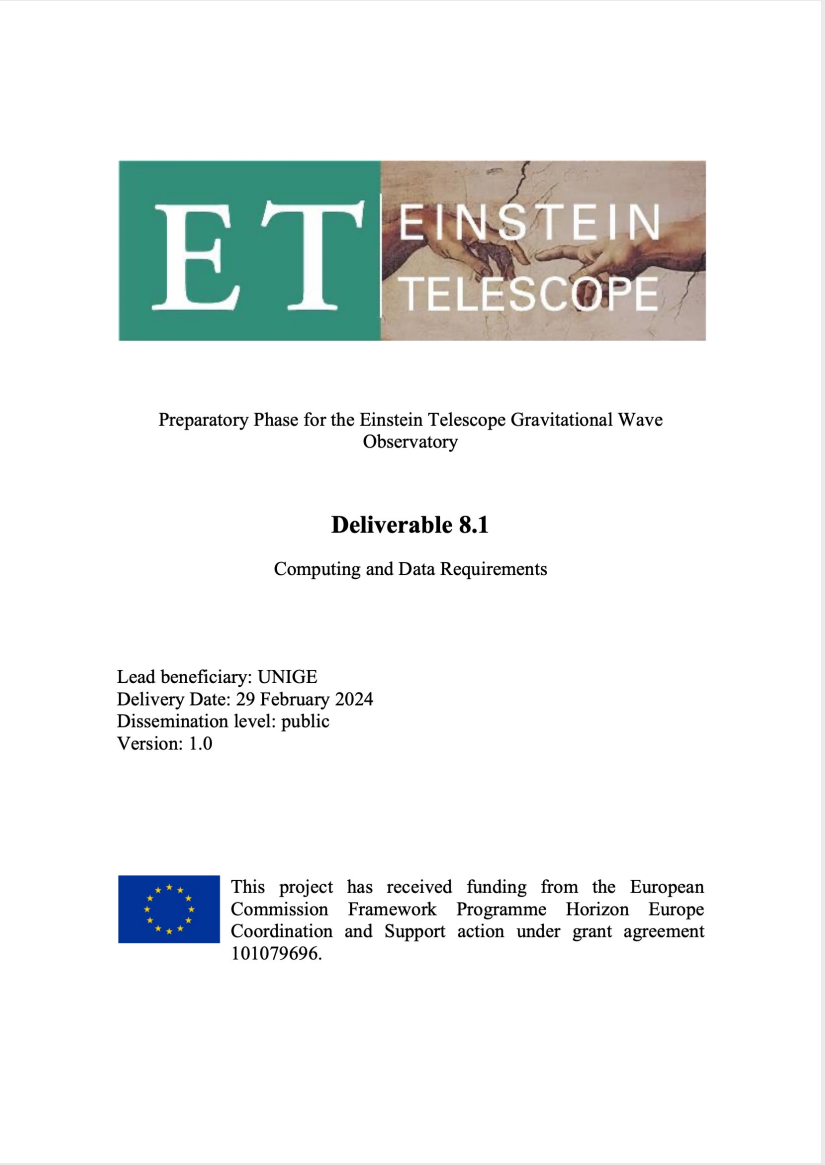

ET EINSTEIN TELESCOPE

Preparatory Phase for the Einstein Telescope Gravitational Wave Observatory

Deliverable 8.1

Computing and Data Requirements

Lead beneficiary: UNIGE
Delivery Date: 29 February 2024
Dissemination level: public
Version: 1.0

This project has received funding from the European Commission Framework Programme Horizon Europe Coordination and Support action under grant agreement 101079696.

WP6: Technologies for multimessenger astronomy together with CTA, KM3NeT, Virgo.

Task 6.1: efficient data processing

- Early robust processing = less to process later = better energy efficiency
- Mostly supervised ML to enable fast/real-time data processing to enhance MM event identification

Task 6.2: sustainable large-scale computing
Involved: CNAF (Daniele Cesini), INFN-Torino/CTLab (SB)

- How to sustainably scale computing to handle large MM event rates and mitigate energy/carbon costs
- Work with large computing centres to study how to scale-up computing for large MM event rates Together with academic partners to bridge the gap between RIs and computing centres

Task 6.3: multimessenger alert tools

Involved: INFN-PG (Giuseppe Greco, task leader)

- How to ensure different research infrastructures can communicate effectively
- Common alert formats, brokers, databases, etc - all while ensuring alerts follow FAIR principle

Frequency Hough Transform analysis on GW continuous sources

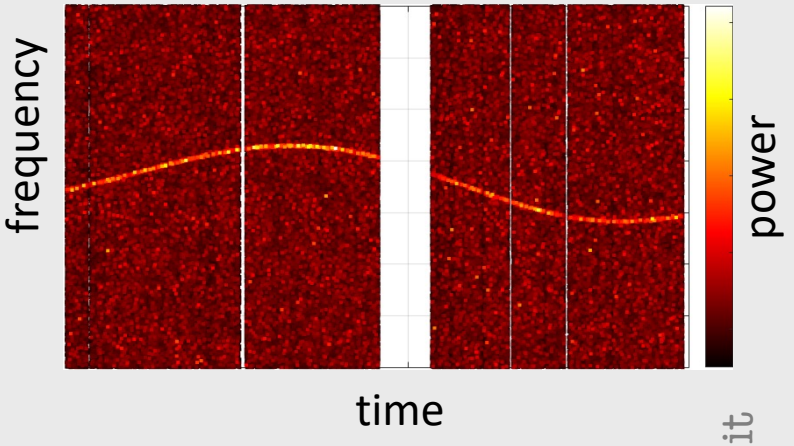

- Frequency Hough is a method for blind, all sky searches of continuous GWs, based on **time-frequency maps of the data processed with a Hough transform**. Computational costs are among the major limitations.
- Einstein Telescope is expected to be able to observe a **huge population of isolated pulsars** as possible GW sources: a big opportunity for all-sky searches.
- With Einstein Telescope, observing runs few-years long are expected. While it contributes to increase the sensitivity of continuous GW searches, it also **increases the volume of the parameter space** and consequently the computational load.
- **Optimizing the algorithm** is crucial: shortening the computing time leaves room to analyze more deeply the data, thus increasing the overall search sensitivity.

Lorenzo.Pierini@roma1.infn.it

ICSC_S2: FLAGSHIP USE CASE 2.3.3

Inference of cosmological and astrophysical population properties from GW observations with and without electromagnetic counterparts

Collaborators: Michele Moresco (Unibo, PI), Matteo Tagliazucchi (PhD, Unibo), Francesco Pannarale (UniRoma 1), Simone Mastrogiovanni (INFN - Roma 1)

Science case: The goal of this UseCase is to **develop, improve, and employ new algorithms to study gravitational waves as cosmological probes** (bright, dark and spectral standard sirens) based on Hierarchical Bayesian Statistics, combining both information of galaxy catalogs and astrophysical properties of the Compact Binary Systems studied (Binary Black Holes, Binary Neutron Stars, Black Hole-Neutron Star binaries) **to be able to handle the data volume of future 3G GW observatories like Einstein Telescope**.

Challenges: Currently, the LVK network identified 90 BBH, and already with current GW detections we are almost at the limit for current data analysis pipeline. In O5 (2027-2029) we expect ~3000 sources, and **with the future Einstein Telescope O(1e5)-O(1e6) GWs**.

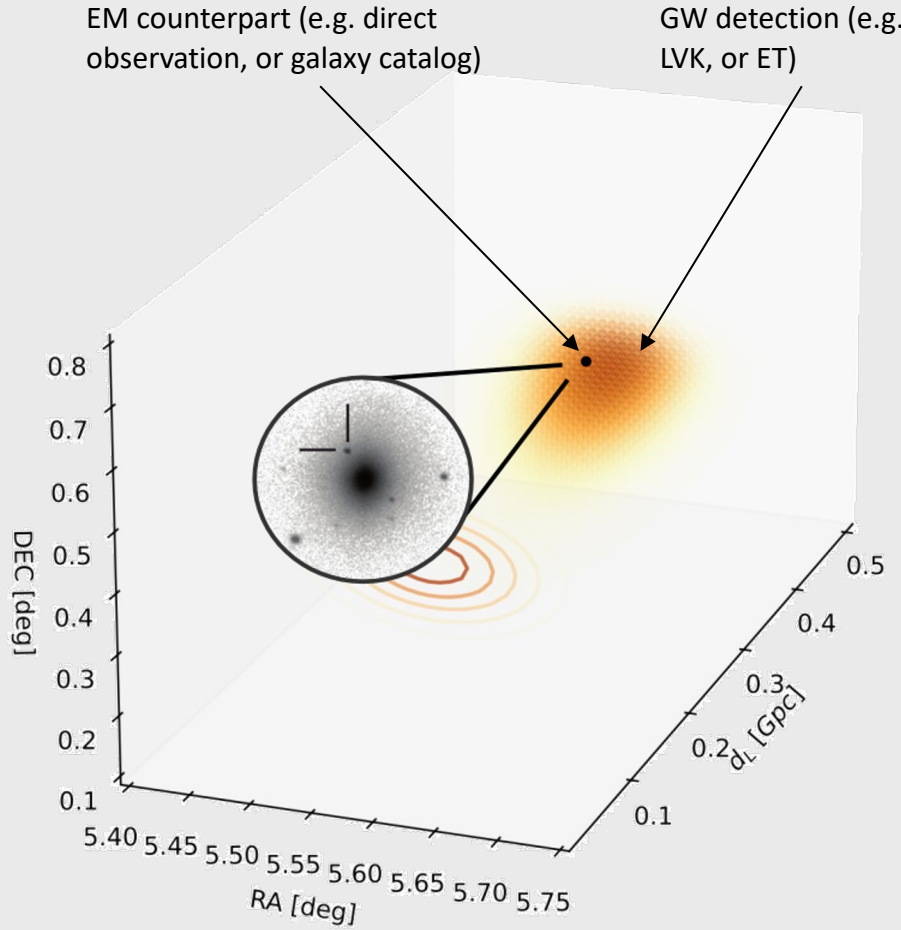

michele.moresco@unibo.it

ICSC_S2: FLAGSHIP USE CASE 2.3.3

Developed codes: CHIMERA (Unibo), icarogw (INFN - Roma 1)

Expected technological advances: We will improve the icarogw and CHIMERA pipelines to accommodate the population fits for thousands of GW sources:

- **porting the codes from CPUs to GPUs** to accelerate the inference pipeline (of at least a factor 10)
- implementation of the codes on High Performance Computing (**HPC**) facilities
- first **mock data challenge** consisting of more than 2000 simulated GW events, to assess possible limitations related to timing or numerical difficulties that we will encounter in the next few years
- potential application to real data

Available resources: Leonardo, ~15k GPU hours and ~200k CPU hours to port and test the pipelines on GPUs, develop and run the MDC; ~ 35 TB of storage space

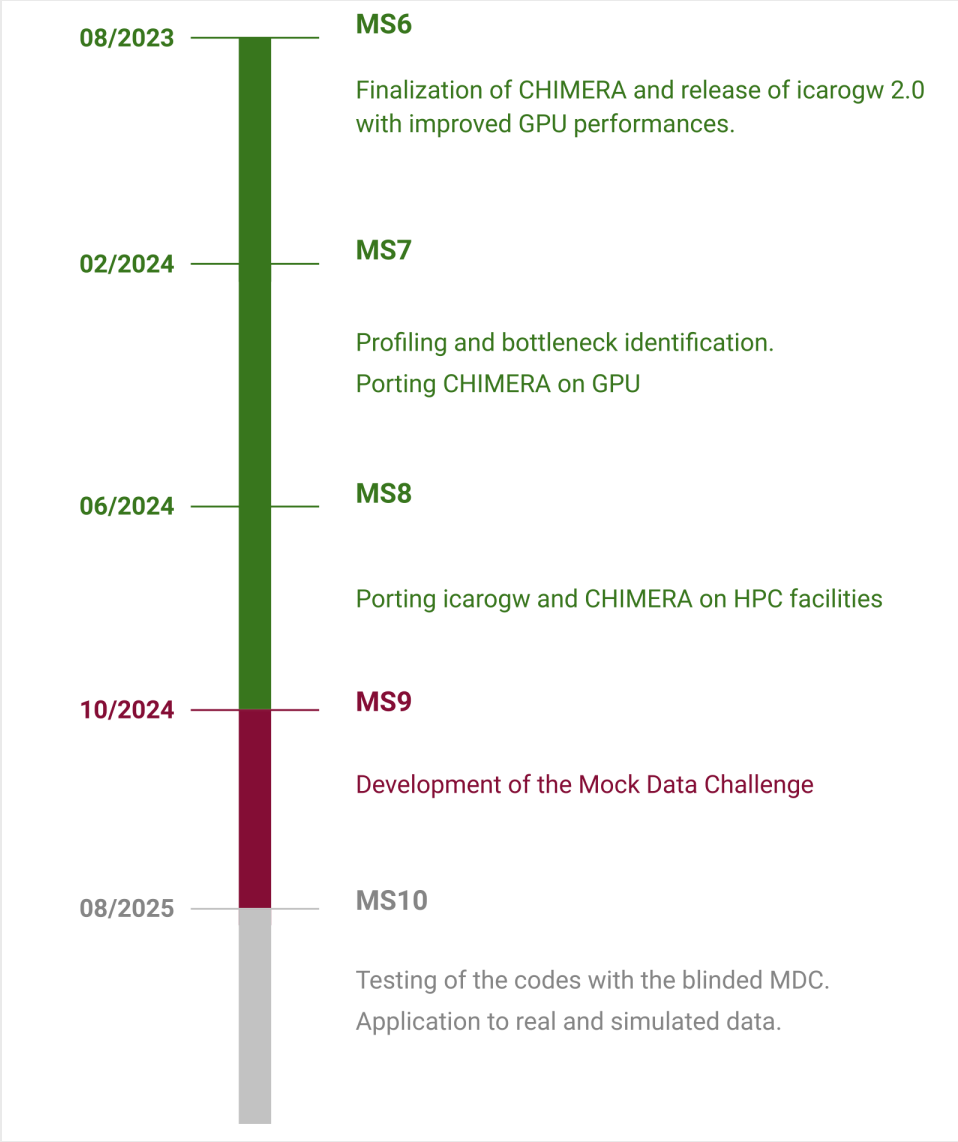

michele.moresco@unibo.it

OSCARS PROPOSALS

- ETAP (Université de Genève)
  - Access to multiple ESCAPE Data Lakes.
  - Rich metadata service integration
  - Access to multiple rich metadata instances
  - A lightweight CRM service monitoring the VRE
- MADDEN (INFN-TO & Université Catholique de Louvain)
  - Multi-RI Data Lake managed with Rucio.
  - Development and test of RucioFS
  - Extend RucioFS to support advanced metadata
- Second OSCARS call (November?)
  - Streaming data for LL?
  - Something IVOA-related?

TECHNOLOGY TRACKING @ETIC

- ETIC TT platform being built in Torino ("TechZoo")
  - Heterogeneous and expandable HPC platform
  - Interoperable with the TeRABIT "HPC Bubble"
  - Access layer via INFNCLOUD, common with similar facility at INFN-BO
  - Usable for code porting, testing, special architectures, accelerators evaluation etc.
  - ...and for regular computing (e.g., numerical relativity)
- Hardware being configured, possibly more coming
- Expect a call for applications in early summer
- See Luca's poster

CONCLUSIONS

- Activities are rampig up exponentially
  - Development of the Computing Model
  - Mitigation of computing needs and development of technologies
  - Support to Mock Data Challenges
- Incidentally: personpower is not!
  - Lots of people with o(10%)
  - But working on that

Questions?