©Johan Jarnestad/The Royal Swedish Academy of Sciences

COMPUTING CHALLENGES IN THE EINSTEIN TELESCOPE PROJECT

Stefano Bagnasco, INFN Torino

The ET Generation | EGO - Cascina
Jun 19, 2024

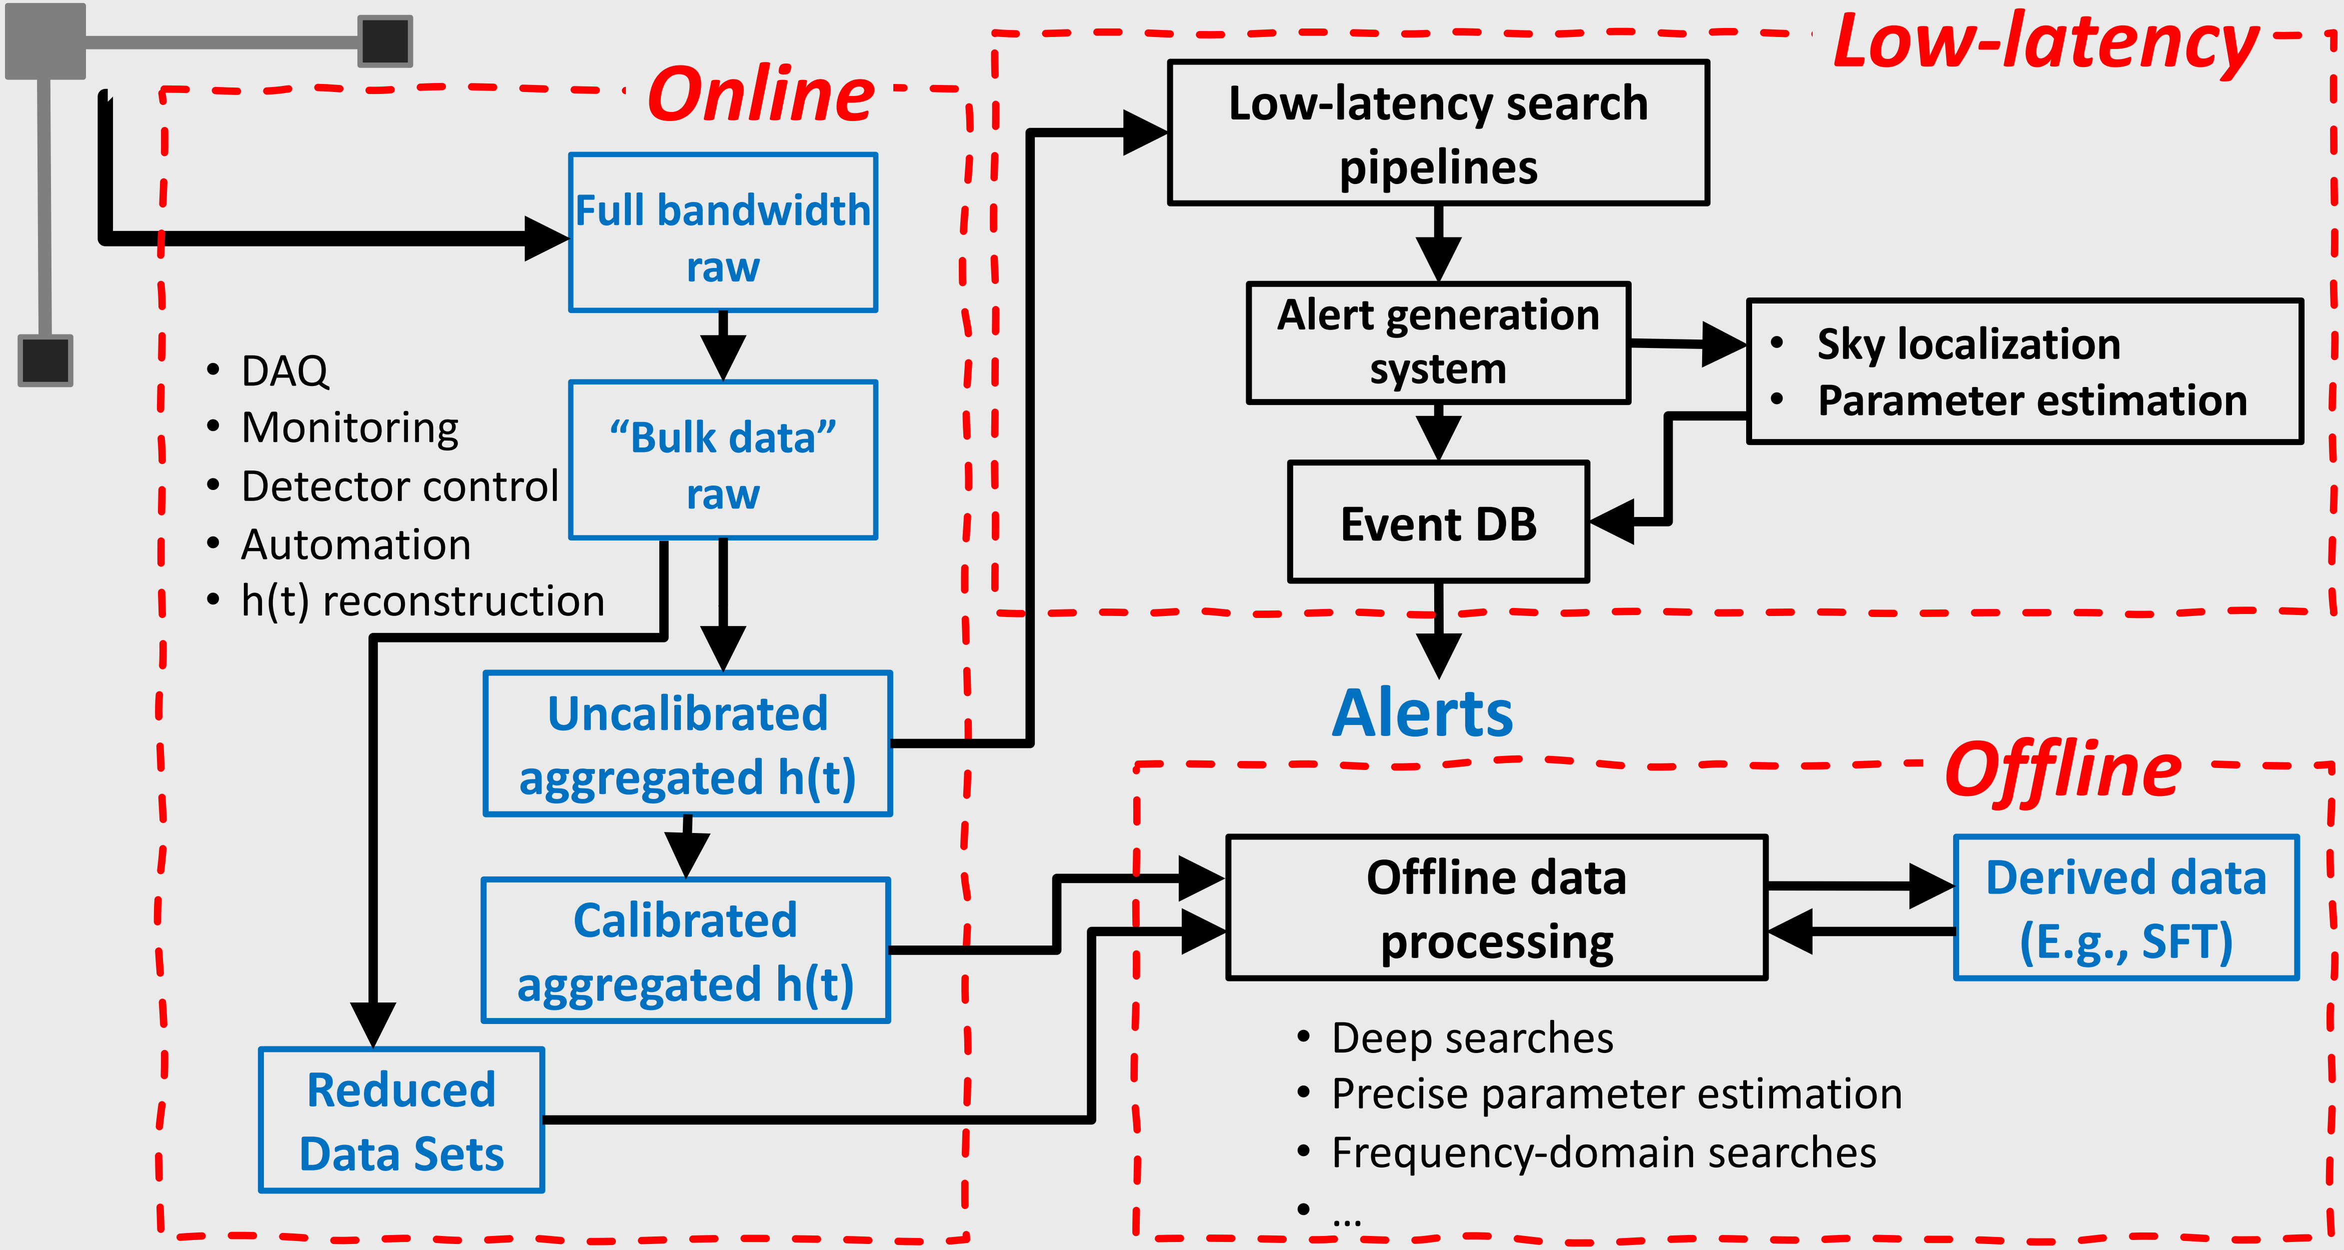

Online

Full bandwidth raw

"Bulk data" raw

- DAQ
- Monitoring
- Detector control
- Automation
- h(t) reconstruction

Uncalibrated aggregated h(t)

Calibrated aggregated h(t)

Reduced Data Sets

Low-latency

Low-latency search pipelines

Alert generation system

- **Sky localization**
- **Parameter estimation**

Event DB

Alerts

Offline

Offline data processing

Derived data (E.g., SFT)

- Deep searches
- Precise parameter estimation
- Frequency-domain searches
- …

BY-NC-SA

THE MANDATORY SLIDE WITH BOXES AND ARROWS

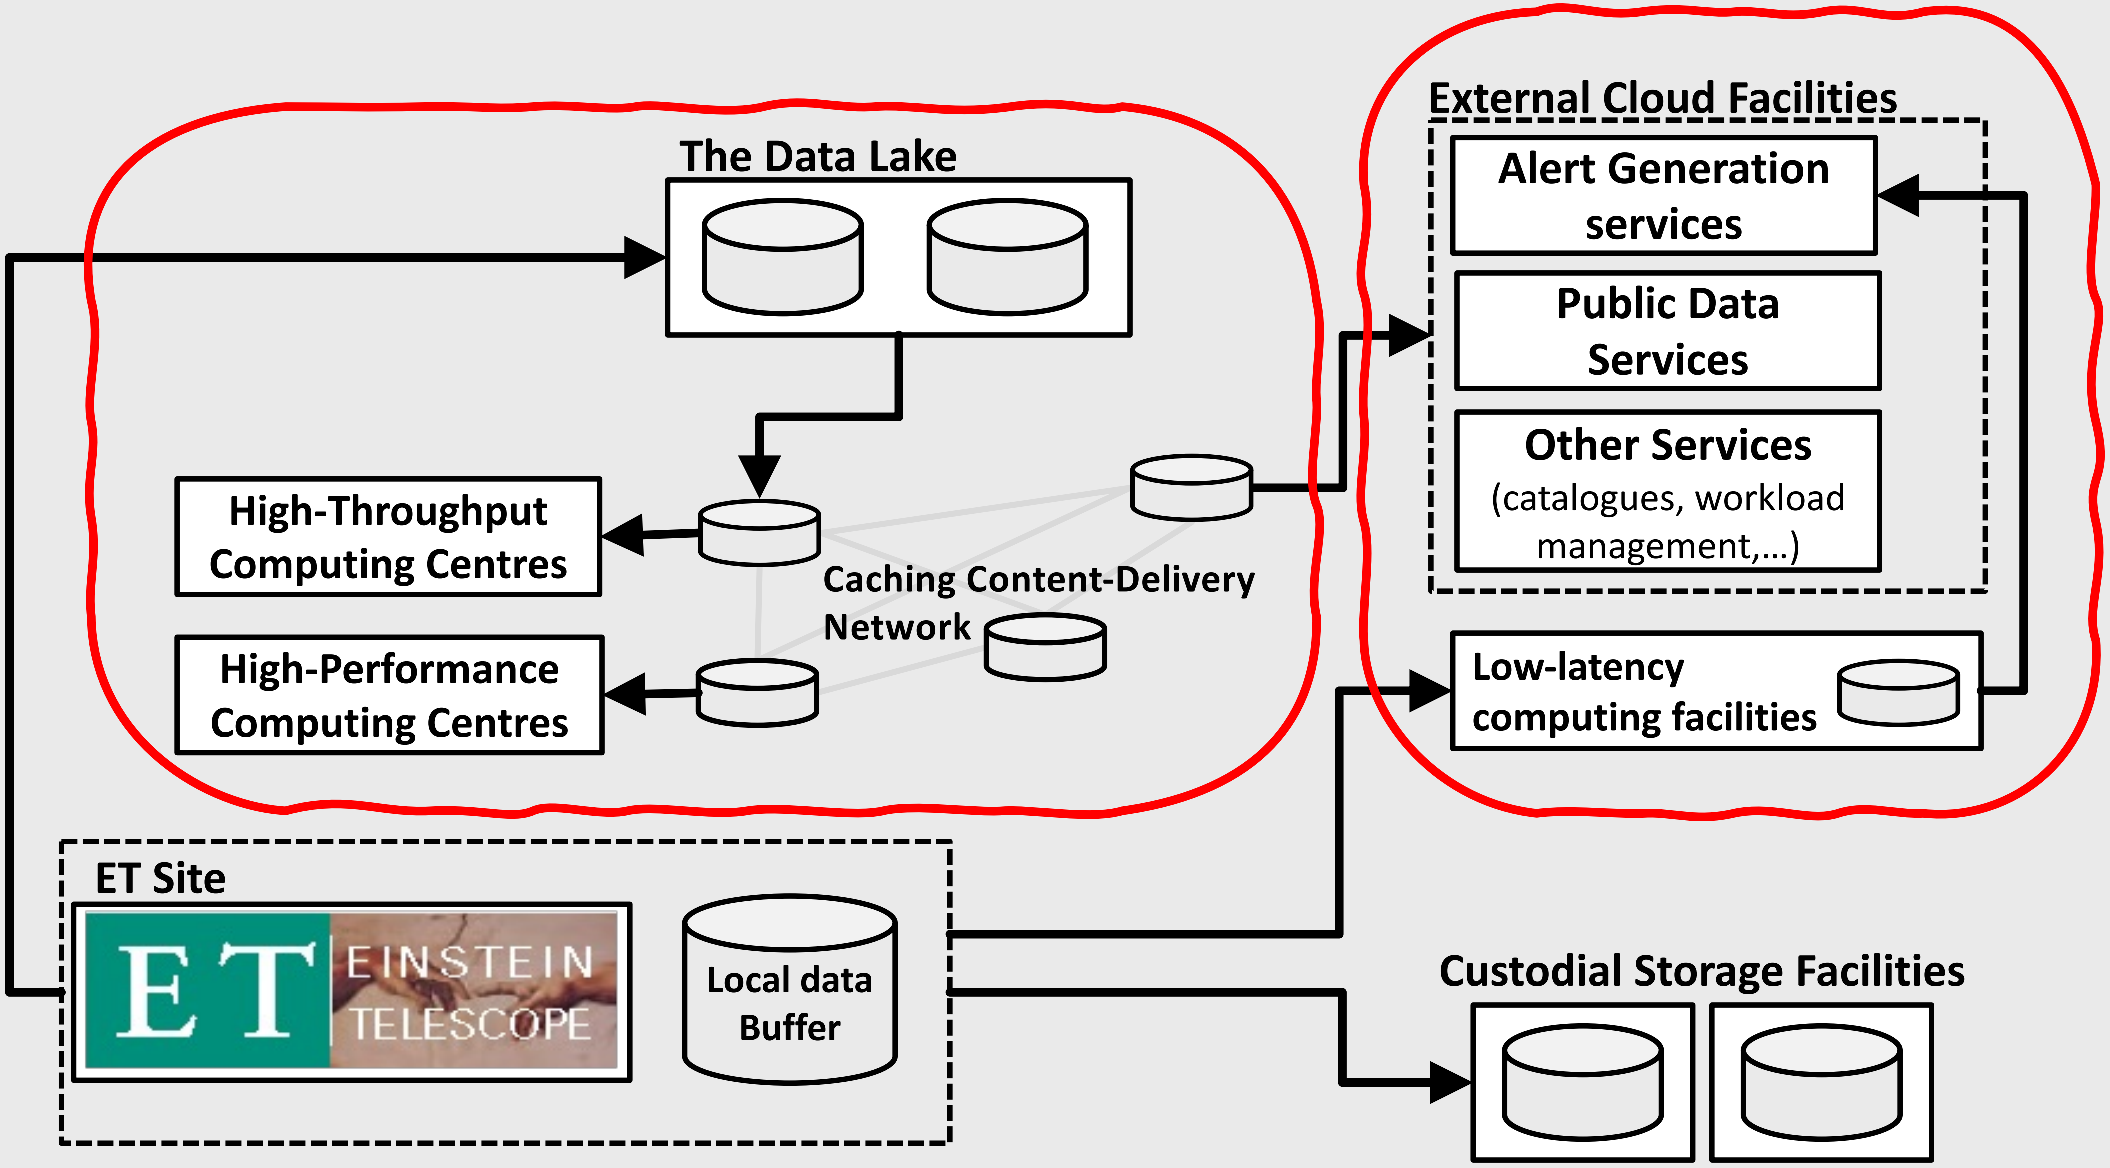

Figure from the ET ESFRI proposal

BY-NC-SA

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

BY-NC-SA

THE ET E-INFRASTRUCTURE BOARD

Chairs: Patrice Verdier (IJCLab), SB (INFN-Torino)

Division 1: Software, frameworks, and data challenge support
Chair: Andres Tanasijczuk (UCLouvain)
OSB Liaison: John Veitch (U. Glasgow), Elena Cuoco (EGO)

Division 2: Services and Collaboration Support
Chair: Antonella Bozzi (EGO)

Division 3: Computing and data model, Resource Estimation
Chair: Gonzalo Merino (PIC)

Division 4: Multimessenger alerts infrastructure
Chair: Steven Schramm (U. Geneva)

TTG: Technology Tracking working Group
Chair: Sara Vallero (INFN-Torino)

BY-NC-SA

WAY FORWARD

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

BY-NC-SA

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 (Virgo+ET) initial R&D program
    - Data Lake (Rucio) for data distribution
    - VRE/REANA for data access and job management
    - Using ET MDCs as testbeds

BY-NC-SA

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

BY-NC-SA

PROJECTS & MORE PROJECTS

- ET-PP
  - WP8 – very good collaboration with eIB
- ICSC_S2
  - "Flagship" use cases development
- M2Tech v2
  - WP6 – see next slides
- ESCAPE/OSCARS proposals
  - MADDEN (INFN-TO & UCLouvain)
  - ETAP (Université de Genève)
  - Streaming data for Low-latency?
- ETIC
  - See slide about CTLab/TechZoo

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.

BY-NC-SA

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?

BY-NC-SA

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

BY-NC-SA

Thanks!

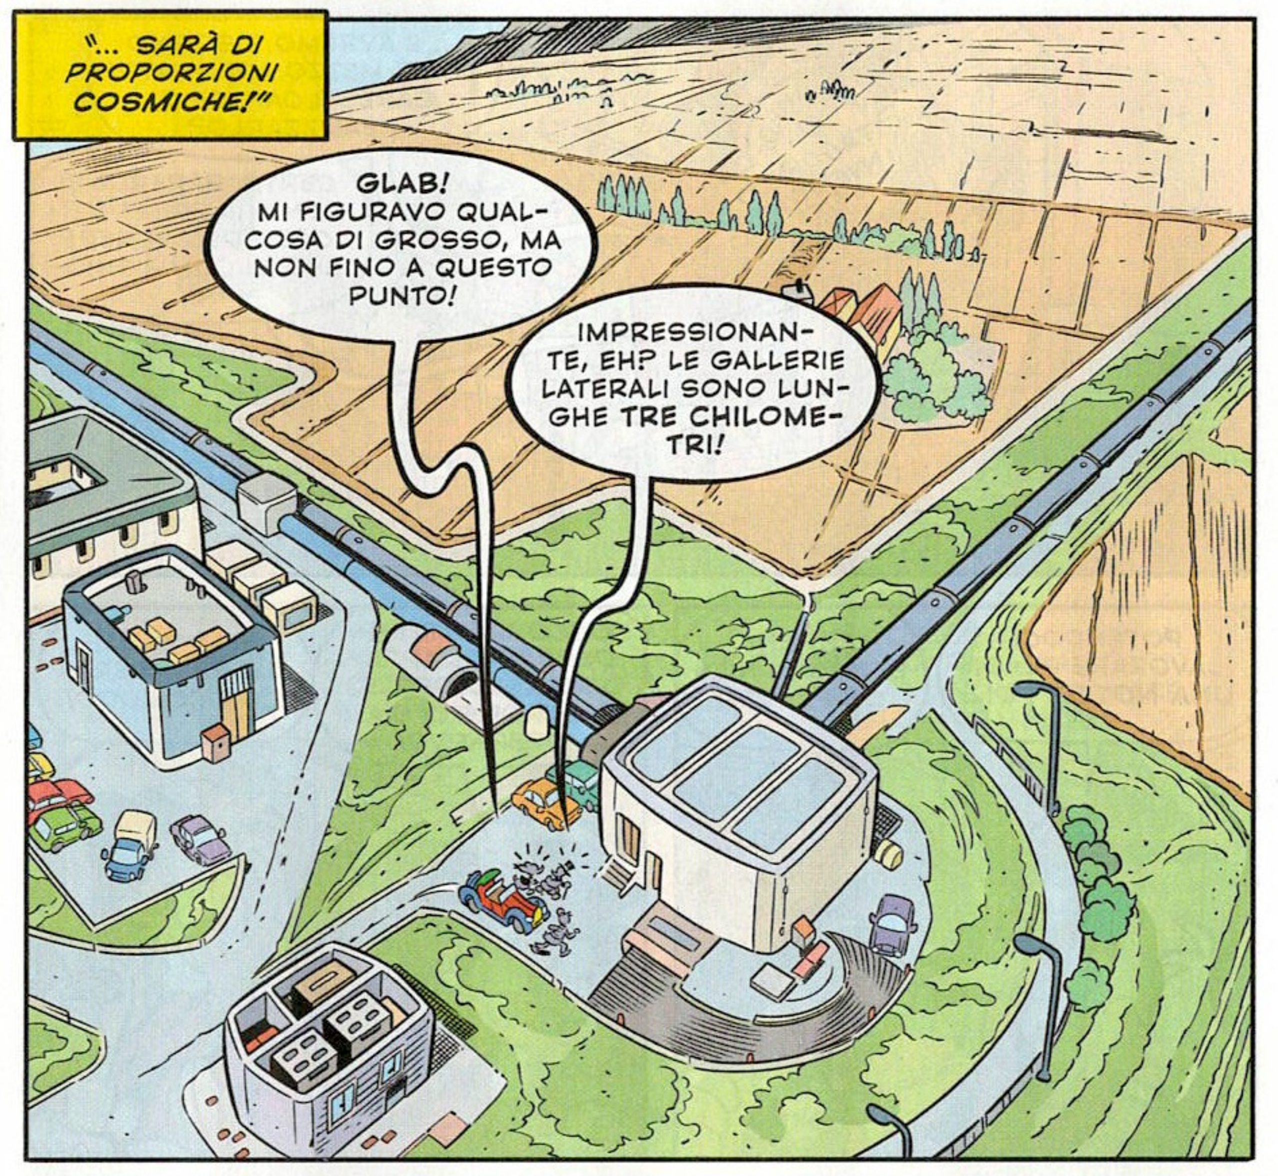

BY-NC-SA

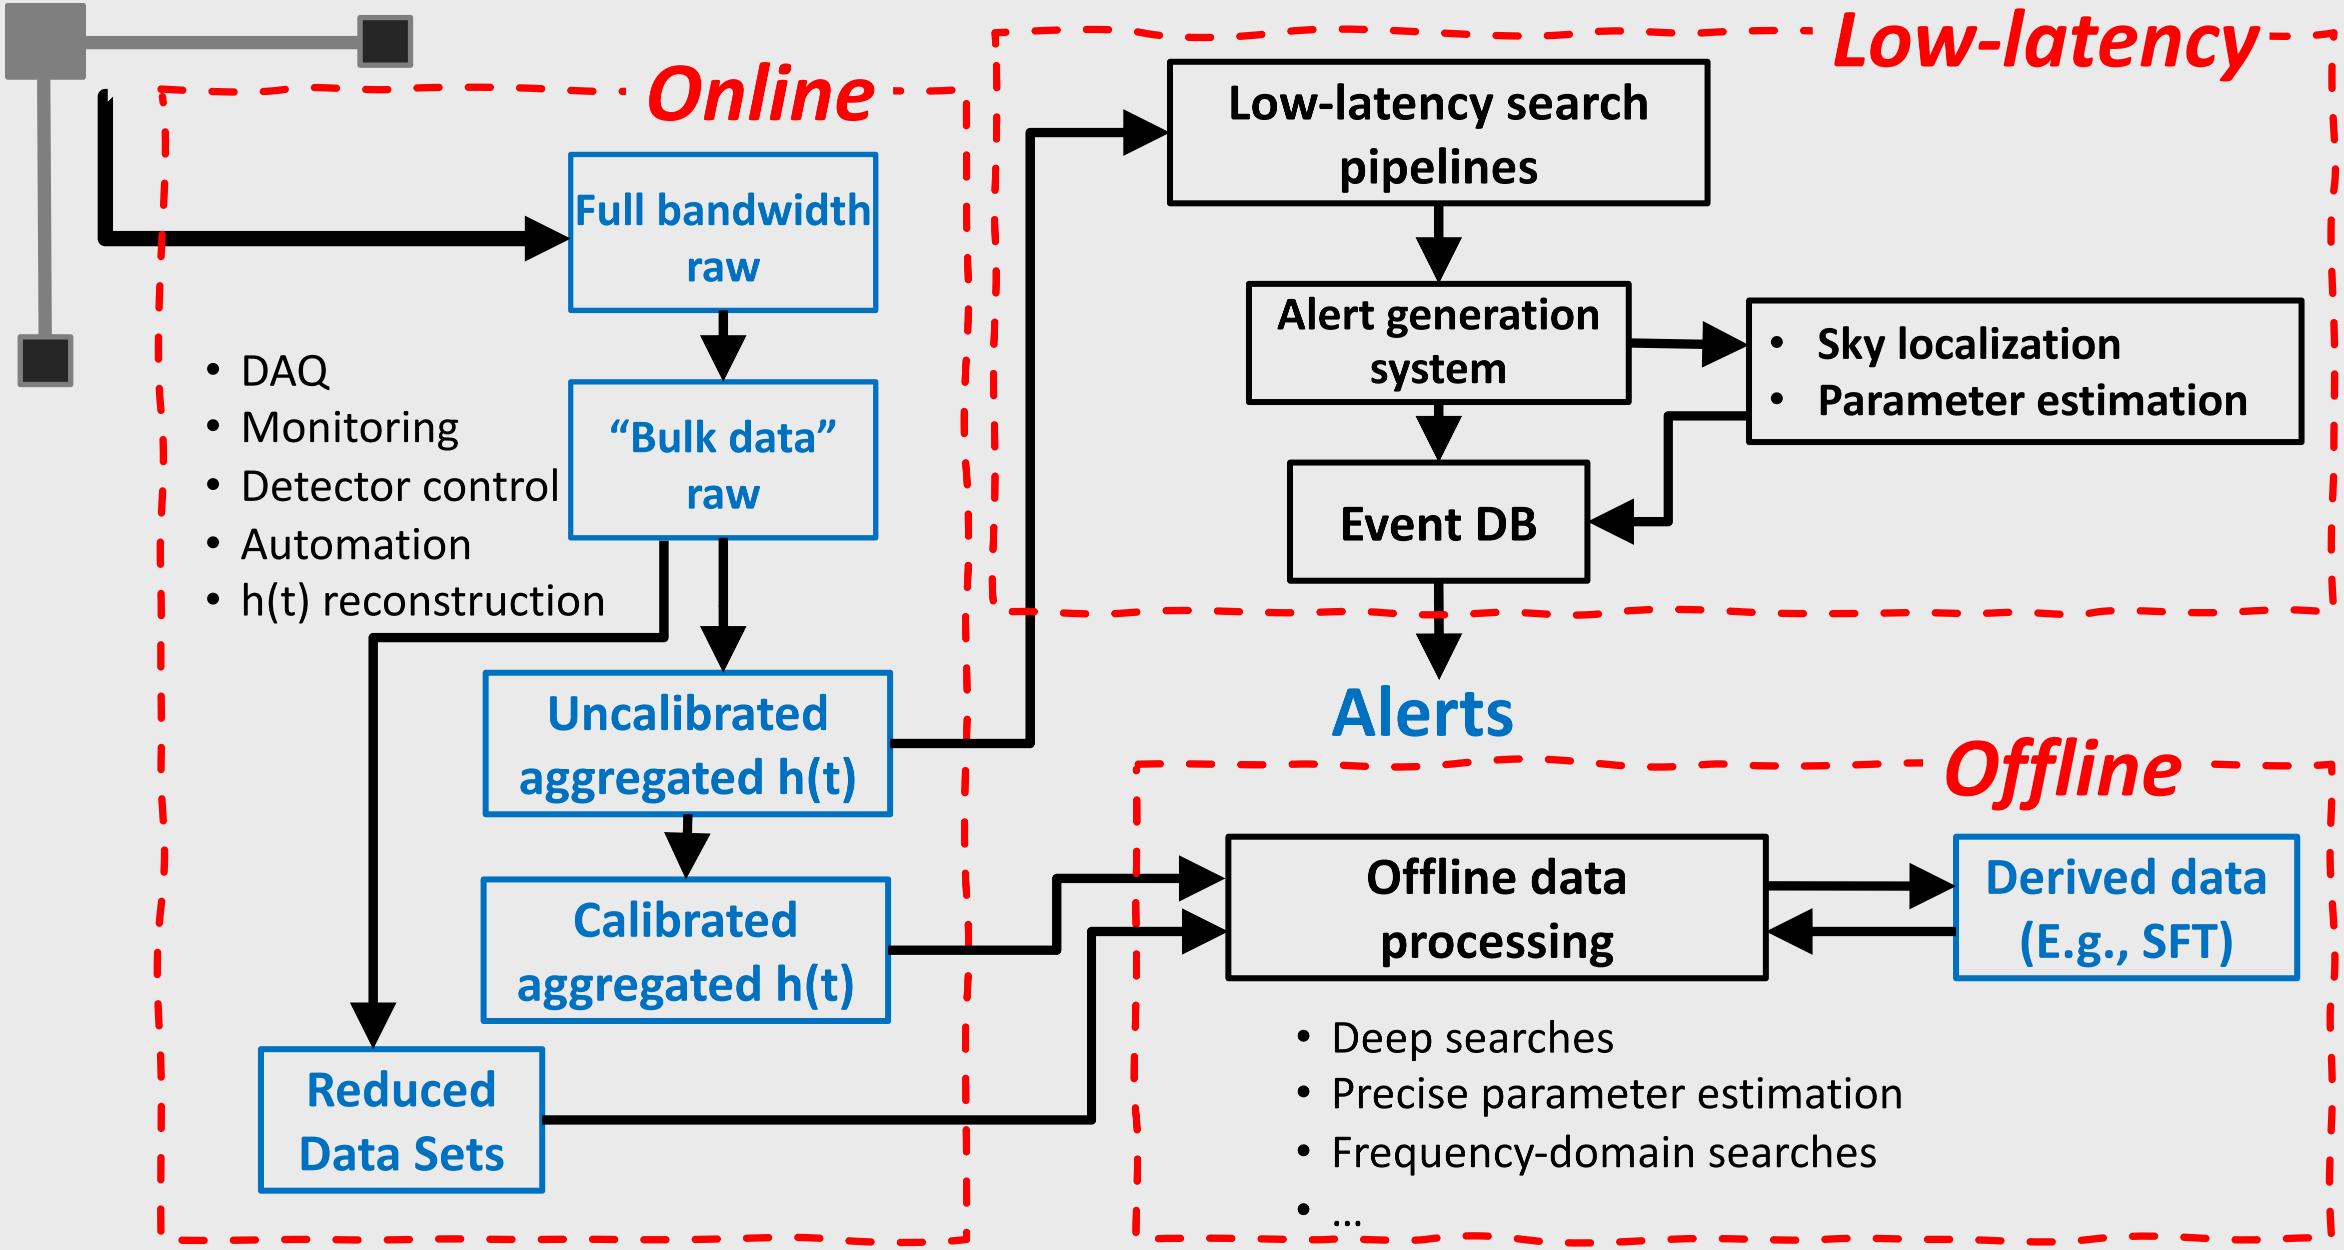

Online

Full bandwidth raw

"Bulk data" raw

Uncalibrated aggregated h(t)

Calibrated aggregated h(t)

Reduced Data Sets

- DAQ
- Monitoring
- Detector control
- Automation
- h(t) reconstruction

Low-latency

Low-latency search pipelines

Alert generation system

- **Sky localization**
- **Parameter estimation**

Event DB

Alerts

Offline

Offline data processing

Derived data (E.g., SFT)

- Deep searches
- Precise parameter estimation
- Frequency-domain searches
- …

BY-NC-SA

COMPLEX OVERALL DATA FLOWS

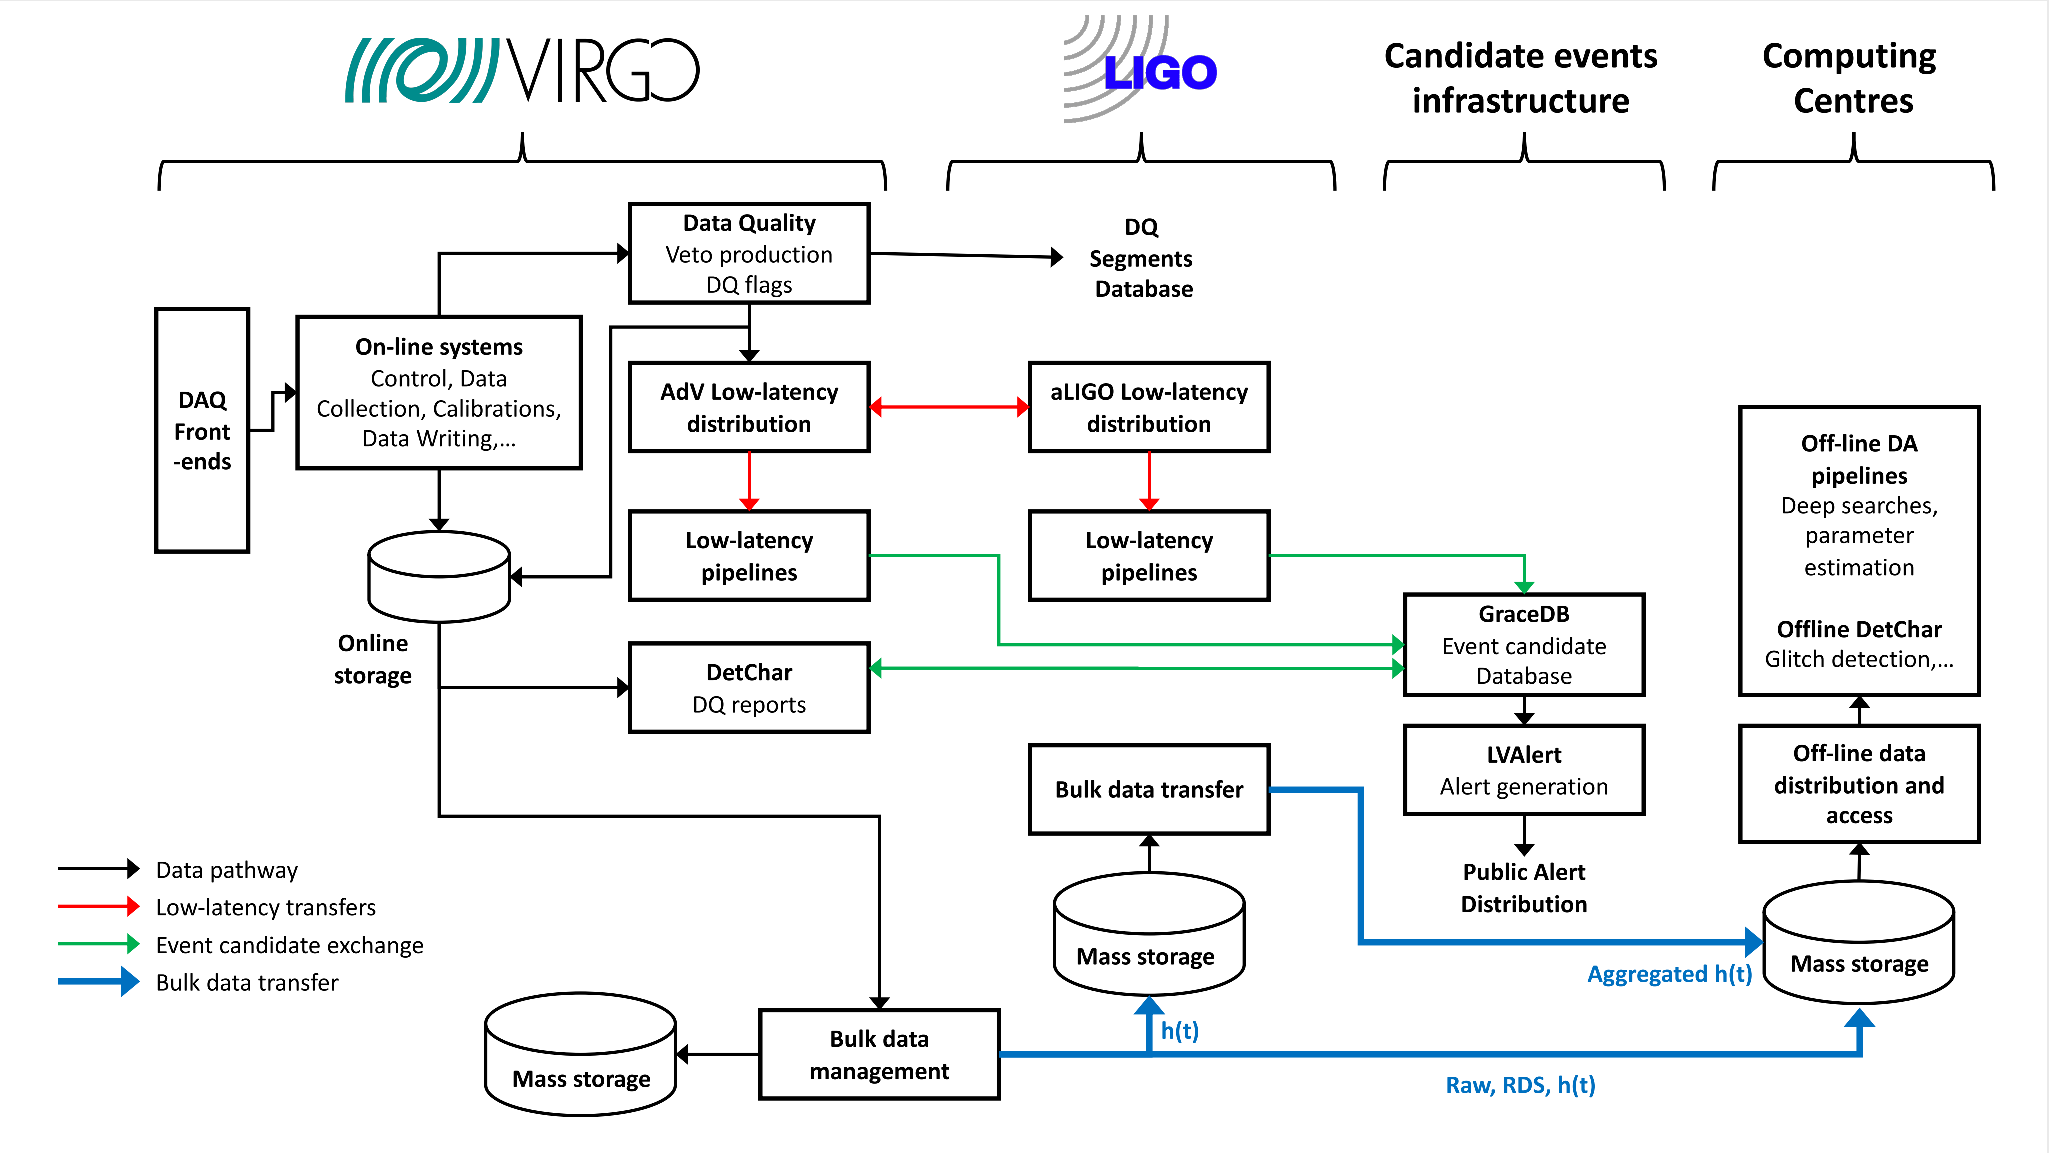

VIRGO

LIGO

Candidate events infrastructure

Computing Centres

DAQ Front -ends

On-line systems
Control, Data Collection, Calibrations, Data Writing,…

Data Quality
Veto production
DQ flags

DQ Segments Database

AdV Low-latency distribution

aLIGO Low-latency distribution

Low-latency pipelines

Low-latency pipelines

Online storage

DetChar
DQ reports

GraceDB
Event candidate Database

LVAlert
Alert generation

Public Alert Distribution

Off-line DA pipelines
Deep searches, parameter estimation

Offline DetChar
Glitch detection,…

Off-line data distribution and access

Bulk data transfer

Mass storage

Mass storage

Mass storage

Bulk data management

Aggregated h(t)

h(t)

Raw, RDS, h(t)

Data pathway
Low-latency transfers
Event candidate exchange
Bulk data transfer

Updated from S.B., *EPJ Web of Conferences* **245**:07050 (2020)

BY-NC-SA

A FEW NUMBERS

- **Raw data size:** about 120TB per month of observation per observatory
  - Includes all control channels from the instrument
  - Transferred to custodial storage for safekeeping
- **“Aggregated” data for analysis:** 10TB/year per observatory
  - Includes the single physics channel and summary “data quality” information
  - Distributed to computing centres for low-latency and offline analysis
  - Published to GWOSC after proprietary period
- **Computing:** nearly 10^9 CPU core hours
  - to process O3 data, both low-latency and offline
  - about 10% of one of the LHC experiments…
  - …or 10.000 years on my Apple M1-based laptop

BY-NC-SA

THE MANDATORY SLIDE WITH BOXES AND ARROWS

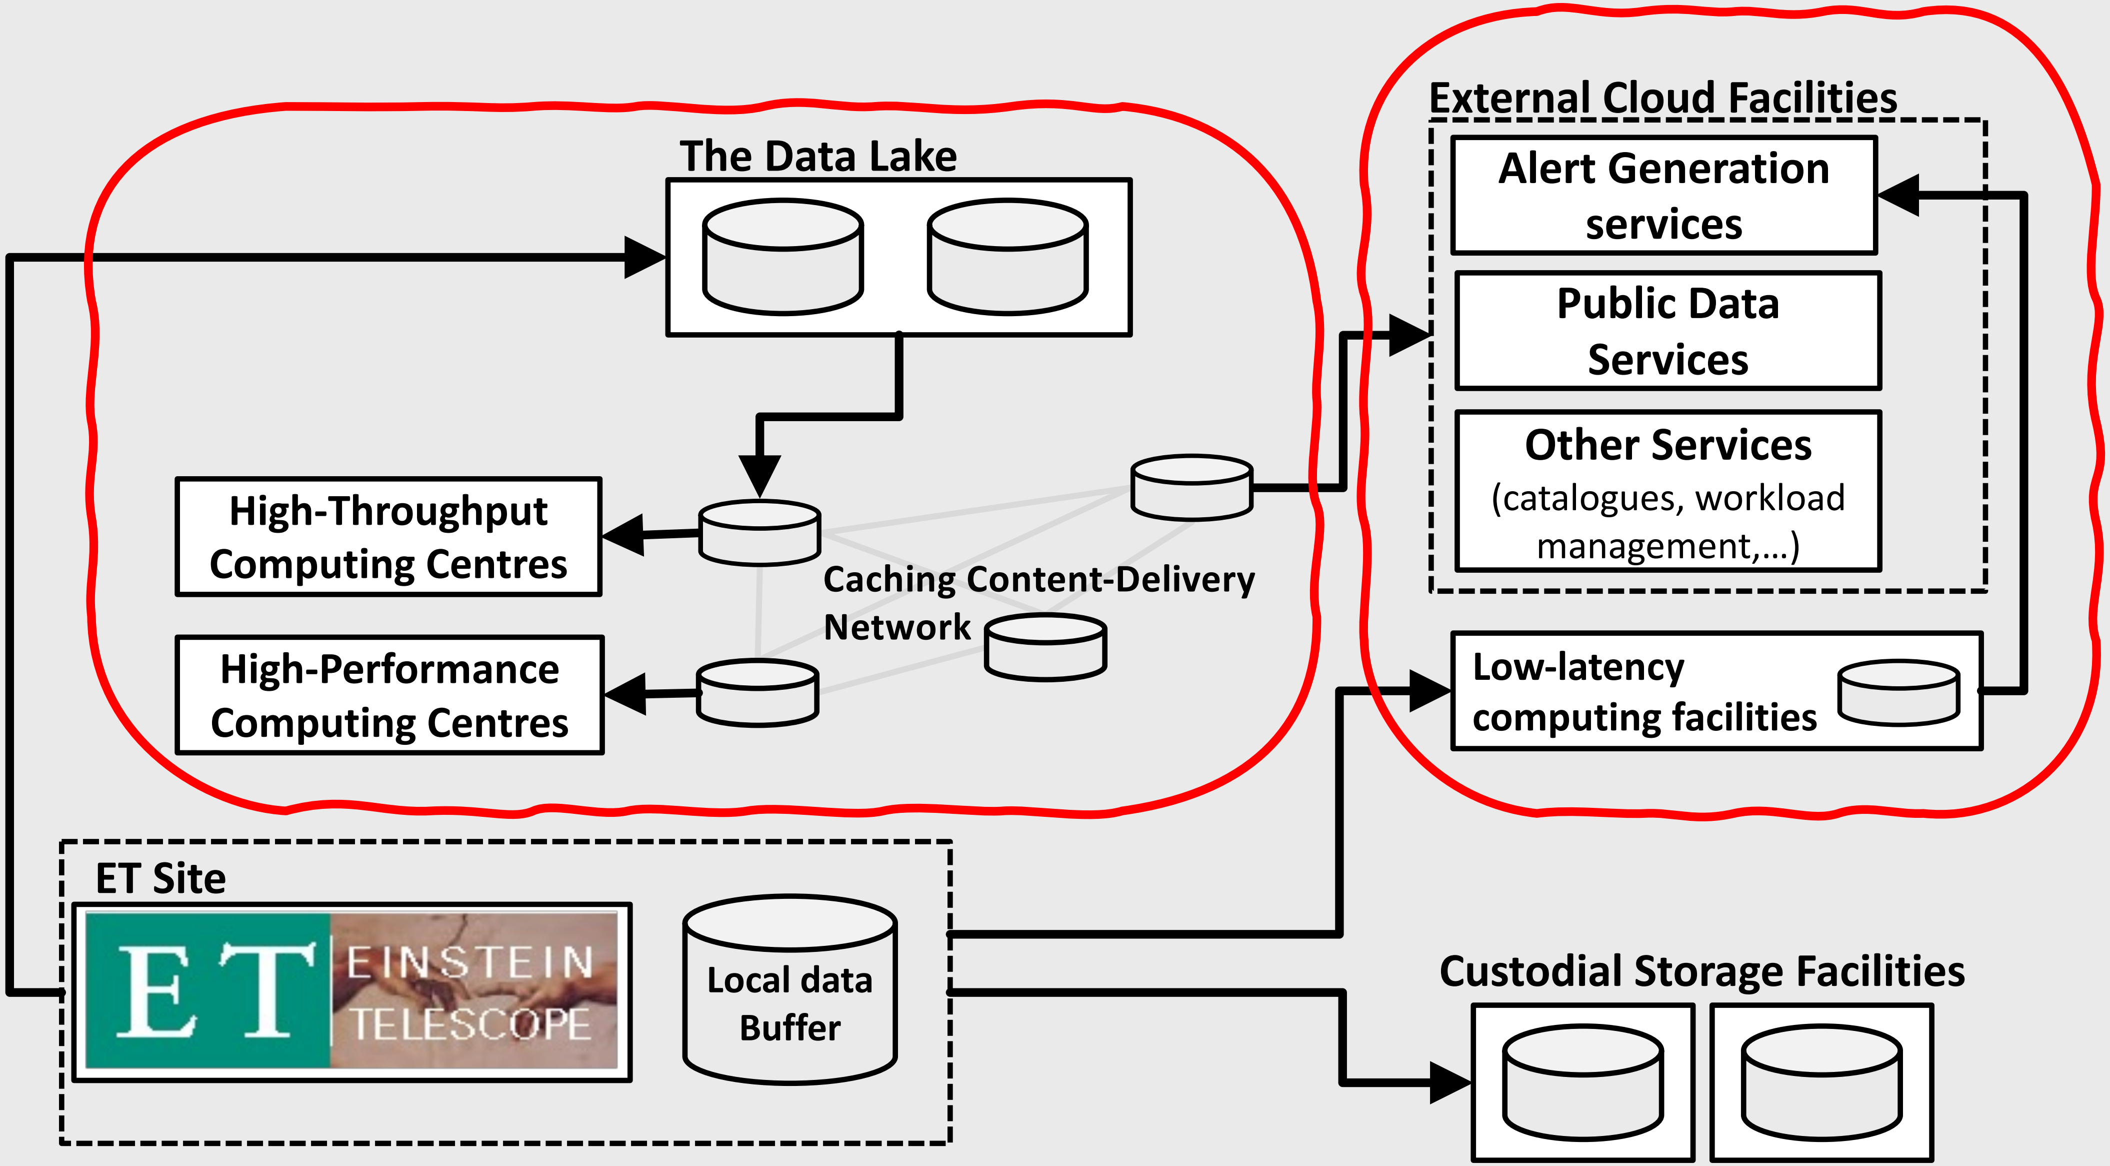

The Data Lake

High-Throughput Computing Centres

High-Performance Computing Centres

Caching Content-Delivery Network

External Cloud Facilities

Alert Generation services

Public Data Services

Other Services (catalogues, workload management,...)

Low-latency computing facilities

ET Site

Local data Buffer

Custodial Storage Facilities

Figure from the ET ESFRI proposal

BY-NC-SA

COMPUTING CHALLENGES

- Data rate (not a challenge)
- Event rate
- Event duration
- Role of AI (with and without hype)
- Reproducibility, scalability, sustainability
- Alert management
- Multimodal analysis

- …and quantum computing?

BY-NC-SA

CHALLENGE 0: DATA (HERE WE'RE LUCKY)

- Raw interferometer data don't grow much with increasing sensitivity
  - It does grow, however, with instrument complexity
  - 6 × interferometers
  - Cryogenics, huge vacuum volume,… mean many more auxiliary channels
- What grows is the amount of useful scientific information embedded in data
  - And the computing power needed to wring it out

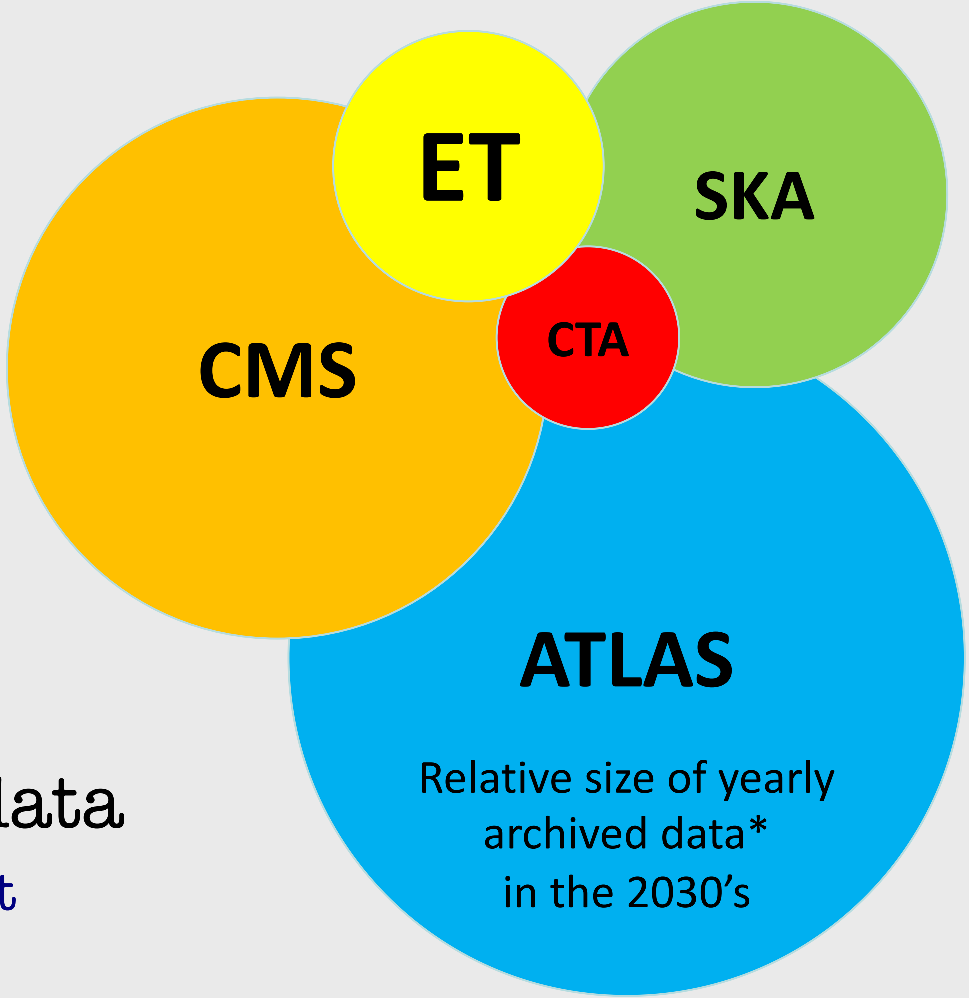

Relative size of yearly archived data*
in the 2030's

*to say nothing of weather forecast, genomics, Earth observation, oil industry, GAFAM and everybody else

BY-NC-SA

CHALLENGE 1: EVENT RATE

- The event rate scales with the third power of the range
  - 1 event per week in O3
  - 1 every couple of days on O4a
- 10^5 events/year
  - one every few minutes
- Standard techniques and strategies will not scale
  - In particular, CBC Parameter Estimation
  - Target: 1/10th of an LHC experiment in Run 5

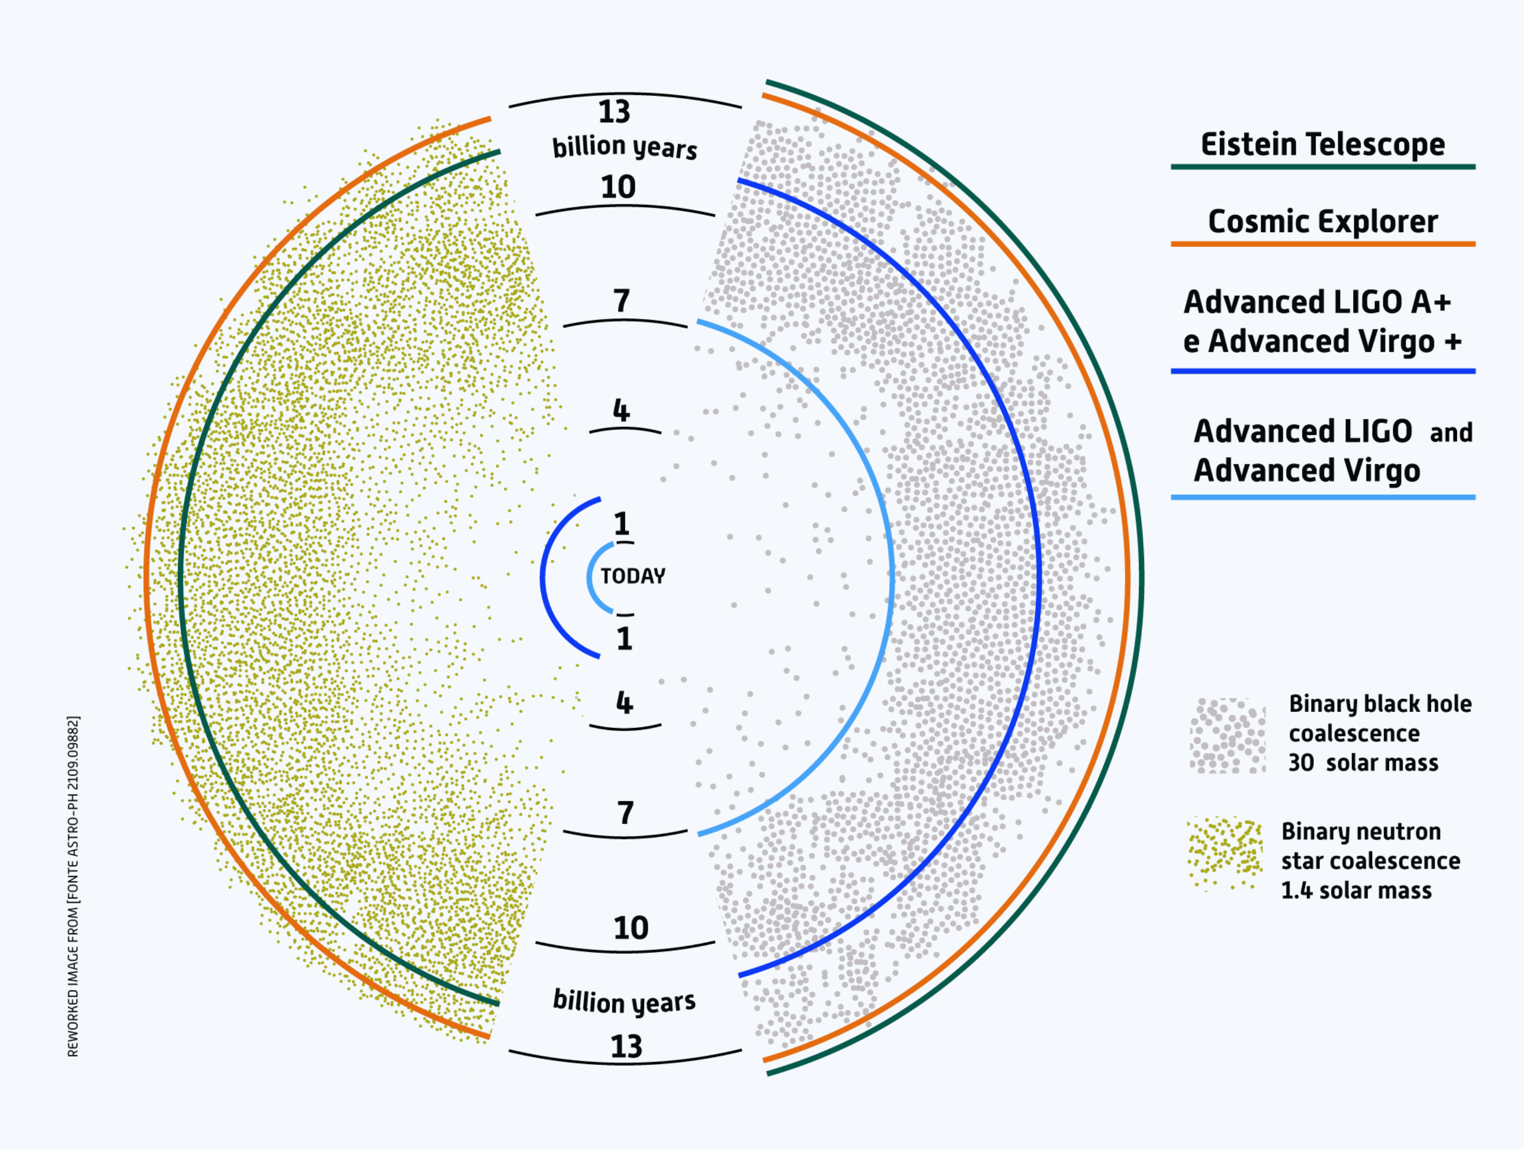

BY-NC-SA

CHALLENGE 2: EVENT DURATION

(few) seconds ⟷

BY-NC-SA

CHALLENGE 2: EVENT DURATION

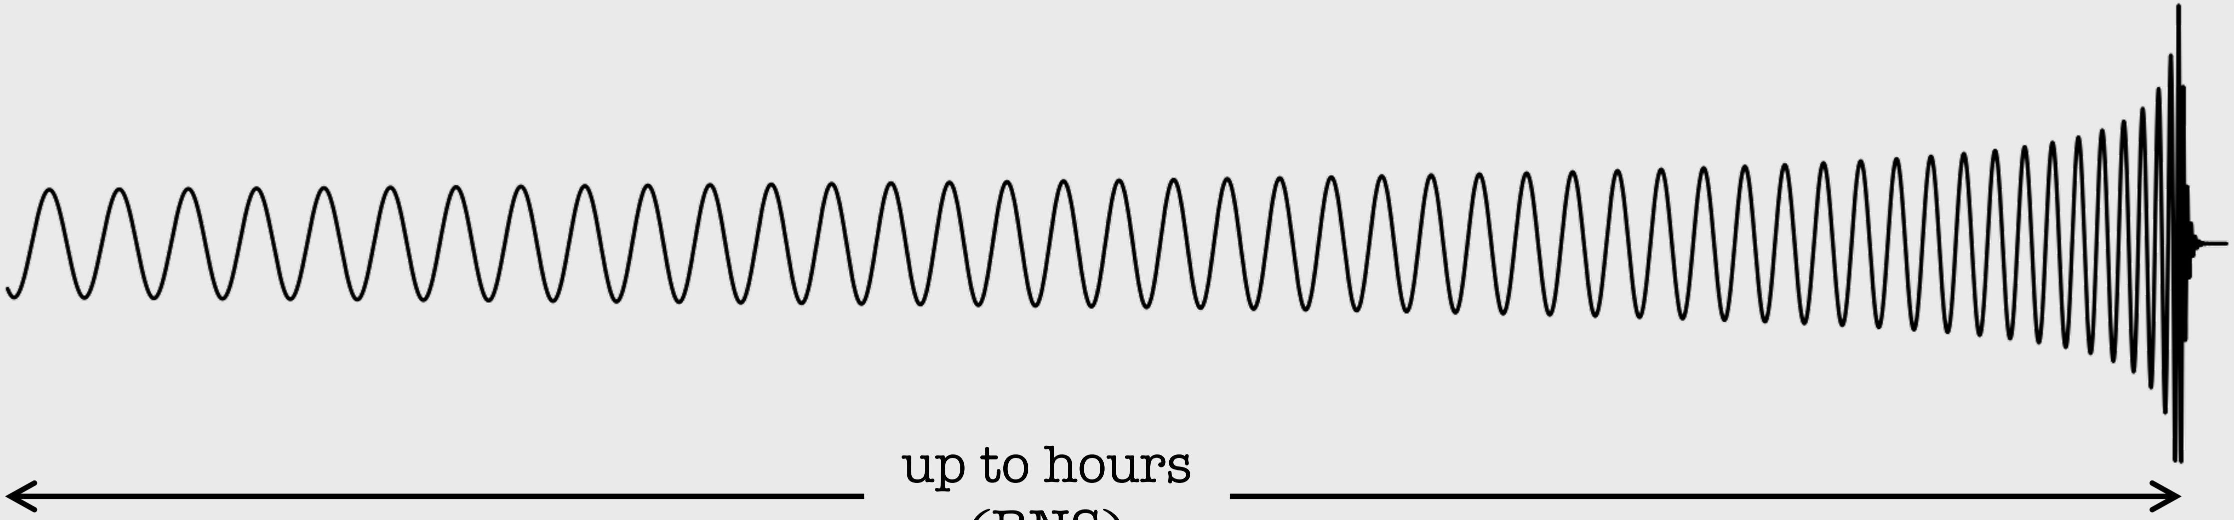

(and larger parameter space to explore)

BY-NC-SA

EARLY WARNING

THE ASTROPHYSICAL JOURNAL LETTERS, 910:L21 (7pp), 2021 April 1 https://doi.org/10.3847/2041-8213/abed54

© 2021. The American Astronomical Society. All rights reserved.

First Demonstration of Early Warning Gravitational-wave Alerts

Ryan Magee[1,2,3,27], Deep Chatterjee[4,5,6,27], Leo P. Singer[7,27], Surabhi Sachdev[2,3,27], Manoj Kovalam[8,9], Geoffrey Mo[10,11], Stuart Anderson[1], Patrick Brady[6], Patrick Brockill[6], Kipp Cannon[12], Tito Dal Canton[13], Qi Chu[8,9], Patrick Clearwater[8,14], Alex Codoreanu[8,14], Marco Drago[15,16], Patrick Godwin[2,3], Shaon Ghosh[17], Giuseppe Greco[18,19], Chad Hanna[2,3,20], Shasvath J. Kapadia[21], Erik Katsavounidis[10,11], Victor Oloworaran[8,9], Alexander E. Pace[2,3], Fiona Panther[8,9], Anwarul Patwary[8,9], Roberto De Pietri[22,23], Brandon Piotrzkowski[6], Tanner Prestegard[6], Luca Rei[24], Anala K. Sreekumar[8,9], Marek J. Szczepańczyk[25], Vinaya Valsan[6], Aaron Viets[6], Madeline Wade[26], Linqing Wen[8,9], and John Zweizig[1]
[1] LIGO, California Institute of Technology, Pasadena, CA 91125, USA
[2] Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA
[3] Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA
[4] Center for Astrophysical Surveys, National Center for Supercomputing Applications, Urbana, IL 61801, USA
[5] Illinois Center for Advanced Studies of the Universe, Department of Physics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA
[6] Leonard E. Parker Center for Gravitation, Cosmology, and Astrophysics, University of Wisconsin-Milwaukee, Milwaukee, WI 53201, USA
[7] Astroparticle Physics Laboratory, NASA Goddard Space Flight Center, Mail Code 661, Greenbelt, MD 20771, USA
[8] Australian Research Council Centre of Excellence for Gravitational Wave Discovery (OzGrav), Australia
[9] Department of Physics, University of Western Australia, Crawley WA 6009, Australia
[10] LIGO Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[11] Department of Physics and Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[12] Research Center for the Early Universe, The University of Tokyo, 113-0033, Japan
[13] Université Paris-Saclay, CNRS/IN2P3, IJCLab, F-91405 Orsay, France
[14] Gravitational Wave Data Centre, Swinburne University, Hawthorn VIC 3122, Australia
[15] Università di Roma “La Sapienza,” I-00185 Roma, Italy
[16] INFN, Sezione di Roma, I-00185 Roma, Italy
[17] Department of Physics and Astronomy, Montclair State University, Montclair, NJ 07043, USA
[18] Università degli Studi di Urbino “Carlo Bo”, I-61029 Urbino, Italy
[19] INFN, Sezione di Firenze, I-50019 Sesto Fiorentino, Firenze, Italy
[20] Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA
[21] International Centre for Theoretical Sciences, Tata Institute of Fundamental Research, Bangalore 560089, India
[22] Dipartimento di Scienze Matematiche, Fisiche e Informatiche, Università di Parma, I-43124 Parma, Italy
[23] 7INFN, Sezione di Milano Bicocca, Gruppo Collegato di Parma, I-43124 Parma, Italy
[24] INFN, Sezione di Genova, I-16146 Genova, Italy
[25] University of Florida, Gainesville, FL 32611, USA
[26] Department of Physics, Kenyon College, Gambier, OH 43022, USA

Received 2021 February 25; accepted 2021 March 10; published 2021 April 6

Abstract

Gravitational-wave observations became commonplace in Advanced LIGO-Virgo’s recently concluded third observing run. 56 nonretracted candidates were identified and publicly announced in near real time. Gravitational waves from binary neutron star mergers, however, remain of special interest since they can be precursors to high-energy astrophysical phenomena like γ-ray bursts and kilonovae. While late-time electromagnetic emissions provide important information about the astrophysical processes within, the prompt emission along with gravitational waves uniquely reveals the extreme matter and gravity during—and in the seconds following—merger. Rapid communication of source location and properties from the gravitational-wave data is crucial to facilitate multimessenger follow-up of such sources. This is especially enabled if the partner facilities are forewarned via an early warning (pre-merger) alert. Here we describe the commissioning and performance of such a low-latency infrastructure within LIGO-Virgo. We present results from an end-to-end mock data challenge that detects binary neutron star mergers and alerts partner facilities before merger. We set expectations for these alerts in future observing runs.

Unified Astronomy Thesaurus concepts: Gravitational waves (678); Gravitational wave astronomy (675); Neutron stars (1108); High energy astrophysics (739)

1. Introduction

The field of gravitational-wave astronomy has exploded in the years following the first direct observation of gravitational waves (GWs) from a binary black hole (BBH) merger (Abbott et al. 2016). Since then, LIGO-Virgo have published 49 candidate events, many of which were identified in low-latency;[28] these include two binary neutron star (BNS) and two neutron star–black hole (NSBH) candidates (Abbott et al. 2020a). The detection of GWs from compact binaries, especially from BBHs, has become routine. GWs from BNS and NSBH mergers, however, remain rare. BNS and NSBH mergers are of special interest due to the possibility of counterpart electromagnetic (EM) signals. For BNS mergers, in particular, it has long been hypothesized that the central engine (post merger) can launch short gamma-ray bursts (SGRBs; Lattimer & Schramm 1976; Lee & Ramirez-Ruiz 2007), kilonovae (Li & Paczynski 1998; Metzger et al. 2010), and radio waves and X-rays post merger (Nakar & Piran 2011; Metzger & Berger 2012). In the special case of the

---

[27] These authors contributed equally to this work.
[28] Some of the 56 have not yet appeared in a LIGO-Virgo publication.

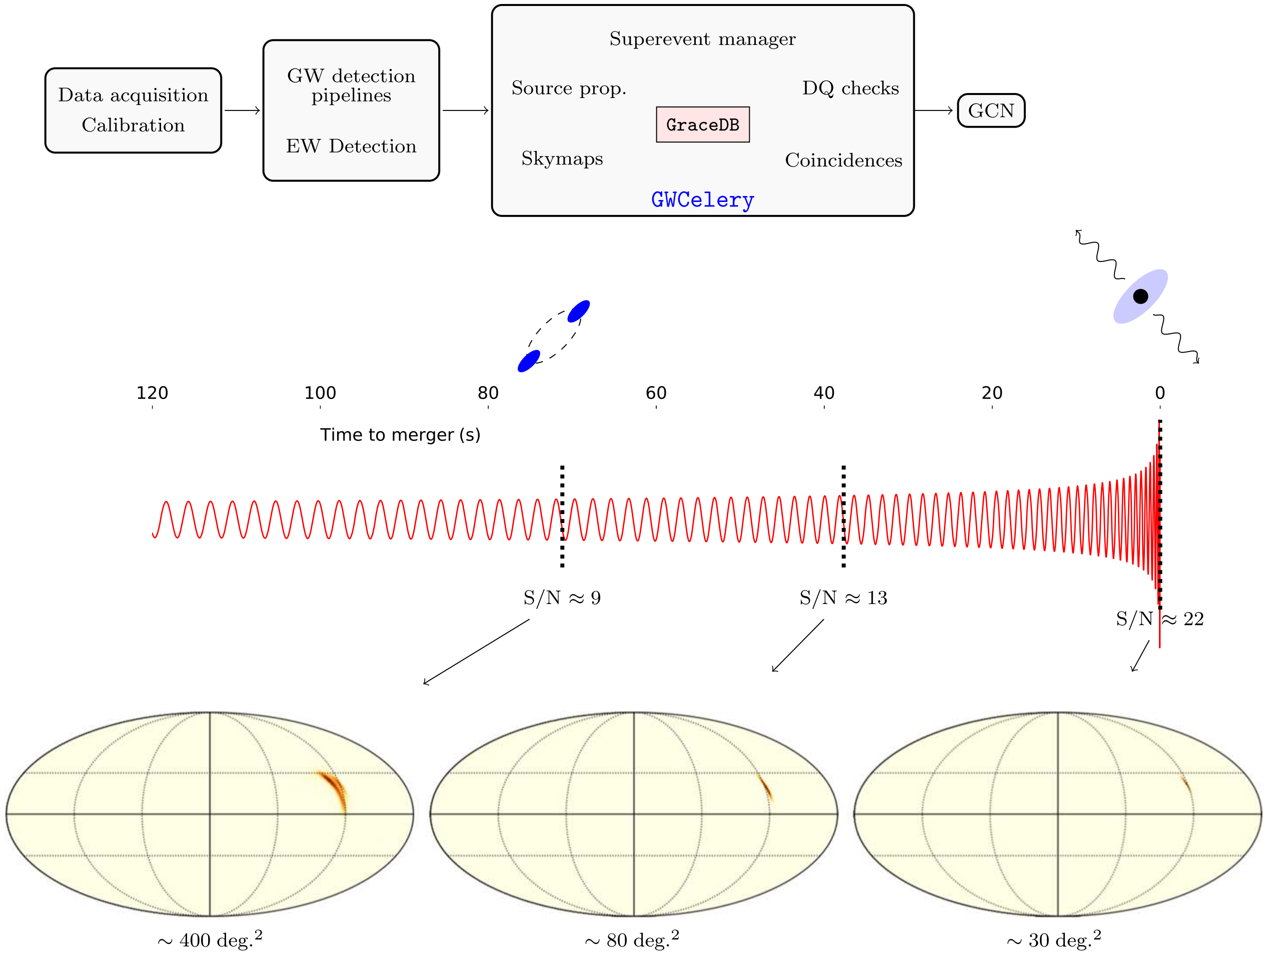

Figure 1. The upper half of the figure illustrates the complete pipeline and interaction of the various (sub)systems, mentioned in Section 2, responsible for disseminating early warning alerts. The waveform evolution with time is shown in the bottom half along with the dependence of the sky-localization area on the cutoff time of the early warning templates and the accumulated S/N during the binary inspiral. The waveforms, time to merger, S/N, and localizations in this figure are qualitative.

BY-NC-SA

CHALLENGE 3: USE AI (WITHOUT THE HYPE)

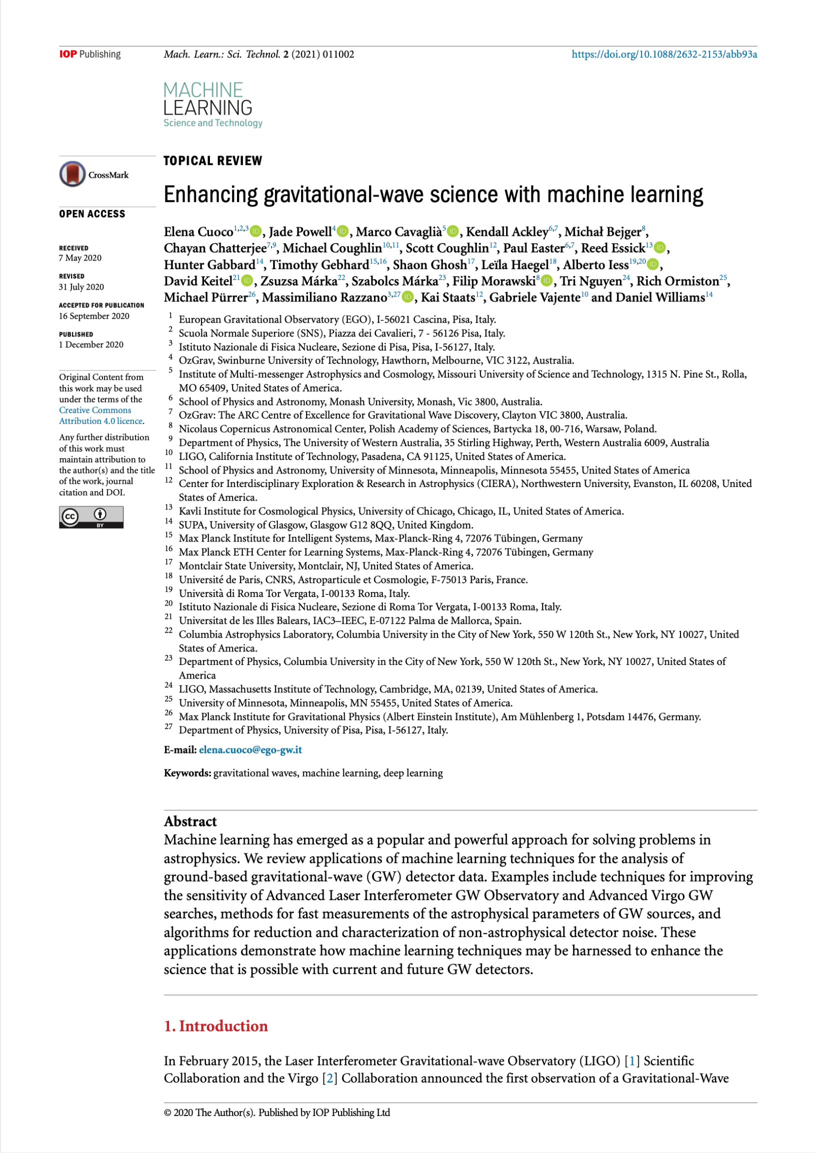

IOP Publishing Mach. Learn.: Sci. Technol. 2 (2021) 011002 https://doi.org/10.1088/2632-2153/abb93a

MACHINE
LEARNING
Science and Technology

CrossMark

OPEN ACCESS

RECEIVED
7 May 2020

REVISED
31 July 2020

ACCEPTED FOR PUBLICATION
16 September 2020

PUBLISHED
1 December 2020

Original Content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence.

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

TOPICAL REVIEW

Enhancing gravitational-wave science with machine learning

Elena Cuoco[1,2,3], Jade Powell[4], Marco Cavaglià[5], Kendall Ackley[6,7], Michał Bejger[8], Chayan Chatterjee[7,9], Michael Coughlin[10,11], Scott Coughlin[12], Paul Easter[6,7], Reed Essick[13], Hunter Gabbard[14], Timothy Gebhard[15,16], Shaon Ghosh[17], Leïla Haegel[18], Alberto Iess[19,20], David Keitel[21], Zsuzsa Márka[22], Szabolcs Márka[23], Filip Morawski[8], Tri Nguyen[24], Rich Ormiston[25], Michael Pürrer[26], Massimiliano Razzano[3,27], Kai Staats[12], Gabriele Vajente[10] and Daniel Williams[14]

[1] European Gravitational Observatory (EGO), I-56021 Cascina, Pisa, Italy.
[2] Scuola Normale Superiore (SNS), Piazza dei Cavalieri, 7 - 56126 Pisa, Italy.
[3] Istituto Nazionale di Fisica Nucleare, Sezione di Pisa, Pisa, I-56127, Italy.
[4] OzGrav, Swinburne University of Technology, Hawthorn, Melbourne, VIC 3122, Australia.
[5] Institute of Multi-messenger Astrophysics and Cosmology, Missouri University of Science and Technology, 1315 N. Pine St., Rolla, MO 65409, United States of America.
[6] School of Physics and Astronomy, Monash University, Monash, Vic 3800, Australia.
[7] OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery, Clayton VIC 3800, Australia.
[8] Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences, Bartycka 18, 00-716, Warsaw, Poland.
[9] Department of Physics, The University of Western Australia, 35 Stirling Highway, Perth, Western Australia 6009, Australia
[10] LIGO, California Institute of Technology, Pasadena, CA 91125, United States of America.
[11] School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, United States of America
[12] Center for Interdisciplinary Exploration & Research in Astrophysics (CIERA), Northwestern University, Evanston, IL 60208, United States of America.
[13] Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL, United States of America.
[14] SUPA, University of Glasgow, Glasgow G12 8QQ, United Kingdom.
[15] Max Planck Institute for Intelligent Systems, Max-Planck-Ring 4, 72076 Tübingen, Germany
[16] Max Planck ETH Center for Learning Systems, Max-Planck-Ring 4, 72076 Tübingen, Germany
[17] Montclair State University, Montclair, NJ, United States of America.
[18] Université de Paris, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France.
[19] Università di Roma Tor Vergata, I-00133 Roma, Italy.
[20] Istituto Nazionale di Fisica Nucleare, Sezione di Roma Tor Vergata, I-00133 Roma, Italy.
[21] Universitat de les Illes Balears, IAC3–IEEC, E-07122 Palma de Mallorca, Spain.
[22] Columbia Astrophysics Laboratory, Columbia University in the City of New York, 550 W 120th St., New York, NY 10027, United States of America.
[23] Department of Physics, Columbia University in the City of New York, 550 W 120th St., New York, NY 10027, United States of America
[24] LIGO, Massachusetts Institute of Technology, Cambridge, MA, 02139, United States of America.
[25] University of Minnesota, Minneapolis, MN 55455, United States of America.
[26] Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, Potsdam 14476, Germany.
[27] Department of Physics, University of Pisa, Pisa, I-56127, Italy.

E-mail: elena.cuoco@ego-gw.it

Keywords: gravitational waves, machine learning, deep learning

Abstract
Machine learning has emerged as a popular and powerful approach for solving problems in astrophysics. We review applications of machine learning techniques for the analysis of ground-based gravitational-wave (GW) detector data. Examples include techniques for improving the sensitivity of Advanced Laser Interferometer GW Observatory and Advanced Virgo GW searches, methods for fast measurements of the astrophysical parameters of GW sources, and algorithms for reduction and characterization of non-astrophysical detector noise. These applications demonstrate how machine learning techniques may be harnessed to enhance the science that is possible with current and future GW detectors.

1. Introduction

In February 2015, the Laser Interferometer Gravitational-wave Observatory (LIGO) [1] Scientific Collaboration and the Virgo [2] Collaboration announced the first observation of a Gravitational-Wave

© 2020 The Author(s). Published by IOP Publishing Ltd

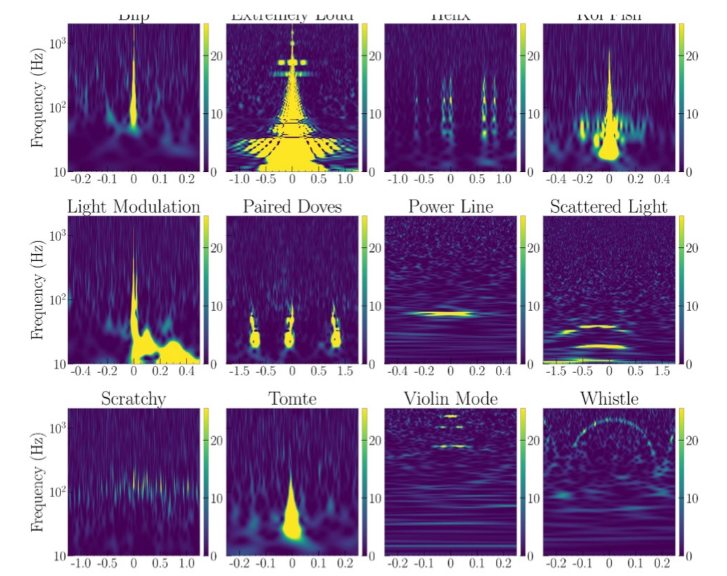

Figure 1. Time-frequency representations of different types of glitches occurring in GW data. (Examples from the GravitySpy project [45]). ML algorithms can help identify the origin of these glitches and increase the sensitivity of GW transient searches.

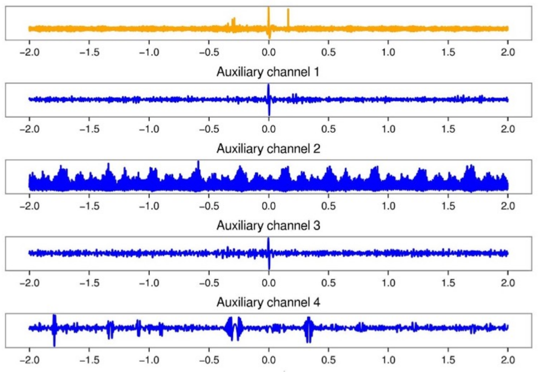

Figure 2. A qualitative illustration of how auxiliary channels can help determine the non-astrophysical nature of detector triggers. The top time series is $h(t)$. The other time series are from detector auxiliary monitoring different sources which are not sensitive to GWs. The spike in the strain time series at $t = 0$ occurs also in auxiliary channels 1 and 3. This may indicate that the trigger is not of astrophysical origin.

BY-NC-SA

CHALLENGE 3: USE AI (WITHOUT THE HYPE)

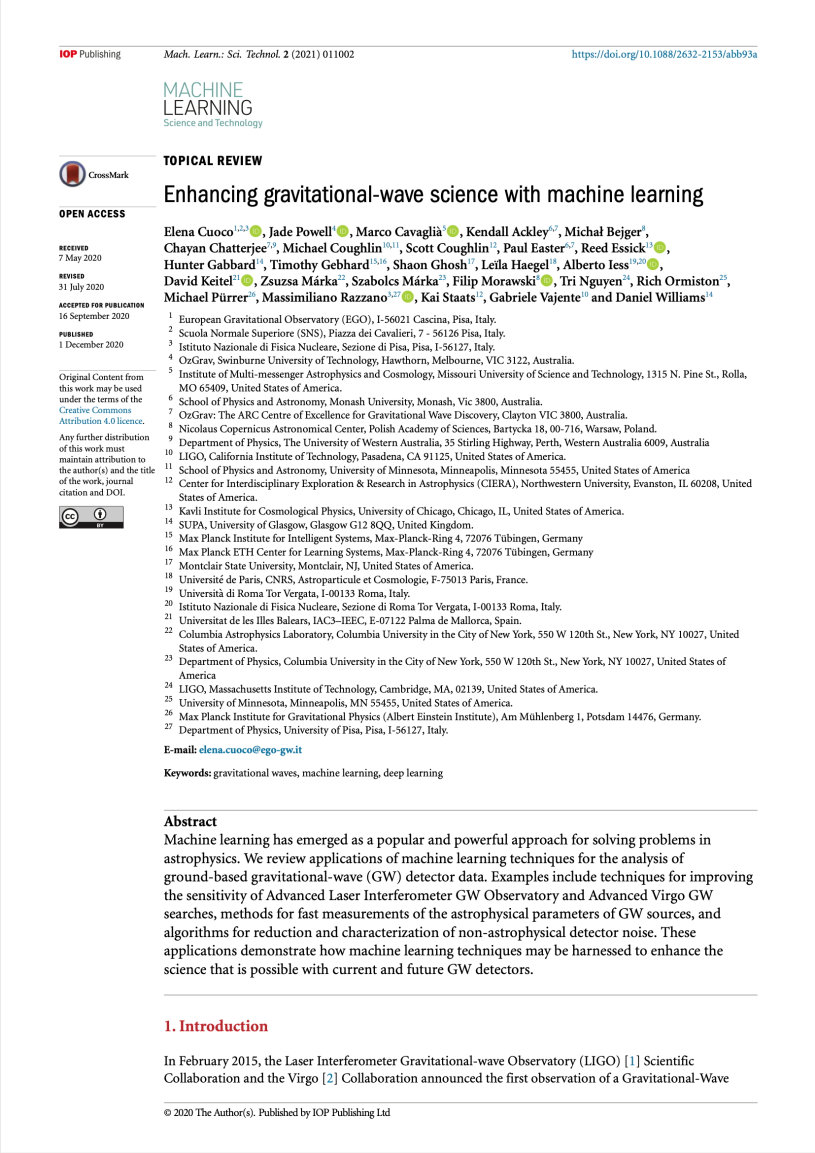

IOP Publishing Mach. Learn.: Sci. Technol. 2 (2021) 011002 https://doi.org/10.1088/2632-2153/abb93a

MACHINE
LEARNING
Science and Technology

TOPICAL REVIEW

CrossMark

OPEN ACCESS

RECEIVED
7 May 2020

REVISED
31 July 2020

ACCEPTED FOR PUBLICATION
16 September 2020

PUBLISHED
1 December 2020

Original Content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence.

Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Enhancing gravitational-wave science with machine learning

Elena Cuoco[1,2,3], Jade Powell[4], Marco Cavaglià[5], Kendall Ackley[6,7], Michał Bejger[8], Chayan Chatterjee[7,9], Michael Coughlin[10,11], Scott Coughlin[12], Paul Easter[6,7], Reed Essick[13], Hunter Gabbard[14], Timothy Gebhard[15,16], Shaon Ghosh[17], Leïla Haegel[18], Alberto Iess[19,20], David Keitel[21], Zsuzsa Márka[22], Szabolcs Márka[23], Filip Morawski[8], Tri Nguyen[24], Rich Ormiston[25], Michael Pürrer[26], Massimiliano Razzano[3,27], Kai Staats[12], Gabriele Vajente[10] and Daniel Williams[14]

[1] European Gravitational Observatory (EGO), I-56021 Cascina, Pisa, Italy.
[2] Scuola Normale Superiore (SNS), Piazza dei Cavalieri, 7 - 56126 Pisa, Italy.
[3] Istituto Nazionale di Fisica Nucleare, Sezione di Pisa, Pisa, I-56127, Italy.
[4] OzGrav, Swinburne University of Technology, Hawthorn, Melbourne, VIC 3122, Australia.
[5] Institute of Multi-messenger Astrophysics and Cosmology, Missouri University of Science and Technology, 1315 N. Pine St., Rolla, MO 65409, United States of America.
[6] School of Physics and Astronomy, Monash University, Monash, Vic 3800, Australia.
[7] OzGrav: The ARC Centre of Excellence for Gravitational Wave Discovery, Clayton VIC 3800, Australia.
[8] Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences, Bartycka 18, 00-716, Warsaw, Poland.
[9] Department of Physics, The University of Western Australia, 35 Stirling Highway, Perth, Western Australia 6009, Australia
[10] LIGO, California Institute of Technology, Pasadena, CA 91125, United States of America.
[11] School of Physics and Astronomy, University of Minnesota, Minneapolis, Minnesota 55455, United States of America
[12] Center for Interdisciplinary Exploration & Research in Astrophysics (CIERA), Northwestern University, Evanston, IL 60208, United States of America.
[13] Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL, United States of America.
[14] SUPA, University of Glasgow, Glasgow G12 8QQ, United Kingdom.
[15] Max Planck Institute for Intelligent Systems, Max-Planck-Ring 4, 72076 Tübingen, Germany
[16] Max Planck ETH Center for Learning Systems, Max-Planck-Ring 4, 72076 Tübingen, Germany
[17] Montclair State University, Montclair, NJ, United States of America.
[18] Université de Paris, CNRS, Astroparticule et Cosmologie, F-75013 Paris, France.
[19] Università di Roma Tor Vergata, I-00133 Roma, Italy.
[20] Istituto Nazionale di Fisica Nucleare, Sezione di Roma Tor Vergata, I-00133 Roma, Italy.
[21] Universitat de les Illes Balears, IAC3–IEEC, E-07122 Palma de Mallorca, Spain.
[22] Columbia Astrophysics Laboratory, Columbia University in the City of New York, 550 W 120th St., New York, NY 10027, United States of America.
[23] Department of Physics, Columbia University in the City of New York, 550 W 120th St., New York, NY 10027, United States of America
[24] LIGO, Massachusetts Institute of Technology, Cambridge, MA, 02139, United States of America.
[25] University of Minnesota, Minneapolis, MN 55455, United States of America.
[26] Max Planck Institute for Gravitational Physics (Albert Einstein Institute), Am Mühlenberg 1, Potsdam 14476, Germany.
[27] Department of Physics, University of Pisa, Pisa, I-56127, Italy.

E-mail: elena.cuoco@ego-gw.it

Keywords: gravitational waves, machine learning, deep learning

Abstract

Machine learning has emerged as a popular and powerful approach for solving problems in astrophysics. We review applications of machine learning techniques for the analysis of ground-based gravitational-wave (GW) detector data. Examples include techniques for improving the sensitivity of Advanced Laser Interferometer GW Observatory and Advanced Virgo GW searches, methods for fast measurements of the astrophysical parameters of GW sources, and algorithms for reduction and characterization of non-astrophysical detector noise. These applications demonstrate how machine learning techniques may be harnessed to enhance the science that is possible with current and future GW detectors.

1. Introduction

In February 2015, the Laser Interferometer Gravitational-wave Observatory (LIGO) [1] Scientific Collaboration and the Virgo [2] Collaboration announced the first observation of a Gravitational-Wave

© 2020 The Author(s). Published by IOP Publishing Ltd

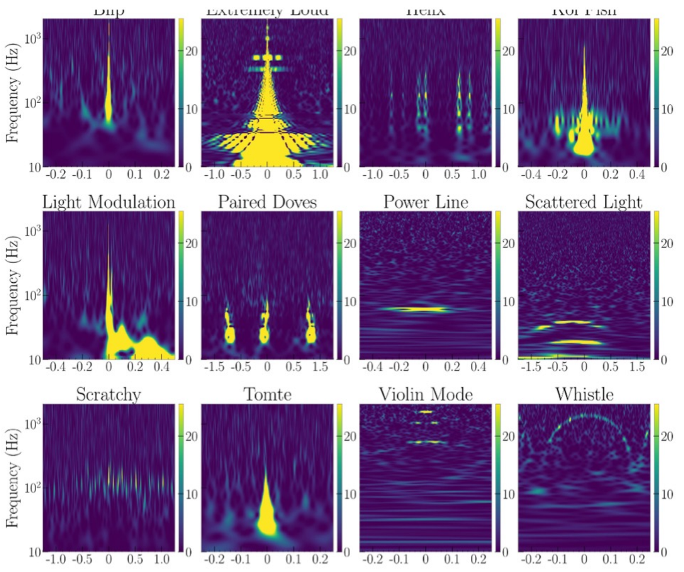

Figure 1. Time-frequency representations of different types of glitches occurring in GW data. (Examples from the GravitySpy project [45]). ML algorithms can help identify the origin of these glitches and increase the sensitivity of GW transient searches.

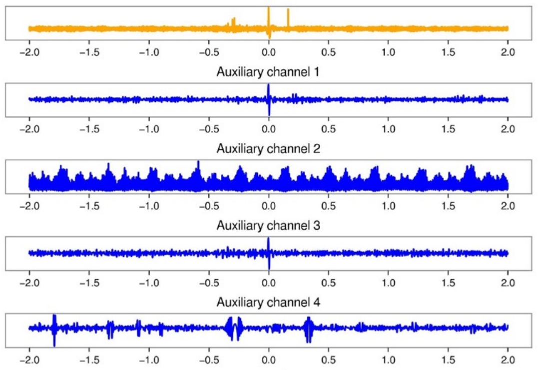

Figure 2. A qualitative illustration of how auxiliary channels can help determine the non-astrophysical nature of detector triggers. The top time series is $h(t)$. The other time series are from detector auxiliary monitoring different sources which are not sensitive to GWs. The spike in the strain time series at $t = 0$ occurs also in auxiliary channels 1 and 3. This may indicate that the trigger is not of astrophysical origin.

interTwin

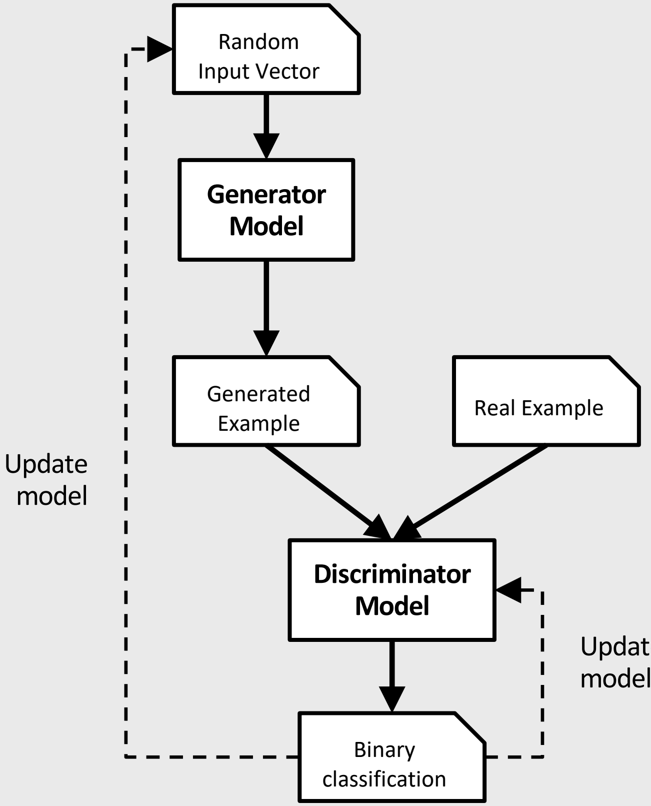

Random Input Vector

Generator Model

Generated Example

Real Example

Discriminator Model

Binary classification

Update model

Update model

BY-NC-SA

CHALLENGE 4: USE AI (WITH SOME HYPE)

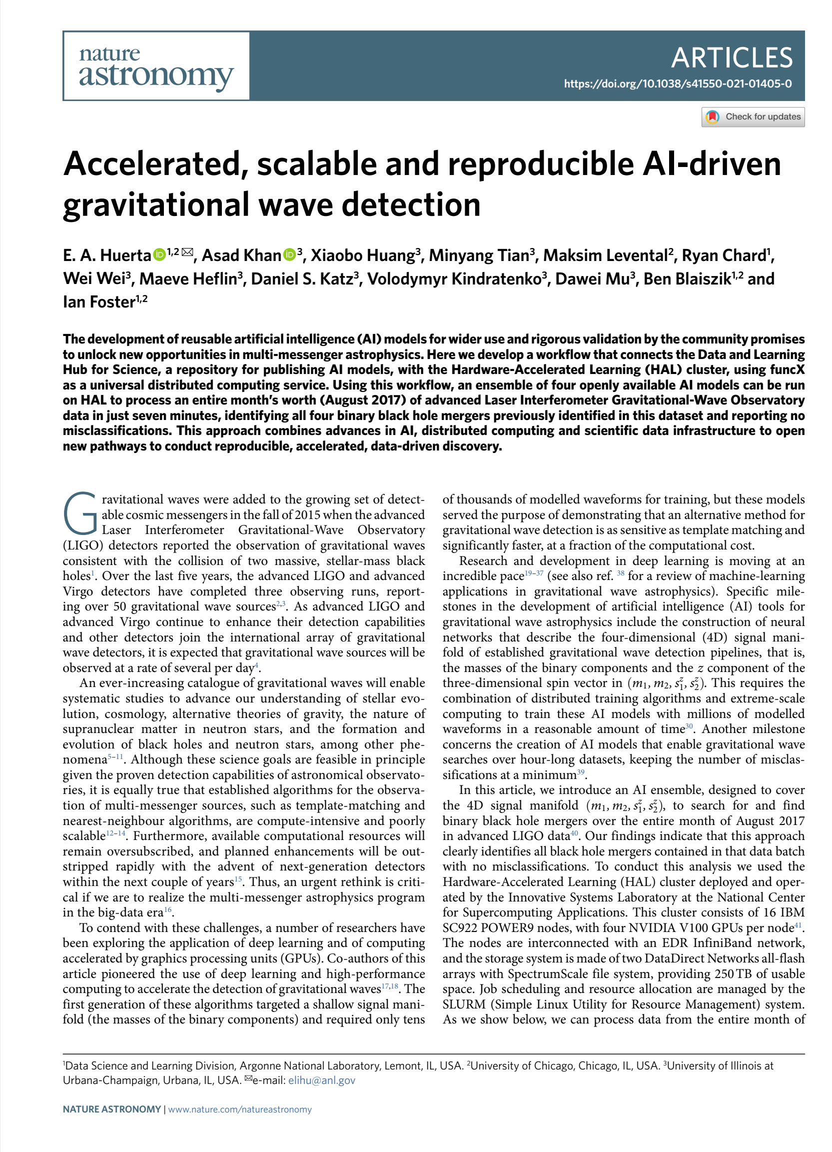

nature astronomy
ARTICLES
https://doi.org/10.1038/s41550-021-01405-0
Check for updates

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

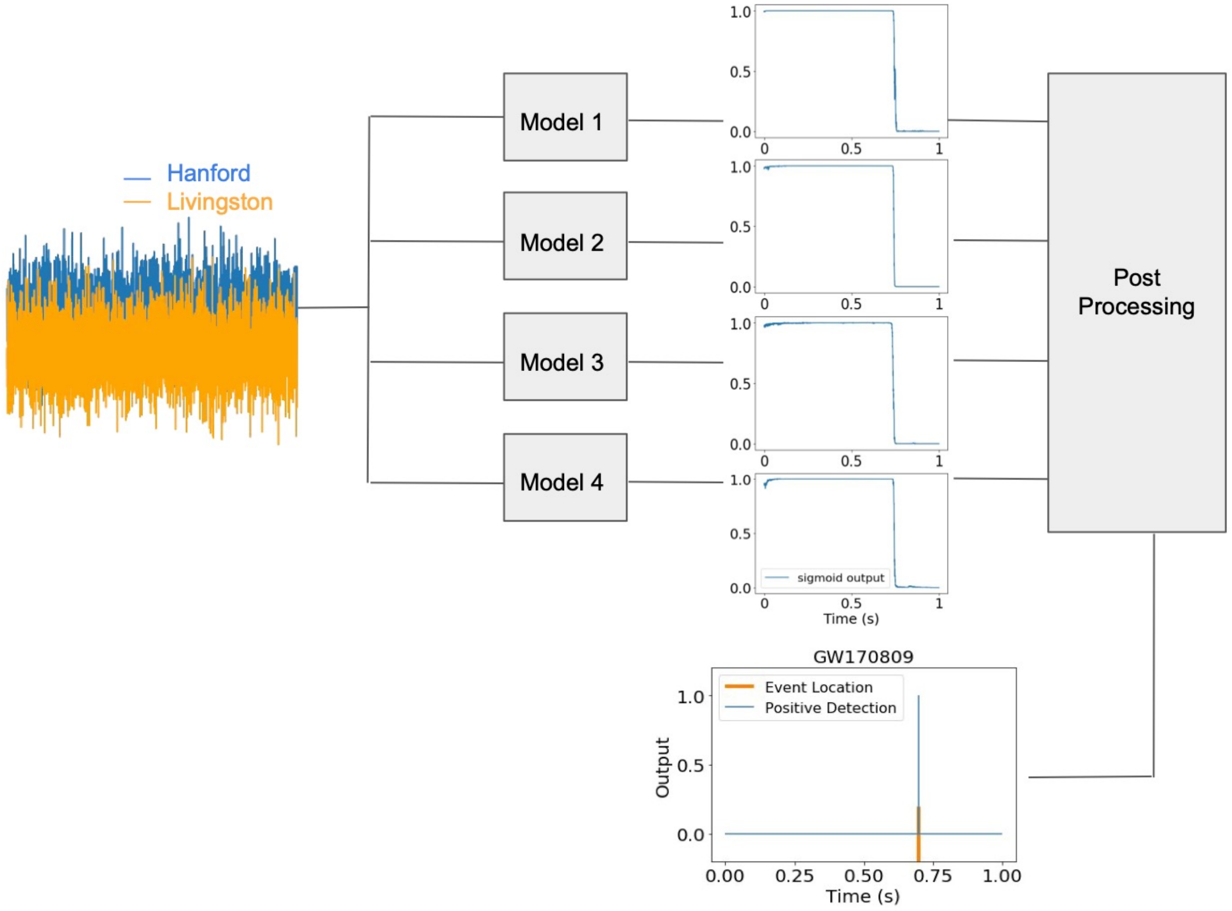

Figure 1. Gravitational wave detection workflow with AI ensemble. Hanford and Livingston gravitational wave data, depicted as blue and orange time-series data on the left, are fed into an AI ensemble of four neural network models. The response of the neural networks to advanced LIGO data is shown to the right of the boxes representing the models. At the post-processing stage, the outputs of the four neural networks are combined. If the outputs of all the models are consistent with the existence of a gravitational wave signal, then the post-processing algorithm indicates a positive detection. The bottom panel showcases a positive detection for the binary black hole merger GW170809.

BY-NC-SA

CHALLENGE 5: REPRODUCIBILITY, SCALABILITY, SUSTAINABILITY

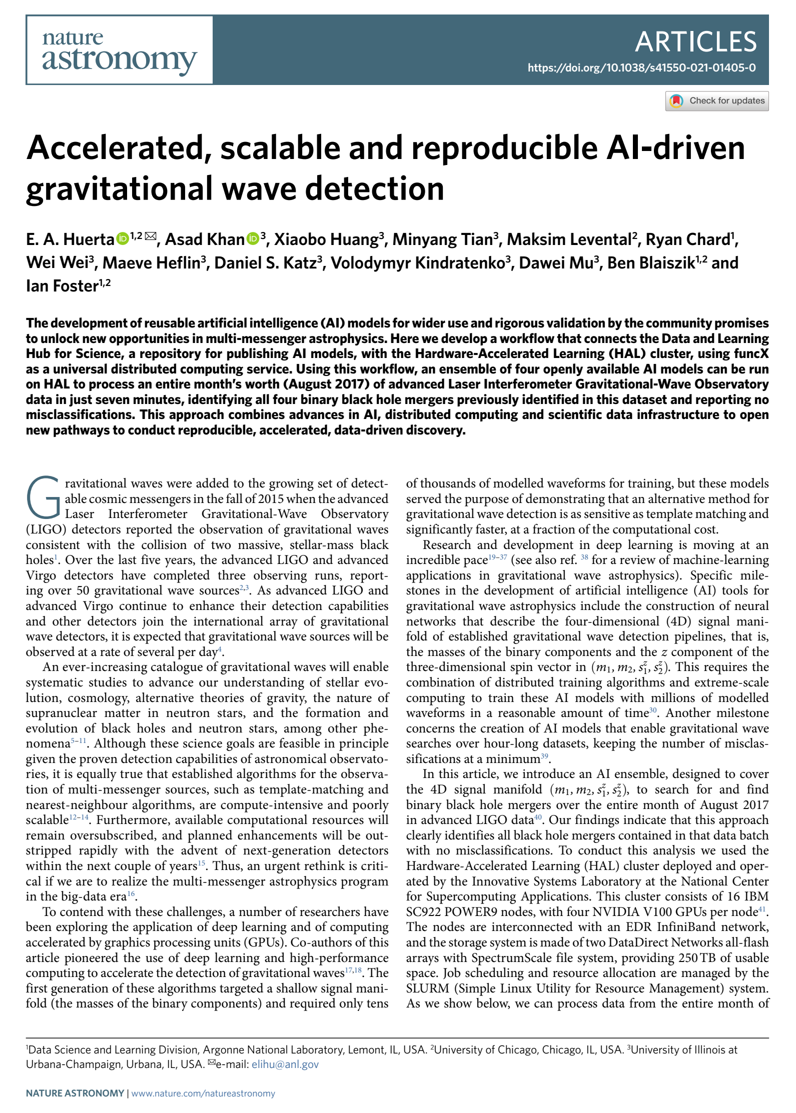

nature astronomy

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

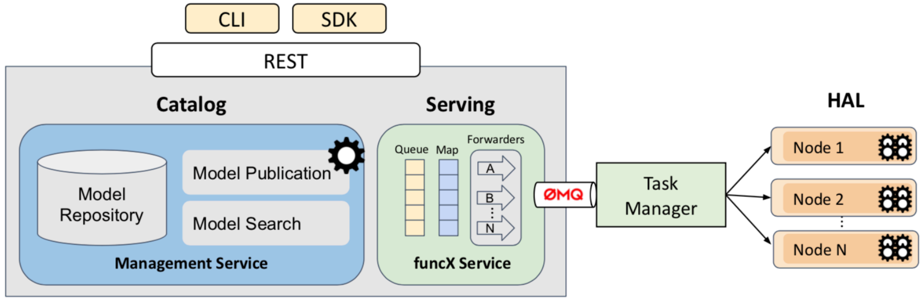

Figure 5. **`DLHub` architecture.** Schematic representation of the cyberinfrastructure resources used to conduct accelerated and reproducible gravitational wave detection on open-source advanced LIGO data. This architecture provides a command line interface (CLI), a Python software development kit (SDK) and a representational state transfer (REST) application programming interface to publish, manage and invoke AI models. The management service coordinates the execution of tasks on remote resources using a ZeroMQ (ØMQ) queue, which sends tasks to registered task managers for execution. This messaging model ensures that tasks are received and executed. `DLHub` supports both synchronous and asynchronous task execution.

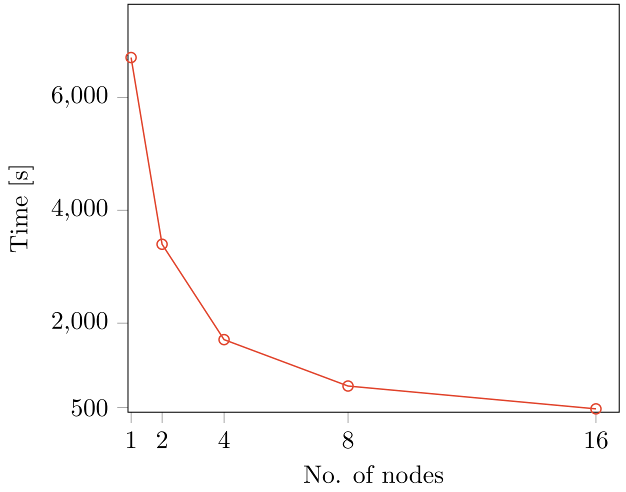

Figure 6. **Throughput of the `DLHub+HAL` architecture.** An AI ensemble of four neural networks, hosted at the `DLHub`, processes advanced LIGO data from an entire month (August 2017) in 7 min using the entire `HAL` cluster that has 64 `NVIDIA` V100 GPUs evenly distributed over 16 nodes.

BY-NC-SA

MEANWHILE, IN THE US...

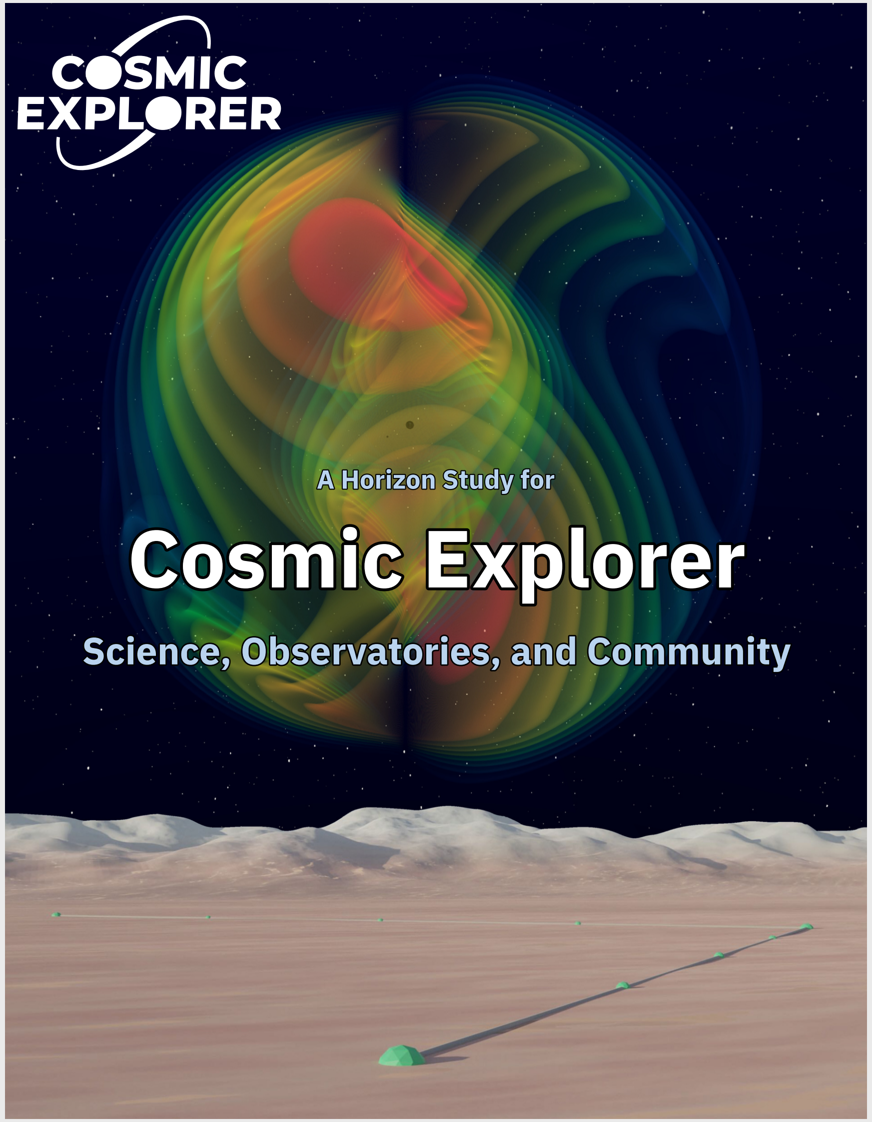

COSMIC EXPLORER

A Horizon Study for

Cosmic Explorer

Science, Observatories, and Community

- Cosmic Explorer: proposed 3G facility in the US
- More conventional (and less risky) surface L-shaped design
- 40km + 20km arms
- Better sensitivity in mid-frequency range (the “bucket”)
- Higher (worse) low-frequency cutoff
- We’ll have a 3G ground-based network!

BY-NC-SA

...AND FARTHER AWAY: LISA

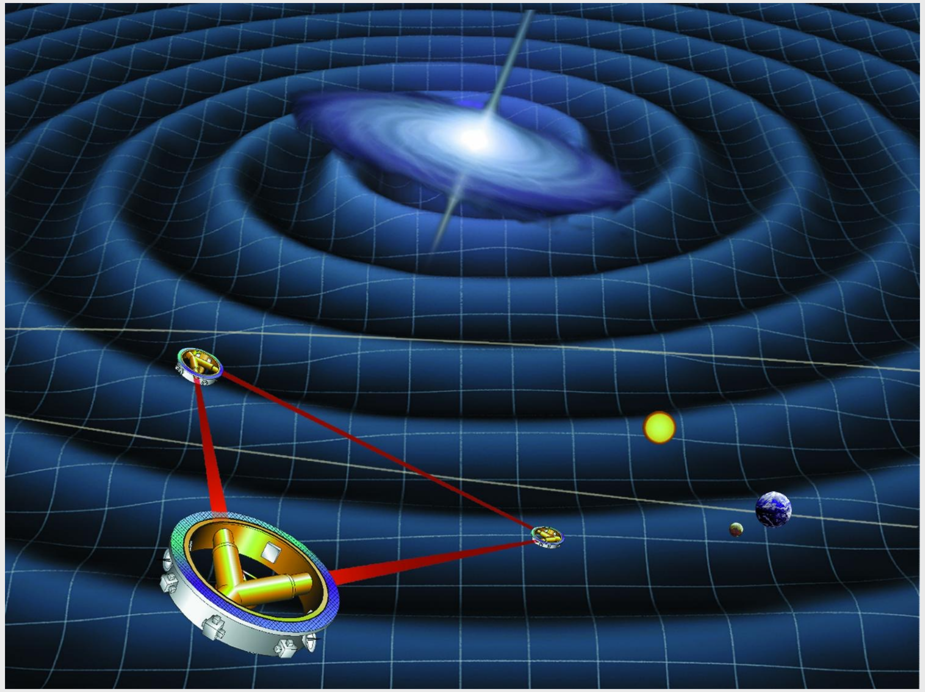

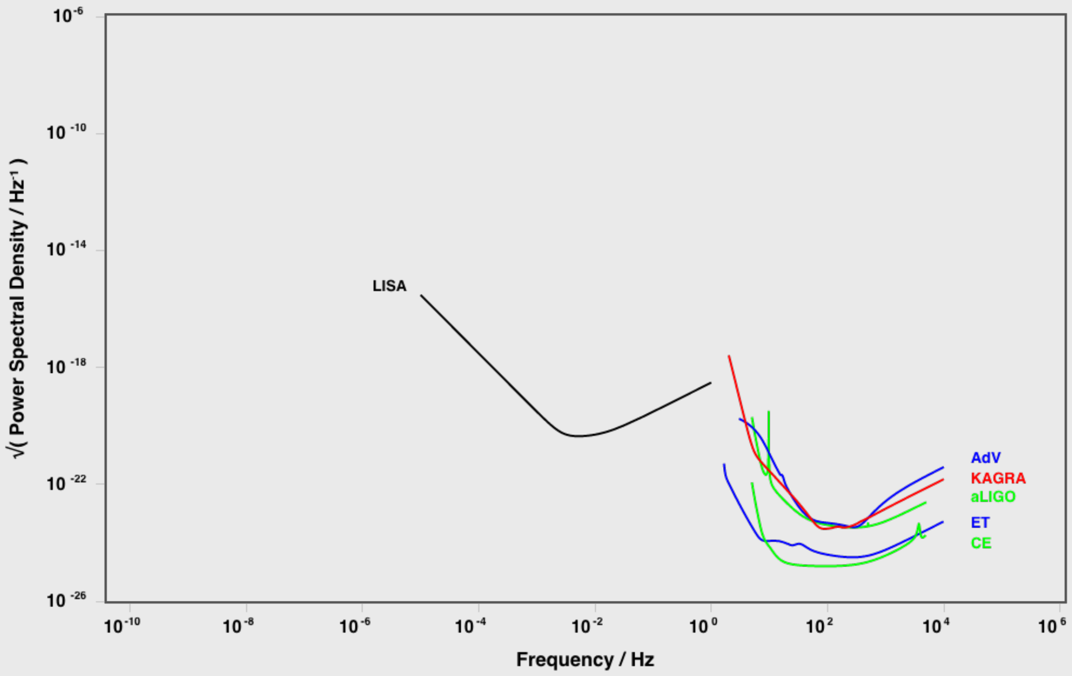

- Space-based (solar orbit) triangular laser interferometer
- 2.5 million km arms, launch planned in 2034
- Very low frequencies: complementary to ground-based interferometers
- Very massive (astrophysical) black holes, very early alerts for BNS, ultra-compact binaries, extreme mass ratio inspirals, precision tests of GR,...

BY-NC-SA

THE MULTIMESSENGER ECOSYSTEM

A large number of existing and future facilities will produce and consume alert triggers. For example, ESFRI-only:

Groud-based optical telescopes, like the Extremely Large Telescope: 5-mirror 39m optical telescope for the ESO on Cerro Armazones, Chile.

Large radiotelescope arrays, like the Square Kilometer Array: huge multi-band radiotelescope arrays in Africa and Australia

Facilities for cosmic ray astronomy, like the Cherenkov Telescope Array: Cherenkov telescopes for highest-energy gamma-ray astronomy, in the Canary Islands and Chile

Neutrino detectors, like KM3NeT: underwater network of neutrino detectors in the Mediterranean

...and many more.

BY-NC-SA

CHALLENGE 6: ALERT MANAGEMENT

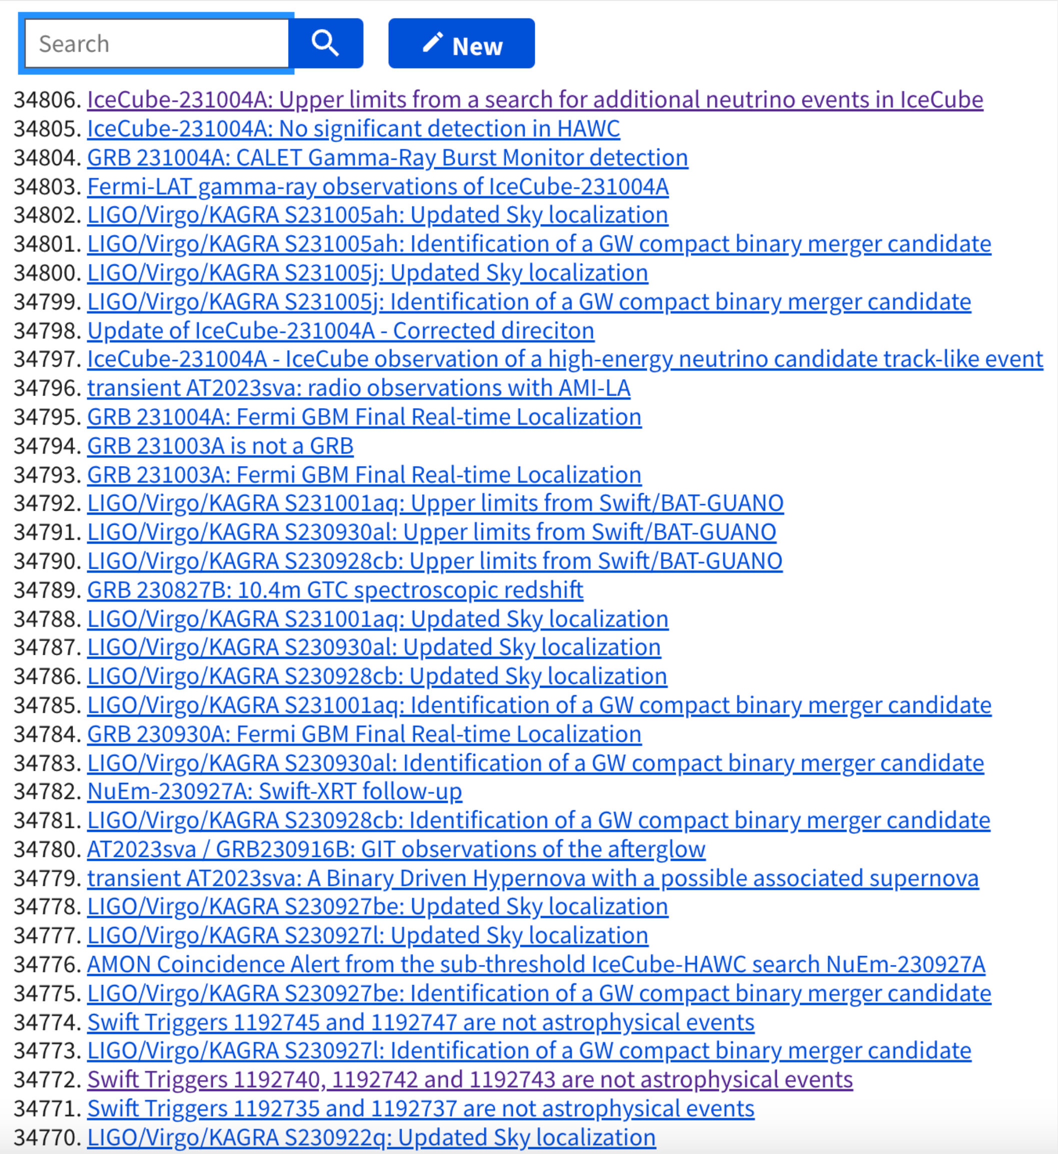

34806. IceCube-231004A: Upper limits from a search for additional neutrino events in IceCube
34805. IceCube-231004A: No significant detection in HAWC
34804. GRB 231004A: CALET Gamma-Ray Burst Monitor detection
34803. Fermi-LAT gamma-ray observations of IceCube-231004A
34802. LIGO/Virgo/KAGRA S231005ah: Updated Sky localization
34801. LIGO/Virgo/KAGRA S231005ah: Identification of a GW compact binary merger candidate
34800. LIGO/Virgo/KAGRA S231005j: Updated Sky localization
34799. LIGO/Virgo/KAGRA S231005j: Identification of a GW compact binary merger candidate
34798. Update of IceCube-231004A - Corrected direciton
34797. IceCube-231004A - IceCube observation of a high-energy neutrino candidate track-like event
34796. transient AT2023sva: radio observations with AMI-LA
34795. GRB 231004A: Fermi GBM Final Real-time Localization
34794. GRB 231003A is not a GRB
34793. GRB 231003A: Fermi GBM Final Real-time Localization
34792. LIGO/Virgo/KAGRA S231001aq: Upper limits from Swift/BAT-GUANO
34791. LIGO/Virgo/KAGRA S230930al: Upper limits from Swift/BAT-GUANO
34790. LIGO/Virgo/KAGRA S230928cb: Upper limits from Swift/BAT-GUANO
34789. GRB 230827B: 10.4m GTC spectroscopic redshift
34788. LIGO/Virgo/KAGRA S231001aq: Updated Sky localization
34787. LIGO/Virgo/KAGRA S230930al: Updated Sky localization
34786. LIGO/Virgo/KAGRA S230928cb: Updated Sky localization
34785. LIGO/Virgo/KAGRA S231001aq: Identification of a GW compact binary merger candidate
34784. GRB 230930A: Fermi GBM Final Real-time Localization
34783. LIGO/Virgo/KAGRA S230930al: Identification of a GW compact binary merger candidate
34782. NuEm-230927A: Swift-XRT follow-up
34781. LIGO/Virgo/KAGRA S230928cb: Identification of a GW compact binary merger candidate
34780. AT2023sva / GRB230916B: GIT observations of the afterglow
34779. transient AT2023sva: A Binary Driven Hypernova with a possible associated supernova
34778. LIGO/Virgo/KAGRA S230927be: Updated Sky localization
34777. LIGO/Virgo/KAGRA S230927l: Updated Sky localization
34776. AMON Coincidence Alert from the sub-threshold IceCube-HAWC search NuEm-230927A
34775. LIGO/Virgo/KAGRA S230927be: Identification of a GW compact binary merger candidate
34774. Swift Triggers 1192745 and 1192747 are not astrophysical events
34773. LIGO/Virgo/KAGRA S230927l: Identification of a GW compact binary merger candidate
34772. Swift Triggers 1192740, 1192742 and 1192743 are not astrophysical events
34771. Swift Triggers 1192735 and 1192737 are not astrophysical events
34770. LIGO/Virgo/KAGRA S230922q: Updated Sky localization

https://gcn.nasa.gov/circulars

eso0717 — Science Release

Controlled by Distant Explosions

VLT Automatically Takes Detailed Spectra of Gamma-Ray Burst Afterglows Only Minutes After Discovery

28 March 2007

Rapid Response Mode Request Received: TELESCOPE PRESET! – @wuves

ATTENTION!

Rapid Response Request Received.

THE TELESCOPE WILL PRESET!

Please follow the instructions below without delay.

Telescope Operator:

The telescope will preset when the countdown reaches zero.

To preset now: press PRESET. If is is unsafe to preset: press STOP.

RA 200645.000

Dec -240000.000

26

Instrument Operator:

Any previous observation is being ended now (shutter closed, reading out).

The Rapid Response Mode OB has been started on a new BOB: the Acquisition is running.

Please execute the rest of the RRM OB WITHOUT DELAY.

Check the e–mail account for the finding chart, and the RRM PSO procedure web page for more information.

https://www.eso.org/public/news/eso0717/

BY-NC-SA

CHALLENGE 7: MULTIMODAL ANALYSIS

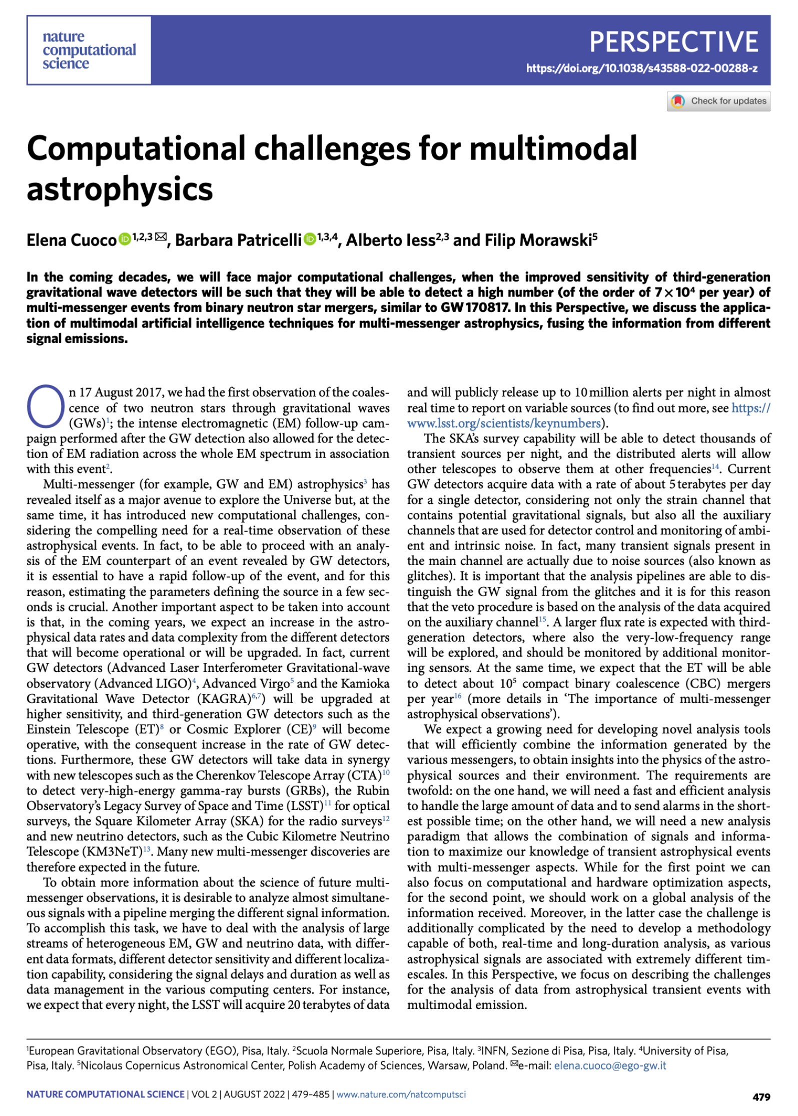

nature computational science

PERSPECTIVE

https://doi.org/10.1038/s43588-022-00288-z

Check for updates

Computational challenges for multimodal astrophysics

Elena Cuoco[1,2,3] ✉, Barbara Patricelli[1,3,4], Alberto Iess[2,3] and Filip Morawski[5]

In the coming decades, we will face major computational challenges, when the improved sensitivity of third-generation gravitational wave detectors will be such that they will be able to detect a high number (of the order of 7 × 10^4 per year) of multi-messenger events from binary neutron star mergers, similar to GW170817. In this Perspective, we discuss the application of multimodal artificial intelligence techniques for multi-messenger astrophysics, fusing the information from different signal emissions.

On 17 August 2017, we had the first observation of the coalescence of two neutron stars through gravitational waves (GWs)[1]; the intense electromagnetic (EM) follow-up campaign performed after the GW detection also allowed for the detection of EM radiation across the whole EM spectrum in association with this event[2].

Multi-messenger (for example, GW and EM) astrophysics[3] has revealed itself as a major avenue to explore the Universe but, at the same time, it has introduced new computational challenges, considering the compelling need for a real-time observation of these astrophysical events. In fact, to be able to proceed with an analysis of the EM counterpart of an event revealed by GW detectors, it is essential to have a rapid follow-up of the event, and for this reason, estimating the parameters defining the source in a few seconds is crucial. Another important aspect to be taken into account is that, in the coming years, we expect an increase in the astrophysical data rates and data complexity from the different detectors that will become operational or will be upgraded. In fact, current GW detectors (Advanced Laser Interferometer Gravitational-wave observatory (Advanced LIGO)[4], Advanced Virgo[5] and the Kamioka Gravitational Wave Detector (KAGRA)[6,7]) will be upgraded at higher sensitivity, and third-generation GW detectors such as the Einstein Telescope (ET)[8] or Cosmic Explorer (CE)[9] will become operative, with the consequent increase in the rate of GW detections. Furthermore, these GW detectors will take data in synergy with new telescopes such as the Cherenkov Telescope Array (CTA)[10] to detect very-high-energy gamma-ray bursts (GRBs), the Rubin Observatory's Legacy Survey of Space and Time (LSST)[11] for optical surveys, the Square Kilometer Array (SKA) for the radio surveys[12] and new neutrino detectors, such as the Cubic Kilometre Neutrino Telescope (KM3NeT)[13]. Many new multi-messenger discoveries are therefore expected in the future.

To obtain more information about the science of future multi-messenger observations, it is desirable to analyze almost simultaneous signals with a pipeline merging the different signal information. To accomplish this task, we have to deal with the analysis of large streams of heterogeneous EM, GW and neutrino data, with different data formats, different detector sensitivity and different localization capability, considering the signal delays and duration as well as data management in the various computing centers. For instance, we expect that every night, the LSST will acquire 20 terabytes of data and will publicly release up to 10 million alerts per night in almost real time to report on variable sources (to find out more, see https://www.lsst.org/scientists/keynumbers).

The SKA's survey capability will be able to detect thousands of transient sources per night, and the distributed alerts will allow other telescopes to observe them at other frequencies[14]. Current GW detectors acquire data with a rate of about 5 terabytes per day for a single detector, considering not only the strain channel that contains potential gravitational signals, but also all the auxiliary channels that are used for detector control and monitoring of ambient and intrinsic noise. In fact, many transient signals present in the main channel are actually due to noise sources (also known as glitches). It is important that the analysis pipelines are able to distinguish the GW signal from the glitches and it is for this reason that the veto procedure is based on the analysis of the data acquired on the auxiliary channel[15]. A larger flux rate is expected with third-generation detectors, where also the very-low-frequency range will be explored, and should be monitored by additional monitoring sensors. At the same time, we expect that the ET will be able to detect about 10^5 compact binary coalescence (CBC) mergers per year[16] (more details in 'The importance of multi-messenger astrophysical observations').

We expect a growing need for developing novel analysis tools that will efficiently combine the information generated by the various messengers, to obtain insights into the physics of the astrophysical sources and their environment. The requirements are twofold: on the one hand, we will need a fast and efficient analysis to handle the large amount of data and to send alarms in the shortest possible time; on the other hand, we will need a new analysis paradigm that allows the combination of signals and information to maximize our knowledge of transient astrophysical events with multi-messenger aspects. While for the first point we can also focus on computational and hardware optimization aspects, for the second point, we should work on a global analysis of the information received. Moreover, in the latter case the challenge is additionally complicated by the need to develop a methodology capable of both, real-time and long-duration analysis, as various astrophysical signals are associated with extremely different timescales. In this Perspective, we focus on describing the challenges for the analysis of data from astrophysical transient events with multimodal emission.

---

[1]European Gravitational Observatory (EGO), Pisa, Italy. [2]Scuola Normale Superiore, Pisa, Italy. [3]INFN, Sezione di Pisa, Pisa, Italy. [4]University of Pisa, Pisa, Italy. [5]Nicolaus Copernicus Astronomical Center, Polish Academy of Sciences, Warsaw, Poland. ✉e-mail: elena.cuoco@ego-gw.it

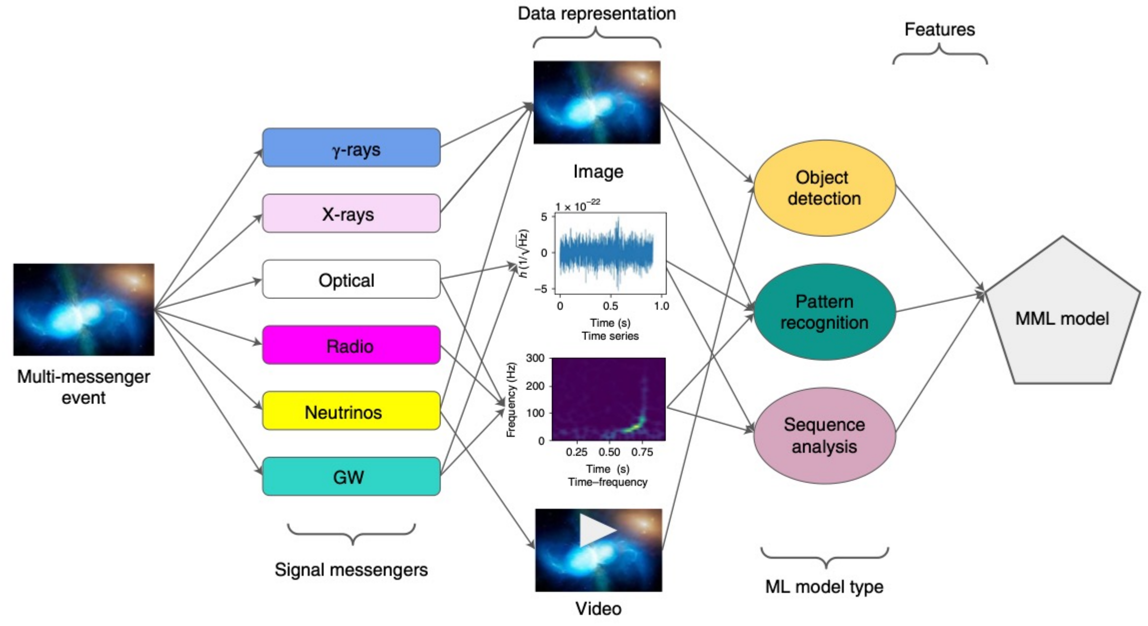

Fig. 3 | The astrophysical event's conceptual MML workflow. A multi-messenger astrophysical transient event can manifest itself via various signal types, including GW, γ-rays, X-rays, optical and radio emission, and neutrinos. Different modalities have their own representation in various domains. We can use the extracted features to perform model prediction in the first stage by using deep learning and ML models. Furthermore, we can analyze the predictions later by combining them in the global MML model.

BY-NC-SA

AND QUANTUM COMPUTING?

OPEN ACCESS
IOP Publishing
Classical and Quantum Gravity
Class. Quantum Grav. **40** (2023) 045001 (15pp)
https://doi.org/10.1088/1361-6382/acafcf

Parameter estimation of gravitational waves with a quantum metropolis algorithm

Gabriel Escrig[1,*], **Roberto Campos**[1,2],
Pablo A M Casares[1] and **M A Martin-Delgado**[1,3]

[1] Departamento de Física Teórica, Universidad Complutense de Madrid, Madrid, Spain
[2] Quasar Science Resources, SL, Madrid, Spain
[3] CCS-Center for Computational Simulation, Universidad Politécnica de Madrid, Madrid, Spain

E-mail: gescrig@ucm.es

Received 22 September 2022; revised 21 November 2022
Accepted for publication 3 January 2023
Published 20 January 2023

CrossMark

Abstract

After the first detection of a gravitational wave in 2015, the number of successes achieved by this innovative way of looking through the Universe has not stopped growing. However, the current techniques for analyzing this type of events present a serious bottleneck due to the high computational power they require. In this article we explore how recent techniques based on quantum algorithms could surpass this obstacle. For this purpose, we propose a quantization of the classical algorithms used in the literature for the inference of gravitational wave parameters based on the well-known quantum walks technique applied to a Metropolis–Hastings algorithm. Finally, we develop a quantum environment on classical hardware, implementing a metric to compare quantum versus classical algorithms in a fair way. We further test all these developments in the real inference of several sets of parameters of all the events of the first detection period GWTC-1 and we find a polynomial advantage in the quantum algorithms, thus setting a first starting point for future algorithms.

Keywords: gravitational waves, quantum walks, Metropolis–Hastings algorithms

(Some figures may appear in colour only in the online journal)

---

[*] Author to whom any correspondence should be addressed.

Original Content from this work may be used under the terms of the Creative Commons Attribution 4.0 licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

1361-6382/23/045001+15$33.00 © 2023 The Author(s). Published by IOP Publishing Ltd Printed in the UK

1

PHYSICAL REVIEW RESEARCH **4**, 023006 (2022)

Quantum algorithm for gravitational-wave matched filtering

Sijia Gao,[*] Fergus Hayes,[†] Sarah Croke, Chris Messenger, and John Veitch
SUPA, School of Physics and Astronomy, University of Glasgow, Glasgow G12 8QQ, United Kingdom

(Received 3 September 2021; accepted 7 March 2022; published 4 April 2022)

Quantum computational devices currently under development have the potential to accelerate data analysis techniques beyond the ability of any classical algorithm. We propose the application of a quantum algorithm for the detection of unknown signals in noisy data. We apply Grover's algorithm to matched filtering, a signal processing technique that compares data to a number of candidate signal templates. In comparison to the classical method, this provides a speedup proportional to the square root of the number of templates, which would make possible otherwise intractable searches. We demonstrate both a proof-of-principle quantum circuit implementation and a simulation of the algorithm's application to the detection of the gravitational wave signal GW150914. We discuss the time complexity and space requirements of our algorithm as well as its implications for the currently computationally limited searches for continuous gravitational waves.

DOI: 10.1103/PhysRevResearch.4.023006

I. INTRODUCTION

Quantum computing holds enormous potential for computational speedup of certain tasks, offering the possibility of solving classically intractable problems, in particular, in quantum chemistry and many-body physics [1,2]. The technology has seen rapid development in the last few years, resulting in processors with 50–100 qubits, and the first demonstrations of clear quantum advantage over classical computation [3,4]. Quantum algorithms (see Ref. [5] for an accessible overview) are being explored for more and more fields of endeavor: for example, finance [6], quantum simulation [7], particle physics [8,9], machine learning [10,11], and, as the technology matures and a new generation of software developers adopt quantum programming languages, it may be anticipated that new and unexpected applications will be discovered. A particularly versatile quantum subroutine is Grover's search algorithm [12], which finds a marked solution in a large unstructured database. Grover's algorithm, one of the earliest proposed quantum algorithms, provides a square-root speedup over classical search. This is less dramatic than the exponential speedup promised by, e.g., Shor's algorithm [13], but can nevertheless provide a significant practical advantage for problems with a large search space. By defining the search space and conditions for a desired solution, Grover's algorithm may be applied to any computational problem with a limited structure and has found use in minimum finding [14], clustering and nearest-neighbor algorithms for supervised and unsupervised learning [15,16], and pattern matching [17–19] to name but a few. In this paper, we propose the use of Grover's search in quantum algorithms for matched filtering, with applications in gravitational wave (GW) astronomy. These algorithms inherit the square root speedup of Grover's search algorithm, an improvement which could enable GW searches currently intractable with state-of-the-art classical techniques.

Matched filtering is a signal processing technique [20] in which an exhaustive search is performed over a bank of templates to find the template that when correlated with the data returns the highest detection statistic [21], making it a natural candidate for a quantum speedup through Grover's algorithm. In GW matched-filtering a geometric definition of distance within the parameter space is defined based on the relative loss in signal-to-noise ratio (SNR) between a template and a potential signal. The required distribution of the templates in the search space are chosen so the distance (or overlap) between adjacent templates is constant throughout the space. Depending on the specific data analysis problem, the number of templates can range up to $\sim 10^{12}$ [22], resulting in a total computational time of $\sim 10^6$ CPU hours. The spacing of templates in the parameter space determines the efficiency of the search but also the overall number of templates, and the sensitivity of searches for certain classes of signals (e.g., continuous wave sources) is currently computationally limited. Thus, even a modest square-root speedup could enable the detection of signals which would be infeasible with classical techniques.

Key to our proposed algorithms is the fact that the potential signals in GW astronomy are well-modeled by general relativity, and the templates may be readily computed as part of the matching procedure. This eliminates the need to preload the database into quantum random access memory (qRAM) [23], and thus avoids hidden complexity associated with this loading step, as well as doubts about the experimental feasibility of constructing qRAM [24–27]. The presented

---

[*] s.gao.2@research.gla.ac.uk
[†] f.hayes.1@research.gla.ac.uk

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

2643-1564/2022/4(2)/023006(24) 023006-1 Published by the American Physical Society

THE ET E-INFRASTRUCTURE BOARD

Chairs: Patrice Verdier (IJCLab), SB (INFN-Torino)

Division 1: Software, frameworks, and data challenge support
Chair: Andres Tanasijczuk (UCLouvain)
OSB Liaison: John Veitch (U. Glasgow), Elena Cuoco (EGO)

Division 2: Services and Collaboration Support
Chair: Antonella Bozzi (EGO)

Division 3: Computing and data model, Resource Estimation
Chair: Gonzalo Merino (PIC)

Division 4: Multimessenger alerts infrastructure
Chair: Steven Schramm (U. Geneva)

TTG: Technology Tracking working Group
Chair: Sara Vallero (INFN-Torino)

BY-NC-SA

WAY FORWARD

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

BY-NC-SA

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 (Virgo+ET) initial R&D program
  - Data Lake (Rucio) for data distribution
  - VRE/REANA for data access and job management
  - Using ET MDCs as testbeds

BY-NC-SA

ET IN ESCAPE

The new ESCAPE Collaboration work programme

ESCAPE CC
Operating the community-based "Competence Center" for EOSC-alignment, train and support, extended outreach, financial model for services and networking with other SCL-CCs

ESCAPE EVSI
R&I for an "European Virtual Institute for Research Software" for advanced technologies

Entities

VRE services

ESCAPE DIOS | Data Infrastructure for Open Science
Access physical & e-infrastructures
Processing & Analysis
Security & Operations

ESCAPE OSSR | Open-source Scientific Software and Service Repository
Aggregator & Integrators
Sharing and Discover
Training & Support

ESCAPE ESAP | ESFRI Science Analysis Platform
Processing & Analysis
Sharing and Discovery
Training & Support

ESCAPE CS | Citizen Science
Sharing and Discovery

ESCAPE VO | Virtual Observatory
Processing & Analysis
Sharing and Discovery
Training & Support

Programmes

ESCAPE COSO
Challenging "Open Science Objectives" by RI commitments in Open Science Projects (OSP) as well as Cross-Cluster Open Science Projects (COSP)

ESCAPE TECH
Bring the FAIRness within technology, R&D and innovation projects as well as explore new "close-to-sensors" low-latency open-data science

ESCAPE CARS
Career development and rewarding for researcher committing in Open Science. Planning, tracking, and assessing scientific knowledge production

ESCAPE SDSS
Building synergies on "Sector Data Spaces" for Society: Green deal, Health, Manufacturing, Education and Skills

24 Giovanni Lamanna 07/07/2023

CC BY-NC-SA

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

BY-NC-SA

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

BY-NC-SA

PROJECTS & MORE PROJECTS

- ET-PP
  - WP8 – very good collaboration with eIB
- ICSC_S2
  - “Flagship” use cases development
- M2Tech v2
  - WP6 – see next slides
- ESCAPE/OSCARS proposals
  - MADDEN (INFN-TO & UCLouvain)
  - ETAP (Université de Genève)
  - Streaming data for Low-latency?
- ETIC
  - See slide about 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.

BY-NC-SA

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

cc BY-NC-SA

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?

BY-NC-SA

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

BY-NC-SA