



Original software publication

DQSEGDB: A time-interval database for storing gravitational wave observatory metadata



Ryan P. Fisher^{a,*}, Gary Hemming^b, Marie-Anne Bizouard^c, Duncan A. Brown^d,
Peter F. Couvares^e, Florent Robinet^f, Didier Verkindt^g

^a CNU, 1 Avenue of the Arts, Newport News, VA 23606, USA

^b European Gravitational Observatory (EGO), I-56021 Cascina, Pisa, Italy

^c Artemis, Université Côte d'Azur, Observatoire Côte d'Azur, CNRS, CS 34229, F-06304 Nice Cedex 4, France

^d Syracuse University, Syracuse, NY 13244, USA

^e LIGO, California Institute of Technology, Pasadena, CA 91125, USA

^f LAL, Univ. Paris-Sud, CNRS/IN2P3, Université Paris-Saclay, F-91898 Orsay, France

^g Laboratoire d'Annecy de Physique des Particules (LAPP), Univ. Grenoble Alpes, Université Savoie Mont Blanc, CNRS/IN2P3, F-74941 Annecy, France

ARTICLE INFO

Article history:

Received 15 March 2020

Received in revised form 4 February 2021

Accepted 9 February 2021

Keywords:

Database 1

Metadata 2

Time segments 3

Gravitational waves 4

LIGO-Virgo 5

ABSTRACT

The Data Quality Segment Database (DQSEGDB) software is a database service, backend application programming interface (API), frontend graphical web interface, and client package used by the Laser Interferometer Gravitational-Wave Observatory, Virgo, GEO600 and the Kamioka Gravitational Wave Detector for storing and accessing metadata describing the status of their detectors. The DQSEGDB has been used in the analysis of all published detections of gravitational waves in the advanced detector era. The DQSEGDB currently stores roughly 600 million metadata entries and responds to roughly 600,000 queries per day with an average response time of 0.317 s.

© 2021 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

Code metadata

Current code version	1.6.1
Permanent link to code/repository used for this code version	https://github.com/ElsevierSoftwareX/SOFTX_2020_125
Code Ocean compute capsule	N/A
Legal Code License	GNU GENERAL PUBLIC LICENSE v3.0
Code versioning system used	Git
Software code languages, tools, and services used	Python, MariaDB, PHP
Compilation requirements, operating environments & dependencies	Scientific Linux 7.5, Python 2.7.5
If available Link to developer documentation/manual	http://ligovirgo.github.io/dqsegdb/
Support email for questions	question@ligo.org

1. Motivation and significance

Gravitational waves (GWs) are disturbances in the metric of space-time that propagate through the Universe and carry information about the astrophysics of sources that generate them. Gravitational waves couple weakly to matter, and so, for current detectors, these sources must be massive objects moving with high accelerations [1]. Although GWs may have very large amplitudes at their origin, they also typically travel extra-galactic distances to reach the Earth. When these waves reach the Earth,

their strength is such that their resulting spacetime perturbation changes measurements of length by 1 part in 10^{20} . This results in an extremely small signal, even if detected with kilometer-scale detectors [2]. The mission of the LIGO Scientific Collaboration and the Virgo Collaboration (LVC) is to detect these weak signals in order to advance our understanding of the Universe. Since 2015, GW detections have shown that Einstein's theory of general relativity holds for colliding black holes and neutron stars [3–5]. These discoveries allow us to estimate the number of binary black hole and binary neutron star mergers in our local Universe [6–8], and have demonstrated that some gamma-ray burst (GRB) events are powered by the coalescence of neutron stars [4,9].

* Corresponding author.

E-mail address: ryan.fisher@cnu.edu (Ryan P. Fisher).

The detection of GWs has been made possible through the colossal effort of thousands of scientists to develop an extremely precise set of interferometers (IFOs) and an ecosystem of computational infrastructure that enables the capture and analysis of the data generated by the IFOs. The Data Quality Segment Database (DQSEGDB) occupies one critical space in this infrastructure. The data analysis (DA) algorithms require information about the state of the IFOs to analyze the observatory data. This requires the definition and distribution of metadata about the data, which we call data quality (DQ) flags. A DQ flag is the name given to a set of metadata that describes a portion of the global status of the detector, operation of the instrument, or quality of the data that may impact its analysis. A category of flags mark the times when the IFOs are operating in an optimal state, thereby indicating which observatory data should be analyzed. Additional flags indicate data that should explicitly not be analyzed, such as when hardware injections are ongoing or when electronics faults cause noise in the GW detection channel. These DQ flags are also called DQ vetoes because they can be used to exclude data from being analyzed [10]. The DQSEGDB is the service used to store and provide access to these flags.

The set of data associated with each flag name is the list of times when the state of that flag was known and the list of times when that state was active or inactive, which are complements within the set of known times. The time periods are contained in a data product known as “segments”, where a segment is a continuous range of time expressed as a half-open GPS time interval $[t_{\text{start}}, t_{\text{end}})$. Within the GW community, the terms DQ segments and DQ flags are often used interchangeably because of this tight relationship. Each flag has a unique name. The flag names are associated with their IFO identifiers, and are combined in the format $[\text{IFO}]:[\text{FLAG-NAME}]$.

1.1. Initial detector databases

The DQSEGDB service and client software were built to replace the aging predecessor services that served the Laser Interferometer Gravitational-Wave Observatory (LIGO) and Virgo collaborations separately during their initial observing runs. These previous services were each able to store hundreds of flags and approximately two million individual DQ segment and metadata entries by the end of the final science runs of initial LIGO and Virgo. The LIGO service relied on an IBM DB2 database instance, and an XML transport layer. The Virgo database (VDB) used a MySQL database, a series of Python management scripts and a PHP-driven web user interface. The VDB used a substantially differing schema to DQSEGDB. Its structure was not normalized to any great degree and was designed to accommodate segment classification in a slightly different manner. The DB explicitly stored times a flag was inactive, where this information can be determined implicitly from the known and active segment information in DQSEGDB.

By the end of its lifetime, the LIGO service had become very slow at the scales of data it was storing, and would often take 10 to 30 min to respond to queries made by GW data analyses. A combination of several factors, including poorly normalized and indexed tables, excessive overall size and memory consumption of the database, relative difficulty in maintaining the DB2 instance over a MySQL approach, an overemphasis on server-side complexity in both query and insert steps including a non-RESTful application programming interface (API), and the slow speed of the internal generation of the XML documents for the data transport layer, all greatly restricted the speed at which the server could respond to queries for large data sets. This consequently severely restricted the usability of the service and indicated a strong need for a replacement for advanced GW detectors, where

the number of flags and number of segments the databases would need to store would grow by factors of hundreds. An increasing number of new software systems were also unable to use the DQ metadata effectively due to the slow response times of the server. Finally, new user requirements pushed for a redesign of the API and database schema. These issues led to the LVC making the decision to pool their resources and to design a new segment database infrastructure. This led to the development of the DQSEGDB software.

2. LVC data landscape and terminology

Each IFO produces one primary data channel, which contains the measurement of the GW strain, and approximately 200,000 channels of auxiliary data that are used to monitor the status of all the hardware and software components used to produce the primary data. This data set constitutes approximately 2 TB per day per IFO. Customized scripts are used to reduce this huge amount of auxiliary data into approximately 1000 DQ flags per LIGO IFO. At the location of the IFOs, a set of real-time processes automatically generate segments for a portion of the total DQ flags. These processes encode the metadata in XML files that each contain information about the status of these flags for 16 s of data. Each of these XML files is about 78 kB in size, which translates to roughly 420 MB of metadata generated per day per IFO. The XML files are then transferred via rsync from each IFO to the DQSEGDB server, which is hosted at the LIGO Laboratory at the California Institute of Technology. The DQSEGDB server then executes all of the code needed to extract the metadata from the XML files, publishes it to the database, and archives the raw XML files.

3. Software development and description

To develop the new DQSEGDB services and software, the collaborations selected a committee of members, who are the authors of this article, representing both the LIGO and Virgo developers and Detector Characterization (DetChar) experts. The committee first developed a set of design principles and then circulated a request for user requirements. The developers then designed an API that would both suit the user needs and speed requirements for the services. After this, the database structure, python application layer and clients were designed and built with a focus on maximizing speed while meeting the use requirements at each step. The remainder of this section will describe additional details about the selected design.

In addition to the requirements that the DQSEGDB service be able to respond rapidly while storing a large amount of metadata, several other design requirements and elements of design philosophy were also met when the new software was written. The database was required to contain both the DQ segments and enough additional metadata to allow the tracking of their provenance. The service was required to allow remote clients to connect via command line or web GUI, and was to provide the metadata within 15 min of its generation. The API was chosen to provide a RESTful set of URLs with a resource-oriented architecture, compatible with multiple programming languages, and restrictive such that data could not be removed from the database. A JSON format was chosen for the returned data, which included an option for all provenance metadata. Additional functionalities were deferred to the client layer to ensure speed at the server.

These design requirements led to the current DQSEGDB software design. The service is split into three major components. The first is the primary database server, which is generally labeled the DQSEGDB. The second is the client software package, which contains both command line tools and a Python package that can

Date	Requests				Avg. response times (s)	
	Get	Patch	Total	Req./s	Get	Patch
31/08/15	13.490	58.311	71.801	0.849	1.99	0.01
01/09/15	13.750	56.195	69.945	0.827	1.86	0.01
01/10/15	18.330	65.670	84.000	0.993	1.876	0.015
14/09/16	37.673	80.857	118.530	1.401	3.584	0.018
22/12/17	42.368	104.538	146.906	1.736	14.228	0.108
01/01/18	41.198	104.573	145.771	1.723	13.036	0.094
01/02/18	36.756	105.753	142.509	1.685	7.962	0.046
10/08/18	41.316	108.166	149.482	1.767	14.192	0.086
10/02/19	36.798	374.346	411.144	4.860	1.285	0.009
10/08/19	48.937	577.708	626.645	7.407	1.27	0.008
10/02/20	39.690	567.588	607.278	7.178	3.238	0.012

Fig. 1. Demonstrating performance stability of the DQSEGDB service. In 2020, the database contains O(100) times more data, and responds to O(5) times more requests per second with nearly identical performance compared to 2016 values.

be used to query the database. This set of tools provides many functions requested by LVC scientists, while satisfying the design requirements listed. The final component is a graphical web interface, which provides a GUI interface to allow collaboration scientists to rapidly access the metadata without needing to write any code.

The DQSEGDB server consists of an Apache layer that calls a custom Python application via the Apache WSGI module. The Python application uses ODBC to communicate with the database. The DQSEGDB database uses the InnoDB engine available within MariaDB. The database hosts the DQ flags, their associated segments, the associated metadata about those segments, and some overall metadata about the data. A normalized schema is used to alleviate the need to store large quantities of text metadata, with information such as flag and flag-version association, originating process, user and interferometer identification all provided with normalized values. In this manner, referential integrity between the various component parts of the database can be properly enforced and maintained. The API provides access to all data associated with a given DQ flag through RESTful URIs, formatted as /dq/IFO/FLAG/VERSION. The data can be downselected based on the information and time interval of interest, using options such as /dq/IFO/FLAG/VERSION/active?s=t1&e=t2. This URI will return all active segments for the given FLAG in the GPS interval [t1, t2).

The DQSEGDB database service is much faster and more stable than its predecessors. Currently, the database contains ≈ 310 M segments and ≈ 480 M segment-summaries in its dedicated tables, along with ≈ 33 M rows providing process-related metadata. Overall, it occupies ≈ 53 GB of disk space. Rates of increase in the amount of data in DQSEGDB vary with time, dependent upon factors such as whether the IFOs are in a scientific data-taking period and publishing frequency. Over the past 2 years, the segment data itself has grown by ≈ 8 GB/year.

The numbers of HTTP GET and PATCH requests over time is reflected in Fig. 1. The system has handled a considerable increase in requests with the passage of time. The table reflects how GET response-times for requests are dependent mainly upon the ways in which these requests are defined by the users. When users build client requests that require the interrogation of data over broad sweeps of time, these will require longer periods of time to resolve and provide a response than a short request. The table also shows the time required to insert new segments into the database with PATCH requests. These include the time required to check the availability of a flag and subsequently the latest available version of the flag, the SQL INSERT of the segments associated to the flag-version, and the INSERT of related process-metadata. This time has remained constantly below the 100 ms level, reaching as low even as 8 ms.

The DQSEGDB software is designed to be used within a more extended infrastructure. A development and a backup server, each containing copies of the production server's data, are used to ensure that the service is always available and operating system software updates may be tested and deployed rapidly. Overall, the system of servers and clients that make up the whole DQSEGDB infrastructure is shown in Fig. 2. The developers also implemented a complex system of monitors to ensure that every part of the service is functioning, from the initial DQ flag XML file generation through to the latency of queryability for newly generated data.

4. Impact

The new DQSEGDB system of servers has been very successful in meeting the needs of the GW community for storing and distributing IFO metadata since 2014. Thanks to the high performance of the DQSEGDB service, nearly all LVC GW searches are using this centralized source of data quality information. It is, thus, also providing a system for careful control and synchronization of the detector status information used by the LVC searches. The DQSEGDB is also used by many automated IFO monitoring processes and many LVC scientists investigating the performance of the IFOs. In particular, the data analyses that concluded in the detection of all GWs thus far have relied on the DQSEGDB infrastructure [4,11–14].

The impact of the DQ information hosted in the DQSEGDB on GW searches is significant, as illustrated in Fig. 3 [10]. The information is used to mitigate systematic noise issues, and thanks to the speed of the DQSEGDB service, the speed of testing of the different choices of DQ flags for use in analyses has been drastically improved.

One example of the types of DQ flags used to remove a significant amount of noise was the “RF45 flag”. This flag indicated times when issues with the electronics that controlled the radio frequency (RF) sidebands used to sense and control LIGO's optical cavities would contaminate the main detection channel with noise that resulted in a significant number of false triggers in DA pipelines [10]. The latency from the time data is collected at the IFO sites to the moment the metadata may be queried by rapid analyses has also been reduced to less than 5 min. This functionality is being used by several “medium latency” analyses that are automatically started in response to external events such as observations of gamma-ray bursts.

The new DQSEGDB service has allowed new collaboration tools that automatically, and very frequently, query the DQ flag metadata to be developed. These services provide many different benefits to the large, 1000+ person LVC. One example, the Summary Page web infrastructure [15] makes heavy use of this




Fig. 2. Schematic view of the DQSEGDB infrastructure. This depicts the set of servers as were deployed for use by the LIGO and Virgo Scientific Collaboration in 2015. DQXML refers to the XML documents that are used to pass data between some layers of the architecture. CIT represents the Caltech network. H1/L1/V1/GEO represent the network at each IFO site, where the raw detector data is converted into segments for flags using the DMT or SegOnline programs. CIT hosts the primary database services described in this article, and receives the data to be published from the IFO site servers as input.




Fig. 3. This image is reproduced with permission from [10], original creator T. J. Massinger. The impact of applying DQ flags from the DQSEGDB to a LVC GW search in the data from Advanced LIGO's first observing run is shown. The initial background events in the search have signal-to-noise ratio (SNR) values reaching above 14. With the DQ applied, the background is reduced below an SNR of 12.5. The upper limit of this background is tied to the limit at which GW events may be detected. The SNR of GW151226 is indicated on the figure to demonstrate that this detection would have been missed without the use of the DQ data. Image License <https://creativecommons.org/licenses/by/3.0/>. Image was cropped from original.

ability. Examples of these pages are available at https://www.gw-open-science.org/detector_status. Many of the plots indicate the state of the interferometers, and all of this metadata is retrieved from the DQSEGDB. These plots are updated on a rolling basis, requiring very frequent queries to the DQSEGDB service that the old service would not have been able to handle. These pages are used by IFO commissioners, data quality investigators, data analysts and the wider astronomical community to easily assess the state of the interferometers and rapidly investigate systematic issues. They have proven invaluable to data and event validation efforts in addition to daily IFO and collaboration operations.

Due to the speed and reliability of the service, additional GW detectors have also begun using this single instance of the DQSEGDB. The GEO600 (GEO) collaboration and the Kamioka Gravitational Wave Detector (KAGRA) collaboration store metadata with this service as well. Thus, the DQSEGDB infrastructure and service is now used by all IFO-based GW detection efforts in the world.

5. Conclusions

The DQSEGDB has been tremendously successful in serving the GW astronomy community. This set of database, backend, frontend and client software has provided rapid access to the DQ segments needed by the LVC for all GW detections made thus far. The speed and reliability of the database combined with its clear, RESTful API has resulted in the design of new tools that enable scientists to more rapidly and easily understand the IFOs in the GW detection network.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors gratefully acknowledge the support of the United States National Science Foundation (NSF), USA for the construction and operation of the LIGO Laboratory and Advanced LIGO as well as the Science and Technology Facilities Council (STFC) of the United Kingdom, the Max-Planck-Society (MPS), and the State of Niedersachsen/Germany for support of the construction of Advanced LIGO and construction and operation of the GEO600 detector. The authors gratefully acknowledge the National Science Foundation, USA Grants PHY-1700765 and PHY-1104371 for the support of this project. Additional support for Advanced LIGO was provided by the Australian Research Council. The authors gratefully acknowledge the Italian Istituto Nazionale di Fisica Nucleare (INFN), the French Centre National de la Recherche Scientifique (CNRS) and the Foundation for Fundamental Research on Matter supported by the Netherlands Organisation for Scientific Research, for the construction and operation of the Virgo detector, their support in all software developments, and the creation and support of the European Gravitational Observatory consortium. The authors also gratefully acknowledge research support from these agencies as well as by the Council of Scientific and Industrial Research of India, the Department of Science and Technology, India, the Science & Engineering Research Board (SERB), India, the Ministry of Human Resource Development, India, the Spanish Agencia Estatal de Investigación, the Vicepresidència i Conselleria d'Innovació, Recerca i Turisme, Spain and the Conselleria d'Educació i Universitat del Govern de les Illes Balears, Spain, the Conselleria d'Educació, Investigació, Cultura i Esport de la Generalitat Valenciana, Spain, the National Science Centre of Poland, the Swiss National Science Foundation (SNSF), the Russian Foundation for Basic Research, the Russian Science Foundation, the European Commission, the European Regional Development Funds (ERDF), the Royal Society, United Kingdom, the Scottish Funding Council, United Kingdom, the Scottish Universities Physics Alliance, United Kingdom, the Hungarian Scientific Research Fund (OTKA), the Lyon Institute of Origins (LIO), France, the Paris Île-de-France Region, the National Research, Development and Innovation Office Hungary (NKFIH), Hungary, the National Research Foundation of Korea, Industry Canada and the Province of Ontario through the Ministry of Economic Development and Innovation, Canada, the Natural Science and Engineering Research Council Canada, the Canadian Institute for Advanced Research, the Brazilian Ministry of Science, Technology, Innovations, and Communications, the International Center for Theoretical Physics South American Institute for Fundamental Research (ICTP-SAIFR), the Research Grants Council of Hong Kong, the National Natural Science Foundation of China (NSFC), the Leverhulme Trust, United Kingdom, the Research Corporation, USA, the Ministry of Science and Technology (MOST), Taiwan and the Kavli Foundation, USA. The authors gratefully for computational resources provided by the NSF, USA, STFC, United Kingdom, INFN, Italy and CNRS, France, and supported by National Science Foundation, USA Grants PHY-1626190, PHY-1700765, PHY-0757058 and PHY-0823459. The authors gratefully acknowledge the contributions of Larne Pekowsky and Ping Wei for the development, support and improvement of the original LIGO segment database and Leone Bosi for the development of the original Virgo database.

Role of funding source

Funding for this project was provided by the National Science Foundation, USA, the European Gravitational Observatory, the Italian Istituto Nazionale di Fisica Nucleare (INFN), the French Centre National de la Recherche Scientifique (CNRS) and the Foundation for Fundamental Research on Matter supported by the Netherlands Organisation for Scientific Research. The funding sources had no involvement in the work described in this article or the writing or submission of the article.

References

- [1] Thorne K. Gravitational radiation. In: Hawking S, Israel W, editors. *Three hundred years of gravitation*. Cambridge: Cambridge University Press; 1987. p. 330–458.
- [2] Creighton JDE, Anderson WG. Gravitational-wave physics and astronomy: An introduction to theory, experiment and data analysis. Weinheim, Germany: Wiley-VCH; 2011. p. 375. URL <http://www.wiley-vch.de/publish/dt/books/ISBN3-527-40886-X>.
- [3] Abbott BP, et al., LIGO Scientific Collaboration, Virgo Collaboration. Observation of gravitational waves from a binary black hole merger. *Phys Rev Lett* 2016;116(6):061102. <http://dx.doi.org/10.1103/PhysRevLett.116.061102>, arXiv:1602.03837.
- [4] Abbott BP, et al., Virgo, LIGO Scientific Collaboration. GW170817: Observation of gravitational waves from a binary neutron star inspiral. *Phys Rev Lett* 2017;119(16):161101. <http://dx.doi.org/10.1103/PhysRevLett.119.161101>, arXiv:1710.05832.
- [5] Abbott BP, et al., LIGO Scientific, Virgo Collaboration. GWTC-1: A gravitational-wave transient catalog of compact binary mergers observed by LIGO and virgo during the first and second observing runs. *Phys Rev* 2019;X9(3):031040. <http://dx.doi.org/10.1103/PhysRevX.9.031040>, arXiv:1811.12907.
- [6] Abbott BP, et al., LIGO Scientific Collaboration, Virgo Collaboration. Upper limits on the rates of binary neutron star and neutron-star-black-hole mergers from advanced ligo's first observing run. 2016, arXiv:1607.07456.
- [7] Abbott BP, et al., LIGO Scientific Collaboration, Virgo Collaboration. Supplement: The rate of binary black hole mergers inferred from advanced LIGO observations surrounding GW150914. 2016, arXiv:1606.03939.
- [8] Abbott BP, et al., LIGO Scientific Collaboration, Virgo Collaboration. The rate of binary black hole mergers inferred from advanced LIGO observations surrounding GW150914. 2016, arXiv:1602.03842.
- [9] Abbott BP, et al., Virgo, Fermi-GBM, INTEGRAL, LIGO Scientific Collaboration. Gravitational waves and gamma-rays from a binary neutron star merger: GW170817 and GRB 170817A. *Astrophys. J.* 2017;848(2):L13. <http://dx.doi.org/10.3847/2041-8213/aa920c>, arXiv:1710.05834.
- [10] Abbott BP, et al., Virgo, LIGO Scientific Collaboration. Effects of data quality vetoes on a search for compact binary coalescences in Advanced LIGO's first observing run. *Classical Quantum Gravity* 2018;35(6):065010. <http://dx.doi.org/10.1088/1361-6382/aaaafa>, arXiv:1710.02185.
- [11] Abbott BP, et al., Virgo, LIGO Scientific Collaboration. Characterization of transient noise in advanced LIGO relevant to gravitational wave signal GW150914. *Classical Quantum Gravity* 2016;33(13):134001. <http://dx.doi.org/10.1088/0264-9381/33/13/134001>, arXiv:1602.03844.
- [12] Abbott BP, et al. GW190425: Observation of a compact binary coalescence with total mass $\sim 3.4M_{\odot}$. 2020, arXiv e-prints arXiv:2001.01761.
- [13] Abbott BP, et al., GROND, SALT Group, OzGrav, DFN, DES, INTEGRAL, Virgo, Insight-Hxmt, MAXI Team, Fermi-LAT, J-GEM, RATIR, IceCube, CAASTRO, LWA, ePESSTO, GRAWITA, RIMAS, SKA South Africa/MeerKAT, H.E.S.S., 1M2H Team, IKI-GW Follow-up, Fermi GBM, Pi of Sky, DWF (Deeper Wider Faster Program), MASTER, AstroSat Cadmium Zinc Telluride Imager Team, Swift, Pierre Auger, ASKAP, VINROUGE, JAGWAR, Chandra Team at McGill University, TTU-NRAO, GROWTH, AGILE Team, MWA, ATCA, AST3, TOROS, Pan-STARRS, NuSTAR, ATLAS Telescopes, BOOTES, CaltechNRAO, LIGO Scientific, High Time Resolution Universe Survey, Nordic Optical Telescope, Las Cumbres Observatory Group, TZAC Consortium, LOFAR, IPN, DLT40, Texas Tech University, HAWC, ANTARES, KU, Dark Energy Camera GW-EM, CALET, Euro VLBI Team, ALMA Collaboration. Multi-messenger observations of a binary neutron star merger. *Astrophys. J.* 2017;848(2):L12. <http://dx.doi.org/10.3847/2041-8213/aa91c9>, arXiv:1710.05833.
- [14] Abbott BP, et al., Virgo, LIGO Scientific Collaboration. GW170814: A three-detector observation of gravitational waves from a binary black hole coalescence. *Phys Rev Lett* 2017;119(14):141101. <http://dx.doi.org/10.1103/PhysRevLett.119.141101>, arXiv:1709.09660.
- [15] Macleod D, Urban AL, Isi M, Massinger T, Paulaitis, Pitkin M, et al. gwpy/gwsumm: 1.0.2. Zenodo; 2019. <http://dx.doi.org/10.5281/zenodo.3590375>.