

GstLAL: A software framework for gravitational wave discovery

Kipp Cannon

RESCEU, The University of Tokyo, Tokyo, 113-0033, Japan

Sarah Caudill

Nikhef, Science Park, 1098 XG Amsterdam, Netherlands

Chiwai Chan

RESCEU, The University of Tokyo, Tokyo, 113-0033, Japan

Bryce Cousins

Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Computational and Data Sciences, The Pennsylvania State University, University Park, PA 16802, USA

Jolien D. E. Creighton

Leonard E. Parker Center for Gravitation, Cosmology, and Astrophysics, University of Wisconsin-Milwaukee, Milwaukee, WI 53201, USA

Becca Ewing

Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

Heather Fong

RESCEU, The University of Tokyo, Tokyo, 113-0033, Japan

Graduate School of Science, The University of Tokyo, Tokyo 113-0033, Japan

Patrick Godwin

Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

Chad Hanna

Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

Department of Astronomy and Astrophysics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Computational and Data Sciences, The Pennsylvania State University, University Park, PA 16802, USA

Shaun Hooper

University of Western Australia, Crawley, Western Australia 6009, Australia

Rachael Huxford

Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

Ryan Magee

Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

Duncan Meacher

Leonard E. Parker Center for Gravitation, Cosmology, and Astrophysics, University of Wisconsin-Milwaukee, Milwaukee, WI 53201, USA

Cody Messick

Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

Soichiro Morisaki

Institute for Cosmic Ray Research, The University of Tokyo, 5-1-5 Kashiwanoha, Kashiwa, Chiba 277-8582, Japan

Debnandini Mukherjee

Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

Hiroaki Ohta

RESCEU, The University of Tokyo, Tokyo, 113-0033, Japan

Alexander Pace

Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

Stephen Privitera

Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik, D-14476 Potsdam-Golm, Germany

Iris de Ruiter

Astronomical Institute, Anton Pannekoek, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, Netherlands

Nikhef, Science Park, 1098 XG Amsterdam, Netherlands

Surabhi Sachdev

Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

LIGO Laboratory, California Institute of Technology, MS 100-36, Pasadena, California 91125, USA

Leo Singer

Astroparticle Physics Laboratory, NASA Goddard Space Flight Center, Mail Code 661, Greenbelt, MD 20771, USA

Divya Singh

Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Gravitation and the Cosmos, The Pennsylvania State University, University Park, PA 16802, USA

Ron Tapia

Department of Physics, The Pennsylvania State University, University Park, PA 16802, USA

Institute for Computational and Data Sciences, The Pennsylvania State University, University Park, PA 16802, USA

Leo Tsukada

RESCEU, The University of Tokyo, Tokyo, 113-0033, Japan

Graduate School of Science, The University of Tokyo, Tokyo 113-0033, Japan

Daichi Tsuna

RESCEU, The University of Tokyo, Tokyo, 113-0033, Japan

Graduate School of Science, The University of Tokyo, Tokyo 113-0033, Japan

Takuya Tsutsui

RESCEU, The University of Tokyo, Tokyo, 113-0033, Japan

Koh Ueno

RESCEU, The University of Tokyo, Tokyo, 113-0033, Japan

Aaron Viets

Leonard E. Parker Center for Gravitation, Cosmology, and Astrophysics, University of Wisconsin-Milwaukee, Milwaukee, WI 53201, USA

Leslie Wade

Department of Physics, Hayes Hall, Kenyon College, Gambier, Ohio 43022, USA

Madeline Wade

Department of Physics, Hayes Hall, Kenyon College, Gambier, Ohio 43022, USA

Abstract

The GstLAL library, derived from Gstreamer and the LIGO Algorithm Library, supports a stream-based approach to gravitational-wave data processing. Although GstLAL was primarily designed to search for gravitational-wave signatures of merging black holes and neutron stars, it has also contributed to other gravitational-wave searches, data calibration, and detector-characterization efforts. GstLAL has played an integral role in all of the LIGO-Virgo collaboration detections, and its low-latency configuration has

enabled rapid electromagnetic follow-up for dozens of compact binary candidates.

Keywords:

Gravitational waves, neutron stars, black holes, multi-messenger astrophysics, data analysis

PACS: 04.30.-w, 04.30.Tv

1. Motivation and significance

Gravitational waves were originally predicted by Einstein in 1916 [1] as a consequence of general relativity, which describes gravity as the warping of space and time caused by mass and energy [2]. Two extremely massive objects orbiting one another e.g., black holes or neutron stars, warp space dynamically and send ripples across the universe that can be observed here on Earth. As they pass by, gravitational waves stretch and squeeze the space around Earth by less than the width of an atom compared to the Earth's diameter. Due to their tiny effect on scientific instruments, gravitational waves were not observed until 100 years after their initial prediction. Technological advances in laser interferometry led to the discovery of gravitational waves from a merging binary black hole in 2015 [3]. This watershed moment was made possible by the Advanced Laser Interferometer Gravitational-wave Observatory (LIGO) [4] and the scientists of the LIGO and Virgo Collaborations.

Advanced LIGO and Advanced Virgo [4, 5] are the currently operating worldwide network of kilometer-scale laser interferometric gravitational wave observatories which have measured gravitational wave signals. These detectors provide a new way to observe our Universe and enable a vast amount of new science. LIGO's observations have already deepened our understanding of the populations of compact objects such as neutron stars and black holes [6, 7], and they have also offered new tests of fundamental physics [8, 9, 10, 11]. The strong gravity regime probed by compact binary mergers is a laboratory for novel tests of general relativity, and a joint observation of gravitational waves along with electromagnetic waves [12] has taught us how matter behaves in the most extreme conditions [13, 14].

The science made possible by LIGO and Virgo is reliant on measuring miniscule changes in the arm lengths of the interferometers known as strain. The perturbations caused by incident gravitational waves manifest themselves as variations in the intensity of laser light output. *Detector calibration* aims to accurately map the intensity of the output to differential changes in the arm length through real-time signal-processing. The calibrated strain




Figure 1: Gravitational wave infographic. Gravitational wave data is time series, audio frequency data that is noise dominated. GstLAL identifies signals consistent with the predictions of general relativity as measured by multiple gravitational wave detectors and assesses the probability that these signals come from merging neutron stars and/or black holes in near real-time.

data contains the encoded properties of the astrophysical systems that produce gravitational waves. The analysis of this data is complicated by the presence of a vast array of transient noise sources. *Detector characterization* aims to quantify departures from stationary noise to identify times where instrumental issues are so severe that the data should not be analyzed or there may be a coupling between environmental sensors (such as seismometers) and the gravitational wave strain data. Once data is calibrated and assessed

for quality, it is analyzed by a host of detection algorithms to identify potential gravitational wave signals. In many cases the signals are invisible to the naked-eye in raw data and discovering them requires sophisticated techniques that may involve checking millions of models against each segment of data. All three of these activities require substantial cyberinfrastructure. The GstLAL software framework [15] was initially designed to support low-latency compact binary searches to facilitate multi-messenger astronomy, but since its conception it has grown to be a key component of the software used to produce accurately calibrated strain data [16], and recently it has contributed to detector characterization efforts [17, 18]. The GstLAL framework is now contributing key cyberinfrastructure to all three of these key aspects of gravitational wave data analysis. This paper will describe how the GstLAL software is used in gravitational-wave searches [19, 20], provide examples, and describe the history and impact of GstLAL on gravitational wave discovery.

2. Software description

Gravitational wave strain data quantifies how the distance between two points will change as a gravitational wave passes. The current gravitational wave observatories are sensitive to changes in strain and measure the stretching and squeezing of space as a function of time. The Advanced LIGO [4] and Advanced Virgo [5] gravitational wave detectors are most sensitive to strain frequencies between 10Hz–10kHz, which is remarkably close to the frequency range of the human ear [21]. For this reason, there is a close connection between the analysis of gravitational wave data and the analysis of audio data. Indeed, techniques such as low pass filtering, high pass filtering, channel mixing, and gating apply equally well to both audio processing and gravitational wave data processing, which provides the motivation for basing GstLAL on Gstreamer [22].

Gstreamer [22] is an open-source, cross-platform multimedia processing framework designed to execute audio and video processing graphs organized into three basic elements: sources, filters and sinks, which are provided by dynamically loaded plugins. A valid Gstreamer graph, called a pipeline, connects elements together ensuring that the capabilities of each element are satisfied. The data are passed along in *buffers* that store both the memory location of the raw data as well as rich metadata. Pipelines can be used to construct complex workflows and scale to thousands of elements. GstLAL combines standard Gstreamer signal processing elements with custom elements to analyze LIGO strain data.

The GstLAL software began development in 2008 through the exploration of novel techniques for filtering gravitational wave data [23]. It derives its




Figure 2: A basic Gstreamer graph. Data starts at a source, “src”, e.g., a file on disk or a network socket, and then is passed through a filter element, “filter 1”, which transforms the data, e.g., by performing a low pass filter. A second data stream starts from “control src” and the output of filter 1 is moderated by a gate controlled by the state of “control src”. The output of the gate is filtered through “filter 2” and sent to a sink which could be another file on disk or a network socket.

name from “Gstreamer wrappings of the LIGO Algorithm Library¹ (LAL)”. GstLAL began to take on its modern form by 2009 and has been since actively developed as open source software. GstLAL currently resides in the LIGO Scientific Collaboration hosted GitLab instance at <https://git.ligo.org/lscsoft/gstlal> [15].

GstLAL is primarily a mix of Python and C with contributions from 75 authors distributed across North America, Europe, Asia and Australia. The master branch currently has over 13,000 commits and 250,000 lines of code. GstLAL is released under the GPLv2 license with 44 distinct releases since 2011 [24]. In 2012, code solely used for gravitational wave searches for compact binaries was split off into its own package: GstLAL-Inspiral. GstLAL-Inspiral has had 45 distinct releases since then. In 2014, code used for gravitational wave burst detection was split into its own package, GstLAL-Burst, with nine distinct releases. And finally, in 2014 code used primarily for LIGO strain data calibration was split into its own package, GstLAL-Calibration, with 60 releases. In addition to tar-ball releases, RedHat and Debian compatible packages were produced for the LIGO Data Grid reference platforms [25].

At present, Docker containers with the full GstLAL/LALSuite software stack are built and distributed using the LIGO Container Registry [26]. The containers are built on top of the CentOS-based Scientific Linux 7, which

¹<https://git.ligo.org/lscsoft/lalsuite>

currently serves as the reference operating system on the LIGO Data Grid. Binary executables are linked against Intel’s high-performance Math Kernel Library (MKL), and compiled to leverage Advanced Vector Extensions. Optimized versions of the GstLAL software stack tuned at the compiler-level to best leverage the native features of the local computing environment are often custom-built by users. Previous studies have demonstrated a $\gtrsim 2$ times increase in overall code throughput as a result of software tuning. GstLAL is also available through CondaForge [27].

2.1. Software Architecture

Prior to 2015, gravitational waves had not been directly observed [3]. Although analysis techniques had been studied for decades [28, 29], the character of gravitational wave data evolved with the interferometers detectors during the initial operation of LIGO and Virgo from 2002–2010 [30, 31, 32, 33, 34, 35]. Therefore, we aimed to make the GstLAL software modular and easy to adapt to the challenges of Advanced LIGO and Virgo data.

The joint observation of gravitational waves and electromagnetic signals, known as multi-messenger astronomy, was a significant goal for Advanced LIGO and Advanced Virgo. In this scenario, gravitational-wave observations were expected to be followed up by observations with telescopes across the electromagnetic spectrum hoping to catch a short-lived transient light source. Discovering gravitational waves quickly is critical because electromagnetic counterparts may quickly fade. GstLAL was designed to offer analysts an extremely short time-to-solution to help ensure that electromagnetic counterparts could be observed quickly.

The key design principles of GstLAL are:

Plugin-based: Libraries within GstLAL provide GStreamer plugins to perform gravitational wave specific signal processing tasks. These are mixed together with stock GStreamer plugins to produce gravitational wave analysis pipelines. Plugins provide elements, which are the building blocks of signal processing workflows. These elements can be ordered in multiple ways with minimal coding effort which allows for quick exploratory work and development of new methods.

Streaming: The goal of GStreamer is to provide ultra low latency signal processing suitable for audio and video playback and editing. GstLAL pipelines typically work with streaming data buffers $\lesssim 1$ s in duration.

Event-driven: GstLAL analysis pipelines are designed to run continuously as data is collected. Each application runs an event loop which controls both application level operations as well as settings within a given

plugin. This allows for the control of program behavior to be altered while the application is running. Dynamic program control is facilitated through embedding the microservices framework Bottle [36]. Simple http APIs push information to the program, retrieve information or alter the program’s behavior.

Scalable: The GstLAL framework is designed both to scale to dozens of cores on a single computer using the multi-threading provided by Gstreamer and to scale to thousands of cores across a compute cluster leveraging HTCondor directed acyclic graph (DAG) scheduling [37].

The GstLAL project is currently comprised of five distinct packages. 1) `gstlal`, 2) `gstlal-ugly`, 3) `gstlal-inspiral`, 4) `gstlal-calibration` and 5) `gstlal-burst`, all of which are described below.

2.2. `gstlal` package

The `gstlal` package curates a set of core plugins, functions and applications that are used by nearly every analysis workflow developed within GstLAL. The `gstlal` package is a dependency for all of the remaining packages which are described in subsequent sections. The `gstlal` package provides Gstreamer elements for Finite-Impulse-Response filtering (`lal_firbank`), $N \rightarrow M$ -channel matrix operations (`lal_matrixmixer`), data whitening (`lal_whiten`), and data gating (`lal_gate`). The `gstlal` package also provides basic python APIs for building Gstreamer pipelines in the module `pipeparts`, basic data access routines in the module `datasource`, and a base class for event handling in the module `simplehandler`.

2.3. `gstlal-ugly` package

The `gstlal-ugly` package is an incubator package for software that is in development. Eventually all `gstlal-ugly` software is migrated to one of the other packages.

2.4. `gstlal-inspiral` package

The primary purpose of the `gstlal-inspiral` package is to house the GstLAL-based search for compact binaries [19, 20], which centers around the application `gstlal_inspiral`. The GstLAL-based compact binary pipeline was created to make near real-time gravitational-wave detections and aimed to one day detect electromagnetically bright systems before coalescence [38, 39].

The GstLAL-based compact binary search is a matched-filter search that incorporates efficient time-domain filtering [23, 40, 41, 42, 43, 44, 38, 45]

of a set of template waveforms that match the gravitational wave signals of merging black holes and neutron stars [46]. LIGO and Virgo detectors are prone to bursts of nonstationary noise called glitches [47] and determining the difference between gravitational waves and glitches is well suited for many classification algorithms. GstLAL Inspiral implements a classification scheme that is a hybrid of hypothesis testing techniques with some elements coming from machine learning approaches. The classifier is an approximate likelihood ratio comprised of many terms which began as a custom implementation of Naive Bayes classification [48] applied to gravitational wave searches [49]. It was realized early in the project that two things were apparent. First, it wasn't practical to treat the classifier as entirely data driven relying purely on training sets. Training sets to adequately classify the full parameter space were too expensive to produce. Second, correlations between some parameters had to be tracked in order to classify well. The first point was addressed by developing semi-analytic models to describe parts of the data [50] and the second point was addressed by factoring the multi-dimensional likelihood ratios into groups of lower dimensional, but not one dimensional, distributions [51, 52].

The GstLAL-based compact binary search has two modes. The first is a near-real-time, “low-latency” mode that discovers and reports compact binaries within tens of seconds of the signal arriving at Earth. The second is an “offline” mode that efficiently processes data in batch jobs where time-to-solution is not as important as computational efficiency and reproducibility. Although both modes share $\gtrsim 95\%$ of the same code, their behavior and design are very different in order to address the differing concerns of real-time vs. batch processing.

The low-latency mode is a collection of typically $\mathcal{O}(1000)$ microservices that communicate (modestly) with one another asynchronously through http using python Bottle, through an Apache Kafka queue, and through a shared file system. Each one of these microservices processes a portion of the nearly 2 million models used in the current low-latency compact binary search. The low-latency workflow is designed to be fault tolerant. If a job dies, another is restarted to take its place. Since information is exchanged asynchronously and there is no guarantee of job success, the behavior in this mode is non-deterministic. In contrast, the offline mode has a fully deterministic execution which can be reproduced to floating point precision. The determinism is imposed by organizing each job in a directed acyclic graph (DAG) using HTCondor.

2.5. `gstlal-burst` package

The GstLAL-based burst package is a collection of utilities intended to search for gravitational-wave sources other than compact binaries as well as non-astrophysical noise transients. One of the recent developments is a pipeline searching for cosmic strings, which are hypothetical objects considered to have formed in the early universe. The pipeline uses time-domain stream-based signal processing algorithms, along with a classification scheme using parameters specific to the search. The algorithms used are mostly in common with the `gstlal-inspiral` package, but simplified due to the smaller number of templates required for matched filtering.

In addition, `gstlal-burst` provides utilities to identify and extract features from non-Gaussian noise transients in near real-time ($O(5s)$) via the Stream-based Noise Acquisition and eXtraction, or SNAX toolkit. The SNAX toolkit also leverages time-domain signal processing, but instead utilizes a sine-Gaussian basis to identify the presence of and extract features from many types of non-Gaussian noise in strain and auxiliary data. Its main data product is multivariate time-series data containing the extracted features, including SNR and phase information as well as the waveform parameters of interest.

2.6. `gstlal-calibration` package

The `gstlal-calibration` package houses the unique software used for calibration of the LIGO strain data. Software in the `gstlal-calibration` package produces the only official LIGO strain data product used in all subsequent analysis. Calibration of LIGO strain data involves standard signal processing and digital filtering techniques in order to derive the differential arm motion observed in the LIGO detectors from the detector’s digital readouts [53, 16, 54]. Many of the signal processing and digital filtering plugins used by the calibration pipeline `gstlal_compute_strain` are housed in the `gstlal` or `gstlal-ugly` packages. A few plugins unique to the calibration process as well as calibration-specific python APIs are housed in the `gstlal-calibration` package.

Much like the GstLAL-based compact binary pipeline, the LIGO calibration pipeline is built to operate in two modes: a “low-latency” mode and an “offline” mode. The low-latency LIGO calibration pipeline operates on hardware located physically at the two LIGO detector sites, LIGO Hanford in Hanford, WA and LIGO Livingston in Livingston, LA. This pipeline produces calibrated LIGO data and a bit-wise state-vector that indicates the fidelity of the calibrated data within $O(5s)$ for each detector at the respective detector sites. The low-latency calibration process involves using a combination of digital filtering performed in the LIGO front-end computers, which

are directly connected to the LIGO detectors and employ the CDS Real-time Code Generator (RCG) core software [55], and further digital filtering and processing performed by the GstLAL calibration software running on non-front-end hardware located at the LIGO detector sites. This two-step process takes advantage of the access the front-end computing system has to the installed detector filters and models and the advanced stream-based filtering techniques housed in the GstLAL software packages [16].

There is often a need to re-calibrate the strain data after the initial low-latency data calibration in order to improve calibration accuracy based on more sophisticated modeling or to remove systematic errors present in the low-latency calibrated strain data [16, 56, 57]. The re-calibration of LIGO data is performed using the offline mode of the `gstlal_compute_strain` pipeline. In this mode, the entire calibration process is performed by software housed in the GstLAL software packages. The offline calibrated data is processed in batch jobs using HTCondor in order to optimize computational efficiency and is completely reproducible to floating-point precision. All analyses derived from LIGO strain data use the calibrated data produced either by the low-latency GstLAL calibration pipeline or the offline GstLAL calibration pipeline.

3. Illustrative Examples

3.1. Example Gstreamer pipeline with GstLAL

The following example was run on a newly instantiated CentOS 7 64 bit virtual machine with miniconda [58] installed by doing:

```

1
2 wget https://repo.anaconda.com/miniconda/Miniconda3-latest-
   Linux-x86_64.sh
3 bash Miniconda3-latest-Linux-x86_64.sh
4 conda create -n myenv python=2.7
5 conda activate myenv
6 conda install -c conda-forge gstlal-inspiral=1.7.3

```

To verify that it works, try:

```

1
2 $ gst-inspect-1.0 gstlalinspiral | head -n4
3 Plugin Details:
4 Name                gstlalinspiral
5 Description          Various bits of the LIGO Algorithm
   Library wrapped in  gstreamer elements
6 Filename             <your miniconda path>/lib/
   gstreamer-1.0/libgstgstlalinspiral.so

```

Next we will try a very simple Gstreamer pipeline that uses two GstLAL elements: `lal_peak` and `lal_nxydump`, along with additional Gstreamer elements to construct a pipeline that generates 10 Hz Gaussian, white noise, finds the peak sample every second and streams the result to the terminal screen as ASCII text. It is possible to construct simple pipelines such as this without any code using the Gstreamer tool `gst-launch` [59]:

```

1
2 $ gst-launch-1.0 audiotestsrc wave=9 ! capsfilter caps=audio/
   x-raw,rate=10 ! lal_peak n=10 ! lal_nxydump ! filesink
   location=/dev/stdout

```

The first element, `audiotestsrc` is a Gstreamer element that can provide many test signals. The `wave=9` property sets it to be unit variance white noise. The second element, `capsfilter` specifies that we want the format of the output to be floating point audio data with a sample rate of 10 Hz. Next, `lal_peak` is the first GstLAL element. In this example it is configured to find the largest absolute value of the signal every 10 sample points. `lal_nxydump` is the second GstLAL element which converts the time-series data to two column ASCII text. Finally, `filesink` dumps the ASCII output to standard out. You should see the following output (with variations caused by the fact that the data is random):

```

1
2 $ gst-launch-1.0 audiotestsrc wave=9 ! capsfilter caps=audio/
   x-raw,rate=10 ! lal_peak n=10 ! lal_nxydump ! filesink
   location=/dev/stdout | head -n 15
3 Setting pipeline to PAUSED ...
4 Pipeline is PREROLLING ...
5 Pipeline is PREROLLED ...
6 Setting pipeline to PLAYING ...
7 New clock: GstSystemClock
8 0.000000000    0
9 0.100000000    0
10 0.200000000    0
11 0.300000000    0
12 0.400000000    0
13 0.500000000    0
14 0.600000000    0.95523328
15 0.700000000    0
16 0.800000000    0
17 0.900000000    0

```

you can see that the maximum sample point was chosen in the 10 sample interval on line 13. Other values are set to 0.

`gst-launch` is useful tool for quickly testing a simple pipeline, or debugging, however it is not suitable for writing large applications with many

elements or situations where program control is exposed dynamically to the user. For building applications, GstLAL relies on the Python bindings for Gstreamer and adds a substantial amount of gravitational wave specific application code written in python. An example of the pipeline above written in the style of GstLAL applications is below.

```

1
2 # boiler plate Gstreamer imports
3 import gi
4 gi.require_version('Gst', '1.0')
5 from gi.repository import GObject, Gst
6 GObject.threads_init()
7 Gst.init(None)
8
9 # Gstlal imports
10 from gstlal import datasource
11 from gstlal import pipeparts
12 from gstlal import simplehandler
13
14 # initialize an event loop, a pipeline and an event handler
15 mainloop = GObject.MainLoop()
16 pipeline = Gst.Pipeline("softwarex_demo")
17 handler = simplehandler.Handler(mainloop, pipeline)
18
19 src = pipeparts.mkaudiotestsrc(pipeline, wave = 9)
20 src = pipeparts.mkcapsfilter(pipeline, src, caps = "audio/x-
    raw, rate=10")
21 src = pipeparts.mkpeak(pipeline, src, n = 10)
22 src = pipeparts.mknxydumpsink(pipeline, src, "/dev/stdout")
23
24 if pipeline.set_state(Gst.State.PLAYING) == Gst.
    StateChangeReturn.FAILURE:
25     raise RuntimeError("pipeline failed to enter PLAYING
    state")
26 mainloop.run()

```

We have found that, using python to procedurally build Gstreamer graphs, we can construct enormous pipelines containing tens of thousands of distinct elements. A prime example of this is our workhorse signal processing pipeline used for discovering compact binary mergers as described in the next section.

3.2. Compact binary searches

`Makefile.softwarex_test` provides an example of the general workflow involved in offline gravitational wave analyses, which primarily rely on the `gstlal` and `gstlal-inspiral` packages. The miniconda installation of GstLAL will run this example in ~ 30 minutes on a single machine. Production

level compact binary searches analyzing data from the three advanced interferometers, on the other hand, take ~ 1 week when distributed over $\mathcal{O}(1000)$ core computing clusters with optimized software builds. The test Makefile is entirely self-contained; the target and dependency relationships and brief comments within describe the workflow. Although the structure presented in this test is linear in nature, full scale searches for compact binaries are heavily parallelized to take advantage of DAG scheduling. The requisite commands to run this test analysis are shown below.

```

1
2 mkdir workflow-test && cd workflow-test
3 export LAL_PATH=${CONDA_PREFIX} GSTLAL_FIR_WHITEN=0 TMPDIR=/
   tmp
4 wget https://dcc.ligo.org/public/0168/P2000195/004/Makefile.
   softwarex_test
5 make -f Makefile.softwarex_test

```

4. Impact

GstLAL has played an integral role in the history of gravitational wave detections, and has participated in gravitational wave searches since S5 [60]. Although GstLAL was designed specifically for near-real-time applications, the low-latency pipeline was prevented from searching for binary black holes at the start of Advanced LIGO’s first observing run (O1). The pipeline was only allowed to use a template bank sensitive to electromagnetically binary neutron stars and neutron star - black hole binaries. The software was fully capable of detecting binary black holes in near-real-time. The LIGO collaboration desired blind binary black hole (BBH) analyses, and since BBH systems are not expected to produce electromagnetic radiation there was no perceived need to detect them in near-real-time. The restriction on the allowed template bank rendered GstLAL unable to detect GW150914 in low-latency, though it was one of two matched filter pipelines used to analyze archival data and validate the event [3]. GW150914 was initially detected in low-latency by the weakly-modeled burst pipeline CWB [61], which demonstrated that there could be no truly blind analysis while low-latency burst pipelines produced alerts. As a result, in late 2015 the GstLAL pipeline was approved to include BBHs with total mass $\lesssim 100M_{\odot}$ in its low-latency configuration. GstLAL quickly demonstrated its ability to recover BBHs in low-latency; it became the first matched-filter pipeline using waveforms based on general relativity to make a near real-time detection of a compact binary with the discovery of GW151226 [62].

Development work between Advanced LIGO’s first and second observing runs focused on enabling single detector discoveries and incorporating

data from Virgo in the analysis. These efforts were rewarded by August 2017 as GstLAL became the first three-detector matched-filter search in the Advanced LIGO era and, more notably, the first (and to date, the only) gravitational-wave detection pipeline to observe a binary neutron star merger in low-latency [63, 64]. Although both Advanced LIGO interferometers and the Advanced Virgo interferometer were operating at the time of GW170817, it was initially observed in a single interferometer. This marked the first single detector observation of a gravitational wave, and the autonomous identification of the candidate enabled rapid offline follow-up of the candidate within the LIGO-Virgo collaboration.

Advanced LIGO’s third observing run (O3) marked the beginning of open public alerts (OPA)[65]. For the first time, candidates with false-alarm-rates below 1 per month² were made public at the time of discovery. Although the first public alert was distributed in error by the collaboration [66], the identification of binary black hole candidate S90408an [67] marked the successful start of the era of automated public alerts. At first only candidates appearing in two or more detectors were approved for automated release, but GW170817 had already demonstrated the importance of single detector searches. Indeed, two weeks into the third observing run GstLAL was the only pipeline to detect GW190425 in near-real-time [63, 68], further highlighting the necessity of single detector searches.

Two months into the observing run, GstLAL became the only pipeline approved to release single detector candidates as OPAs. This was a high risk, high reward endeavor. Matched filter searches have traditionally been able to suppress the background by demanding coincidence across interferometers; single detector candidates do not benefit from this effect and can therefore be more susceptible to short term noise transients. Unfortunately, this resulted in several retractions throughout O3 as the GstLAL team worked on ways to mitigate the effects of noise transients in single detectors. By the end of the second half of the observing run, pipeline tuning had reduced the rate of retractions.

GstLAL has contributed to all gravitational-wave discoveries published by the LIGO Scientific Collaboration [68, 6, 69, 70, 71, 72, 62, 3], but it has also contributed to searches for as yet undetected sources. Sub-solar mass and intermediate mass black holes both pose problems for conventional models of stellar evolution, and GstLAL has directly contributed to searches of both [73, 74, 75, 76]. Although these searches have not yet yielded any

²after applying a trials factor to account for the number of concurrently running searches.

detections, the null results have been able to place strict limits on the abundance of such objects and have also provided the tightest limit to date on a primordial black hole model of the dark matter.

5. Conclusions

The GstLAL library has significantly impacted the progress of gravitational-wave astrophysics, not only via compact binary searches, but also through contributions to detector calibration and characterization efforts. The low-latency GstLAL based inspiral pipeline was instrumental in the first multimessenger discovery with gravitational waves, and strives to lead the march towards more remarkable observations with ground-based interferometers.

6. Conflict of interest

No conflict of interest exists: we wish to confirm that there are no known conflicts of interest associated with this publication and there has been no significant financial support for this work that could have influenced its outcome.

Acknowledgments

Funding for this work was provided by the National Science Foundation through awards: PHY-1454389, OAC-1642391, PHY-1700765, OAC-1841480, PHY-1607178, and PHY-1847350. Funding for this project was provided by the Charles E. Kaufman Foundation of The Pittsburgh Foundation. Computations for this research were performed on the Pennsylvania State University’s Institute for Computational and Data Sciences Advanced CyberInfrastructure (ICDS-ACI) and VM hosting. We are grateful for computational resources provided by the Leonard E Parker Center for Gravitation, Cosmology and Astrophysics at the University of Wisconsin-Milwaukee. Computing support was provided by the LIGO Laboratory through National Science Foundation grant PHY-1764464. GstLAL relies on many other open source software libraries; we gratefully acknowledge the development and support of NumPy [77], SciPy [78], PyGTK [79], PyGST [80], Bottle [36], Kafka [81], Fftw3F [82], Intel MKL [83], GLib2 [84], GNU Scientific Library [85], and GWpy [86]. The authors gratefully acknowledge the LIGO-Virgo-Kagra collaboration for support, review, and valuable critiques throughout various stages of development of the GstLAL library. We are especially thankful for collaborations within the Compact Binary Coalescence working group.

References

- [1] A. Einstein, Approximative Integration of the Field Equations of Gravitation, *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* 1916 (1916) 688–696.
- [2] A. Einstein, The Field Equations of Gravitation, *Sitzungsber. Preuss. Akad. Wiss. Berlin (Math. Phys.)* 1915 (1915) 844–847.
- [3] B. P. Abbott, et al., Observation of Gravitational Waves from a Binary Black Hole Merger, *Phys. Rev. Lett.* 116 (6) (2016) 061102. [arXiv:1602.03837](#), [doi:10.1103/PhysRevLett.116.061102](#).
- [4] J. Aasi, et al., Advanced LIGO, *Class. Quant. Grav.* 32 (2015) 074001. [arXiv:1411.4547](#), [doi:10.1088/0264-9381/32/7/074001](#).
- [5] F. Acernese, et al., Advanced Virgo: a second-generation interferometric gravitational wave detector, *Class. Quant. Grav.* 32 (2) (2015) 024001. [arXiv:1408.3978](#), [doi:10.1088/0264-9381/32/2/024001](#).
- [6] B. P. Abbott, et al., 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. X* 9 (3) (2019) 031040. [arXiv:1811.12907](#), [doi:10.1103/PhysRevX.9.031040](#).
- [7] B. P. Abbott, et al., Binary Black Hole Population Properties Inferred from the First and Second Observing Runs of Advanced LIGO and Advanced Virgo, *Astrophys. J.* 882 (2) (2019) L24. [arXiv:1811.12940](#), [doi:10.3847/2041-8213/ab3800](#).
- [8] B. P. Abbott, et al., A gravitational-wave standard siren measurement of the Hubble constant, *Nature* 551 (7678) (2017) 85–88. [arXiv:1710.05835](#), [doi:10.1038/nature24471](#).
- [9] B. P. Abbott, et al., Tests of General Relativity with GW170817, *Phys. Rev. Lett.* 123 (1) (2019) 011102. [arXiv:1811.00364](#), [doi:10.1103/PhysRevLett.123.011102](#).
- [10] B. P. Abbott, et al., Tests of General Relativity with the Binary Black Hole Signals from the LIGO-Virgo Catalog GWTC-1, *Phys. Rev. D* 100 (10) (2019) 104036. [arXiv:1903.04467](#), [doi:10.1103/PhysRevD.100.104036](#).

- [11] B. Abbott, R. Abbott, T. Abbott, S. Abraham, F. Acernese, K. Ackley, C. Adams, R. Adhikari, V. Adya, C. Affeldt, et al., A gravitational-wave measurement of the hubble constant following the second observing run of advanced ligo and virgo, arXiv preprint arXiv:1908.06060.
- [12] B. P. Abbott, et al., Multi-messenger Observations of a Binary Neutron Star Merger, *Astrophys. J.* 848 (2) (2017) L12. arXiv:1710.05833, doi:10.3847/2041-8213/aa91c9.
- [13] B. P. Abbott, et al., GW170817: Measurements of neutron star radii and equation of state, *Phys. Rev. Lett.* 121 (16) (2018) 161101. arXiv:1805.11581, doi:10.1103/PhysRevLett.121.161101.
- [14] M. Nicholl, et al., The Electromagnetic Counterpart of the Binary Neutron Star Merger LIGO/VIRGO GW170817. III. Optical and UV Spectra of a Blue Kilonova From Fast Polar Ejecta, *Astrophys. J.* 848 (2) (2017) L18. arXiv:1710.05456, doi:10.3847/2041-8213/aa9029.
- [15] GstLAL software: git.ligo.org/lscsoft/gstlal.
- [16] A. Viets, et al., Reconstructing the calibrated strain signal in the Advanced LIGO detectors, *Class. Quant. Grav.* 35 (9) (2018) 095015. arXiv:1710.09973, doi:10.1088/1361-6382/aab658.
- [17] R. Essick, P. Godwin, C. Hanna, E. Katsavounidis, R. Vaulin, L. Blackburn, iDQ: Statistical Inference for Non-Gaussian Noise with Auxiliary Degrees of Freedom in Gravitational-Wave Detectors, In preparation.
- [18] P. Godwin, C. Hanna, R. Essick, , In preparation.
- [19] C. Messick, et al., Analysis Framework for the Prompt Discovery of Compact Binary Mergers in Gravitational-wave Data, *Phys. Rev. D* 95 (4) (2017) 042001. arXiv:1604.04324, doi:10.1103/PhysRevD.95.042001.
- [20] S. Sachdev, et al., The GstLAL Search Analysis Methods for Compact Binary Mergers in Advanced LIGO's Second and Advanced Virgo's First Observing Runs arXiv:1901.08580.
- [21] H. Olson, *Music, Physics and Engineering*, Dover Books, Dover Publications, 1967.
URL <https://books.google.com/books?id=RUDTFBbb7jAC>
- [22] Gstreamer : <https://gstreamer.freedesktop.org/>.

- [23] K. Cannon, A. Chapman, C. Hanna, D. Keppel, A. C. Searle, A. J. Weinstein, Singular value decomposition applied to compact binary coalescence gravitational-wave signals, *Phys. Rev. D* 82 (2010) 044025. [arXiv:1005.0012](#), [doi:10.1103/PhysRevD.82.044025](#).
- [24] Ligo software: <http://software.ligo.org>.
- [25] Ligo packages: <http://software.ligo.org/lscsoft/>.
- [26] Container registry - lscsoft/gstlal: https://git.ligo.org/lscsoft/gstlal/container_registry.
- [27] Gstlal inspiral :: Anaconda cloud: <https://anaconda.org/conda-forge/gstlal-inspiral>.
- [28] L. S. Finn, D. F. Chernoff, Observing binary inspiral in gravitational radiation: One interferometer, *Phys. Rev. D* 47 (1993) 2198–2219. [arXiv:gr-qc/9301003](#), [doi:10.1103/PhysRevD.47.2198](#).
- [29] B. Allen, W. G. Anderson, P. R. Brady, D. A. Brown, J. D. E. Creighton, FINDCHIRP: An Algorithm for detection of gravitational waves from inspiraling compact binaries, *Phys. Rev. D* 85 (2012) 122006. [arXiv:gr-qc/0509116](#), [doi:10.1103/PhysRevD.85.122006](#).
- [30] B. Abbott, et al., Analysis of LIGO data for gravitational waves from binary neutron stars, *Phys. Rev. D* 69 (2004) 122001. [arXiv:gr-qc/0308069](#), [doi:10.1103/PhysRevD.69.122001](#).
- [31] B. Abbott, et al., Search for gravitational waves from galactic and extragalactic binary neutron stars, *Phys. Rev. D* 72 (2005) 082001. [arXiv:gr-qc/0505041](#), [doi:10.1103/PhysRevD.72.082001](#).
- [32] B. Abbott, et al., Search for gravitational waves from binary inspirals in S3 and S4 LIGO data, *Phys. Rev. D* 77 (2008) 062002. [arXiv:0704.3368](#), [doi:10.1103/PhysRevD.77.062002](#).
- [33] B. P. Abbott, et al., Search for Gravitational Waves from Low Mass Binary Coalescences in the First Year of LIGO’s S5 Data, *Phys. Rev. D* 79 (2009) 122001. [arXiv:0901.0302](#), [doi:10.1103/PhysRevD.79.122001](#).
- [34] B. P. Abbott, et al., Search for Gravitational Waves from Low Mass Compact Binary Coalescence in 186 Days of LIGO’s fifth Science Run, *Phys. Rev. D* 80 (2009) 047101. [arXiv:0905.3710](#), [doi:10.1103/PhysRevD.80.047101](#).

- [35] J. Abadie, et al., Search for Gravitational Waves from Low Mass Compact Binary Coalescence in LIGO's Sixth Science Run and Virgo's Science Runs 2 and 3, *Phys. Rev. D* 85 (2012) 082002. [arXiv:1111.7314](#), [doi:10.1103/PhysRevD.85.082002](#).
- [36] Bottle: Python web framework: <https://bottlepy.org/docs/dev/>.
- [37] Htcondor high throughput computing: <https://research.cs.wisc.edu/htcondor/>.
- [38] K. Cannon, et al., Toward Early-Warning Detection of Gravitational Waves from Compact Binary Coalescence, *Astrophys. J.* 748 (2012) 136. [arXiv:1107.2665](#), [doi:10.1088/0004-637X/748/2/136](#).
- [39] S. Sachdev, et al., An early warning system for electromagnetic follow-up of gravitational-wave events [arXiv:2008.04288](#).
- [40] K. Cannon, C. Hanna, D. Keppel, A. C. Searle, Composite gravitational-wave detection of compact binary coalescence, *Phys. Rev. D* 83 (2011) 084053. [arXiv:1101.0584](#), [doi:10.1103/PhysRevD.83.084053](#).
- [41] K. Cannon, C. Hanna, D. Keppel, Efficiently enclosing the compact binary parameter space by singular-value decomposition, *Phys. Rev. D* 84 (2011) 084003. [arXiv:1101.4939](#), [doi:10.1103/PhysRevD.84.084003](#).
- [42] K. Cannon, C. Hanna, D. Keppel, Interpolating compact binary waveforms using the singular value decomposition, *Phys. Rev. D* 85 (2012) 081504. [arXiv:1108.5618](#), [doi:10.1103/PhysRevD.85.081504](#).
- [43] K. Cannon, J. Emberson, C. Hanna, D. Keppel, H. Pfeiffer, Interpolation in waveform space: enhancing the accuracy of gravitational waveform families using numerical relativity, *Phys. Rev. D* 87 (4) (2013) 044008. [arXiv:1211.7095](#), [doi:10.1103/PhysRevD.87.044008](#).
- [44] R. Smith, K. Cannon, C. Hanna, D. Keppel, I. Mandel, Towards Rapid Parameter Estimation on Gravitational Waves from Compact Binaries using Interpolated Waveforms, *Phys. Rev. D* 87 (12) (2013) 122002. [arXiv:1211.1254](#), [doi:10.1103/PhysRevD.87.122002](#).
- [45] L. Tsukada, K. Cannon, C. Hanna, D. Keppel, D. Meacher, C. Messick, Application of a Zero-latency Whitening Filter to Compact Binary Coalescence Gravitational-wave Searches, *Phys. Rev. D* 97 (10) (2018) 103009. [arXiv:1708.04125](#), [doi:10.1103/PhysRevD.97.103009](#).

- [46] D. Mukherjee, et al., The GstLAL template bank for spinning compact binary mergers in the second observation run of Advanced LIGO and Virgo [arXiv:1812.05121](https://arxiv.org/abs/1812.05121).
- [47] B. P. Abbott, et al., Characterization of transient noise in Advanced LIGO relevant to gravitational wave signal GW150914, *Class. Quant. Grav.* 33 (13) (2016) 134001. [arXiv:1602.03844](https://arxiv.org/abs/1602.03844), doi:10.1088/0264-9381/33/13/134001.
- [48] H. Zhang, The optimality of naive bayes, *AA* 1 (2) (2004) 3.
- [49] K. C. Cannon, A Bayesian coincidence test for noise rejection in a gravitational-wave burst search, *Class. Quant. Grav.* 25 (2008) 105024. doi:10.1088/0264-9381/25/10/105024.
- [50] C. Hanna, et al., Fast evaluation of multi-detector consistency for real-time gravitational wave searches, *Phys. Rev. D* 101 (2) (2020) 022003. [arXiv:1901.02227](https://arxiv.org/abs/1901.02227), doi:10.1103/PhysRevD.101.022003.
- [51] K. Cannon, C. Hanna, J. Peoples, Likelihood-Ratio Ranking Statistic for Compact Binary Coalescence Candidates with Rate Estimation [arXiv:1504.04632](https://arxiv.org/abs/1504.04632).
- [52] K. Cannon, C. Hanna, D. Keppel, Method to estimate the significance of coincident gravitational-wave observations from compact binary coalescence, *Phys. Rev. D* 88 (2) (2013) 024025. [arXiv:1209.0718](https://arxiv.org/abs/1209.0718), doi:10.1103/PhysRevD.88.024025.
- [53] B. P. Abbott, et al., Calibration of the Advanced LIGO detectors for the discovery of the binary black-hole merger GW150914, *Phys. Rev. D* 95 (6) (2017) 062003. [arXiv:1602.03845](https://arxiv.org/abs/1602.03845), doi:10.1103/PhysRevD.95.062003.
- [54] D. Tuyenbayev, S. Karki, J. Betzwieser, C. Cahillane, E. Goetz, K. Izumi, S. Kandhasamy, J. S. Kissel, G. Mendell, M. Wade, A. J. Weinstein, R. L. Savage, Improving LIGO calibration accuracy by tracking and compensating for slow temporal variations, *Classical and Quantum Gravity* 34 (1) (2016) 015002. doi:10.1088/0264-9381/34/1/015002. URL <https://doi.org/10.1088/0264-9381/34/1/015002>
- [55] R. Bork, K. Thorne, D. Barker, J. Hanks, Real-time code generator (rcg) version 3.0 release notes T1600055.

- [56] C. Cahillane, J. Betzwieser, D. A. Brown, E. Goetz, E. D. Hall, K. Izumi, S. Kandhasamy, S. Karki, J. S. Kissel, G. Mendell, R. L. Savage, D. Tuyenbayev, A. Urban, A. Viets, M. Wade, A. J. Weinstein, Calibration uncertainty for advanced ligo's first and second observing runs, *Phys. Rev. D* 96 (2017) 102001. doi:10.1103/PhysRevD.96.102001. URL <https://link.aps.org/doi/10.1103/PhysRevD.96.102001>
- [57] L. Sun, E. Goetz, J. S. Kissel, J. Betzwieser, S. Karki, A. Viets, M. Wade, D. Bhattacharjee, V. Bossilkov, P. B. Covas, L. E. H. Datrier, R. Gray, S. Kandhasamy, Y. K. Lecoecuche, G. Mendell, T. Mistry, E. Payne, R. L. Savage, A. J. Weinstein, S. Aston, A. Buikema, C. Cahillane, J. C. Driggers, S. E. Dwyer, R. Kumar, A. Urban, Characterization of systematic error in Advanced LIGO calibration, arXiv e-prints (2020) arXiv:2005.02531arXiv:2005.02531.
- [58] Miniconda: <https://docs.conda.io/en/latest/miniconda.html>.
- [59] gst-launch-1.0: <https://gststreamer.freedesktop.org/documentation/tools/gst-launch.html?gi-language=c>.
- [60] M. Wade, Gravitational-wave Science with the Laser Interferometer Gravitational-wave Observatory, Ph.D. thesis, University of Wisconsin-Milwaukee (2015). arXiv:LIGO-P1500068. URL <https://dcc.ligo.org/LIGO-P1500068/public>
- [61] S. Klimenko, et al., Method for detection and reconstruction of gravitational wave transients with networks of advanced detectors, *Phys. Rev. D* 93 (4) (2016) 042004. arXiv:1511.05999, doi:10.1103/PhysRevD.93.042004.
- [62] B. P. Abbott, et al., GW151226: Observation of Gravitational Waves from a 22-Solar-Mass Binary Black Hole Coalescence, *Phys. Rev. Lett.* 116 (24) (2016) 241103. arXiv:1606.04855, doi:10.1103/PhysRevLett.116.241103.
- [63] V. C. LIGO Scientific Collaboration, GCN 24168. [link]. URL <https://gcn.gsfc.nasa.gov/gcn3/24168.gcn3>
- [64] V. C. LIGO Scientific Collaboration, GCN G298048. [link]. URL <https://gcn.gsfc.nasa.gov/other/G298048.gcn3>
- [65] V. C. LIGO Scientific Collaboration, GCN 24045. [link]. URL <https://gcn.gsfc.nasa.gov/gcn3/24045.gcn3>

- [66] V. C. LIGO Scientific Collaboration, GCN 24109. [link].
URL <https://gcn.gsfc.nasa.gov/gcn3/24109.gcn3>
- [67] V. C. LIGO Scientific Collaboration, GCN 24069. [link].
URL <https://gcn.gsfc.nasa.gov/gcn3/24069.gcn3>
- [68] B. P. Abbott, et al., GW190425: Observation of a Compact Binary Coalescence with Total Mass $\sim 3.4M_{\odot}$ [arXiv:2001.01761](#).
- [69] B. P. Abbott, et al., GW170817: Observation of Gravitational Waves from a Binary Neutron Star Inspiral, *Phys. Rev. Lett.* 119 (16) (2017) 161101. [arXiv:1710.05832](#), doi:10.1103/PhysRevLett.119.161101.
- [70] B. P. Abbott, et al., GW170814: A Three-Detector Observation of Gravitational Waves from a Binary Black Hole Coalescence, *Phys. Rev. Lett.* 119 (14) (2017) 141101. [arXiv:1709.09660](#), doi:10.1103/PhysRevLett.119.141101.
- [71] B. P. Abbott, et al., GW170608: Observation of a 19-solar-mass Binary Black Hole Coalescence, *Astrophys. J.* 851 (2) (2017) L35. [arXiv:1711.05578](#), doi:10.3847/2041-8213/aa9f0c.
- [72] B. P. Abbott, et al., GW170104: Observation of a 50-Solar-Mass Binary Black Hole Coalescence at Redshift 0.2, *Phys. Rev. Lett.* 118 (22) (2017) 221101, [Erratum: *Phys. Rev. Lett.* 121, no. 12, 129901 (2018)]. [arXiv:1706.01812](#), doi:10.1103/PhysRevLett.118.221101, 10.1103/PhysRevLett.121.129901.
- [73] B. P. Abbott, et al., Search for Subsolar-Mass Ultracompact Binaries in Advanced LIGO's First Observing Run, *Phys. Rev. Lett.* 121 (23) (2018) 231103. [arXiv:1808.04771](#), doi:10.1103/PhysRevLett.121.231103.
- [74] B. P. Abbott, et al., Search for intermediate mass black hole binaries in the first and second observing runs of the Advanced LIGO and Virgo network, *Phys. Rev. D* 100 (6) (2019) 064064. [arXiv:1906.08000](#), doi:10.1103/PhysRevD.100.064064.
- [75] B. P. Abbott, et al., Search for Subsolar Mass Ultracompact Binaries in Advanced LIGO's Second Observing Run, *Phys. Rev. Lett.* 123 (16) (2019) 161102. [arXiv:1904.08976](#), doi:10.1103/PhysRevLett.123.161102.
- [76] B. P. Abbott, et al., All-Sky Search for Short Gravitational-Wave Bursts in the Second Advanced LIGO and Advanced Virgo Run, *Phys. Rev. D*

- 100 (2) (2019) 024017. arXiv:1905.03457, doi:10.1103/PhysRevD.100.024017.
- [77] T. Oliphant, NumPy: A guide to NumPy, USA: Trelgol Publishing, [Online; accessed *today*] (2006–).
URL <http://www.numpy.org/>
- [78] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. Carey, Í. Polat, Y. Feng, E. W. Moore, J. Vand erPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, S. . . Contributors, SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, Nature Methods 17 (2020) 261–272. doi:<https://doi.org/10.1038/s41592-019-0686-2>.
- [79] Pygtk software: <https://wiki.python.org/moin/PyGtk>.
- [80] Erik, L. Saldyt, Rob, J. Gross, tjproct, kmrudin, T. L. Scholten, msarovar, kevincyoung, R. Blume-Kohout, pyIonControl, D. Nadlinger, pygstio/pygsti: Version 0.9.9.1 (Feb. 2020). doi:10.5281/zenodo.3675466.
URL <https://doi.org/10.5281/zenodo.3675466>
- [81] Apache kafka software: <https://kafka.apache.org/>.
- [82] M. Frigo, S. G. Johnson, The design and implementation of fftw3, Proceedings of the IEEE 93 (2) (2005) 216–231.
- [83] Intel math kernel library: <https://software.intel.com/content/www/us/en/develop/tools/math-kernel-library.html>.
- [84] Glib2 software: <https://github.com/GNOME/glib>.
- [85] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, GNU Scientific Library - Reference Manual, Third Edition, for GSL Version 1.12 (3. ed.), 2009.
- [86] D. Macleod, A. L. Urban, S. Coughlin, T. Massinger, M. Pitkin, paulaltin, J. Areeda, E. Quintero, T. G. Badger, L. Singer, K. Leinweber, gwpy/gwpy: 1.0.1 (Jan. 2020). doi:10.5281/zenodo.3598469.
URL <https://doi.org/10.5281/zenodo.3598469>

Current code version

nr.	code metadata description	
c1	current code version	1.7.3
c2	permanent link to code/repository used for this code version	https://git.ligo.org/lscsoft/gstlal
c3	code ocean compute capsule	N/A
c4	legal code license	gnu general public license
c5	code versioning system used	git
c6	software code languages, tools, and services used	python, c, sqlite
c7	compilation requirements, operating environments & dependencies	scientific linux 7, python 2.7
c8	if available link to developer documentation/manual	https://lscsoft.docs.ligo.org/gstlal/
c9	support email for questions	gstlal-discuss@ligo.org

Table 1: code metadata (mandatory)