GW Cloud: Executive Summary

Paul D. Lasky,[1, 2, *] Vivien Raymond,[3] and Eric Thrane[1, 2, †]

[1] *Monash Centre for Astrophysics, School of Physics and Astronomy, Monash University, VIC 3800, Australia*
[2] *OzGrav: The ARC Centre of Excellence for Gravitational-Wave Discovery, Clayton, VIC 3800, Australia*
[3] *School of Physics and Astronomy, Cardiff University, Cardiff, CF243AA, Wales, UK*[‡]

Over the past three years, the field of astrophysics has been revolutionized by the discovery of gravitational waves from merging black holes and inspiralling neutron stars. The raw data from these discoveries is publicly available[1]. However, in order to extract the most interesting scientific results, the data must be analyzed using specialized Bayesian inference tools [e.g 1–3]. The technical nature of Bayesian inference has limited the impact of publicly available data. We propose GW Cloud, a public database for the Bayesian inference data products derived from gravitational-wave data. GW Cloud will enable breakthroughs by making gravitational-wave astronomy more accessible, facilitating research by scientists with diverse skill sets and interests. It will enhance the science output of gravitational-wave astronomers by consolidating the Bayesian inference data products used for so much of our research into an easy-to-use, reliable database. There are four components to GW Cloud: the database, workflow management, enabling new science, and outreach.

Workflow management—GW Cloud will provide a portal through which Bayesian inference jobs are launched, allocate appropriate high-performance computing or cloud resources to run each job, and record the results using a standardized format in the GW Cloud database. Jobs will be queued on the LIGO Data Grid resources. If/when additional resources are available—for example, the Open Science Grid[2]—GW Cloud will send containerized jobs to external clusters. The portal can be accessed by a simple graphical user interface or from the command line. There will be a public portal (with restricted resources) and a LIGO/Virgo portal (with unlimited resource access). GW Cloud will provide administrators and users with high-level views of inference jobs in order to allocate resources and optimize workflows. The portal will provide users with helpful information such as estimated run time, issue warnings guided by analytics of previous submissions, and flag submissions similar to existing database entries.

Database—For each job, the database will record output including posterior samples, Bayesian evidence, weights, and all the metadata necessary to reproduce the submission. Database entries can be marked as private, LIGO/Virgo only, or public. The database will include a search engine to retrieve results based on variety of a tags and/or attributes. For example, users could search for inferences from GW170729 obtained using a specific waveform approximant, or they could search for all inferences that returned posterior samples with distances less than 100 Mpc. The database will include inferences from both real data and simulated data.

Enabling new science—We anticipate that GW Cloud will sometimes have spare computing power that is not needed for user submissions. The GW Cloud queuing system will use these spare cycles to facilitate new science, for example, running inference jobs on all available data in order to detect and characterize the astrophysical background from unresolved binary black holes [4], to enhance the significance of marginal detections [5], and to pursue targets of opportunity such as the follow-up of a new detections.

Outreach—GW Cloud will be accessible both as a public and private portal, with the latter having restricted access for LIGO/Virgo collaboration members. The public portal will include access to state-of-the-art inference tools, publicly available data, and the results of previously executed inference jobs tagged as public. The public portal can be used by professional astronomers outside of LIGO/Virgo as well as citizen scientists.

GW Cloud is a database and workflow-management tool for astrophysical inference of gravitational-wave data. It enables new science and expands the community of scientists who can engage with gravitational-wave astronomy while reshaping the way that specialists do gravitational-wave astronomy. Over time, we will build up Bayesian inference data products for all times of LIGO/Virgo data. This will enable a new understanding of weak gravitational-wave

---

[1] https://www.gw-openscience.org/about/
[2] https://opensciencegrid.org

signals, probing deeper into the cosmos than previously possible. This LIGO document LIGO-L1900080.

* Electronic address: paul.lasky@monash.edu
† Electronic address: eric.thrane@monash.edu
‡ Electronic address: vivien.raymond@astro.cf.ac.uk

[1] J. Veitch et al., Phys. Rev. D **91**, 042003 (2015).
[2] Biwer et al., PASP **131**, 024503 (2019).
[3] G. Ashton et al., arXiv e-prints (2018), 1811.02042.
[4] R. Smith and E. Thrane, Phys. Rev. X **8**, 021019 (2018).
[5] M. Isi, R. Smith, S. Vitale, T. Massinger, J. Kanner, and A. Vajpeyi, Phys. Rev. D **98**, 042007 (2018).