Draft 30 July 2019

Gravitational-Wave Data Analysis Computing Challenges in the 3G Era

Ian Bird, Peter Couvares, Ed Porter, Josh Willis

Abstract

Third generation ground-based GW detectors will require not only R&D for new detector technologies. Equally important will be investments in computing infrastructure -- including but not limited to the development of vastly more efficient algorithms to search for transient GW signals and extract their underlying astrophysical parameters, assessments of needed cyberinfrastructure resources, and trade studies of computational 'service provider' models. Here, we describe the major computing challenges facing the GW community in the 3G era and present recommendations for addressing them.

Introduction

To frame and motivate this section of the report, we begin with background on the existing 2G detector scientific collaborations and an overview of the computing models and methods currently employed. The Advanced LIGO/Advanced Virgo collaboration (LVC) is composed of three Gravitational Wave (GW) interferometers located in Hanford (WA), Livingston (LA) and Pisa (Italy). In September 2015, the LVC began a series of advanced era detector runs, with the nomenclature "O#". O1 ran from September 2015 to January 2016, and as well as the first ever detection of GWs, the run ended with the detection of three binary black hole (BBH) mergers. O2 ran from December 2016 until the end of August 2017. As well as the detection of a number of other BBH mergers, O2 saw the first ever detection of a merger of two neutron stars (BNS). O3 began on April 1st 2019, and will run for one year. It is further expected that the Japanese interferometer KAGRA will join the O3 run in late-2019.

From a computing perspective, the biggest challenge in the transition from O1 to O2 was the increased computational power needed for both the search and parameter estimation phases. In the search (detection) phase, the template banks increased in size to accommodate a larger range of masses. In the parameter estimation phase, while the computational cost of each run remained almost the same as in O1, the number of sources, as well as the number of exploratory runs that were needed for the BNS merger, caused the computational cost to explode. In addition, unforeseen and computationally-intensive analyses were needed to measure the Hubble-Lemaitre constant H0, test the validity of GR and to constrain the internal physics of neutron stars.

In its third observing run (O3), the LIGO-Virgo collaboration estimates its data analysis computing requirements at ~300 million CPU core-hours per year, to perform some 80 astrophysical searches, follow-up activities, and detector characterization. The 10 most demanding of these analyses comprise about 90% of the demand, with a long tail of the remaining 70. Most of this computing consists of pleasingly parallel High Throughput

Computing (HTC) for “deep” offline searches; ~10% is for low-latency data analysis needed to generate rapid alerts for multi-messenger (electromagnetic, neutrino) followup. Very little high-performance parallel computing is required, with the exception of Numerical Relativity simulations, which is not included in this assessment.

Currently ~90% of this computing is provided by dedicated LIGO-Virgo clusters, and ~10% by external shared computing resources; however, growth of the dedicated resources has flattened while the shared component is growing. This growth of shared, external computing resources presents new distributed computing and data access challenges. This is something the gravitational-wave data analysis community shares in common with HL-LHC computing.

The LVC has carefully investigated the use of parallel GPU and MIC architectures for its most compute-intensive searches; CUDA GPUs have been the most successful and cost-effective, and were deployed at scale for the first time in O3.

Currently the LVC makes little use of commercial cloud resources; there are no major technical obstacles (the cloud looks similar to other shared resources); however, the logistics and cost-effectiveness of funding and managing metered computing are still being understood.

LIGO/Virgo generates ~20TB per IFO per observing year of h(t) strain data used by most analyses, and ~1PB per IFO per observing year of raw data (all channels, full sample rate). The scale of gravitational-wave data analysis is no longer “big data” by 2019 standards — but data access remains non-trivial in a distributed HTC environment nonetheless.

The demand for data analysis computing in the 3G era will be driven by the high number of detections (up to hundreds per day), the expanded search parameter space (0.1 solar masses to 1000+ solar masses) and the subsequent PE follow-up required to extract the optimal science goals. Given that signals will now last hours to days in the detector (as opposed to seconds/minutes in the 2G era), these extended signals may strain physical memory abilities. On top of this, as the 3G network will be an international effort, a large effort will be needed to develop an appropriate and scalable computing infrastructure, including data access and transfer protocols, and storage and management of software tools, that has sustainable development, support and management processes.

3G Computing Resources

In terms of the processing hardware needed for development, simulation, testing, and production data analysis of gravitational-wave data, there appear two major challenges for the 3G era, relative to the 2G era: CBC detection and CBC parameter estimation (PE).

CBC Detection

At present, the majority of GW detections are made using a matched-filter based template bank. This requires covering the parameter search space with many hundreds of thousands of theoretical waveform models (or templates). Application of existing methods are unlikely to meet the science requirements of the 3G era, as the cost of matched-filter searches for CBC signals grows dramatically with improvements in frequency sensitivity. Both the number and duration of waveforms needed increases substantially for the lower frequency limits afforded by the 3G detector designs. Thankfully, the dominant costs are independent of predicted rates and rely on parallel algorithms with independent tasks that are horizontally scalable. However, given the increased duration and higher number of signals, it is almost certain that current methods will not be sufficiently capable in the 3G era.

A non-negligible amount of search pipeline R&D will be needed to prepare for 3G detection. Assuming template banks are still applicable in the 3G era, a number of questions need to be answered:

- How many templates will be needed to cover 1-1000(+) solar mass events?
- Will sub-solar-mass events (0.1-1 solar masses) need to be included in the template bank?
- How will template banks cope if there are overlapping mergers in the data set?
- How closely do templates need to be spaced, or inversely, how much SNR are we prepared to lose to detect events?
- What are the science impacts of massively reducing the number of templates at the cost of some sensitivity?

3G detectors will observe vastly more gravitational waves than 2G detectors, including some with tremendous signal-to-noise ratios and some that remain in the detector's sensitive frequency band over much longer times. Extracting all of the science encoded in these observations will require next-generation computational techniques, as today's approaches would be prohibitively expensive. Even with these new methods, the large number of observations to process will likely require a distributed system to analyze all of it, rather than the centralized LIGO Data Centers used today.

A major foreseeable problem is how to answer these questions when the ground-based GW data analysis community is constantly involved in development to prepare for and exploit current science runs. One potential way around this problem, and to bring similar expertise to bear on identifying and interpreting gravitational-wave detections for third-generation detector data on Earth, would be to initiate a series of mock data challenges in the same spirit as The Laser Interferometer Space Antenna (LISA) Mock Data Challenges. Such an endeavour would require establishing dialog among 2G, 3G, and LISA data analysis experts as these data challenges are created and run.

LISA has faced similar data analysis challenges with some overlap with 3G data analysis, including analyzing extraordinarily loud signals and signals that remain in band for much longer than they do in LIGO (in LISA’s case, this can be months or even years). Since 2005 (and most recently in 2018 [1]), the LISA community has begun building up techniques to address computational challenges in analyzing its data, using "Mock Data Challenges." Each challenge begins with publicly released simulated LISA noise, with different simulated signals hidden inside. Anyone in the community is then welcome to build and apply computational tools that attempt to find and interpret the hidden signals. In particular, the data challenges spurred the development of stochastic search techniques, because a LIGO-style template-bank approach would not be computationally feasible.

Although some initial investigations have been performed, large and well-organized 3G mock-data challenges with broad participation from 2G data analysis experts will be critical for gaining the necessary understanding of 3G data analysis computational challenges and costs. They will drive development of the necessary new techniques, and they will demonstrate the feasibility of using a distributed (instead of centralized) approach.

CBC Parameter Estimation

The cost of CBC Parameter Estimation using Bayesian inference may present an even larger challenge. Improvements in detector sensitivity translate into much higher detection rates, and PE costs scale with the number of detections. Furthermore, they rely on some inherently sequential algorithms (e.g., Monte Carlo Markov Chain, or “MCMC”), whose convergence properties vary from algorithm to algorithm. To make statements on the parameters of a system, a certain number of statistically independent samples from the posterior distribution are required. The current 2G runtime to achieve this number of samples varies not only from algorithm to algorithm, but also from source to source, with the fastest analysis taking days, and the longest taking weeks. This latency is unacceptable in the 3G era.

The development of new, more convergent (and thus more rapid and efficient) algorithms is a field of research in its own right, with the typical timescale for the development of a new algorithm usually being on the order of a decade. If a decision is made to develop more convergent algorithms “in-house”, this will exacerbate the problem identified above with respect to the development of new search algorithms. The development of such an algorithm would have to be an interdisciplinary effort involving GW scientists, experts in statistical inference and professional programmers. The adaptation of advanced algorithms designed for other domains will almost certainly require extensive modification for GW astronomy applications. Again, this is something that would take a number of years to accomplish.

The scientific requirements for PE in the 3G era are as yet undefined. In the 2G era, the LVC performed multiple, computationally-intensive PE analyses on every viable CBC candidate, and yet were constrained by available human effort as much as by computing. In addition, the computing costs of development, testing, and exploration of PE codes were greater than that of the final PE runs used for publication. Approximately 1 million CPU core-hours of computing (on a reference Intel Xeon E5-2670 2.6Ghz CPU) were used for parameter estimation of each signal in the LVC’s second observing run (O2), compared to ~10k CBC signals per year expected for 3G IFOs; a 3 orders-of-magnitude increase.

In an era of multiple signals per day, it is unclear if it will be feasible or desirable to run a deep PE analysis on every CBC candidate. It might be that the science goals demand different degrees of PE investigation (and therefore computing resources) for different candidates. The answer will obviously be driven by the computing efficiency of PE codes (and the degree of automation possible in the PE process), but the range of scientific scenarios is not obvious due to the many remaining uncertainties:

- what is the lower limit of PE investigation needed to realize our most basic 3G scientific goals, and what is its present computing cost;
- what is the upper limit to PE investigation we can currently achieve by applying more computing power, given the expected signal quality;
- where between these two limits do we get the maximum scientific benefit as a function of cost; and do we get diminishing returns for additional computing after a certain degree of precision?
- Are there obvious “tiers” of investigation and cost which we should target for different kinds of candidate signals?
- How should we define those tiers given the science we hope to extract from the data?

These questions need to be explored iteratively by GW scientists over the next decade in close collaboration with statistical inference and computing experts, as new computing efficiencies are realized and new scientific goals are identified and translated into PE requirements.

Burst and Continuous-Wave Detection

Detection of unmodeled burst sources, a major component of 2G gravitational-wave data analysis computing demand, should be straightforward in the 3G era. Burst searches should scale more or less linearly with the # of IFOs and their observing time. Unlike CBC, there is no scaling with detector sensitivity or low-frequency cut-off.

It is less problematic for Continuous-Wave (pulsar) searches to scale in the 3G era, simply because they have always had to cope with many orders of magnitude less computing than optimal; as a result, their computational challenges will remain fundamentally similar to those faced in the 2G era.

Hardware Resources Needed

A naive scaling of current matched-filter CBC searches and PE to 3G design sensitivities will require many (3+) orders of magnitude more CPU and RAM than 2G. Moore’s Law alone will not deliver the required performance increase as it is slowing already due to both cost management by a shrinking set of hardware producers and — barring a breakthrough — the physical limits of silicon transistors. CERN’s outlook is a 5-10%/year performance/cost improvement at constant cost, with very large uncertainties. This does not include the coprocessor/accelerator (GPU, FPGA, etc.) market, but there are no panaceas there; the performance improvements that are being delivered by both increased CPU parallelism and co-processors are becoming more complicated to exploit in software (with task and/or data parallelism). Many existing codes won’t benefit fully.

Something has to give: barring an unexpected breakthrough, we may need to be clever about how and where in parameter space our searches can be less sensitive, to not miss what we care about most. This could be done, e..g., through matched-filtering with sparser template banks and/or less faithful waveforms in the most expensive and/or less scientifically valuable regions of the parameter space, reduced-order modeling, waveform compression, greater reliance on unmodeled burst searches, and/or machine learning. We don't yet know the right mix of approaches for 3G-scale analysis. We may to identify new "sweet spots" of sensitivity vs. cost given our science goals.

If the GW science we can realize will be limited by the efficiency of search and PE codes, targeted investments in data analysis development and optimization in advance of the 3G era may pay outsized scientific dividends. These investments should be driven by better understanding the detection and PE requirements needed for specific discoveries.

3G Data Analysis Software Infrastructure

Given that the hardware resources for 3G will be orders of magnitude larger than 2G, how will these resources need to be organized and managed? A centralized "walled garden" of homogeneous, dedicated GW clusters likely to be inefficient and prohibitively expensive. A diverse worldwide GW collaboration is unlikely to select a single external computing provider (commercial or otherwise). 3G computing optimization may require specialized hardware for different searches.

Improved software infrastructure (aka "cyberinfrastructure") will be needed to make a more complex, dynamic network of resources to scientists usable and robust, including distributed scheduling, identity management and access control infrastructure, cybersecurity tools and services, data analysis software development and testing tools, resource accounting and reporting tools, low-latency alert and coordination infrastructure, as well as public data and code release infrastructure.

We expect increasing heterogeneity and complexity of computing platforms in 3G era:

- of processing hardware; due to the opportunities for cost savings, data analyses may need to support multiple generations of CPUs, GPUs, MIC platforms and treat them each as distinct platforms. "Lowest common denominator" code capable of running on any platform may not be efficient enough.
- of providers — resources will include those internal to the project, partners & collaborators, institutional, regional/national, commercial, volunteer.
- of target operating systems and software environments. Containerization, etc. are tools to help mitigate this complexity but carry their own costs and aren't a complete solution.
- of batch/queueing systems.
- of storage and network interfaces and capabilities.
- of policies for identity and access management, workflow prioritization.
- of accounting models and accounting systems
- of motivations and expectations — mutual scientific/strategic interest, public or scientific recognition, financial or other compensation, etc. — not all of which will be spelled out in an MOU, SLA, or contract.

3G Computing - Human Effort

One major challenge to computing In the 2G era has been uncertain and discontinuous funding streams for computing labor embedded in the collaboration groups outside the IFO laboratories. The need to professionalize software development and engineering and to support increasingly complex computing environments demands more full-time professional computing expertise side-by-side with collaboration scientists (vs. part-time volunteer/service work by scientists).

Many of these are not strict software development or IT roles that can be outsourced beyond the project — they are hybrids of research computing, consulting, software engineering, and distributed systems development and administration roles. These roles benefit enormously from institutional (project) “memory” — projects pay dearly in time, money and quality when experience and relationships are lost.

It is also hard to recruit and retain career professionals on overlapping 1-3 year awards. It can also be hard to find funding for this work, which is not always “transformative” science in and of itself, but needed to enable transformative science. This is an old problem but will become more acute with the increasing need for these computing roles in the 3G era.

These concerns are common with other science communities, and could help to drive a common effort to recognise software and computing as scientific career tracks, as well as collaborate in a common lobby for more focussed investment in software.

Recommendation: in concert with HEP and Astronomy communities, we recommend that funding agencies provide more stable funding streams, and recognized roles for hybrid domain science / Computer Science staff dedicated to the computing challenges discussed in this report.

3G GW Community Services

Increasing heterogeneity and complexity of computing platforms drives the need for better software engineering and testing; additional organizational expertise and effort in optimization, distributed computing (architecture, engineering, support), and computing management; better tools, services, and processes for sustainable optimization; better education and consulting for scientists/developers who are not first and foremost software engineers; automated testing for diverse hardware platforms and environments; and more complex deployment, orchestration, instrumentation, and accounting of DA workflows.

Data Analysis Optimization Support

The 3G era provides several challenges and opportunities for algorithmic and software engineering, including:

- Scientific Prioritization and Scoping
- Estimation and Benchmarking of Computational Costs

- Optimization of Data Analysis Methods and Algorithms
- Optimization of Code Implementation and Libraries
- Advanced Compiler Optimizations
- Workflow Management Optimizations
- Development, Testing, and Simulation Process Optimizations
- 3G Computing Resource Scheduling Optimizations
- Resource "Supply" Optimizations (make more cycles available)
- Workflow Portability Optimizations (expand usable resources)
- Hardware Procurement Informed by Code Performance
- Scientific Reviews that include Computational Efficiency
- Documentation, Training, Collaboration and External Engagement on Optimization

Collaborations with non-GW computing efforts

The HEP community has a working group looking into the evolution of the cost of hardware (CPU, disk, tape) in preparation for the HL-LHC upgrades on the 2026 timescale. CERN and other institutes have tracked costs over more than 15 years, and use recent trends to set expectations for the evolution over the coming years. As noted above, simple "Moore's law" predictions no longer seem to be valid and are outweighed by market forces and cost management by the relatively few hardware producers. It is a fact that there are now only a handful of chip fabrication companies, only 2 HDD manufacturers and that tape drive technology is a monopoly with only IBM building drives and only 2 media producers.

For the past several years the HEP community has used as guidance a 20% per year improvement in performance for CPU and disk storage capacity for a constant cost, to predict the affordable capacity of complete server systems (not the raw disk space or raw CPU performance). For CPU this applies to standard x86-like processors. For tape storage capacity the capacity increase per year was estimated to be somewhat more, around 25%/yr.

In the last 2 years however we have observed a worrying trend, that there are large fluctuations in the market prices of key components (in particular RAM), and shortages of supply, and that the effective growth is now only around 10% per year for constant cost. In addition all of these markets are now in the hands of very few large manufacturers, who control the markets, and thus prices. The upshot of this is that it is currently very difficult to make a realistic estimate of the price/performance evolution for the coming years. This is a significant concern when trying to understand the overall cost of computing and storage for future experiments.

Recommendation: Identify (and if necessary, establish) the right forums for the GW community, HEP, and Astro communities to regularly and systematically discuss the overlap in our computing problems, share information, and plan specific technical collaborations where useful.

For example:

- Coordination of Distributed Scheduling, Computing Security, and Identity and Access Management efforts between LHC-HL and LIGO/Virgo.
- Computing optimization approaches, techniques and tools.
- Software engineering tools and technologies.

Some of this interaction already happening, but the GW community should expand those efforts and have a stronger voice in forums where these matters are discussed. The 3G gravitational-wave community cannot afford to reinvent wheels.

The high-luminosity LHC experiments face similarly-daunting computational scaling problems over the same timeframe, but they have an order of magnitude larger starting requirements — they are helping to blaze a trail in optimization and distributed computing infrastructure which we can and should follow and collaborate on.

CERN and the Worldwide LHC Computing Grid project (WLCG) are reviewing the LHC computing models and infrastructure for the future high luminosity runs of the LHC which increase the data volumes and compute requirements by many factors. At the same time it has become clear that the global facilities that support WLCG will also host other large-data science experiments, such as SKA, LSST, CTA, and 3G GW experiments. It is clear that a common international data management and computing infrastructure capable of supporting all of these would be mutually beneficial to the experiments, the computing facilities, and the funding agencies. Some current regional efforts (e.g., the ESCAPE project) are encouraging steps in this direction.

Being explored are Exabyte-scale capabilities for data management, able to serve data in an efficient way to heterogeneous and globally distributed compute resources including HTC clusters, HPC clusters, and commercial clouds. We foresee the requirement to support a variety of processors ranging from CPU, accelerators (e.g., GPU, MIC), FPGA and other innovative architectures.

CERN/WLCG is basing its strategy on 3 areas of investment:

- A federated data cloud (“data lake”) capable of the long term curation and serving of exabyte-scale data, complemented with tools and services to allow experiments to manage that data and content delivery tools to efficiently serve the data to remote compute resources. This data lake would also provide an effective mechanism for hosting and serving data to the broader community including public open access, potentially with the inclusion of commercial interests.
- A significant investment in application software skills. The ability to make use of the full set of available resources in coming years requires development of software skills and the ability to move applications to appropriate compute architectures in an agile way as they evolve. The infrastructure will also depend on a clearinghouse of software tools and services for science communities. The HEP community has initiated the HEP Software Foundation (HSF) as the vehicle through which to address the long term software investment and also to provide the locus for a software tools clearinghouse. It

is not necessarily specific to HEP, and could also be strengthened with the inclusion of other science communities.
- Networks to support the infrastructure and the use of appropriate technologies to manage data management both within the data lake and in serving data externally. WLCG has developed an overlay network (LHCONE) that the National Research and Education Networks use to manage the science data traffic. This concept has proven extremely versatile and is now used by many more science communities.

In addition to the above infrastructure developments, CERN has pioneered aspects of open access and open data, including the Zenodo publishing platform and the CERN open data portal, as well as the storage infrastructure underpinning these. These could be interesting components of a future service for the GW community. CERN also provides “CERNBox” - a DropBox-like open source solution that is widely used in the HEP community.

To complete the picture, CERN has the “openlab” concept, which is a vehicle for collaboration with industry, disconnected from any expectation of procurement. This provides a mechanism to fund and staff projects of joint interest between CERN and various industrial partners. In recent years this has expanded to include other research institutes. This is another avenue of potential technology transfer to the 3G GW community, enabling the investigation of new computing technologies in collaboration with the companies, and to provide access to their expertise.

From the CERN viewpoint collaborating with the GW community on some or all of these aspects and the related technologies is of great interest and could aid in developing a broad toolset useful to a broad scientific base. In the recent European Strategy for Particle Physics meeting (May 2019), a broad collaboration of HEP with the Gravitational Wave community was seen as strategic, and very much encouraged from both communities.

More efficient techniques for background estimation may reduce CBC search costs by a factor of a few; machine learning-based searches may live up to their promise in culling signals from non-Gaussian noise. Even a diminished Moore’s law could help improve the performance of sequential codes by a factor of a few over the next decade. Parallel algorithms will track Moore’s Law closer than sequential algorithms.

LIGO’s recent experience has shown that the return-on-investment of dedicated, “full-stack” computing optimization by a team of GW+computing experts is overwhelmingly positive, in terms of saved costs for expensive data analysis codes. The GW community should plan on making this investment for 3G. More efficient data analysis codes needed for 3G will require more software development effort. This will require time from busy GW scientists and/or hiring more computing specialists. An investment in computing specialists can free up scientists to do science.

Long-term software development costs may be higher for rapidly-evolving parallel hardware platforms (GPU, MIC, AVX512, etc.) than for traditional CPUs, given that parallel programming interfaces are less stable targets. Single-threaded CPU codes have worked on new hardware with minimal modifications for 30+ years. Data analysis on more distributed, non-dedicated computing grid/cloud platforms will require ongoing computing infrastructure investment. This will require time from busy GW scientists and/or hiring more computing specialists.

3G computing labor costs may go up by a factor of a few related to to existing 2G computing models. Notably, there may be potential and interest in broader collaboration in this area, e.g. with HL-LHC (and other HEP like DUNE) as well as SKA.

Conclusion and Summary Recommendations

The GW community must make smart data analysis computing investments well in advance of, as well as during, the 3G era. Without major innovation, 3G data analysis (search and PE) naively requiring a 1000x increase in computing demand may need to be implemented on 10-100x faster hardware. Current algorithms will not scale; new techniques have not yet been developed and **there is not yet a broad consensus on how difficult it may be**. Data analysis innovations, and expert code optimization effort for the most expensive search and PE codes will be needed to bridge this gap.

We need to plan for increased investment in data analysis software development and computing infrastructure effort compared to the 2G era, or or potentially sacrifice important science opportunities.

Where current methods cannot scale, we need to understand computing cost vs. scientific benefit and make smart tradeoffs to achieve our goals with available computing resources. Without focused planning and funding, it will be difficult to drive the gravitational-wave community to address this problem while it is dealing with the many detections of O4 and O5, and computing is less of a concern.

Specific recommendations:

- R&D on 3G search pipeline development and scientific analysis should become a higher priority of the GW community.
- Initiate R&D on 3G PE methods, in collaboration with experts in statistical inference and computing methods.
- Organize new mock-data challenges for ground-based 3G computing, similar to ones undertaken by the LISA community, to better understand these challenges and identify the most promising solutions.
- Identify (and if necessary, establish) the right forums for the GW community and HEP community to regularly discuss the overlap in our computing problems, share information, and plan specific technical collaborations where useful.
- In concert with HEP and Astronomy communities, recommend that funding agencies provide more stable funding streams, and recognized roles for hybrid domain science / Computer Science staff dedicated to the computing challenges discussed in this report.

Suggested other recommendations:

- Pursue collaborative efforts with HL-LHC, SKA, and other relevant communities on investment in advanced software techniques, to address performance and adaptability to the landscape of changing computing architectures
- Technology tracking, perhaps in collaboration with industry (CERN openlab as vehicle or model), could include research in how to adapt codes to be reasonably efficient across many architectures
- Collaborate on cyberinfrastructure (e-infrastructures in Europe) together with HL-LHC, and others - OSG, European Open Science Cloud, and others. Broker partnerships and involvement in defining directions to support the science use cases.