

GPU coprocessors as a service for deep learning inference in high energy physics

Jeffrey Krupa¹, Kelvin Lin², Maria Acosta Flechas³, Jack Dinsmore¹, Javier Duarte⁴, Philip Harris¹, Scott Hauck², Burt Holzman³, Shih-Chieh Hsu², Thomas Klijsma³, Mia Liu³, Kevin Pedro³, Natchanon Suaysom², Matt Trahms², Nhan Tran^{3,5}

¹ Massachusetts Institute of Technology, Cambridge, MA 02139

² University of Washington, Seattle, WA, 98195

³ Fermi National Accelerator Laboratory, Batavia, IL 60510

⁴ University of California San Diego, La Jolla, CA 92093

⁵ Northwestern University, Evanston, IL 60208

Abstract. In the next decade, the demands for computing in large scientific experiments are expected to grow tremendously. During the same time period, CPU performance increases will be limited. At the CERN Large Hadron Collider (LHC), these two issues will confront one another as the collider is upgraded for high luminosity running. Alternative processors such as graphics processing units (GPUs) can resolve this confrontation provided that algorithms can be sufficiently accelerated. In many cases, algorithmic speedups are found to be largest through the adoption of deep learning algorithms. We present a comprehensive exploration of the use of GPU-based hardware acceleration for deep learning inference within the data reconstruction workflow of high energy physics. We present several realistic examples and discuss a strategy for the seamless integration of coprocessors so that the LHC can maintain, if not exceed, its current performance throughout its running.

Submitted to: *Mach. Learn.: Sci. Technol.*

1. Introduction

The detectors at the CERN Large Hadron Collider (LHC) have enormous data rates, with a current aggregate of 100 Tb/s and plans to exceed over 1 Pb/s. The challenge of processing this data continues to be one of the most critical elements in the execution of the LHC physics program. A three-tiered approach is utilized to process LHC data, where at each tier, the data rate is reduced by roughly two orders of magnitude, resulting in a manageable final data rate of 10 Gb/s. Due to the high initial rate and restrictions coming from the high radiation collision environment, the first tier of computing consists of specialized hardware that utilizes field-programmable gate arrays (FPGAs) and application-specific integrated circuits (ASICs). The second tier, the high-level trigger (HLT), consists of a CPU-based computing cluster on-site at the LHC. The third tier, performing complete “offline” event processing, consists of a globally distributed CPU-based computing grid.

The first decade of LHC running has led to an extensive set of scientific results. These results include the discovery of the Higgs boson [1–3] and, more recently, strong constraints on the nature of dark matter [4–6]. To contend with these strong dark matter constraints, physicists have been forced to re-think their approach to searching for dark matter and, generically, new physics models. This has led to the development of light dark matter models [7]. These models often predict signatures that could be produced at the LHC but would be discarded in the early tiers of data reduction. To enable the search for these particles, it is imperative to improve the quality of LHC data reconstruction at all tiers of processing. Additionally, over the next decade, the LHC will progressively increase the beam intensity, resulting in more data generated by the detectors. As a consequence, the demands for computing will increase proportionally to sustain the current level of physics output. Figure 1 shows the expected computing needs over the next decade. To contend with the high-luminosity upgrade of the LHC (HL-LHC), a large increase is needed starting from 2026. These demands outpace the expected growth of CPU performance. As a consequence, the LHC needs a computing solution at least to sustain the current computing performance and, potentially, to exceed it.

With the end of Dennard scaling in the late 2000s, processor technology has undergone several changes. These changes have included the adoption of multicore processors and the rise of alternative processing architectures, or coprocessors, such as graphics processing units (GPUs), FPGAs, and ASICs. With the rise of deep learning (DL), these alternatives have become increasingly appealing due to the inherent parallelism in both DL algorithms and in these coprocessors. The gains from using coprocessors can be substantial, with improvements in algorithmic latency exceeding multiple orders of magnitude. Given the scale of developments related to DL, future growth in processor technology is increasingly leaning towards heterogeneous systems in which combinations of CPUs, GPUs, FPGAs, and ASICs are all deployed, with each designed to solve specific tasks. However, high energy physics (HEP) experiments have thus far undertaken only limited use of alternative processors within the HLT and offline computing grids, despite common use of machine learning (ML). HEP experiments have historically relied on ML as a way to improve the overall quality of the data and to separate small signals from enormous backgrounds [11]. DL approaches have enhanced both the performance and flexibility of ML techniques. In light of this, the LHC has been quick to adopt DL techniques to improve the quality of data analysis. This includes core components such as low-level detector energy reconstruction, electron and photon reconstruction, and quark and gluon identification. The increasing deployment of these algorithms is starting to comprise a significant portion of the overall computing budget. The goal of this study is to enable the use of these algorithms in online and offline data processing tiers, in the context of the LHC experiments’ increasing data rates. Our approach does this by offloading the computational burden of these algorithms to GPUs while making minimal changes to existing CPU-based workflows.

To achieve this, we move existing work a step further by exploiting the “as-a-service” paradigm. In this paper, we design a prototypical framework for LHC computing as a service. We develop DL algorithms to replace domain-specific algorithms, to solve a variety of physics problems through DL inference. We then transfer the algorithm to a coprocessor on an independent (local or remote) server and re-configure the existing CPU nodes to communicate with this server through asynchronous and non-blocking inference requests. With the inference task offloaded

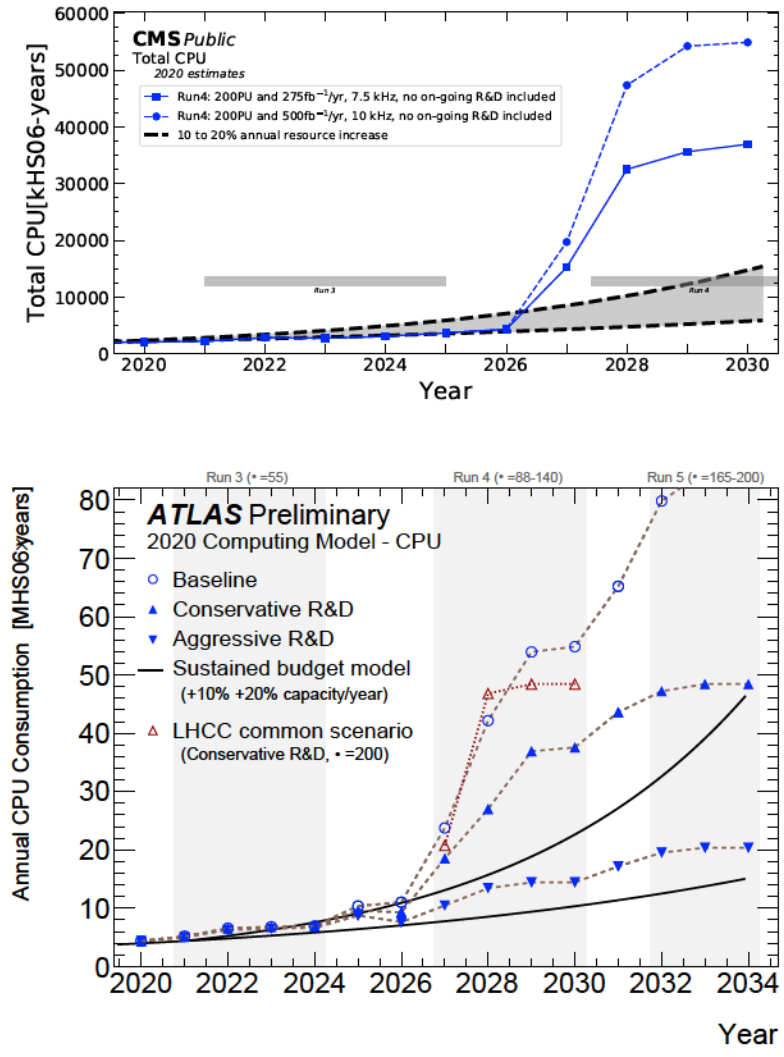


Figure 1. Computing needs projections for the CMS (upper) and ATLAS (lower) experiments at the LHC [8–10]. For ATLAS, the computing needs estimate is shown under the baseline computing model and under different R&D scenarios, while for CMS, the estimate is shown without including ongoing R&D in different collider conditions. Computing growth assuming a 10–20% CPU budget increase is also shown. MHS06-years (kHS06-years) stands for 10^6 (10^3) HEPSPEC06 per year, a standard CPU performance metric for high energy physics. The different LHC operating runs are also indicated. A large upgrade to the collider and all LHC detectors will occur between Run 3 and Run 4, starting in 2026.

as a request to the server, the CPU is free to perform the rest of the necessary computing within the event.

Deploying GPUs as a service (GPUaaS) is a natural way to incorporate alternative coprocessors that has several advantages over a direct-connection approach. In particular, deploying GPUaaS increases hardware cost-effectiveness by reducing the number of GPUs required to achieve the same throughput. This is possible because each GPU can service many more CPUs than a direct-connection paradigm would allow. It is nondisruptive to the existing LHC computing model by offloading the specific algorithms with minimal client-side re-configuration (see Section 3). It facilitates seamless integration and scalability of heterogeneous coprocessors (such as GPUs and FPGAs), as suited for optimal algorithmic performance. Finally, by exploiting existing open-source, widely-adopted frameworks that have been optimized for fast GPU-based DL, this approach can be adapted quickly to different tasks at the LHC and beyond.

In this paper, we present several examples of integrating GPUaaS into LHC workflows. We consider three ML-based algorithms that span a variety of LHC computing applications. We integrate these algorithms into both online and offline LHC workflows with GPUaaS and we benchmark them to evaluate the impact of GPUaaS on the operation of the HLT and the offline computing grid. Based on our results, we propose a model for incorporating GPUs and other coprocessors into LHC computing workflows.

The remainder of this paper is organized as follows. In Section 2, we briefly review related work. In Section 3, we provide an overview of the current LHC computing model and the as-a-service computing model and we derive metrics that maximize the cost-effectiveness of coprocessors in LHC workflows. Section 4 describes the three ML-based algorithms to be deployed. We describe our configuration of the servers in Google Cloud and LHC data centers and evaluate the limitations of each site in Section 5. We also measure several performance-related quantities relevant for full-scale LHC reconstruction as a service, including hardware throughput, network bottlenecks, and scaling with number of GPUs. In Section 6, we determine the hardware and networking requirements for maximizing the throughput of these algorithms at scale for LHC computing as a service. Finally, we conclude in Section 7.

2. Related Work

Researchers across HEP have investigated the use of GPUs in detector reconstruction. At the LHC, the focus has been on implementations for the HLT where faster computing times lead directly to increased throughput. For offline computing at the LHC, GPUs have not been considered since this would require a larger redesign of the LHC computing grid. The work in this paper is different from previous approaches in that we employ GPUaaS, which allows for use in both the HLT and offline workflows without a large redesign of the existing LHC computing grid. Additionally, we utilize DL algorithms, which allow for the use of existing GPU compiler frameworks to quickly obtain optimized code. We stress that this is the first instance of GPU usage for offline computing at the LHC, the first usage of GPUaaS at the LHC, and the first use of DL with GPU acceleration for reconstruction of physics objects at the LHC.

Outside of detector reconstruction, DL algorithms are extensively used in HEP in the later stages of data analysis. In this context, training of DL algorithms is almost exclusively performed on GPUs. Additionally, the use of GPUs for DL has

led to several new HEP data analysis frameworks that exploit GPUs, including the hepaccelerate framework [12], GooFit [13], and pyhf [14].

Within the context of online computing, GPUs were first integrated into the 180 compute nodes of the HLT workflow of the ALICE experiment at the LHC to perform charged particle reconstruction. They were part of the ALICE operations during 2010–2013 and 2015–2018 [15]. More recently, GPUs are being considered for the HLT of the LHCb experiment [16, 17]. This HLT system relies completely on charged particle tracking algorithms. Additionally, within the CMS experiment, GPUs have been explored for charged particle reconstruction through the use of cellular automata [18], and through the use of an accelerated pattern recognition algorithm [19]. Beyond charged particle tracking, algorithms for GPUs and FPGAs have been developed (but not implemented) for real time processing of ring imaging Cherenkov detectors for the LHCb HLT [20, 21]. Beyond the LHC experiments, GPU algorithms have been developed for the trigger readout of the mu3e experiment [22], and the dark matter experiment NA62 [23]. These algorithms are planned to run in the next round of data taking for each experiment, starting in 2021. In all instances, GPUs have been considered in the context of direct connection to CPUs via PCI Express. Additionally, the algorithms on GPUs presented in the aforementioned works did not use DL.

Within the context of offline computing, GPU use in HEP has remained limited. For LHC reconstruction it has never been considered. In neutrino physics, GPUs have been used for simulation of the propagation of Cherenkov light signatures for the IceCube experiment [24]. The experiment recently performed a large-scale test of a GPU-only simulation of neutrino signatures, using over 50,000 GPU cores for a period of 20 minutes [24, 25]. While the study was able to utilize a large number of cores for a single type of algorithm used in the simulation of the IceCube detector, it did not constitute a full HEP reconstruction workflow consisting of a broad set of algorithms that perform many different tasks.

The offline and online reconstruction software for large LHC experiments consists of several million lines of CPU code. Rewriting this code to run on GPUs, for example using CUDA, would be prohibitively costly and in some cases would likely lead to substantially worse performance. In this paper, we present, for the first time, an alternative model whereby only algorithms with substantial speedups are ported to GPUs, with each GPU serving many CPU nodes. We demonstrate that GPUaaS can be integrated within full LHC workflows and can produce significant overall algorithmic speed improvements. A similar model for the utilization of CPUs and FPGAs within the LHC workflow was presented in Ref. [26] using the Services for Optimized Network Inference on Coprocessors (SONIC) framework [27]. The study exploited the Microsoft Brainwave service [28] and demonstrated a decrease in deep neural network (DNN) inference time by nearly 3 orders of magnitude when using an FPGA compared to a CPU. This paper extends SONIC to support GPUaaS, demonstrating a viable model for fast and nondisruptive integration of GPUs into the LHC workflow.

3. As-a-service computing for LHC physics

The current LHC computing model shown in Figure 2. In typical LHC event reconstruction, data is processed sequentially event-by-event, possibly on multiple threads on the CPU. However, if certain algorithms are significantly accelerated by

the use of coprocessors (as shown in Ref. [26]), a modified scenario with coprocessing as a service can be considered. In this model, a single coprocessor can serve hundreds of CPU processing elements. The CPUs are executing numerous different algorithms of the full event reconstruction, whereas the inference server is executing a single algorithm very efficiently. To benefit from this type of computing model, there must be a sufficiently large acceleration such that the overhead of offloading this processing onto a separate server does not further increase the reconstruction latency. To explain when this is the case, we first review the reconstruction model at the LHC and then discuss how as-a-service computing can be implemented within the LHC reconstruction workflow.

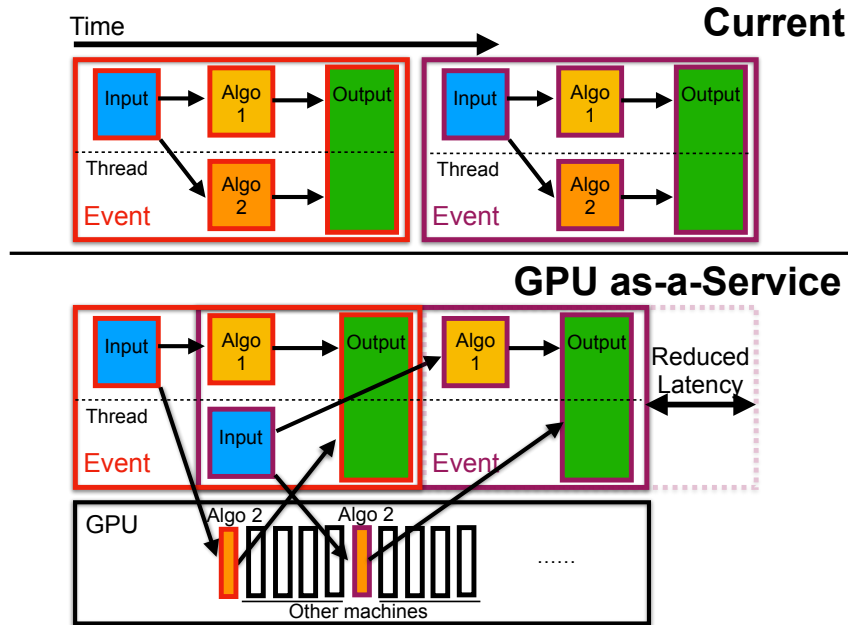


Figure 2. Diagram comparing the traditional LHC production model on CPUs (upper) with the GPUaaS approach (lower). Each block represents a module within the reconstruction framework. For the GPUaaS approach, algorithm 2 is run on the GPU, which allows the processing of the second event (outlined in purple) to run concurrently with the first event (outlined in red).

3.1. LHC Reconstruction

Detectors at the LHC are general-purpose devices with millions of channels, each of which records information from a particle passing through or decaying within it. Event reconstruction involves optimally combining these the individual signals from different detector channels to form the set of observable particles, including their energy, momentum, and type, for each event. This collection of particles is then used to infer the underlying physics process. For example, an event containing a Higgs boson decaying to a bottom quark-antiquark pair can lead to roughly 100 particles and we can use the aggregate properties of these particles to infer the presence of a Higgs boson. The variety of particles with different signatures in each detector (physics objects) that may be present in any given collision leads to a large number

of different reconstruction algorithms that must be run on each event as each physics object typically has its own reconstruction algorithm. This, in turn, leads to a large codebase that is written entirely for CPUs.

Parallelization of the reconstruction algorithms that create particles is possible by splitting the reconstruction into separate geometric regions and reconstructing the individual particles within that region. Further parallelization is possible through the separate reconstruction of the individual detectors before they are aggregated into particles. The current reconstruction aims to exploit possible parallelization by compartmentalizing separate reconstruction algorithms into modules that can be run in parallel. No single algorithm dominates the overall computing time, but some fundamental tasks, such as tracking and clustering detector hits to form particles, are the most computationally intensive. The potential for parallelization has only partly been realized through standard CPU optimizations, such as auto-vectorization. In this work, large-scale batching of the reconstruction to allow for algorithm-level parallelism is achieved through the use of DL algorithms on the GPU.

3.2. As-a-service computing

To apply DL under the as-a-service paradigm, we choose an algorithm that has a significant speedup when using a GPU. We then take this algorithm and set up a GPU inference server using the NVIDIA Triton Inference Server [29]. This package uses a custom gRPC-based communication protocol [30], and it supports load-balancing between multiple GPUs hosted together in a single server. Inference requests can be made for models from various ML frameworks, and multiple models can be loaded on the same server.

Software frameworks in HEP are typically written in C++. The software framework for the CMS experiment, CMSSW [31], uses task-based multithreading enabled by tbb [32]. This facilitates asynchronous, non-blocking calls to external resources using a feature called ExternalWork [33]. This is the most efficient way to utilize coprocessors as a service because the CPU running the experiment software can perform other tasks while the service call finishes. For this paper, we have taken advantage of these features by extending the CMSSW version of the SONIC software to perform remote gRPC calls to GPUs via the Triton Inference Server. In the SONIC approach, only the client code needs to be provided in the software framework. This minimizes the maintenance burden, as the client code has just two responsibilities: converting between the experiment and server data formats, and making the call to the server via the chosen communication protocol. All the details of the model architecture, any optimizations, and even the choice of coprocessors can be decided on the server without any change in the client. This setup enables the modified reconstruction workflow depicted in Figure 2.

By extending the SONIC framework to handle the gRPC calls utilized by the NVIDIA Triton Inference Server, the new client code uses a standard interface such that the user-developed software to convert between experiment and server data formats remains independent. Beyond the specification of the remote server protocol and location within a global configuration file, the user code remains completely intact, and switching between the remote FPGA calls, remote GPU calls, and other local calls are done seamlessly through a configuration file.

3.3. Metrics for optimization

To determine the cost-effectiveness of deploying a given algorithm as a service, we compose a simplified heuristic. We assume a computing model similar to those used by LHC experiments, which schedules modules to run during the service request, as in Figure 2. We introduce the GPU-to-CPU replacement ratio $F_{\text{GPU}}^{\text{eq}}$ to maintain the same throughput:

$$F_{\text{GPU}}^{\text{eq}} = \frac{X - S - L}{Y}, \quad (1)$$

where X is the algorithm processing time on the CPU, S is the overhead time due to input/output packaging in SONIC, and $L = f(Y + T)$ is the rescheduling time as a function of the algorithm processing time on the GPU Y and the packet transfer time T . For instance, a value of $F_{\text{GPU}}^{\text{eq}} = 32$ implies that one GPU can replace 32 CPUs at no cost in the overall event throughput. The optimal value of $F_{\text{GPU}}^{\text{eq}}$ depends on the demands of the system design, as well as the algorithm- and software-dependent values for both S and L . Time spent in data transfer or queuing on the server plays a small role in total throughput because of the asynchronous, non-blocking call employed in SONIC.

We use $F_{\text{GPU}}^{\text{eq}}$ as a guide to contextualize our results for GPU acceleration for each of the different scenarios studied. It is derived provided that no substantial bottlenecks are present in the software infrastructure, and further studies will refine this model. In the following sections, we explore the GPU speedups utilizing the SONIC framework for algorithms with various $F_{\text{GPU}}^{\text{eq}}$ values. We discuss the discovered bottlenecks and present a path towards a realistic implementation of GPUaaS at the LHC. An interesting potential extension of these studies would be to systematically investigate the relationship between $F_{\text{GPU}}^{\text{eq}}$ and throughput and latency.

4. Algorithms

To investigate the scalability of deploying DL as a service for LHC experiments, we study three distinct algorithms. Together, these algorithms span LHC computing, from low-level tasks of local detector energy reconstruction to high-level tasks of offline object identification. They also exhibit a range of speedups on coprocessors. Each algorithm performs as well as a CPU reference algorithm at resolving physical quantities. We then accelerate these algorithms with GPUaaS in realistic LHC workflows.

While the emphasis in this paper is on DL algorithms because optimized GPU implementations already exist, many LHC algorithms are currently not ML-based and likely will remain that way in the future. Nonetheless, many of these tasks have been shown to benefit significantly in computational performance if deployed on coprocessors with custom implementations [34]. The technology we develop for the ML algorithms as a service is flexible and its extension to non-ML algorithms is straightforward.

4.1. Hadron Calorimeter Reconstruction

The simplest algorithm that we study is called Fast Calorimeter Learning (FACILE), a deep neural network consisting of 2,000 parameters. This algorithm was trained on simulated collisions at the LHC using generator-level information to reconstruct

the energy deposited by particles in the hadron calorimeter (HCAL) of the CMS experiment. FACILE uses 15 inputs that contain information about raw charge deposits, geometry, and the gain of the HCAL channel in question. FACILE consists of batch normalization layers and dense layers with rectified linear unit (ReLU) activation functions. The layers consist of 30, 20, 10, 5, and 3 neurons each, respectively.

The HCAL is a core component of LHC experiments and a prototypical subdetector for which to implement ML as-a-service reconstruction for several reasons.

First, good resolution in the HCAL is important for sensitive measurements in particle physics, such as events with a Higgs boson decaying to bottom quarks. We find that local and global objects reconstructed with FACILE have as good or better resolution compared to the nominal algorithm that does not use ML. Second, the nominal HCAL reconstruction algorithm in CMS requires 60 ms of CPU time, accounting for approximately 15% of the online computing budget [35]. FACILE offers a significant improvement in computing performance when operated as a service by reducing the CPU time to less than 7 ms, resulting in an estimated $F_{\text{GPU}}^{\text{eq}} = 27$, which we verify experimentally. The remote time (including 2 ms of GPU latency) is largely eliminated by the asynchronous, non-blocking ExternalWork feature employed by SONIC. Finally, by exploiting GPU performance for large batch sizes, FACILE offers enhanced physics potential by reconstructing all 16,000 HCAL channels in parallel with little added latency, instead of reconstructing only the highest energy channels.

In terms of physics and computational performance, FACILE is well-suited for both online and offline applications. We deploy it in both settings. For instance, we perform a high-bandwidth test designed to emulate, for the first time at scale, a realistic LHC online computing system with coprocessors as a service.

4.2. Electron Regression

DeepCalo is a midsize convolutional neural network (CNN) trained for electron energy regression for the ATLAS detector. It operates at a higher, more abstract level compared to FACILE since it reconstructs the energy from an entire region of a calorimeter subdetector. Compared to nominal techniques, DeepCalo improves energy resolution and robustness against pileup, both of which are important for the HL-LHC [36]. The model is trained on electrons reconstructed from a Monte Carlo simulation of collisions spanning a wide range of energies. Each collision deposits energy in the electromagnetic calorimeter (ECAL) cells. These energy deposits are encoded as a 56×11 pixel image with 4 channels that represents a 2D patch of the detector of width 0.175 in η and 0.270 in ϕ . The 4 channels represent 4 separate layers of the ECAL and each pixel value represents the amount of energy deposited at that location and in that layer in η and ϕ . Using these images, DeepCalo estimates the energy of the electron.

DeepCalo is composed of 1.8 million parameters. The first component of the CNN consists of 5 convolutional blocks. The first block performs a 5×5 convolution, followed by batch normalization and a leaky ReLU activation function. Each subsequent convolutional block performs a 2×2 maximum pooling, followed by two instances of a sub-block consisting of a 3×3 convolution, batch normalization, and a leaky ReLU activation function. The final component of DeepCalo consists of three fully-connected layers, with the last layer producing a prediction for the electron energy.

In this study, we deploy DeepCalo as a service on GPU coprocessors for offline

reconstruction. In the offline test, we maximize the event throughput and compare the performance to on-site, CPU-based implementations. When deployed as a service, DeepCalo shows significant performance gains by reducing the latency per event from 75 to 1.5 ms and, when optimized, to 0.1 ms. This yields an estimated $F_{\text{GPU}}^{\text{eq}} = 50$ (750) before (after) optimization.

4.3. Top Quark Tagging

ResNet-50 is a CNN composed of 23 million parameters, 49 convolutional layers of 7×7 , 3×3 , and 1×1 convolutions with “skip connections,” and 1 fully-connected layer, which predicts 1000 class probabilities for natural images [37]. In earlier studies [26], the ResNet-50 CNN architecture was re-purposed to identify events containing top quarks (top quark tagging). In addition, CMS has implemented similar CNN-based top quark tagging algorithms for offline reconstructions [38]. Another study [39] showed that ResNet-50 could be modified to perform top quark tagging with performance rivaling leading ML algorithms. Of the three algorithms, ResNet-50 is the most complex and has the longest latency on GPU, and we estimate $F_{\text{GPU}}^{\text{eq}} = 150$. We choose it to enable benchmarking of a CPU-prohibitive algorithm as a service. In particular, ResNet-50 has a CPU latency of the order of seconds, which is prohibitively high for use even in offline reconstruction scenarios.

In Ref. [26], we observed a speedup by orders of magnitude by deploying ResNet-50 as a service on FPGA coprocessors. In this study, we extend our earlier studies by deploying ResNet-50 as a service on GPU coprocessors in LHC workflows. This enables top quark tagging to be performed in offline reconstruction. ResNet-50 also serves as a prototypical large benchmark algorithm comparable in burden to other major tasks in LHC computing, such as tracking. The specifications and GPU utilization of the three algorithms are summarized in Table 1.

Table 1. Summary of the specifications of each algorithm, including model parameters, GPU memory usage, and GPU utilization. The memory usage and GPU utilization are quoted at the point of maximum GPU throughput.

Algorithm	Batch size	Architecture type	Trainable parameters	Number of layers	GFLOP per batch	GPU memory usage [GB]	GPU utilization [%]
FACILE	16,000	Dense	2k	5	0.032	1	20
DeepCalo	5	Convolutional	2M	13	0.43	2.6	40
ResNet-50	10	Convolutional	23M	50	39	12	95

5. GPU Performance Studies

For online computing, we integrate FACILE into the HLT, the second tier of CMS data acquisition. For offline computing, we consider all three algorithms in stand-alone workflows. The client is implemented based on SONIC in CMSSW with real inputs and conditions. We quantify the hardware and networking requirements to run these algorithms as a service in LHC computing. To achieve this, we measure the coprocessor throughput (in events processed per second), the number of servers and GPUs required to service a given number of clients, and the network bandwidth limitations (arising from on-premises and external sources) in both LHC computing clusters and on the Google Cloud Platform.

We focus on achieving a hardware-efficient deployment of the algorithms by monitoring server properties and GPU utilization. We also measure how throughput scales with the number of GPUs by deploying many GPUs on a single server with a customized Google Kubernetes engine setup (as described in Section 5.3). Finally, we investigate various optimizations for the GPUs to further increase the throughput. Ultimately, in Section 6, we apply our findings to determine the hardware and networking requirements to perform full-scale LHC computing with coprocessors as a service.

5.1. Online Computing

To study the use of coprocessors as a service in online computing, we run the full CMS HLT with local HCAL reconstruction performed by FACILE as a service. FACILE is particularly well-suited for an online computing application because the algorithm it replaces is responsible for 15% of the HLT latency per event. In this study, the clients are deployed as jobs running single-thread HLT instances on virtual machines in Google Cloud using the HEPCloud framework [40–42]. HEPCloud deploys jobs submitted on batch systems to CPU instances created dynamically at a cloud computing site. The jobs are synchronized by adding a waiting period such that each job begins processing information only when all jobs are ready. This ensures that all jobs send calls to the GPU server during the same time period, enabling an accurate measurement of GPU and network throughput. Since FACILE has a small GPU latency (2 ms) compared to the HLT (500 ms), it proved essential to run on CPUs absent of other jobs for a realistic emulation of the current system of dedicated HLT cores. The cloud enabled this by reducing systematic uncertainties arising from shared CPUs on-premises. The server was deployed at the same site and consisted of a Google compute instance with either 1 or 4 NVIDIA Tesla V100 GPUs. This client-server configuration realistically emulates a fraction of the dedicated HLT CPU farm at CERN with the addition of as-a-service computing.

The results of this test are shown in Figure 3. Each client is allotted 7,000 simulated LHC benchmark timing events. The timing distribution for the HLT running FACILE as a service is shown in the top panel for servers with 1 or 4 GPUs in red and blue violins. The average time to run the nominal HLT algorithm locally on the CPU is shown in a dotted black line. For fewer than 500 clients, a decrease of approximately 10% in the total time is observed with FACILE as a service with 1 GPU when compared to the nominal algorithm. This largely eliminates the CPU burden of HCAL reconstruction. Since the HLT farm at CERN operates under latency restrictions, this demonstrates an opportunity to increase the throughput of the current trigger system by 10%, or alternatively, partitioning 10% of the existing machines to be used for other tasks. An increase in aggregate HLT latency occurs only above 300 clients for a single GPU, and above 1,000 clients for 4 GPUs. This increase represents the point where GPU throughput limitations begin to dominate, indicating that at least 300 HLT instances can be serviced by a single GPU without penalty. This slightly exceeds our expectation of 180 HLT instances, based on our computation of $F_{\text{GPU}}^{\text{eq}} = 27$ divided by the 15% CPU time fraction, but confirms the overall scaling. As a result, we conclude that operating reconstruction as a service is more efficient than having GPUs directly connected to CPUs. We explore this further by describing a scale design in Section 6. We note that the long tails in the figure are caused by scheduler assignments where fewer jobs are run on certain machines,




Figure 3. The distribution of the total time (including median and whiskers) to run the high level trigger with the HCAL reconstruction performed with FACILE as a service (upper). Servers with 1 and 4 GPUs are shown as red and blue violins. The average time taken to process the same events using the nominal HCAL reconstruction is shown as a dotted black line. High level trigger throughput running FACILE for servers with 1 and 4 GPUs in red and blue markers, respectively (lower).

leading to improved throughput for a small number of jobs, but a negligible effect in overall throughput.

The HLT throughput with FACILE as a service is shown in the bottom panel of Figure 3. For the single GPU server (red circles), the throughput asymptotes above 300 simultaneous processes, while for the 4 GPU server (blue triangles), it does not yet asymptote even up to 1,000 simultaneous processes.

5.2. Offline Reconstruction

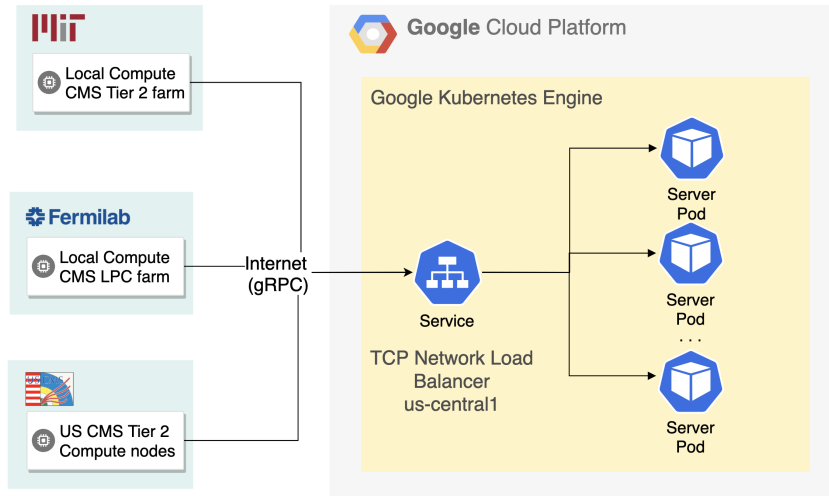


Figure 4. Diagram of the architecture for large-scale processing using GPU as a service. In this scenario, a GPU server within Google Cloud (right) is used to serve many offline computing centers processing LHC data (left). The calls are sent remotely over the internet as gRPC requests.

In the offline computing scenario, a single GPU service can be used by several remote computing clusters at the same time as depicted in Figure 4. We investigate the use of FACILE, DeepCalo, and ResNet-50 for LHC offline computing by executing a dedicated workflow process for each model. Our tests assume a benchmark LHC computing throughput of 5,000 events per second, and we estimate the coprocessors necessary to attain this. The processing of each model includes realistic input, formatting, and output steps. To emulate a realistic global offline computing scenario with CPU workers, as shown in Figure 4, we deploy clients to CPU clusters at MIT and Fermilab. These CPUs send gRPC requests over the internet to servers in Google Cloud’s `us-central1-a` zone in Council Bluffs, Iowa. While not shown here, we repeated these same tests on-premises going from on-site CPUs to the GPU with Google Cloud and we observed nearly identical throughput saturation and networking effects to the tests observed when going from a remote location to the same GPU within Google Cloud. This implies that communication over distance is reliable at the network bandwidths of interest.

5.2.1. FACILE The throughput of FACILE as a service is shown for different numbers of clients and GPUs in Figure 5. A server with a single GPU is found




Figure 5. The throughput of FACILE operated as a service in events per second versus the number of simultaneous clients. The marker indicates the mean and the error bars indicate the standard deviation. The throughput tests are repeated with 1, 4, 8, and 24 GPUs connected to the server, represented by red circles, orange triangles, green diamonds, and pink hexagons, respectively. The network bandwidth from the transfer of inputs from the client to the server is shown on the right vertical axis.

to saturate at a throughput of 500 inferences per second for a V100 GPU. This limit is due to the hardware latency and occurs above 50 clients. The V100 GPU is used because it offers a 10–20% gain over other GPU models.

As we increase the number of GPUs on the server using a customized Google Kubernetes Engine setup (see 5.3.2, we find that the throughput scales linearly and with high efficiency. Servers with 4 and 8 GPU saturate at approximately 2,000 and 4,000 inferences per second, as shown in Figure 5, respectively. Therefore, the LHC throughput requirement can be satisfied by a single 10 GPU server. This indicates that the Google Kubernetes Engine employed here is an efficient way to increase the throughput. The 24 GPU server test with 2,000 clients, in pink, is designed to probe the limit on network bandwidth between the LHC clusters and Google Cloud. This test becomes limited by a network bottleneck of unknown origin and we observe a peak bandwidth exceeding 70 Gb/s. The number of clients at which saturation occurs also scales with the number of GPUs; for example, the 4 GPU server in orange does not saturate until nearly 500 clients. We note that we were not able to plot out the entire throughput distributions due to the expense of each test.

The throughput is highly sensitive to the server configuration. Initially, we deployed a server with a single 4 CPU ingress node handing off the request to nodes with GPUs. These tests proved to be limited to a throughput of 1,500 inferences per second (12 Gb/s) regardless of client number, indicating there was a bandwidth limitation at the destination rather than between MIT and Google Cloud. As a result, we iteratively reconfigured our server to deploy multiple machines behind a




Figure 6. The throughput of DeepCalo as a service on an NVIDIA V100 GPU in events per second versus the number of simultaneous clients. The markers indicate the mean throughput and the error bars indicate the standard deviation. With server side optimizations, the batch size is configured dynamically to prefer a batch of 250 or 500 (as explained in Section 5.3.1) and to use five concurrent model executions. The network bandwidth from the transfer of inputs from the client to the server is shown on the right vertical axis.

load balancer, as described in Section 5.3.2.

5.2.2. DeepCalo As DeepCalo performs an image classification task, we expect it to be computationally bound rather than bandwidth limited. In our studies, we investigate the application of DeepCalo in offline reconstruction, which is throughput limited. We evaluate the performance of running DeepCalo as a service by running up to 1,000 clients on-premises on Fermilab’s computing cluster and deploying a GPU server in Google cloud. We set the batch size to 5 because this is the approximate number of electrons expected per reconstructed collision in a realistic LHC scenario.

We consider the case of a single NVIDIA V100 GPU server deployed on Google Cloud. The results are shown in Figure 6. For a batch size of 5, the throughput increases rapidly until 20 simultaneous clients, and it saturates at 680 events per second between 20 and 50 simultaneous clients. At 200 simultaneous clients, the utilization of the GPU saturates at 45% and the bandwidth peaks at 270 Mb/s, suggesting that the batch size is limiting GPU utilization. Further optimizations to DeepCalo, namely dynamic batching, are discussed in Section 5.3.1.

In a previous study, the latency on four Xeon E5-2698 CPUs was found to be 15 ms per electron—or 75 ms for an event of 5 electrons [36]. With our baseline GPU performance, we compute 680 events per second or 1.5 ms per event. We also observe a factor of 50 improvement in the throughput.

5.2.3. *ResNet-50* ResNet-50 is deployed as a service with clients on Fermilab’s computing cluster and servers with 1, 4, and 8 NVIDIA Tesla T4 GPUs in Google Cloud. The throughput obtained using ResNet-50 are shown in Figure 7. A single GPU server saturates at 25 (batch 10) inferences per second, at about 10 clients. We find that the throughput scales linearly with the number of GPUs, although with approximately 85% efficiency, slightly lower than for FACILE.




Figure 7. Throughput of ResNet-50 as a service in events per second versus the number of simultaneous clients. The markers indicate the mean throughput and the error bars indicate the standard deviation. The throughput tests are repeated with 1, 4, and 8 GPUs connected to the server, represented by yellow triangles, gold squares, and orange circles, respectively. The network bandwidth from the transfer of inputs from the client to the server is shown on the right vertical axis.

5.3. Optimizations

5.3.1. *Dynamic Batching* Dynamic batching is a feature of the Triton Inference Server that serves to increase both the throughput and hardware efficiency. It performs an added server-side queue of requests from the client until an optimal batch size is reached. The performance gains of dynamic batching are also related to the “instance groups,” or simultaneous executions of a model. This poses an optimization problem between the dynamic batch size and model concurrency. The use of dynamic batching is particularly interesting in that it circumvents the LHC computing paradigm of splitting computations on an event-by-event basis. Here, multiple events can be processed simultaneously within a single computation, without redesigning the computing model. We stress that this type of scheduling is *only* beneficial when GPUs are servicing a large number of parallel processes.

In the initial DeepCalo measurements, we found that the choice of batch size limited the utilization and throughput of the GPU. In our studies, we found that a

low number of model execution instances and a high dynamic batch size yielded the best throughput result. Figure 6 shows the throughput gains when using dynamic batching. At 200 simultaneous clients, the throughput shows no signs of saturation; at this point, the throughput is about 4,200 events per second, representing a factor of 6 improvement. When extended to 1,000 simultaneous clients, the throughput reaches 9,800 events per second, representing a factor of 14 gain in throughput, and the GPU utilization increases to 80%. At 1,000 clients, the bandwidth peaks at 3.9 Gb/s, which is roughly the same bandwidth limit observed with FACILE. On the other hand, dynamic batching for FACILE yielded no gain in throughput. We conclude that the most significant gains of dynamic batching are found for large models that naturally operate with small batch sizes.

5.3.2. Server optimization and monitoring We performed tests on many different combinations of computing hardware, which provided us with a deeper understanding of networking limitations within Google Cloud and on-premises data centers. Even though the Triton Inference Server does not consume significant CPU power, the number of CPU cores provisioned for the node did have an impact on the maximum ingress bandwidth reached in our early tests.

To scale the GPU throughput in flexibly, we deployed a Google Kubernetes Engine cluster for server side workloads. The cluster was configured using a Deployment and ReplicaSet [43], which control Pods, groups of one or more containers with shared storage and network and a specification to run the containers [44], and their resource requests. Additionally, a load balancing service was deployed which distributed incoming network traffic among the Pods. We implemented Prometheus-based monitoring [45] of overall system health and inference statistics. All metrics were visualized through a Grafana [46] instance, also deployed in the same cluster.

We note a Pods contents are always co-located and co-scheduled, and run in a shared context within Kubernetes Nodes [44]. We kept the Pod to Node ratio at 1:1 throughout the studies, with each Pod running an instance of Triton Inference Server (v20.02-py3) from the NVIDIA Docker repository.

It can be naively assumed that a small instance `n1-standard-2` with 2 vCPUs, 7.5 GB of memory, and different GPU configurations (1, 2, 4, 8) would be able to handle the workload. However, Google Cloud imposes a hard limit on network bandwidth per virtual CPU (vCPU). After performing several tests, we found that horizontal scaling would allow us to increase our ingress bandwidth since Google Cloud imposes a hard limit on network bandwidth at 2 Gb/s per vCPU up to a theoretical maximum of 16 Gb/s for each virtual machine [43].

Given these parameters, the ideal setup for optimizing ingress bandwidth was to provision multiple Pods on 16-vCPU machines with fewer GPUs per Pod. For GPU-intensive tests, we took advantage of having a single point of entry through Kubernetes load balancing and provisioning multiple identical Pods, where the sum of the GPUs attached to each Pod is the total GPU requirement.

5.3.3. Future optimizations Throughout these tests, we monitored the GPU utilization. ResNet-50 largely saturated GPU utilization, whereas FACILE and DeepCalo used 20% and 45% of the GPU, respectively. This indicates the throughput is batch limited. Throughput and GPU utilization can be optimized using dynamic batching for DeepCalo, as described above, but not for FACILE. Follow-up studies

will investigate optimizations for small models like FACILE.

6. Scaling

We now apply our findings to determine the GPU resources, networking and compute resources required for a modified LHC computing system in which algorithms with large speedups on coprocessors are deployed as a service. To assess the amount needed, we compute the required resources for the scaling of the three algorithms in either the HLT or offline computing cases. As a benchmark, we use the plateau performance numbers measured per GPU for each of the algorithms. We summarize these in Table 2. As an estimate of the total amount of required resources, we make some assumptions for the typical latency and throughput that we would expect for the online reconstruction and offline reconstruction. We emphasize that these numbers are approximate and used here for illustrative purposes.

Table 2. Summary of the three algorithms in terms of replacement with respect to a CPU ($F_{\text{GPU}}^{\text{req}}$), inferences per second, individual packet size (Mb), and network bandwidth per GPU (Gb/s). These numbers are quoted for the desired batch per event. The asterisk (*) identifies an algorithm using dynamic batching.

Algorithm	$F_{\text{GPU}}^{\text{req}}$	Inf./s [Hz]	Size/batch [Mb]	Throughput/GPU [Gb/s]
FACILE	27	500	7.7	3.9
DeepCalo	50	680	0.4	0.3
DeepCalo*	750	9,800	0.4	3.9
ResNet50	150	25	48.2	1.2

A typical LHC HLT consists of 1,000 servers, each with 32 cores, for a total of 32,000 CPU cores. This system is designed to process events at a rate of 100,000 events per second. The goal of the HLT is decide whether the event is sufficiently interesting to preserve for further reconstruction. As a consequence, the HLT performs a sequential, tiered reconstruction and filtering of the event and immediately rejects the event if it fails any filter in the sequence in order to prevent further reconstruction [47, 48]. With an emulated HLT, we find that a single GPU running FACILE can serve 300–500 different HLT nodes while preserving a 10% reduction in the per event throughput. This means that roughly 100 GPUs are needed to serve FACILE for the whole HLT. Moreover, if 100 GPUs were added to the system, 10% or 3,200 CPU cores could be removed from the system. We emphasize that FACILE represents only the tip of the iceberg in the use of deep learning in low-level reconstruction at the LHC. The algorithm uses less than a tenth of the GPU memory, and its GFLOP is less than a tenth of the other algorithms (see Table 1). As a result, it can be extended to carry out sophisticated reconstruction tasks (thus merging multiple algorithms) without large increases in latency.

While there is a significant reduction in the number of processing cores, there is an increase in network usage. With each GPU, an additional network bandwidth of 3.9 Gb/s is required. A server of 10 GPUs would thus require 40 Gb/s while simultaneously serving 100 HLT servers (3,200 cores). While this bandwidth is large,

it is already attainable. A setup of one 10 GPU server, serving 100 HLT servers, would be a logical design for the system that could be implemented with existing technology.

Lastly, we consider the option of adapting the DeepCalo algorithm or ResNet-50 to run in the HLT with our benchmark batch values. We expect that a 4 GPU server with a bandwidth of 16 Gb/s would be sufficient to run DeepCalo for the whole HLT system. ResNet-50 with the default batch size of 10 images would require 1,600 GPUs with a total added bandwidth of 1.9 Tb/s. The large number of GPUs and the large bandwidth would require significant and costly modifications to the design of the current HLT, making it impractical. However, a ResNet-50 implementation using a batch size of 1, or equivalently an inference rate below the expectation for a batch size of 1 by applying the algorithm only to some events, would result in a comparable number of GPU servers to FACILE. This conclusion meshes well with the fact that high-energy top quark candidates are relatively rare so it may be reasonable to run a ResNet-50 top quark tagging algorithm once or less per event.

Offline computing at the scale of a single LHC experiment consists of computing clusters that provide a total of roughly 150,000 CPU cores. Event reconstruction times are on the order of 30 s per event, yielding a throughput of 5,000 reconstructed events per second. If we were to run FACILE in this system, a single server of 10 GPUs operating with a bandwidth of 39 Gb/s would be able to sustain the full reconstruction load. Applying the same scaling for DeepCalo, we find that 1 GPU would be sufficient to run the reconstruction for a whole LHC experiment. Lastly, for ResNet-50 at a batch size of 10, a setup of 200 GPUs with 240 Gb/s bandwidth would be sufficient to support 150,000 cores. In this case, with ResNet-50 applied to all reconstructed events, a realistic scenario would consist of 10 separate GPU servers, each running with 24 Gb/s bandwidth. In contrast, for ResNet-50 inference on CPUs with a batch size of 10, the reconstruction time per event would increase by 18 s. This would require a 60% increase in the computing clusters or an additional 90,000 cores to sustain the same throughput.

In the context of LHC algorithms, DL algorithms are being developed at all levels of the detector reconstruction. Algorithms that run on aggregate event properties are comparable in size to ResNet-50. The full collection of particles in a collision after reconstruction is found to be on the order of 1,000 particles per event [49]. Given the number of particles, a computation of the event size ranges from 0.5–2.5 Mb, which is less than size of single event requests performed in tests with FACILE. Consequently, we observe that data rates and throughput for future LHC algorithms are comparable to that of the results presented with FACILE. Therefore, at no added complexity in networking or design, the framework presented in this paper can be extended for algorithms designed for the future particle reconstruction at the LHC.

7. Conclusion

We have demonstrated a core framework that enables the use of deep learning (DL) as a service with direct applications to the processing of LHC data. Our framework, Services for Optimized Network Inference on Coprocessors (SONIC), utilizes gRPC to perform asynchronous, non-blocking calls to a GPU server. Our server infrastructure can address both small and large scale use cases. With our infrastructure, we have tested three algorithms that span a large space of ML model complexities and batch sizes. Together, these algorithms serve as benchmarks for a wide array of LHC reconstruction tasks. In each case, we have measured the algorithm throughput

and demonstrated comparable throughput for as-a-service computing both remotely and on-site. We fully integrated a DL algorithm called Fast Calorimeter Learning (FACILE) for LHC reconstruction in the high-level trigger (HLT), the second tier of the LHC data processing and filter, and we found that this algorithm can lead to a 10% overall reduction in the computing demands. This latency reduction is almost equivalent to removing the hadron calorimeter reconstruction latency from the HLT entirely, and it matches the expected optimal performance when performing standalone algorithmic tests. Finally, we demonstrated the use of FACILE, DeepCalo, and ResNet-50 for offline reconstruction. A server implementation in the cloud was found to operate at data rates and inference times comparable to the demands set by LHC offline reconstruction. This is a concrete validation of the SONIC framework, demonstrating the viability of coprocessors as a service on representative scales for LHC computing.

While our focus was on accelerating DL algorithms with GPUs, this work can be applied to any algorithm that can be implemented on a GPU and appropriately integrated into a GPU server. This work is largely agnostic to the hardware and software implementations of the algorithm and can be adapted for other types of coprocessors and other scientific experiments. As DL accelerator tools are constantly evolving and improving, we expect the speedups observed in this paper to become even larger.

From our studies, we delineated certain considerations for designing an optimal system with GPUs-as-a-service (GPUaaS) for the LHC. In particular, an optimized scheduling framework is needed to ensure that remote operations of algorithms incur minimal losses in performance. Additionally, sufficient bandwidth is needed to ensure that the full performance of the accelerator servers can be achieved for both remote and local as-a-service operation. Finally, both a load balancer and an optimized and flexible server infrastructure are needed to ensure robust operation. In this paper, we have demonstrated that all of these demands can be met with existing resources.

In the context of physics performance, our results lead to direct performance improvements that can be implemented immediately in the LHC computing model. In particular, we found that: (1) DL inference as a service can enable a significant increase in event throughput, (2) algorithms with complexities not previously attainable can be operated in the LHC reconstruction workflow, (3) optimized versions of algorithms can be implemented without disrupting the existing computing model, and (4) simultaneous multi-event processing is achievable in the reconstruction workflow. Concurrently with these studies, an extensive suite of new DL techniques for LHC reconstruction have been developed [50–63]. The current work will enable the integration of these algorithms into the LHC computing model in a seamless and computationally efficient way.

We would like to stress that this work represents the beginning of developments in coprocessor computing both at the LHC and other large scale experiments. This work and related studies are encouraging for physicists in other fields, such as neutrino physics, gravitational wave detection, and astrophysics, to pursue a similar computing model. As a consequence, we believe that this work may lead to a paradigm shift in the scientific computing model, enabling us to meet the enormous scientific computing demands in the next decade.

Acknowledgments

P. H., and D. R. are supported by NSF grants #1934700, #1931469, and the IRIS-HEP grant #1836650. J. K. is supported by NSF grant #190444. Cloud credits for this study were provided by Internet2 managed Exploring Cloud to accelerate Science (NSF grant #190444). Additional support on networking was provided by Doug Burton, Dan Speck, and Aaron Strong of the Burwood group. Also, we thank Emma Levett from MIT for additional networking support. P. H. would like to thank Matthew Harris for discussion on GCP networking capabilities. M. A. F, B. H., T. K., M. L., K. P., N. T. are supported by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the U.S. Department of Energy, Office of Science, Office of High Energy Physics. K. P. is partially supported by the High Velocity Artificial Intelligence grant as part of the Department of Energy High Energy Physics Computational HEP sessions program.

References

- [1] ATLAS Collaboration, “Observation of a new particle in the search for the standard model Higgs boson with the ATLAS detector at the LHC”, *Phys. Lett. B* **716** (2012) 1, doi:10.1016/j.physletb.2012.08.020, arXiv:1207.7214.
- [2] CMS Collaboration, “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”, *Phys. Lett. B* **716** (2012) 30, doi:10.1016/j.physletb.2012.08.021, arXiv:1207.7235.
- [3] CMS Collaboration, “Observation of a new boson with mass near 125 GeV in pp collisions at $\sqrt{s} = 7$ and 8 TeV”, *J. High Energy Phys.* **06** (2013) 081, doi:10.1007/JHEP06(2013)081, arXiv:1303.4571.
- [4] CMS Collaboration, “Search for new physics in final states with an energetic jet or a hadronically decaying W or Z boson and transverse momentum imbalance at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **97** (2018) 092005, doi:10.1103/PhysRevD.97.092005, arXiv:1712.02345.
- [5] CMS Collaboration, “Search for high mass dijet resonances with a new background prediction method in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *J. High Energy Phys.* **05** (2020) 033, doi:10.1007/JHEP05(2020)033, arXiv:1911.03947.
- [6] CMS Collaboration, “Search for low mass vector resonances decaying into quark-antiquark pairs in proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Phys. Rev. D* **100** (2019) 112007, doi:10.1103/PhysRevD.100.112007, arXiv:1909.04114.
- [7] J. Alexander et al., “Dark sectors 2016 workshop: Community report”, in *Dark Sectors 2016 Workshop*. 8, 2016. arXiv:1608.08632.
- [8] HEP Software Foundation Collaboration, “A roadmap for HEP software and computing R&D for the 2020s”, *Comput. Softw. Big Sci.* (2019) 7, doi:10.1007/s41781-018-0018-8, arXiv:1712.06982.
- [9] CMS Collaboration, “CMS offline and computing public results”, 2020. <https://twiki.cern.ch/twiki/bin/view/CMSPublic/CMSOfflineComputingResults>.
- [10] ATLAS Collaboration, “ATLAS Computing and Software Public Results”, 2020. <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ComputingandSoftwarePublicResults>.
- [11] D. Guest, K. Cranmer, and D. Whiteson, “Deep learning and its application to LHC physics”, *Ann. Rev. Nucl. Part. Sci.* **68** (2018) 161, doi:10.1146/annurev-nucl-101917-021019, arXiv:1806.11484.
- [12] J. Pata and M. Spiropulu, “Processing collider data with GPU-accelerated kernels”, arXiv:1906.06242.
- [13] H. Schreiner et al., “GooFit 2.0”, *J. Phys. Conf. Ser.* **1085** (2018) 042014, doi:10.1088/1742-6596/1085/4/042014, arXiv:1710.08826.
- [14] ROOT Collaboration, “HistFactory: A tool for creating statistical models for use with RooFit and RooStats”, Technical Report CERN-OPEN-2012-016, 2012.
- [15] ALICE Collaboration, “Real-time data processing in the ALICE high level trigger at the LHC”, *Comput. Phys. Commun.* **242** (2019) 25, doi:10.1016/j.cpc.2019.04.011, arXiv:1812.08036.

- [16] R. Aaij et al., “Allen: A high level trigger on GPUs for LHCb”, *Comput. Softw. Big Sci.* **4** (2020) 7, doi:10.1007/s41781-020-00039-7, arXiv:1912.09161.
- [17] D. Vom Bruch, “Real-time data processing with GPUs in high energy physics”, *J. Instrum.* **15** (2020) C06010, doi:10.1088/1748-0221/15/06/C06010, arXiv:2003.11491.
- [18] D. Funke et al., “Parallel track reconstruction in CMS using the cellular automaton approach”, *J. Phys. Conf. Ser.* **513** (2014) 052010, doi:10.1088/1742-6596/513/5/052010.
- [19] F. Pantaleo et al., “Development of a phase-II track trigger based on GPUs for the CMS experiment”, in *2015 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, p. 7581775. 2016. doi:10.1109/NSSMIC.2015.7581775.
- [20] G. Lamanna, “Almagest, a new trackless ring finding algorithm”, *Nucl. Instr. Methods Phys. Res. A* **766** (2014) 241, doi:10.1016/j.nima.2014.05.073.
- [21] C. Fäerber, R. Schwemmer, J. Machen, and N. Neufeld, “Particle identification on an FPGA accelerated compute platform for the LHCb upgrade”, *IEEE Trans. Nucl. Sci.* **64** (2017) 1994, doi:10.1109/TNS.2017.2715900.
- [22] Mu3e Collaboration, “Online data reduction using track and vertex reconstruction on GPUs for the Mu3e experiment”, *Eur. Phys. J. Web Conf.* **150** (2017) 00013, doi:10.1051/epjconf/201715000013.
- [23] R. Ammendola et al., “Real-time heterogeneous stream processing with NaNet in the NA62 experiment”, *J. Phys. Conf. Ser.* **1085** (2018) 032022, doi:10.1088/1742-6596/1085/3/032022.
- [24] D. Chirkin et al., “Photon propagation using GPUs by the IceCube neutrino observatory”, in *15th International Conference on eScience, eScience 2019, San Diego, CA, USA, September 24-27, 2019*, p. 388. IEEE, 2019. doi:10.1109/eScience.2019.00050.
- [25] I. Sfligoi, F. Würthwein, B. Riedel, and D. Schultz, “Running a pre-exascale, geographically distributed, multi-cloud scientific simulation”, in *High Performance Computing*, p. 23. 2020. arXiv:2002.06667. doi:10.1007/978-3-030-50743-5_2.
- [26] J. Duarte et al., “FPGA-accelerated machine learning inference as a service for particle physics computing”, *Comput. Softw. Big Sci.* **3** (2019) 13, doi:10.1007/s41781-019-0027-2, arXiv:1904.08986.
- [27] K. Pedro, “SonicCMS”. [software] version v5.1.0 (accessed 2020-04-22) <https://github.com/hls-fpga-machine-learning/SonicCMS>, 2020. v5.1.0 (accessed 2020-04-22). <https://github.com/hls-fpga-machine-learning/SonicCMS>.
- [28] A. Caulfield et al., “A cloud-scale acceleration architecture”, in *2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, p. 1. IEEE, October, 2016. doi:10.1109/MICRO.2016.7783710.
- [29] NVIDIA, “Triton Inference Server”. [software] version v1.8.0 (accessed 2020-02-17) <https://docs.nvidia.com/deeplearning/sdk/triton-inference-server-guide/docs/index.html>, 2019. v1.8.0 (accessed 2020-02-17). <https://docs.nvidia.com/deeplearning/sdk/triton-inference-server-guide/docs/index.html>.
- [30] Google, “gRPC”. [software] version v1.19.0 (accessed 2020-02-17) <https://grpc.io/>, 2018. v1.19.0 (accessed 2020-02-17). <https://grpc.io/>.
- [31] CMS Collaboration, “CMS physics: technical design report Volume 1: Detector performance and software”, CMS Technical Design Report CERN-LHCC-2006-001, 2006.
- [32] Intel, “Thread Building Blocks”. [software] version 2018.U1 (accessed 2019-02-11) <https://www.threadingbuildingblocks.org>, 2018. 2018.U1 (accessed 2019-02-11). <https://www.threadingbuildingblocks.org>.
- [33] A. Bocci et al., “Bringing heterogeneity to the CMS software framework”, in *24th International Conference on Computing in High Energy and Nuclear Physics (CHEP 2019)*. 2019. arXiv:2004.04334. Submitted to *Eur. Phys. J. Web Conf.*
- [34] M. Rovere et al., “CLUE: A fast parallel clustering algorithm for high granularity calorimeters in high energy physics”, (2020). arXiv:2001.09761.
- [35] A. Massironi, V. Khristenko, and M. DAlfonso, “Heterogeneous computing for the local reconstruction algorithms of the CMS calorimeters”, *J. Phys. Conf. Ser.* **1525** (2020) 012040, doi:10.1088/1742-6596/1525/1/012040.
- [36] F. Faye, “Energy reconstruction of electrons and photons using convolutional neural networks”, Master’s thesis, University of Copenhagen, 2019.
- [37] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition”, in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, p. 770. IEEE, 2016. arXiv:1512.03385. doi:10.1109/CVPR.2016.90.
- [38] CMS Collaboration, “Identification of heavy, energetic, hadronically decaying particles using machine-learning techniques”, *J. Instrum.* **15** (06, 2020) P06005,

- doi:10.1088/1748-0221/15/06/P06005, arXiv:2004.08262.
- [39] A. Butter et al., “The machine learning landscape of top taggers”, *SciPost Phys.* **7** (2019) 014, doi:10.21468/SciPostPhys.7.1.014, arXiv:1902.09914.
- [40] B. Holzman et al., “HEPCloud, a new paradigm for HEP facilities: CMS Amazon Web Services investigation”, *Comput. Softw. Big Sci.* **1** (2017) 1, doi:10.1007/s41781-017-0001-9, arXiv:1710.00100.
- [41] M. Altunay et al., “Intelligently-automated facilities expansion with the HEPCloud decision engine”, in *2018 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*, p. 352. 2018. arXiv:1806.03224. doi:10.1109/CCGRID.2018.00053.
- [42] P. Mhashilkar et al., “HEPCloud, an elastic hybrid HEP facility using an intelligent decision support system”, *Eur. Phys. J. Web Conf.* **214** (2019) 03060, doi:10.1051/epjconf/201921403060.
- [43] Google LLC, “Compute Engine Documentation - Concepts - Virtual machine instances - Machine types”, 2020. <https://cloud.google.com/compute/docs/machine-types>.
- [44] Kubernetes Authors, “Documentation - Concepts - Workloads - Pods - Pods”, 2020. <https://kubernetes.io/docs/concepts/workloads/pods/pod/>.
- [45] Prometheus Authors, “Prometheus”, 2020. <https://prometheus.io/>.
- [46] Grafana Labs, “Grafana”, 2020. <https://grafana.com/>.
- [47] A. Perrotta, “Performance of the CMS High Level Trigger”, *J. Phys. Conf. Ser.* **664** (2015) 082044, doi:10.1088/1742-6596/664/8/082044.
- [48] CMS Collaboration, “CMS Trigger Performance”, *Eur. Phys. J. Web Conf.* **182** (2018) 02037, doi:10.1051/epjconf/201818202037.
- [49] CMS Collaboration, “Measurement of charged particle spectra in minimum-bias events from proton-proton collisions at $\sqrt{s} = 13$ TeV”, *Eur. Phys. J. C* **78** (2018) doi:10.1140/epjc/s10052-018-6144-y, arXiv:1806.11245.
- [50] CMS Collaboration, “A deep neural network to search for new long-lived particles decaying to jets”, doi:10.1088/2632-2153/ab9023, arXiv:1912.12238. Accepted by *Mach. Learn.: Sci. Technol.*
- [51] E. M. Metodiev, B. Nachman, and J. Thaler, “Classification without labels: Learning from mixed samples in high energy physics”, *J. High Energy Phys.* **10** (2017) 174, doi:10.1007/JHEP10(2017)174, arXiv:1708.02949.
- [52] B. Nachman and D. Shih, “Anomaly detection with density estimation”, *Phys. Rev. D* **101** (2020) 075042, doi:10.1103/PhysRevD.101.075042, arXiv:2001.04990.
- [53] A. Andreassen, B. Nachman, and D. Shih, “Simulation assisted likelihood-free anomaly detection”, *Phys. Rev. D* **101** (2020) 095004, doi:10.1103/PhysRevD.101.095004, arXiv:2001.05001.
- [54] J. H. Collins, K. Howe, and B. Nachman, “Anomaly detection for resonant new physics with machine learning”, *Phys. Rev. Lett.* **121** (2018) 241803, doi:10.1103/PhysRevLett.121.241803, arXiv:1805.02664.
- [55] J. H. Collins, K. Howe, and B. Nachman, “Extending the search for new resonances with machine learning”, *Phys. Rev. D* **99** (2019) 014038, doi:10.1103/PhysRevD.99.014038, arXiv:1902.02634.
- [56] M. Farina, Y. Nakai, and D. Shih, “Searching for new physics with deep autoencoders”, *Phys. Rev. D* **101** (2020) 075021, doi:10.1103/PhysRevD.101.075021, arXiv:1808.08992.
- [57] T. Heimel, G. Kasieczka, T. Plehn, and J. M. Thompson, “QCD or What?”, *SciPost Phys.* **6** (2019) 030, doi:10.21468/SciPostPhys.6.3.030, arXiv:1808.08979.
- [58] S. R. Qasim, J. Kieseler, Y. Iiyama, and M. Pierini, “Learning representations of irregular particle-detector geometry with distance-weighted graph networks”, *Eur. Phys. J. C* **79** (2019) 608, doi:10.1140/epjc/s10052-019-7113-9, arXiv:1902.07987.
- [59] S. Farrell et al., “Novel deep learning methods for track reconstruction”, in *4th International Workshop Connecting The Dots 2018 (CTD2018) Seattle, Washington, USA, March 20-22, 2018*. 2018. arXiv:1810.06111.
- [60] X. Ju et al., “Graph Neural Networks for Particle Reconstruction in High Energy Physics Detectors”, in *Machine Learning and the Physical Sciences Workshop at the 33rd Annual Conference on Neural Information Processing Systems*. 2019. arXiv:2003.11603.
- [61] H. Qu and L. Gouskos, “ParticleNet: Jet tagging via particle clouds”, *Phys. Rev. D* **101** (2020) 056019, doi:10.1103/PhysRevD.101.056019, arXiv:1902.08570.
- [62] E. A. Moreno et al., “JEDI-net: a jet identification algorithm based on interaction networks”, *Eur. Phys. J. C* **80** (2020) 58, doi:10.1140/epjc/s10052-020-7608-4, arXiv:1908.05318.
- [63] E. A. Moreno et al., “Interaction networks for the identification of boosted $H \rightarrow b\bar{b}$ decays”,

