

# Modeling a Large Data-Acquisition Network in a Simulation Framework

Tommaso Colombo<sup>\*†</sup>, Holger Fröning<sup>†</sup>, Pedro Javier García<sup>‡</sup> and Wainer Vandelli<sup>\*</sup>

<sup>\*</sup>Physics Department, CERN, Geneva, Switzerland

Email: {Tommaso.Colombo,Wainer.Vandelli}@cern.ch

<sup>†</sup>Institut für Technische Informatik (ZITI), Universität Heidelberg, Mannheim, Germany

Email: holger.froening@ziti.uni-heidelberg.de

<sup>‡</sup>Departamento de Sistemas Informáticos, Universidad de Castilla-La Mancha, Albacete, Spain

Email: PedroJavier.Garcia@uclm.es

**Abstract**—The ATLAS detector at CERN records particle collision “events” delivered by the Large Hadron Collider. Its data-acquisition system identifies, selects, and stores interesting events in near real-time, with an aggregate throughput of several 10 GB/s. It is a distributed software system executed on a farm of roughly 2000 commodity worker nodes communicating via TCP/IP on an Ethernet network. Event data fragments are received from the many detector readout channels and are buffered, collected together, analyzed and either stored permanently or discarded. This system, and data-acquisition systems in general, are sensitive to the latency of the data transfer from the readout buffers to the worker nodes. Challenges affecting this transfer include the many-to-one communication pattern and the inherently bursty nature of the traffic. In this paper we introduce the main performance issues brought about by this workload, focusing in particular on the so-called TCP incast pathology. Since performing systematic studies of these issues is often impeded by operational constraints related to the mission-critical nature of these systems, we focus instead on the development of a simulation model of the ATLAS data-acquisition system, used as a case study. The simulation is based on the well-established OMNeT++ framework. Its results are compared with existing measurements of the system’s behavior. The successful reproduction of the measurements by the simulations validates the modeling approach. We share some of the preliminary findings obtained from the simulation, as an example of the additional possibilities it enables, and outline the planned future investigations.

## I. INTRODUCTION

ATLAS [1] is a high-energy physics experiment installed at CERN, Geneva, Switzerland. Its scientific program is mainly focused on the discovery and study of rare particle physics phenomena. The experiment’s detectors observe proton-proton collision events delivered by the Large Hadron Collider (LHC) at a design frequency of 40 MHz. Each event corresponds to 1-2 MB of data, which are used to reconstruct the physical processes produced by the collisions. If all collision events were to be read out and acquired, ATLAS would produce 80 TB/s.

The sheer amount of data produced by ATLAS cannot be all stored for later analysis. Instead, as with many other large scale experiments, a real-time data-acquisition and data-selection system is necessary. These systems are usually implemented with a mix of custom hardware and software running on commercial off-the-shelf (COTS) hardware, and their performance and reliability have a strong impact on the experiment as a whole. Failures and poor performance result

in the permanent loss of extremely valuable experimental data.

Due to the mission-critical nature of these systems, a systematic study of their performance envelope is often impeded by operational constraints, such as system availability requirements or limited opportunities of performing hardware or system software modifications. A simulation model can thus be a worthwhile alternative, assuming that it is accurate enough to reliably reproduce the key traits of the system.

This publication reports on the development of a simulation of the ATLAS data-acquisition system, used as a case study. The most significant source of performance degradation in data-acquisition networks is described and measurements showing its significance are presented. The paper introduces the basic assumptions underlying the simulation model and presents the validation tests undertaken to ensure that the model can reliably reproduce the behavior of the real system. Finally, some preliminary insights gained through the simulation are reported and the future investigations enabled by the simulation are outlined.

## II. BACKGROUND: THE ATLAS TRIGGER AND DATA-ACQUISITION SYSTEM

The ATLAS Trigger and Data-Acquisition (TDAQ) system [2] is responsible for the selection of interesting collision events (*triggering*, in high-energy physics jargon) reducing the initial frequency of 40 MHz to ~1 kHz of stored events. This requires an overall trigger rejection factor of the order of  $10^4$  against unremarkable events, while retaining potential candidates containing new physics processes, such as Higgs boson decays. The TDAQ system, outlined in Figure 1, is based on the combination of a hardware-based first stage and a software-based second stage.

The first stage, called Level-1 trigger, is a synchronous pipelined electronics system, with a guaranteed maximum latency of 2.5  $\mu$ s. It selects events using coarse-grained data from a subset of the experiment’s detectors. It triggers the readout at a maximum rate of 100 kHz. As part of the selection process, it identifies regions-of-interest (RoI): parts of the detector that recorded interesting signals that are used as a seed for the second stage.

When an event is accepted by the Level-1, its data fragments (1860 fragments of variable sizes around 1 kB, corresponding to different regions of the experiment) are readout by the detector electronics. Each fragment is pushed via a custom point-to-point optical link into a specific hardware

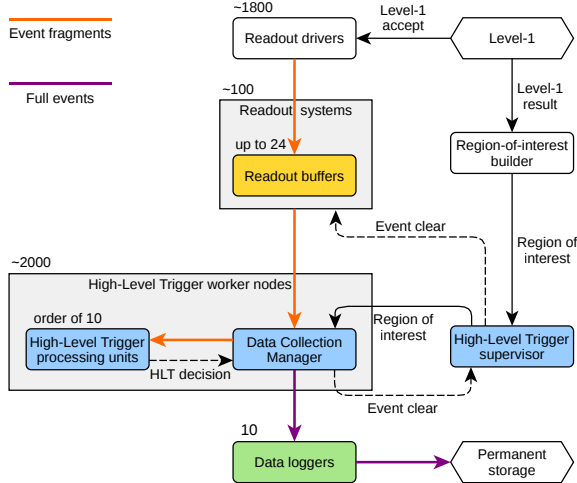


Figure 1. Logical message flow in the ATLAS Trigger and Data-Acquisition system (the experiment is the input of the Readout Drivers and the Level-1).

buffer in one of the  $\sim 100$  “readout system” nodes, which act as the coupling between the first stage’s custom electronics and the second stage’s COTS hardware.

The second stage, called High-Level Trigger (HLT), is a distributed software system running on around 2000 Linux PCs interconnected by an Ethernet network. Each HLT worker node hosts one HLT “processing unit” per CPU core and a single “data-collection manager”, which handles communications with the rest of the system on behalf of the node’s processing units. A central scheduler, the HLT “supervisor”, receives regions-of-interest information from the Level-1, and assigns it to one of the available HLT processing units. Using the regions-of-interest as starting points, the processing unit incrementally retrieves and analyses event fragments, until a decision can be taken. The event can be rejected even without analyzing all its fragments, thus limiting the fraction of data to be retrieved. Fragments of rejected events are deleted from the readout systems buffers, while accepted events are transferred to one of the Data Logger nodes for storage. The rate of events that can be accepted is mostly determined by the availability of computing and storage resources for subsequent data analysis, and is foreseen to be around 1 kHz for the 2015–2018 data taking period.

The physical layout of the Data-Acquisition and High-Level Trigger system is represented in Figure 2. The core of the system consists of two large network routers with a maximum capacity of several hundred 10GbE ports. Readout systems are directly connected to both core routers with 4 10GbE links. HLT worker nodes are organized in racks of at most 40 nodes. Each node in a rack is connected to an aggregation switch with a GbE link. The rack switches have 10 GbE links to both core routers. For the rationale behind this design, please refer to [3].

### III. PERFORMANCE ISSUES IN DATA-ACQUISITION NETWORKS

#### A. Traffic pattern

Data-acquisition networks have to deal with a particularly problematic traffic pattern:




Figure 2. Network architecture of the ATLAS Data-Acquisition and High-Level Trigger system.




Figure 3. Visualization of the potential network congestion issues.

- The communication pattern is many-to-one.
- Data are transmitted in multiple bursts, rather than a smooth flow.

These two characteristics do not reflect a design error, but rather the nature of data acquisition itself. The goal is to gather together data fragments from different components of the experiment, hence the many-to-one communication. The data are transferred right after the experiment has generated them, hence the burstiness.

In the case of ATLAS, event data are striped over all the readout systems, since each node buffers data from a specific region of a detector. A single HLT processing unit usually requests fragments from multiple readout systems at the same time. As the fragments are already available in the readout systems’ buffers and are sent as soon as the request reaches the readout systems, many nodes will start sending fragments at the same time to the same destination, thus creating instantaneous network congestion (see Figure 3).

#### B. The TCP incast pathology

As mentioned in Section I, the ATLAS data-acquisition system is based on Ethernet. The messaging among applications relies on the TCP protocol. The combination of TCP, a lossy link layer such as Ethernet, and the traffic pattern described above is subject to a well-known TCP pathology called *incast* [4]: first observed in data-center storage networks, it occurs “when a client simultaneously receives a short burst of data from multiple sources, overloading the switch buffers

associated with its network link such that all original packets from some sources are dropped”.

TCP has both an active packet-drop detection mechanism, in which the receiver detects the lost packet and causes the sender to react very quickly, and a passive mechanism, which is based on retransmission timeouts (*RTO*) on the sender side [5]. The *RTO* value is calculated starting from the estimated round-trip time [6]. However, to protect against spurious retransmissions and interference with other TCP features such as delayed acknowledgments, the *RTO* has a fixed minimum, specified as 1 s in [6], with lower values used in actual implementations (200 ms in Linux, 30 ms in FreeBSD). These values are orders of magnitude larger than the round-trip time of a data-center network, which is usually in the sub-millisecond range. In addition, the active mechanism requires at least some of the original packets to reach the destination for the packet loss to be detected. When this is not the case, the delivery of packets incurs the large delay caused by the *RTO*.

### C. Impact and mitigation

The data-collection latency (i.e., the latency of the data transfer between the data sources and the processing nodes) is critical to the performance of a data-acquisition system as a whole. In the case of ATLAS, each of the HLT processing units operates exclusively on the single event assigned to it, with the HLT selection proceeding iteratively starting from the region-of-interest identified by the Level-1, collecting data fragments incrementally as needed. A processing unit is blocked while it waits for the data fragments to be collected, which means that the data-collection latency effectively translates into lost CPU time. With average event processing times of the order of 100 ms, waiting for a TCP *RTO* is very wasteful.

The TCP incast problem is well studied in literature, with many solutions being proposed. In general the solutions point in the direction of either specialized hardware (i.e. switches with huge buffers), or alterations of the TCP implementation. Due to the operational constraints mentioned in Section I, none of these can be readily deployed in the ATLAS data-acquisition system, as they would require significant modifications to the existing hardware or to core system software (specifically the kernel TCP implementation).

A less invasive, but less general, solution is application-level traffic shaping. In ATLAS, smoothing the rate of data requests generated by a HLT node can alleviate the network congestion by controlling the maximum size of the traffic burst from the readout systems [7]. Obviously such a smoothing mechanism imposes a trade-off: excessive smoothing can increase the data-collection time by unnecessarily delaying the requests for data, whereas insufficient smoothing will not eliminate packet drops.

The ATLAS data-acquisition software currently uses a credit-based traffic-shaping algorithm, implemented in the Data-Collection Manager running on every HLT node. Its basic rules are as follows.

- Each HLT worker node has a fixed number of credits available. All the HLT processing units on a node share these credits.

- Each data request from a processing unit to a readout system uses as many credits as the number of fragments it asks for.
- Each response returns the credits used by the corresponding request.
- If all available credits are used, further requests are blocked until the necessary credits become available.

The number of fragments in a request gives a rough estimation of the size of the corresponding response. Therefore, this algorithm effectively limits the maximum burst size of data transfers directed to the same HLT node. However, it relies on the assumptions that all event fragments are similar in size, and that this size is known beforehand. These assumptions are reasonable for ATLAS under normal operating conditions, but not necessarily for other systems or scenarios.

## IV. MEASUREMENTS

Detecting buffer overflows in an Ethernet network is a relatively simple task: the Linux kernel provides the total number of packet retransmissions that have occurred in a TCP connection, and some commercial switch models report cumulative counts of dropped packets per port via SNMP. In-depth analysis of the root causes is not as straightforward: important metrics, such as switch buffer occupancies, are either not available at all or too coarse-grained to be useful. This, combined with a protocol with sophisticated congestion control algorithms, such as TCP, leads to an extremely intricate landscape. On the other hand, generating synthetic traffic patterns, which can be tightly controlled and are known in advance, proves instrumental in reducing the complexity down to a manageable level. The tests described in the following paragraphs employ this approach.

### A. Test set-up

The measurements were performed using some of the hardware available in the ATLAS data-acquisition system. At the time of these tests, the infrastructure was still under consolidation, so the test set-up is slightly different from the one described in Section I. In particular, readout systems were not directly connected with 10GbE links to the core routers. They were instead connected with GbE links to an intermediate aggregation switch with a 10GbE uplink to each core. The intermediate switch was kept under-subscribed so that no congestion could appear there.

Given these constraints, the following test set-up, shown in Figure 4, was chosen:

- 10 readout system groups:
  - Each group consists of 16 nodes connected to one switch (160 total).
  - 12 event fragments of 1.1 kB are served by each node (1920 total fragments, 2112 kB full event size).
- 1 HLT rack:
  - It consists of 39 PCs connected to one switch.
  - Each PC hosts 24 HLT processing units (936 total).

This configuration was chosen because it provides a realistic model of the expected network buffer usage in the final system topology. In particular, the congestion at the rack-level switch is well represented. The small amount of HLT nodes in




Figure 4. Test set-up.

use with respect to the complete system does not prejudice the usefulness of this set-up: the network congestion phenomena individually affect the output buffers of the network ports that are actually in use. The obtained results should scale reliably with the higher number of HLT racks in the complete system.

The core routers use input-buffering with a peculiar implementation of switch-level virtual output queuing: input ports are grouped in modules of 8 ports at most and each module maintains multiple, distinct queues to every output port on the router. The routers are equipped with deep buffers: each module has a packet memory of 1.5 GB.

The top-of-rack switches use output-buffering. Different models were tested. This paper focuses on two representative models:

- Switch 1: a switch with per-port dedicated buffers of 750 kB, of which ~600 kB are available for standard-priority packets
- Switch 2: a switch with two buffers of 12 MB each, shared by 32 ports each, with a per-port limit of 8 MB

Among the synthetic traffic pattern tested, this paper reports on the so-called full event building pattern:


- Events are assigned by the HLT supervisor to HLT processing units at a constant rate.
- The processing units immediately collect all fragments of assigned events, process them for a fixed amount of time, and ask for a new assignment.

It should be noted that this pattern is the harshest in terms of generated traffic bursts: since the data corresponding to an event are collected all at once, the maximum burst size corresponds to the event size. On the other hand, the fixed-size fragments enable the simple traffic-shaping algorithm to operate most efficiently.


The assignment rate parameter for the presented tests was selected with the goal of utilizing a sizable portion of the available bandwidth of the HLT rack uplinks, while maintaining a comfortable margin to avoid the effects of link saturation. The chosen 750 Hz assignment rate corresponds to a total throughput of  $\sim 13$  Gb/s. Due to performance constraints, the only policy that the supervisor can use when choosing the processing unit to which an event is assigned is first-come first-served (FCFS), i.e. events are assigned to processing units in order of arrival of their assignment requests.

## B. Results

The results of the measurements are shown in Figure 5. The total data-collection time per event is influenced both by the network conditions and by the traffic-shaping mechanism.



(a)



(b)

Figure 5. Data-collection latency as a function of the number of traffic-shaping credits, for (a) switch 1 and (b) switch 2. The test conditions are detailed in Section IV-A. The bullets represent the average values. The horizontal box lines represent the first quartile, the median, and the third quartile. The box whiskers represent the first and the 99th percentile.

With too few traffic-shaping credits available, the large data-collection time is due to collection inefficiency because the HLT nodes cannot fully utilize the network bandwidth. This is analogous to a TCP connection with a congestion window smaller than the bandwidth-delay product. With too many traffic-shaping credits the top-of-rack switch buffers are overflowed, triggering the TCP incast pathology. Packet drops are not an issue in the core router due to the deep buffers it offers.

In the following sections, these measurements are used as the baseline for verifying that the simulation models developed correctly reproduce the system behavior. For a more in-depth look at similar tests on the ATLAS data-acquisition system, please refer to [7]. One key consideration is worth reporting: it is possible to estimate the minimum possible value for the data-collection latency to use as a reference point. For this purpose, the time it takes for data requests to reach the readout systems and the time it takes for the readout systems to prepare the data can be neglected<sup>1</sup>: the dominant component is the time it takes to transmit the data fragments on the slowest link in the path. In this set-up, this corresponds to the amount of time it takes to transmit a full event on the GbE link from the rack-level switch to the HLT node, which is  $\sim 18$  ms. It is worth noting that while the credit-based traffic-shaping algorithm is successful in keeping the data-collection latencies under control when configured correctly, it also prevents the average data-collection latency

<sup>1</sup>Measurements show that both time intervals are smaller than 0.5 ms.

from reaching that minimum, as evidenced by the latency distributions shown in the figures.

## V. SIMULATION MODEL DEVELOPMENT

### A. Simulation framework

The simulation presented in this paper uses the OMNeT++ discrete event simulation<sup>2</sup> framework [8]. OMNeT++ is free for non-commercial use and its source code is available. It was chosen for two main reasons: its wide acceptance in the academic community and its ease of use for the purposes of modeling computer networks. In OMNeT++ simulations are composed of modules, defined in the declarative NED language, which communicate exchanging messages via module-to-module channels. So-called simple modules are the active components of the simulation and are implemented in C++, leveraging the class hierarchy provided by the simulation framework. Modules can be grouped together to form compound modules and networks. The development of network simulations is aided by built-in support for physical channels, with latency, transmission delay, and message loss properties.

The simulation model described here is based on the INET Framework of OMNeT++. INET is a protocol model library which includes detailed implementations of all network layers, from the MAC layer onwards. In particular, the Ethernet, IP, and TCP implementations are used.

### B. Hosts and applications

Data-acquisition applications are implemented on top of the standard INET host model, as shown in Figure 6. Applications interact only with the TCP module. A useful feature of INET’s TCP module is that it supports a data transfer mode that preserves application-level message boundaries: e.g. if an application sends a 1 MB message the receiver application will receive the same message, after TCP has completely finished simulating the transmission of 1 MB over the connection. This greatly simplifies the modeling of message-based applications, like those in the ATLAS data-acquisition software.

In order to simulate the test system presented in Section IV-A, four applications need to be modeled: the readout systems, the HLT supervisor, the data-collection manager, and the HLT processing units.

In the real system, the HLT supervisor, which assigns events to processing units, and the readout systems, which serve event data fragments, are data-driven applications: their behavior is dependent on that of the experiment and the Level-1 trigger. In principle, this would require models of those applications to follow traces recorded on the real system. However, this is not necessary when trying to reproduce the synthetic traffic patterns described in Section IV-A. The readout systems become trivial server applications, responding to fragment requests with a configurable delay and response size. The supervisor instead is reduced to a periodic scheduler. The processing units send a message to the supervisor when they are available, i.e. when they are ready to start processing another event. The supervisor stores this information and uses

<sup>2</sup>In this particular sentence, the word “event” refers to simulation events. To avoid confusion, throughout the rest of the paper, “event” will only be used in its high-energy physics meaning, i.e. “collision event”.




Figure 6. Example of a model a host: HLT working node with 2 processing units.

it to assign events to processing units at a configurable global rate, by sending assignment messages.

As already mentioned, OMNeT++ simple models are implemented using the C++ programming language. Since the same language is used throughout the ATLAS data-acquisition software, the application-level code that is relevant to the simulation model can be ported to the simulation environment with minimal changes. This ensures a bug-for-bug compatible reproduction of the applications’ behavior within the model. This approach is used for modeling the processing units, which generate the data requests for an event, and the per-node data-collection managers, which act as proxies between the node’s processing units and the readout systems. In particular, a processing unit can simulate per-event iterative collection and processing of data: after it receives an event assignment, it can request data from the readout systems with a configurable pattern, emulate processing by waiting for a configurable amount of time, and repeat this process several times before considering the event fully processed and asking for another one from the supervisor. Just like in the real system, the processing units of a worker node do not interface directly with the TCP module: their communications are mediated by the per-node data-collection manager. Its most relevant functions in the context of this model are the mapping of each data request from the processing units to messages to multiple readout systems the enforcement of the traffic-shaping algorithm described in Section III-C.

### C. Network switches

For the purposes of this simulation, the most relevant aspect of the network hardware is packet buffering. The internal architecture of the switches is not modeled in detail. It is assumed that the switch speedup is high enough to prevent input head-of-line blocking, and to make the packetization delay of the switch cells negligible. With these assumptions, switches are modeled as follows. For each switch port, an INET Ethernet MAC module acts as the interface between the physical transmission channel and the switch. It sends incoming frames to an ideal frame relay unit, which maintains the switch’s MAC address table and instantaneously forwards frames towards their destination port (or broadcasts them if the destination is not yet present in the address table). One or more “packet droppers” intercept frames between the relay unit and the switch output queues. These modules selectively drop frames, depending on the total buffer space available

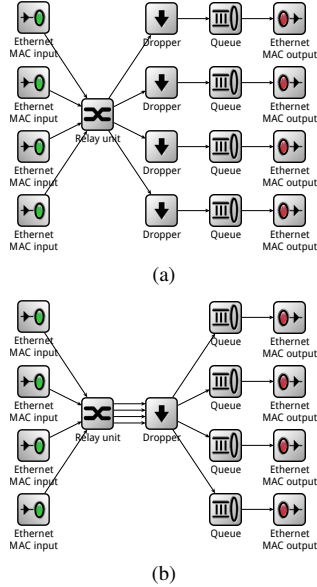


Figure 7. Models of output buffered switches with (a) dedicated per-port buffers and (b) shared buffers.

in the queues that are connected to their outputs, effectively defining the switch buffering scheme. Frames that were not dropped are stored in the switch’s output queues, waiting for the Ethernet MAC to pull them from the queue when the transmission channel is ready.

Two basic buffering schemes are considered: dedicated and shared. In the dedicated buffers model, shown in Figure 7a, there is a dropper for every output port, effectively modeling tail-drop queues. In the shared buffer model, shown in Figure 7(b), a single dropper guards all the output ports. The two basic schemes can naturally be integrated to model more complex architectures, e.g. with buffer space limits both on a per-switch basis and on a per-port basis.

#### D. Complete model

The components described in the previous sections are assembled to create a model of the test system described in Section IV-A. The observed simulation run-time depends on the rate of events scheduled by the supervisor. The relation between simulated time  $t$  and simulation runtime  $T$  is approximately given by:  $T = t \cdot r / R$ , where  $r$  is the supervisor assignment rate, and  $R$  is a constant which is roughly equal to 5 Hz when running the simulation on a modern CPU. As an example, this gives a runtime of ca. 100 minutes for simulating 30 s of a system running at 1 kHz.

## VI. ANALYSIS AND COMPARISON OF MEASURED AND SIMULATED RESULTS

### A. Validation

The model outlined in Section V has some obvious approximations, the most important one being the simplified switch architecture. This is not only a design choice: details on the architecture of commercial network equipment are extremely scarce and, when available, are geared towards marketing

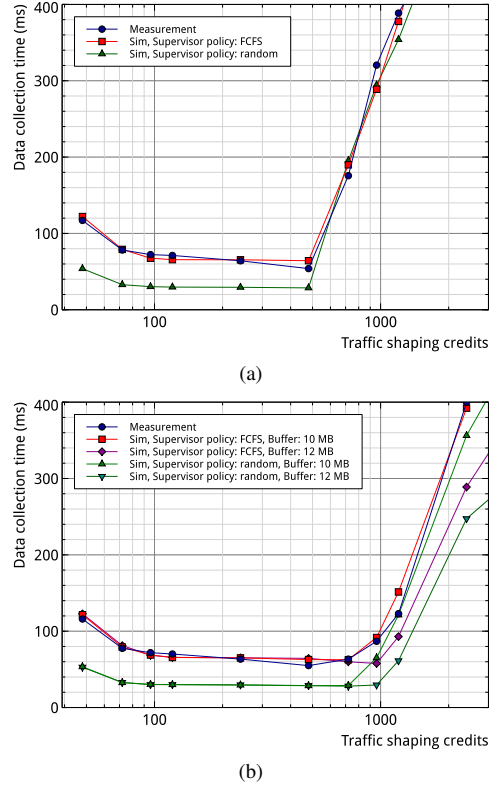


Figure 8. Comparison of measured and simulated data-collection latencies for different settings of the traffic-shaping algorithm, using (a) switch 1 and (b) switch 2.

rather than engineering. Another potential source of inaccuracies in the simulation is the TCP congestion avoidance algorithm. The most commonly used protocol model libraries only provide the “traditional” TCP variants: Tahoe, Reno, Vegas and New Reno. Support for more recent variants such as TCP CUBIC (the algorithm currently in use in Linux) relies on models built by adapting the Linux kernel TCP stack for use in the simulation. To avoid this extra complexity, for the results presented here New Reno was selected. While none of the mentioned algorithms can effectively prevent the incast problem, their impact on other facets of the simulation might render it unreliable. For these reasons, it is crucial to validate the model against the measured behavior of the system, before it is used to draw conclusions on scenarios that cannot easily be tested in practice.

As explained in Section III-C, the key application performance metric is the average data-collection latency per event. Therefore, the focus here is on comparing the measured data presented in Section IV-B with the simulated latencies, as shown in Figure 8. As a significant example of the additional investigations enabled by the simulation, two settings for the HLT supervisor event assignment policy are considered. The first one, FCFS, more closely corresponds to the actual implementation. The second policy consists in choosing a random processing unit out of all of those that requested an assignment, without regard for the order in which they did so.




Figure 9. Comparison of measured and simulated data-collection latency distributions, using Switch 1, at 480 traffic-shaping credits.

The most important feature to reproduce is the onset of the TCP incast pathology, evidenced by the abrupt increase in data-collection latency as the traffic-shaping algorithm allows an excessive number of data requests to be in-flight. The model is successful in this. When simulating Switch 1, the onset point is identical at 480 traffic-shaping credits, corresponding to a maximum per-node burst size of  $\sim 550$  kB. This is expected, as the switch has 600 kB per-port output buffers. When simulating Switch 2, the incast onset points of measurement and simulations are close, but not identical, with the simulation being more forgiving. This can be explained by an overly optimistic choice of buffering parameters in the simulation. Taking the manufacturer’s parameters at face value, the switch’s shared buffer memory is 12 MB. However, some of that memory is reserved for quality-of-service purposes and cannot be used by standard-priority packets. Indeed, with a simulated buffer size of 10 MB, the incast onset point coincides in simulation and measurement.

The model of the traffic-shaping algorithm benefits from the code sharing with its actual implementation. As a consequence, the simulation is particularly accurate in the region of the parameter space where the latency is heavily influenced by the shaping, i.e. where the algorithm is effective in preventing packet drops. Further confirmation of the accuracy of the simulation comes from comparing the distributions of the data-collection times, rather than just their mean values. As an example, the two superimposed histograms in Figure 9 show the measured and simulated distributions for Switch 1 at 480 traffic-shaping credits, i.e. the highest setting not suffering from incast. The histograms are in very good agreement and demonstrate the effect of the traffic-shaping algorithm. A sizable portion of the events are fully collected within 20 ms, which is compatible with the minimum latency estimated in Section IV-B (18 ms), presumably because all their fragments were collected when no other events were competing for the same credits. The other events need to wait for enough credits to become available, hence the long tail of the distribution.

### B. Preliminary findings

With the consistency of the simulation model with the real system measurements reasonably established, the focus can shift to more simulated scenarios that could not easily be enacted in practice.

One such scenario is applying a different event-scheduling policy in the HLT supervisor. As mentioned in Section IV-A performance limitations currently prevent event-scheduling policies more complex than FCFS from being implemented

in the supervisor. The simulation is not affected by such issues, so an alternative policy, such as randomly choosing a processing unit out of the available ones, can be modeled.

The results are shown in Figure 8. The change in policy from FCFS to random leads to a very significant reduction of the data-collection latency, especially in the region of the parameter space where the traffic-shaping algorithm prevails. The explanation for this difference lies in the mapping of processing units to nodes. While the traffic-shaping credits limit is applied on a per-node basis, the event scheduling (and associated data collection) happens on a per-processing-unit basis. In the particular set-up used, there are 24 units per node (see Section IV-A). The random policy reduces the probability that two or more events will be assigned in a short interval to units hosted by the same node. Units on different nodes do not compete for the same credit pool, ultimately resulting in lower average data-collection latency.

The simulation also enables studying the effects of hardware parameters that cannot so easily be modified in practice. One of these is the size of the packet buffers in the top-of-rack switches. The simulation can be used to determine the relation between the amount of memory and the data-collection latency, with the workload described in Section IV-A and with the traffic-shaping algorithm disabled.

The results for a switch with dedicated per-port memory are shown in Figure 10a. When employing the FCFS assignment policy, buffers of at least 4.8 MB effectively prevent packet drops and the latency can reach its lowest value (slightly lower than 20 ms, compatible with the minimum latency estimated in Section IV-B). However, with the random assignment policy, the lowest latency is only reached with 8 MB buffers. This discrepancy, and in general the better performance of FCFS in this scenario, can be explained. If two events are consecutively assigned to two processing units on the same worker node, the data for the second event incurs a larger delay, since the buffer of the worker node’s switch port still contains data from the first event, leading to queuing delays or overflow. This effect changes the order of the supervisor’s FCFS queue: over time, entries in the queue referring to different units on the same worker node distance themselves. The available switch buffer memory is therefore used more efficiently, and less of it is necessary to prevent packet drops.

The results for a switch with shared memory are shown in 10b. In this scenario, both with the FCFS and random assignment policies, packet drops are prevented with a buffer size of at least 26 MB. However, for smaller buffer sizes, the random assignment policy performs better than FCFS. This is explained by the fact that, in a switch with fully shared memory, packet drops are not directly related to the occupancy of their output queue. Therefore, the beneficial effect described in the above paragraph does not apply, and the random assignment policy ensures a more uniform usage of the switch outputs, thus reducing packet drops.

## VII. RELATED WORK

Characteristics of other large-scale data-acquisition systems can be found for example in [9] for the LHCb experiment and [10] for the ALICE experiment. Details of the CMS experiment’s InfiniBand-based system are in [11].

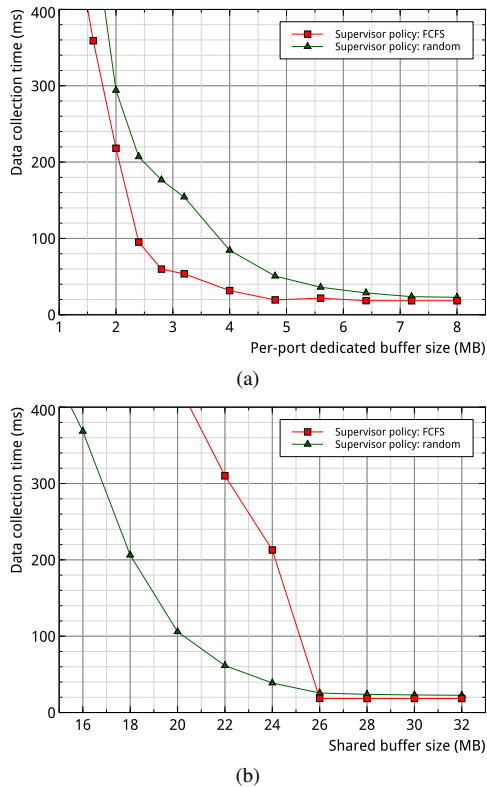


Figure 10. Average data-collection latency (with no traffic-shaping) as a function of the top-of-rack switch buffer size, using (a) dedicated per-port buffers and (b) shared buffers.

Alternative general-purpose network simulation frameworks include the open-source ns-3 [12] and the commercial SteelCentral NetModeler (formerly OPNET Modeler). Ns-3 includes a good-quality port of the Linux TCP implementation, which could improve the accuracy of the simulation in reproducing TCP behavior.

The TCP incast pathology received a considerable amount of attention in the academic community. See [13] for a review. The application of some of the proposed solutions to a small prototype of the ATLAS data-acquisition network is reported on in [14].

## VIII. CONCLUSION AND FUTURE DIRECTIONS

The simulation model presented in this paper was proven capable of reproducing the key performance traits of a complex data-acquisition system such as the ATLAS TDAQ system. Already at this verification stage, comparing the simulated and measured results yields useful indications that can drive optimizations and further development of the system. The approach employed in the development of the model, aiming at keeping complexity at a minimum without sacrificing accuracy, can be extended to systems with similar traffic patterns.

Using this work as a starting point, more speculative trials can be undertaken. In particular, three categories of solutions to the incast pathology can be modeled: more sophisticated application-level traffic shaping, alterations of the transport protocol itself (requiring kernel modifications in the real

system) and of the link layer (requiring hardware changes). The second category includes solutions such as reducing TCP's minimum RTO limit or experimental incast-aware TCP variants. The third category mainly refers to alternative link-layer technologies such as Infiniband or recent additions to the Ethernet standards such as IEEE 802.1Q Congestion Notification.

## ACKNOWLEDGMENTS

This work has been partially supported by the Ministry of Economy and Competitiveness of Spain (project TIN2012-38341-C04) and the European Commission (ERDF).

## REFERENCES

- [1] ATLAS Collaboration, "The ATLAS experiment at the CERN large hadron collider," *J. Instrumentation*, vol. 3, no. 08, p. S08003, Aug. 2008.
- [2] "ATLAS high-level trigger, data-acquisition and controls," CERN, Geneva, Technical Design Report ATLAS-TDR-016 CERN-LHCC-2003-022, 2003.
- [3] M. E. Pozo Astigarraga, "Evolution of the ATLAS Trigger and Data Acquisition System," *J. Phys.: Conf. Ser.*, vol. 608, no. 1, p. 012006, 2015.
- [4] A. Phanishayee *et al.*, "Measurement and analysis of TCP throughput collapse in cluster-based storage systems," in *Proc. of the 6th USENIX Conference on File and Storage Technologies*, ser. FAST'08. Berkeley: USENIX Association, 2008, pp. 12:1–12:14.
- [5] M. Allman, V. Paxson, and E. Blanton, "TCP congestion control," RFC 5681, IETF, Sep. 2009.
- [6] V. Paxson, M. Allman, J. Chu, and M. Sargent, "Computing TCP's retransmission timer," RFC 6298, IETF, Jun. 2011.
- [7] T. Colombo, "Data-flow Performance Optimisation on Unreliable Networks: the ATLAS Data-Acquisition Case," *J. Phys.: Conf. Ser.*, vol. 608, no. 1, p. 012005, 2015.
- [8] A. Vargas, "OMNeT++," in *Modeling and Tools for Network Simulation*, K. Wehrle, J. Gross, and M. Günes, Eds. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 35–59.
- [9] F. Alessio *et al.*, "The LHCb Data Acquisition during LHC Run 1," *J. Phys.: Conf. Ser.*, vol. 513, no. 1, p. 012033, 2014.
- [10] F. Carena *et al.*, "The ALICE data acquisition system," *Nucl. Instruments and Methods in Physics Research A*, vol. 741, pp. 130–162, 2014.
- [11] T. Bawej *et al.*, "Boosting Event Building Performance using Infiniband FDR for the CMS Upgrade," *Proc. Sci.*, vol. TIPP2014, p. 190, 2014.
- [12] "The ns-3 network simulator." [Online]. Available: <http://www.nsnam.org>
- [13] Y. Zhang and N. Ansari, "On Architecture Design, Congestion Notification, TCP Incast and Power Consumption in Data Centers," *IEEE Commun. Survey Tutorials*, vol. 15, no. 1, pp. 39–64, 2013.
- [14] G. Jereczek, G. Lehmann-Miotto, and D. Malone, "Analogues between tuning TCP for data acquisition and datacenter networks," presented at IEEE Int. Conf. Comm., 2015.