

Modeling a Large Data-Acquisition Network in a Simulation Framework

Tommaso Colombo^{*†}, Holger Fröning[†], Pedro Javier García[‡] and Wainer Vandelli^{*}

^{*}Physics Department, CERN, CH-1211 Geneva 23, Switzerland

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

[†]Institut für Technische Informatik (ZITI), Universität Heidelberg, B6 26, 68131 Mannheim, Germany

holger.froening@ziti.uni-heidelberg.de

[‡]Departamento de Sistemas Informáticos, Universidad de Castilla-La Mancha, Avda. de España, 02071 Albacete, Spain

PedroJavier.Garcia@uclm.es

Abstract—The ATLAS detector at CERN records particle collision “events” delivered by the Large Hadron Collider. Its data-acquisition system is a distributed software system that 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 the 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 the Large Hadron Collider (LHC) at a design frequency of 40 MHz. Each event corresponds to 1 to 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 com-

mercial 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 loss of potentially unique 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 μ 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


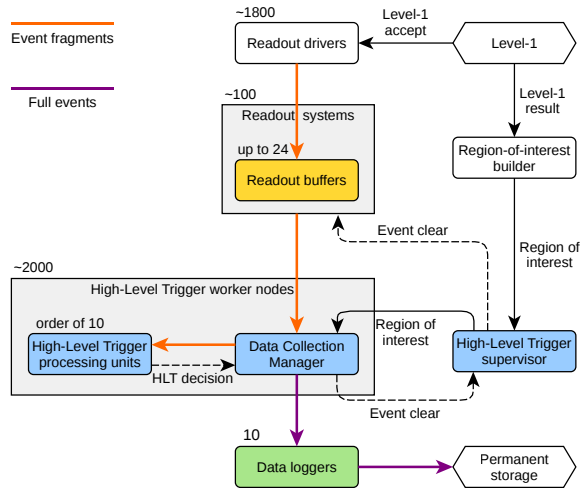


Figure 2. Network architecture of the ATLAS Data-Acquisition and High-Level Trigger system.

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).

85 detector that recorded interesting signals that are used as a
 86 seed for the second stage.
 87 When an event is accepted by the Level-1, its data fragments
 88 (1860 fragments of variable sizes around 1 kB, corresponding
 89 to different regions of the experiment) are readout by the
 90 detector electronics. Each fragment is pushed via a custom
 91 point-to-point optical link into a specific hardware buffer in
 92 one of the ~ 100 “readout system” nodes, which act as the
 93 coupling between the first stage’s custom electronics and the
 94 second stage’s COTS hardware.
 95 The second stage, called High-Level Trigger (HLT), is a
 96 distributed software system running on around 2000 Linux
 97 PCs interconnected by an Ethernet network. Each HLT worker
 98 node hosts one HLT “processing unit” per CPU core and a
 99 single “data-collection manager”, which handles communi-
 100 cations with the rest of the system on behalf of the node’s
 101 processing units. A central scheduler, the HLT “supervisor”,
 102 receives regions-of-interest information from the Level-1, and
 103 assigns it to one of the available HLT processing units. Using
 104 the regions-of-interest as starting points, the processing unit
 105 incrementally retrieves and analyses event fragments, until a
 106 decision can be taken. The event can be rejected even without
 107 analyzing all its fragments, thus limiting the fraction of data to
 108 be retrieved. Fragments of rejected events are deleted from the
 109 readout systems buffers, while accepted events are transferred
 110 to one of the Data Logger nodes for storage. The rate of events
 111 that can be accepted is mostly determined by the availability of
 112 computing and storage resources for subsequent data analysis,
 113 and is foreseen to be around 1 kHz for the 2015–2018 data
 114 taking period.

121 Each node in a rack is connected to an aggregation switch with
 122 a GbE link. The rack switches have 10 GbE links to both core
 123 routers. For the rationale behind this design, please refer to
 124 [3].

125
 126
 127
 128
 129
 130
 131
 132

III. PERFORMANCE ISSUES IN DATA-ACQUISITION NETWORKS

A. Traffic pattern

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

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

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

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

B. The TCP incast pathology

148 As mentioned in Section I, the ATLAS data-acquisition sys-
 149 tem is based on Ethernet. The messaging among applications
 150 relies on the TCP protocol. The combination of TCP, a lossy
 151 link layer such as Ethernet, and the traffic pattern described
 152 above is subject to a well-known TCP pathology called *incast*
 153 [4]: first observed in data-center storage networks, it occurs
 154 “when a client simultaneously receives a short burst of data
 155 from multiple sources, overloading the switch buffers associ-
 156 ated with its network link such that all original packets from
 157 some sources are dropped”.

115 The physical layout of the Data-Acquisition and High-Level
 116 Trigger system is represented in Figure 2. The core of the
 117 system consists of two large network routers with a maximum
 118 capacity of several hundred 10GbE ports. Readout systems are
 119 directly connected to both core routers with 4 10GbE links.
 120 HLT worker nodes are organized in racks of at most 40 nodes.




Figure 3. Visualization of the potential network congestion issues.

159 TCP has both an active packet-drop detection mechanism, in
 160 which the receiver detects the lost packet and causes the sender
 161 to react very quickly, and a passive mechanism, which is based
 162 on retransmission timeouts (*RTO*) on the sender side [5]. The
 163 *RTO* value is calculated starting from the estimated round-trip
 164 time [6]. However, to protect against spurious retransmissions
 165 and interference with other TCP features such as delayed
 166 acknowledgments, the *RTO* has a fixed minimum, specified as
 167 1 s in [6], with lower values used in actual implementations
 168 (200 ms in Linux, 30 ms in FreeBSD). These values are orders
 169 of magnitude larger than the round-trip time of a data-center
 170 network, which is usually in the sub-millisecond range. In
 171 addition, the active mechanism requires at least some of the
 172 original packets to reach the destination for the packet loss to
 173 be detected. When this is not the case, the delivery of packets
 174 incurs the large delay caused by the *RTO*.

175 C. Impact and mitigation

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

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

197 A less invasive but less general solution is application-level
 198 traffic shaping. In ATLAS, smoothing the rate of data requests
 199 generated by a HLT node can alleviate the network congestion
 200 by controlling the maximum size of the traffic burst from the

201 readout systems [7]. Obviously such a smoothing mechanism
 202 imposes a trade-off: excessive smoothing can increase the
 203 data-collection time by unnecessarily delaying the requests for
 204 data, whereas insufficient smoothing will not eliminate packet
 205 drops.

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

- 210 • Each HLT worker node has a fixed number of credits
- 211 available. All the HLT processing units on a node share
- 212 these credits.
- 213 • Each data request from a processing unit to a readout
- 214 system uses as many credits as the number of fragments
- 215 it asks for.
- 216 • Each response returns the credits used by the correspond-
- 217 ing request.
- 218 • If all available credits are used, further requests are
- 219 blocked until the necessary credits become available.

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

228 IV. MEASUREMENTS

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

244 A. Test set-up

245 The measurements were performed using some of the
 246 hardware available in the ATLAS data-acquisition system.
 247 At the time of these tests, the infrastructure was still under
 248 consolidation, so the test set-up is slightly different from the
 249 one described in Section I. In particular, readout systems
 250 were not directly connected with 10GbE links to the core
 251 routers. They were instead connected with GbE links to an
 252 intermediate aggregation switch with a 10GbE uplink to each
 253 core. The intermediate switch was kept under-subscribed so
 254 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 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 ~ 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. 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




Figure 4. Test set-up.

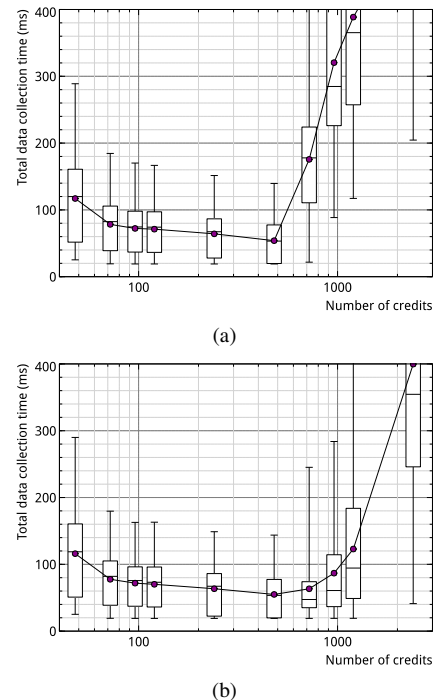


Figure 5. Data-collection latency as a function of the number of traffic-shaping credits, for switch 1 (a) and switch 2 (b). The test conditions are detailed in Section IV-A. The bullets represent the average values. The horizontal box lines represent the first quartile, the median, and the third quartile.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

326 correctly reproduce the system behavior. For a more in-depth
327 look at similar tests on the ATLAS data-acquisition system,
328 please refer to [7]. One key consideration is worth reporting:
329 it is possible to estimate the minimum possible value for the
330 data-collection latency to use as a reference point. For this
331 purpose, the time it takes for data requests to reach the readout
332 systems and the time it takes for the readout systems to prepare
333 the data can be neglected¹: the dominant component is the time
334 it takes to transmit the data fragments on the slowest link in
335 the path. In this set-up, this corresponds to the amount of
336 time it takes to transmit a full event on the GbE link from
337 the rack-level switch to the HLT node, which is ~18 ms.
338 It is worth noting that while the credit-based traffic-shaping
339 algorithm is successful in keeping the data-collection latencies
340 under control when configured correctly, it also prevents the
341 average data-collection latency from reaching that minimum,
342 as evidenced by the latency distributions shown in the figures.

343 V. SIMULATION MODEL DEVELOPMENT

344 A. Simulation framework

345 The simulation presented in this paper uses the OMNeT++
346 discrete event simulation² framework [8]. OMNeT++ is free
347 for non-commercial use and its source code is available. It
348 was chosen for two main reasons: its wide acceptance in the
349 academic community and its ease of use for the purposes of
350 modeling computer networks. In OMNeT++ simulations are
351 composed of modules, defined in the declarative NED lan-
352 guage, which communicate exchanging messages via module-
353 to-module channels. So-called simple modules are the active
354 components of the simulation and are implemented in C++,
355 leveraging the class hierarchy provided by the simulation
356 framework. Modules can be grouped together to form com-
357 pound modules and networks. The development of network
358 simulations is aided by built-in support for physical channels,
359 with latency, transmission delay, and message loss properties.

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

365 B. Hosts and applications

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

¹Measurements show that both time intervals are smaller than 0.5 ms.

²In 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".

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

As already mentioned, OMNeT++ simple models are im-
plemented 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 units 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.

376 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




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

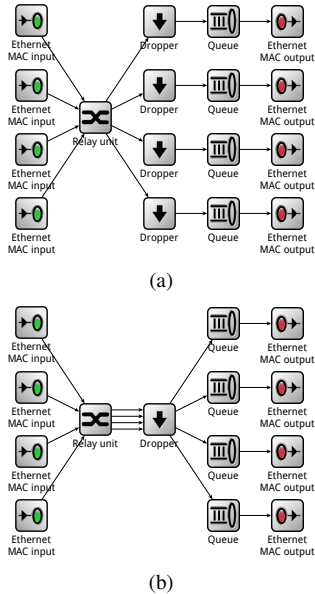


Figure 7. Models of output buffered switches with dedicated per-port buffers (a) and shared buffers (b).

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 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 be 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 7b, 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.

451 D. Complete model

452 The components described in the previous sections are
 453 assembled to create a model of the test system described in
 454 Section IV-A. The observed simulation run-time depends on
 455 the rate of events scheduled by the supervisor. The relation
 456 between simulated time t and simulation runtime T is ap-
 457 proximately given by: $T = t \cdot r/R$, where r is the supervisor
 458 assignment rate, and R is a constant which is roughly equal to
 459 5 Hz when running the simulation on a modern CPU. As an
 460 example, this gives a runtime of ca. 100 minutes for simulating
 461 30 s of a system running at 1 kHz.

462 VI. ANALYSIS AND COMPARISON OF MEASURED AND 463 SIMULATED RESULTS

464 A. Validation

465 The model outlined in Section V has some obvious ap-
 466 proximations, the most important one being the simplified
 467 switch architecture. This is not only a design choice: details
 468 on the architecture of commercial network equipment are
 469 extremely scarce and, when available, are geared towards
 470 marketing rather than engineering. Another potential source of
 471 inaccuracies in the simulation is the TCP congestion avoidance
 472 algorithm. The most commonly used protocol model libraries
 473 only provide the “traditional” TCP variants: Tahoe, Reno,
 474 Vegas and New Reno. Support for more recent variants such
 475 as TCP CUBIC (the algorithm currently in use in Linux) relies
 476 on models built by adapting the Linux kernel TCP stack for
 477 use in the simulation. To avoid this extra complexity, for the
 478 results presented here New Reno was selected. While none
 479 of the mentioned algorithms can effectively prevent the incast
 480 problem, their impact on other facets of the simulation might
 481 render it unreliable. For these reasons, it is crucial to validate
 482 the model against the measured behavior of the system, before
 483 it is used to draw conclusions on scenarios that cannot easily
 484 be tested in practice.

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

496 The most important feature to reproduce is the onset of
 497 the TCP incast pathology, evidenced by the abrupt increase in
 498 data-collection latency as the traffic-shaping algorithm allows
 499 an excessive number of data requests to be in-flight. The model
 500 is successful in this. When simulating Switch 1, the onset point
 501 is identical at 480 traffic-shaping credits, corresponding to a
 502 maximum per-node burst size of ~550 kB. This is expected, as
 503 the switch has 600 kB per-port output buffers. When simulat-
 504 ing Switch 2, the incast onset points of measurement and sim-
 505 ulations are close, but not identical, with the simulation being

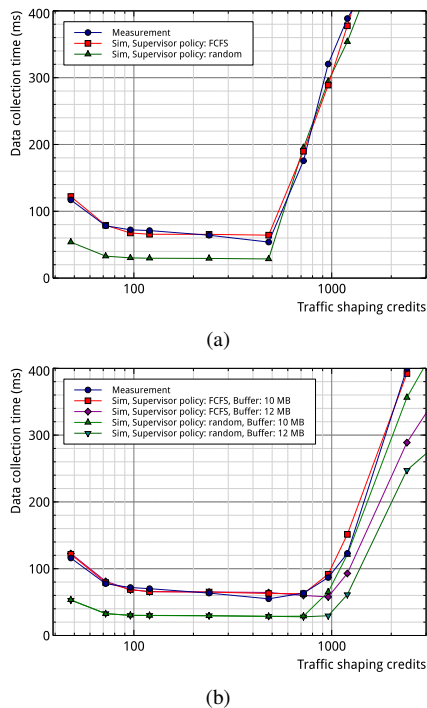


Figure 8. Comparison of measured and simulated data-collection latencies for different settings of the traffic-shaping algorithm, using switch 1 (a) and switch 2 (b).




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

506 more forgiving. This can be explained by an overly optimistic
 507 choice of buffering parameters in the simulation. Taking the
 508 manufacturer’s parameters at face value, the switch’s shared
 509 buffer memory is 12 MB. However, some of that memory is
 510 reserved for quality-of-service purposes and cannot be used
 511 by standard-priority packets. Indeed, with a simulated buffer
 512 size of 10 MB, the incast onset point coincides in simulation
 513 and measurement.

514 The model of the traffic-shaping algorithm benefits from
 515 the code sharing with its actual implementation. As a conse-
 516 quence, the simulation is particularly accurate in the region of
 517 the parameter space where the latency is heavily influenced by
 518 the shaping, i.e. where the algorithm is effective in preventing
 519 packet drops. Further confirmation of the accuracy of the
 520 simulation comes from comparing the distributions of the data-
 521 collection times, rather than just their mean values. As an
 522 example, the two superimposed histograms in Figure 9 show
 523 the measured and simulated distributions for Switch 1 at 480
 524 traffic-shaping credits, i.e. the highest setting not suffering

525 from incast. The histograms are in very good agreement and
 526 demonstrate the effect of the traffic-shaping algorithm. A
 527 sizable portion of the events are fully collected within 20 ms,
 528 which is compatible with the minimum latency estimated in
 529 Section IV-B (18 ms), presumably because all their fragments
 530 were collected when no other events were competing for the
 531 same credits. The other events need to wait for enough credits
 532 to become available, hence the long tail of the distribution.

533 B. Preliminary findings

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

538 One such scenario is applying a different event-scheduling
 539 policy in the HLT supervisor. As mentioned in Section IV-A
 540 performance limitations currently prevent event-scheduling
 541 policies more complex than FCFS from being implemented in
 542 the supervisor. The simulation is not affected by such issues, so
 543 an alternative policy, such as randomly choosing a processing
 544 unit out of the available ones, can be modeled.

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

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

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

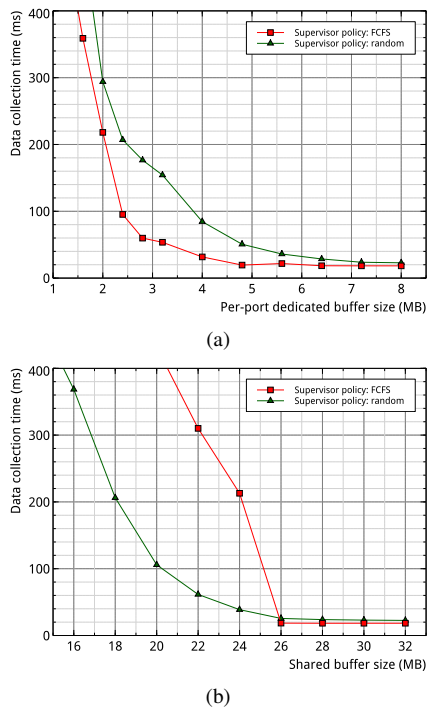


Figure 10. Average data-collection latency (with no traffic-shaping) as a function of the top-of-rack switch buffer size, using dedicated per-port buffers (a) and shared buffers (b).

580 entries in the queue referring to different units on the same
 581 worker node distance themselves. The available switch buffer
 582 memory is therefore used more efficiently, and less of it is
 583 necessary to prevent packet drops.

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

595 VII. RELATED WORK

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

600 Alternative general-purpose network simulation frameworks
 601 include the open-source ns-3 [12] and the commercial Steel-
 602 Central NetModeler (formerly OPNET Modeler). Ns-3 in-
 603 cludes a good-quality port of the Linux TCP implementation,
 604 which could improve the accuracy of the simulation in repro-
 605 ducing TCP behavior.

606 The TCP incast pathology received a considerable amount
 607 of attention in the academic community. See [13] for a review.

608 The application of some of the proposed solutions to a small
 609 prototype of the ATLAS data-acquisition network is reported
 610 on in [14].

611 VIII. CONCLUSION AND FUTURE DIRECTIONS

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

621 Using this work as a starting point, more speculative trials
 622 can be undertaken. In particular, three categories of solutions
 623 to the incast pathology can be modeled: more sophisticated
 624 application-level traffic shaping, alterations of the transport
 625 protocol itself (requiring kernel modifications in the real
 626 system) and of the link layer (requiring hardware changes).
 627 The second category includes solutions such as reducing
 628 TCP's minimum RTO limit or experimental incast-aware TCP
 629 variants. The third category mainly refers to alternative link-
 630 layer technologies such as Infiniband or recent additions
 631 to the Ethernet standards such as IEEE 802.1Q Congestion
 632 Notification.

633 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," 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.