Measuring Performance Under Failures in the LHCb Data Acquisition Network

Eloise Stein [1,2], Flavio Pisani [2], Tommaso Colombo [2] and Cristel Pelsser [1,3]

[1]University of Strasbourg [2]CERN [3]UCLouvain

Introduction

Data acquisition (DAQ) systems play a crucial role in the collection of scientific data [1, 2, 3]. Such systems, for instance, are deployed at experiments along the Large Hadron Collider (LHC) to collect fragmented data from various sensors and assemble them to reconstruct each particle collision event. This process, illustrated in Figure 1, is known as *Event Building*.

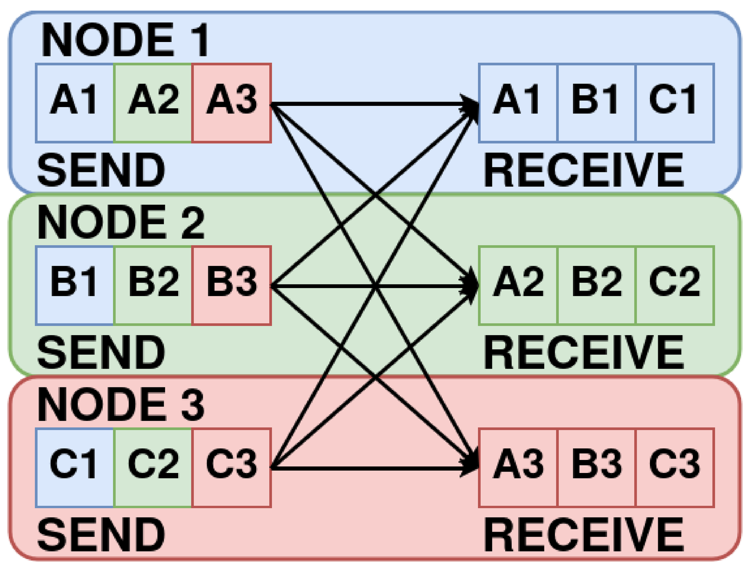

Figure 1. The Event Building process with 3 nodes.

Background

The resulting network traffic pattern of the Event Building process is a continuous succession of all-to-all exchanges. This exchange requires a bandwidth proportional to n^2, where n is the number of nodes. Given that no network topology provides such high bandwidth, the collective all-to-all exchange is typically divided into multiple phases to distribute the necessary bandwidth over time as illustrated in Figure 2.

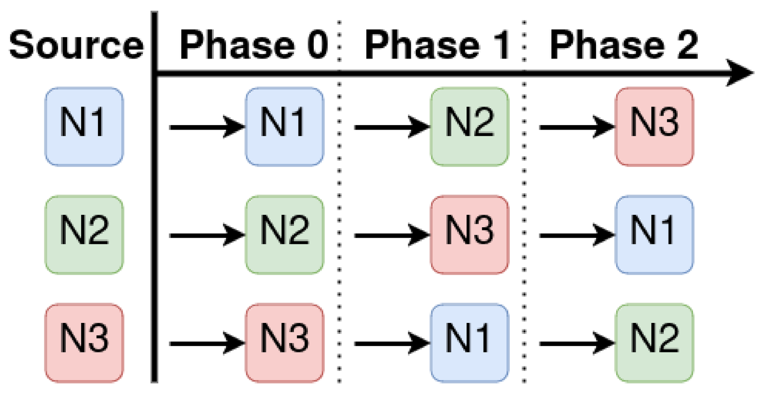

Figure 2. The all-to-all exchange between 3 nodes divided into 3 phases.

The all-to-all collective exchange can be *synchronized*, meaning that the DAQ application ensures that all computers complete their data exchange before moving to the next phase. Another approach involves a simpler *unsynchronized* all-to-all, where no nodes wait for the others to finish their exchange.

Main contributions

- Statistics on the duration, frequency, and underlying causes of link failures in the LHCb DAQ network.
- Throughput measurement of the synchronized and unsynchronized all-to-all in terms of scalability and fault-tolerance.
- Design recommendations.

Measurements of experienced failures in the LHCb DAQ network

Link failures in DAQ networks are common. During the month of March, the Event Builder network experienced 1933 failures, out of which 549 were attributed to performance tests, while 1384 were real network failures. The causes of the real failures were either hardware problems or dirty optical connectors.

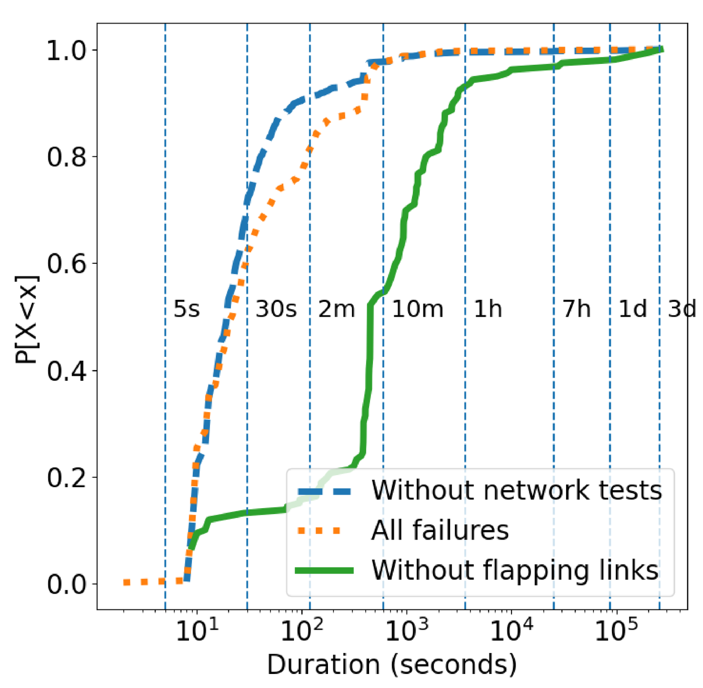

(a) Cumulative distribution of failure duration, which includes all failures, failures observed excluding test periods and treats flapping links as a single failure.

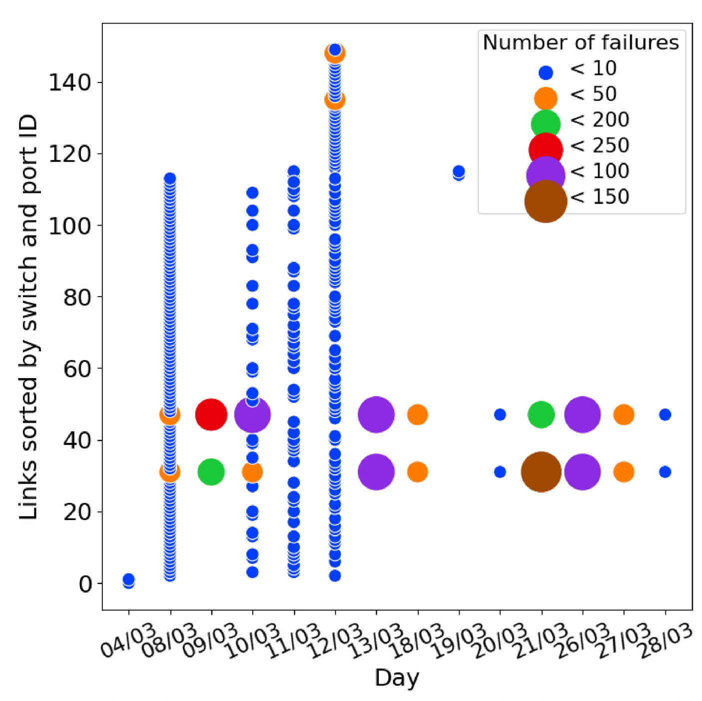

(b) Distribution of failures for each link and day of March 2024. The links on the y axis are sorted by switch and port ID.

Figure 3. Duration and frequency of failures observed during a one-month period.

Scalability of synchronized and unsynchronized all-to-all exchange

The all-to-all with synchronization exhibits a better throughput than the one without synchronization. As the system grows in size, the difference increases. When all servers are included, the throughput of the unsynchronized event builder is only 67% of its synchronized counterpart.

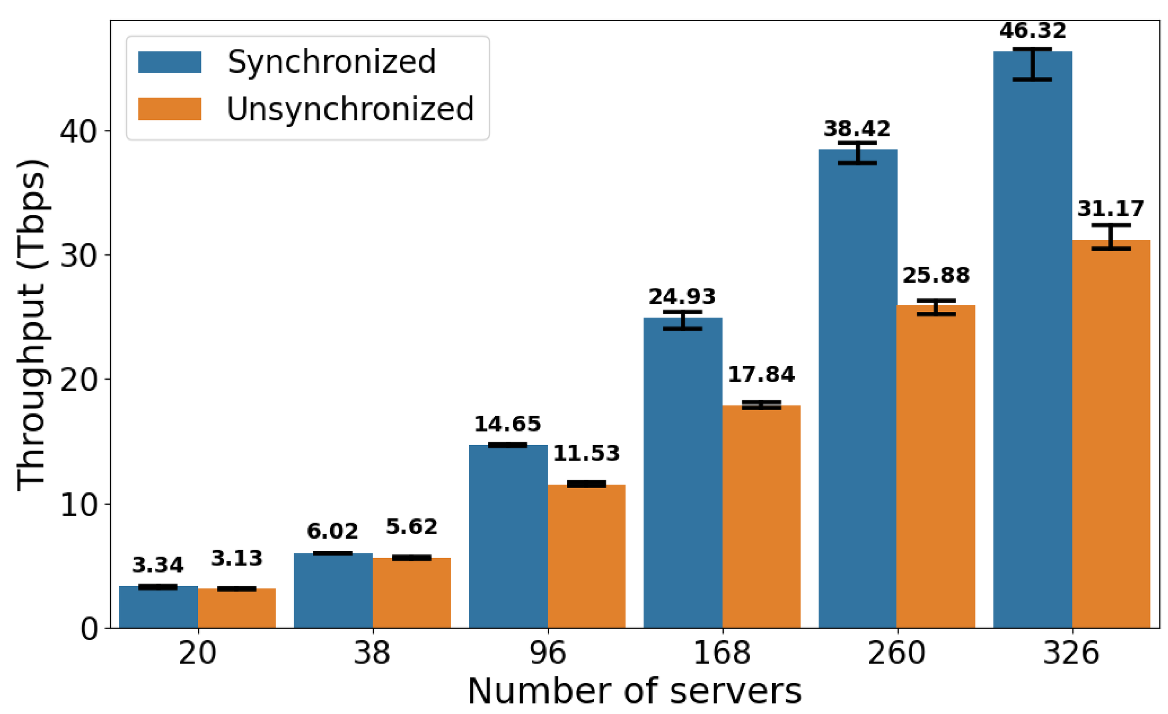

Figure 4. Scalability of the synchronized and unsynchronized all-to-all applications. The error bars represent the minimum, mean and maximum values.

Throughput achieved under link failures

Synchronized exchange shows lower performance than unsynchronized exchange in the event of failures.

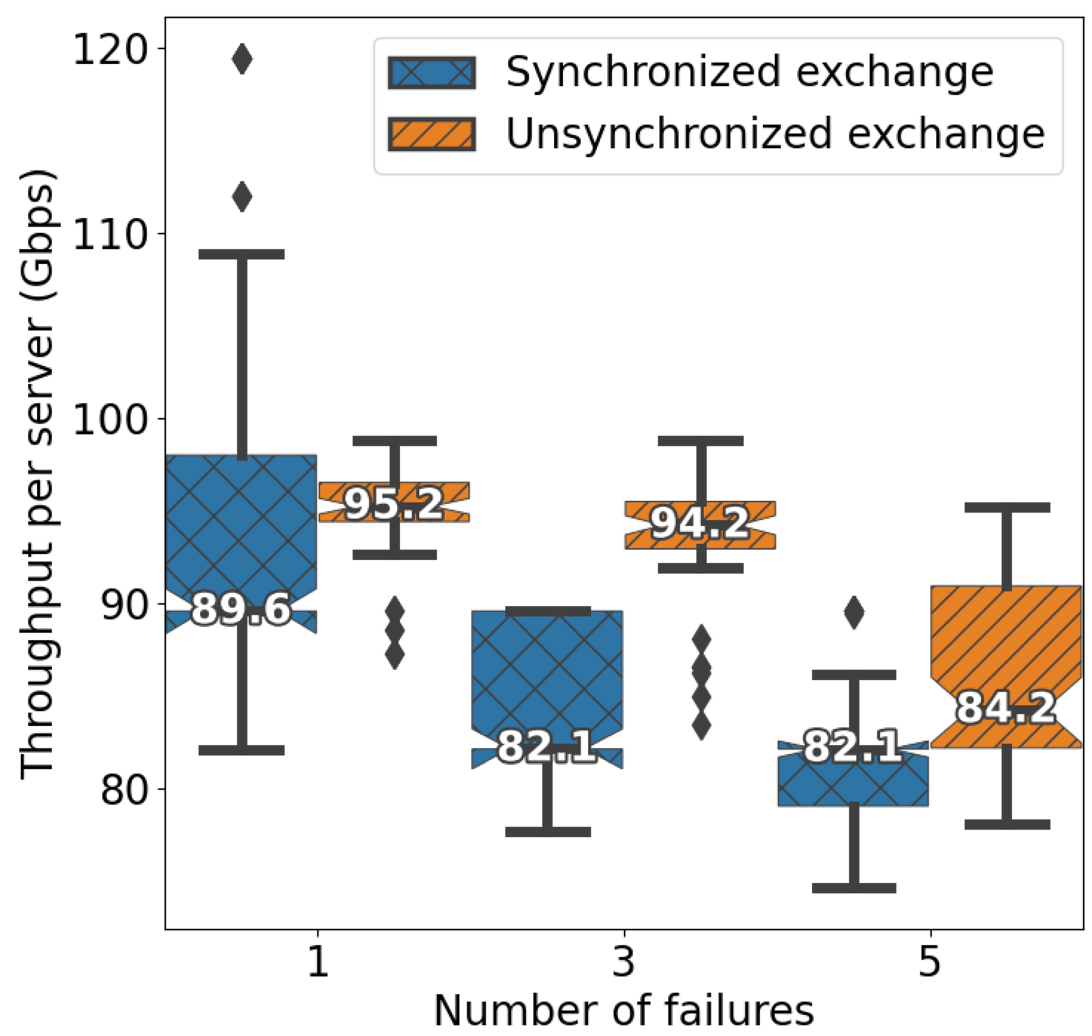

Figure 5. Synchronized and unsynchronized all-to-all exchange throughput as a function of the number of failures. Failure scenarios are randomly generated 10 times with 1, 3 or 5 simultaneous failures on the Event Builder network. The boxplots represent the minimum, 25th percentile, median, 75th percentile, and maximum values. Outliers are also depicted.

Design recommendations

In our measurements, we demonstrate that the unsynchronized all-to-all exchange performs better than the synchronized one in the event of failures. However, the synchronized all-to-all exchange demonstrates better performance during normal operation. Therefore, an optimized configuration of the all-to-all exchange would be to use the synchronized all-to-all and then move to the unsynchronized exchange if a failure occurs.

References

[1] N. e. a. Belyaev, "High performance computing system in the framework of the higgs boson studies," tech. rep., ATL-COM-SOFT-2017-089, 2017.

[2] G. Jereczek, G. Lehmann Miotto, and D. Malone, "Analogues between tuning tcp for data acquisition and datacenter networks," in *2015 IEEE International Conference on Communications (ICC)*, pp. 6062–6067, 2015.

[3] T. e. a. Bawej, "The new cms daq system for run-2 of the lhc," *IEEE Transactions on Nuclear Science*, vol. 62, no. 3, pp. 1099–1103, 2015.