

High speed adaptive rack-scale fabrics

Omer S. Sella
University of Cambridge
Omer.Sella@cl.cam.ac.uk

Andrew W. Moore
University of Cambridge
Andrew.Moore@cl.cam.ac.uk

Noa Zilberman
University of Cambridge
Noa.Zilberman@cl.cam.ac.uk

1 ABSTRACT

Rack-scale systems contain thousands of densely packed connected components. While a data center may accommodate a fully provisioned network, rack-scale systems demand a more compact and versatile network that would even up within a heavily populated system. Unless the critical path between communicating hosts is made faster, distributed rack-scale applications cannot scale. We present adaptive rack-scale fabrics, an architecture that uses *Physical Layer Primitives*, coupled with a *Closed Ring Control*. The resulting fabric uses pre-fetching techniques, but at the physical layer of the interconnect, to optimize performance within strict power-budget limitations.

2 MOTIVATION

Rack-scale systems do not necessarily follow the cpu-board-centric architecture that traditional racks use [1]. Instead of using regular server blades, we strip down the components and redesign according to the relevant metric - NVMe for fast storage, significant amount of DRAM for caching etc. This leads to a layout of hundreds and even thousands of interconnected nodes in a single rack. The meaning is that within a single rack we find a network as sophisticated and complex as in a data center, only much more constrained. In particular two problems arise: latency and power consumption.

Figure 1 shows the latency a packet experiences by traversing multiple hops through layer 2 cut-through switches. It also shows that the delay due to the media, (e.g., fiber) is negligible relative to the use of packet switching. The conclusion is that in the scale of a rack, it is packet switching that prevents distributed rack-scale applications from scaling. As an example, consider a MapReduce operation that requires transmission from all nodes. Since a reducer has to wait for data from all mappers, the slowest link pulls down the performance of an entire system.

Power budget is also a constraint, since rack-scale systems inherit the power budget of a traditional rack, and is factored into our proposed architecture as shown in figure 2. We give an overview of adaptive rack-scale fabrics, a network architecture that extends the physical layer and provides a control framework for it. Three key points of the architecture are:

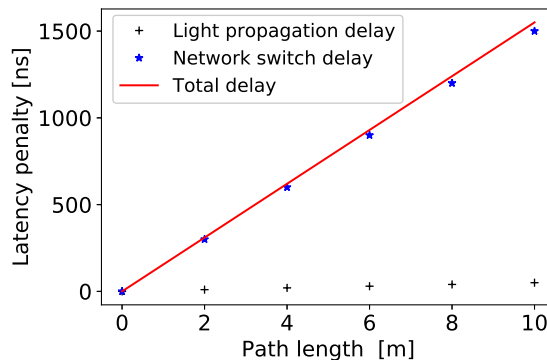


Figure 1: The latency due to propagation of packets in the media vs. the latency due to packet traversing a layer 2 state-of-the-art cut through switch. We assume a switch every 2 meters. In the scale of a rack, i.e.: a connection spanning few meters, the latency due to packet switching is dominant, and hence is bottlenecking scalability.

- Backwards compatibility - No restructuring of the network layer is needed. In particular, existing applications benefit from the architecture with no required change.
- Media agnostic - the specific underlying media is irrelevant. We only expect it to provide some subset of the *Physical Layer Primitives* that we define.
- Forward compatibility - future structuring of the network layer would be able to better utilize our architecture.

3 PROPOSED ARCHITECTURE

Configurable interconnect has seen many advances in recent years. Both on the optics side as in [2], as well as the electrical side as in [3]. While these solutions are different in the underlying media (optics vs. electrical) as well as in configuration times, they could be treated as functionalities that were added to the (already existing) physical layer. We place these extensions to the physical layer under a single framework, which we call *Physical Layer Primitives (PLP)*. In turn, these *PLP* are orchestrated by a control mechanism, that also schedules flows according to the availability of *PLP*'s. The control part of the architecture, called *Closed Ring*

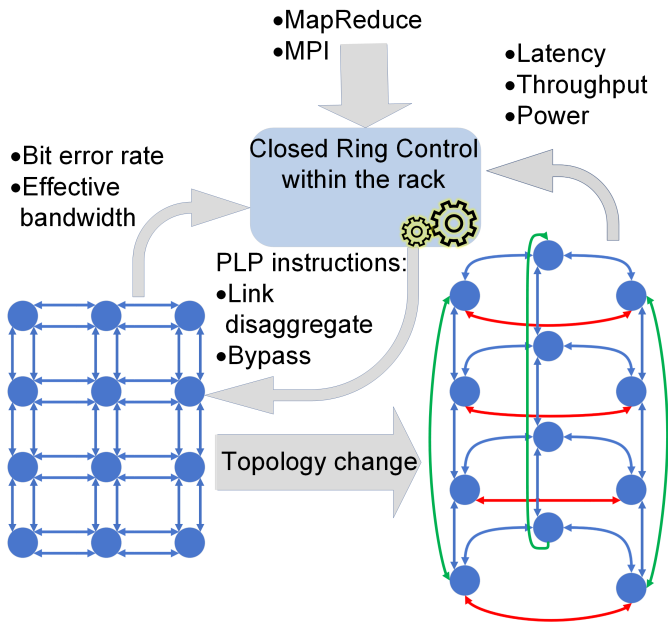


Figure 2: An example of the adaptive rack-scale network operation. Initially, the rack is configured using a grid topology of two lanes per link. Internal indications such as bit error rate and bandwidth are fed to the *Close Ring Control - CRC*. The *CRC* issues commands to the *Physical Layer Primitives - PLP*, which result in a torus topology running at one lane per link. New indications such as latency and power, are fed back to the in-rack *CRC* for further action.

Control (CRC), uses feedback from the interconnect such as latency, power consumption etc., to tag each link with a cost function. In this way, both routing as well as changes to the topology, are subject to the tools of control theory.

3.1 Physical Layer Primitives - PLP

We assume that a physical link is made up from physical lanes. The canonical example is a 100Gbps link that is made from four 25Gbps physical links, but different wavelengths under wavelength division multiplexing is an equivalent example. Looking at [3] and [2], we can identify several *Physical Layer Primitives*, and in addition draw new ones:

- (1) Link breaking / bundling - separating a link of N lanes into two links of k and N-k lanes and vice versa.
- (2) High speed bypass - connecting two links at the lowest possible physical level.
- (3) Turning a link on or off.
- (4) Adaptive forward error correction.
- (5) Per-lane statistics such as: bit error rate, latency, and effective bandwidth.

3.2 Closed Ring Control - CRC

Allocation of the *PLP* in the underlying fabric requires a control and scheduling entity. The *Closed Ring Control*, or *CRC* uses per-link price tags, with respect to metrics such as latency, congestion, link health etc. to allocate *PLP*'s and schedule flows. The problem that arises in all reconfigurable fabrics is finding the minimum flow size for which reconfiguration is worth the cost. This could be formulated as a convex optimization problem and solved distributively by the *CRC*. Further insights on rapid provisioning and reconfiguration, as well as traffic engineering for virtual switching can be found in Andromeda [4]. Figure 2 shows a *CRC* embedded in the rack. Upon receiving per-link statistics directly from the fabric, the *CRC* issues *PLP* instructions to improve the target metric, e.g: latency, by reducing the amount of switching logic that a packet has to go through.

4 EVALUATION

Since rack-scale systems contain hundreds to thousands of connected nodes, a simulation is used to evaluate the solution. We chose omnet++ as our simulation framework. To be certain that a large scale simulation is sound and credible, we begin with a small scale simulation verified by a hardware proof of concept (POC). We intend to use the NETFPGA SUME platform [5] for the hardware POC. Once the small scale simulation is validated, the POC will be integrated into the large scale simulation.

5 ACKNOWLEDGEMENTS

This work was partly funded by Microsoft Research through its PhD Scholarship Programme, the Leverhulme Trust (ECF-2016-289) and the Isaac Newton Trust.

REFERENCES

- [1] G. S. Zervas, F. Jiang, Q. Chen, V. Mishra, H. Yuan, K. Katrinis, D. Syrivelis, A. Reale, D. Pneumatikatos, M. Enrico *et al.*, "Disaggregated compute, memory and network systems: A new era for optical data centre architectures," in *Optical Fiber Communication Conference*. Optical Society of America, 2017, pp. W3D-4.
- [2] M. Ghobadi, R. Mahajan, A. Phanishayee, N. Devanur, J. Kulkarni, G. Ranade, P.-A. Blanche, H. Rastegarfar, M. Glick, and D. Kilper, "Projector: Agile reconfigurable data center interconnect," in *Proceedings of the 2016 ACM SIGCOMM Conference*. ACM, 2016, pp. 216-229.
- [3] V. Shrivastav, A. Valadarsky, H. Ballani, P. Costa, K. S. Lee, H. Wang, R. Agarwal, and H. Weatherspoon, "Shoal: A lossless network for high-density and disaggregated racks," Tech. Rep., 2017.
- [4] M. Dalton, D. Schultz, J. Adriaens, A. Arefin, A. Gupta, B. Fahs, D. Rubinstein, E. C. Zermeno, E. Rubow, J. A. Doucaer *et al.*, 2018.
- [5] N. Zilberman, Y. Audzevich, G. A. Covington, and A. W. Moore, "Netfpga sume: Toward 100 gbps as research commodity," *IEEE micro*, vol. 34, no. 5, pp. 32-41, 2014.