RETROSPECTIVE: Energy Proportional Datacenter Networking

Michael R. Marty[1], Philip Wells[1], Hong Liu[1], and Dennis Abts[2]

[1]Google, [2]Nvidia

I. CONTEXT

In 2007, Barroso and Hoelze made the case for energy proportional computing [3]. Their key observation was that servers spent most of their time at low to moderate utilizations of 10 to 50 percent and exhibited poor energy efficiency at these levels. Meanwhile, energy consumption from servers was rapidly growing and Google was expanding its construction of warehouse-scale datacenters.

Around the same time, Google formed an Advanced Development Lab, located in Madison, WI, seeded with industry veterans from Cray Inc. and Sun Microsystems, a University of Wisconsin Professor, and recent PhD graduates. With other colleagues from Google's Platforms division in Silicon Valley, we were chartered with finding and making "big bets" in hardware technologies that we could deploy to our warehouse-scale datacenters. This was before the era that custom silicon was popular, thus many of us converged around bigger bets in networking technologies using off-the-shelf building blocks.

Google had already made substantial innovations in datacenter networking by deploying Clos-based cluster fabrics with custom switches composed of merchant Ethernet-based silicon [29]. The impact this had on our computing infrastructure and software engineering was enormous, as flat bandwidth across 10,000-machine clusters freed us from the inefficiencies of structuring software around much smaller per-rack localities– it was easier and more efficient to schedule compute onto machines and streamlined the development of fault-tolerant distributed systems software. Our appetite for more capable datacenter networking was whetted and we sought to further innovate in this area by pursuing more efficient ways to substantially increase bandwidth and reduce latency – but without a corresponding (or even super-linear) increase in cost.

Meanwhile the high-performance networking (HPC) community had recently published network designs based on novel topologies like the flattened butterfly [15] and the dragonfly. Compared to a folded Clos, these topologies exploited high-radix switch chips and packaging locality to gain substantial cost efficiency. We realized that it was possible to reap the benefits of these novel topologies with off-the-shelf switches targeted towards HPC (notably Infiniband products from Mellanox). Some key issues we focused on included adaptive routing algorithms and deadlock avoidance, topology design to fit Google's datacenter environment, control-plane software, physical design, and integration into applications.

Power consumption of our proposed datacenter network designs was *not* a first-order consideration for our internal efforts to deploy a new innovative cluster fabric. Rather than the typical Google publication that describes a productionized system after-the-fact, we decided to put on our "academic hats" in a side-project to combine the work we were doing with novel network topologies with the emerging idea of energy proportional computing.

II. ENERGY PROPORTIONAL DATACENTER NETWORKS

The premise of the paper was that *if* we largely solved energy proportionality for server machines, *if* server utilization remained low, and *if* we continued to deploy higher-bandwidth fabrics, then pursuing energy proportionality in the datacenter networking equipment was a meaningful optimization. Perhaps a stretch in retrospect, but something we felt was plausible and timely. We started with the topology we were investigating at the time– the flattened butterfly– and made the case that it was naturally more power efficient by using fewer switch chips and links for comparable performance. We then explored the idea of exploiting the dynamic power range of link speed, motivated by actual measurements from the Mellanox Infiniband switch chips that were capable of varying link speeds from 2.5 Gb/s up to 40 Gb/s with energy proportionality. We evaluated the broader potential through simulation using actual workload traces from Google datacenters. Finally we discussed future ideas to dynamically change topologies and evolve switch designs.

We did make an error in the paper in how we framed what factor was held constant in topology comparisons, thereby incorrectly implying that a flattened butterfly has comparable bisection bandwidth to a folded Clos with half the switch chips. What was actually held constant in the topology comparison was not the bisection bandwidth, but rather the bisection bandwidth needed to support the same uniform random traffic pattern. As described in the prior Kim et al. paper presenting the flattened butterfly [15], a folded Clos only achieves 50% link utilization to support a uniform random traffic pattern whereas a flattened butterfly can provision half the bisection bandwidth to support the same traffic pattern. In retrospect, we do not believe this error changes the high-level points of the paper. This discovery of this mistake highlights the imperfect nature of academic research. The recent movement towards artifact evaluation and reproducibility in our field is a large step forward in strengthening the scientific process and in the sifting and winnowing of truth[1]. However the community could potentially benefit from some kind of corrections forum like seen in other journal-based fields.

As we reflect on the impact of the work 13 years later, energy proportional computing remains as important as ever with the growing compute costs (exacerbated by the end of Moore's Law and rise of compute-hungry AI), along with the constant struggle to increase utilization while maintaining latency targets. Yet we have made progress in substantially increasing machine utilization, in part, from our amazing Borg compute environment [31]. Thus with higher utilizations and ever-increasing compute density from servers and GPUs/TPUs, the power consumed by our datacenter networking equipment remains less than 10% of total power [27] which diminishes the need for energy proportionality in the network itself.

III. Direct-connect Topologies

We encountered significant operational challenges with direct-connect topologies, such as the flattened butterfly explored in our work. In datacenter environments, the ability to gradually expand serving capacity and make live changes to clusters, including removing and upgrading server racks, is of utmost importance. However, a direct-connect topology complicates the process because the fabric switches are deployed alongside servers or server racks, carrying traffic associated with other server racks. This makes it more difficult to scale and evolve clusters without causing operational disruptions. On the other hand, an indirect topology, like a folded Clos, connects top-of-rack (ToR) switches to pre-deployed spine switches. This arrangement enables easy addition and removal of server racks with minimal impact.

After publication, we opted for a "hybrid" approach to a flattened butterfly. We designed the first stage of the fabric to connect ToRs to the second stage, which was organized as a direct-connect flattened butterfly. This approach allowed us to take advantage of the scalability and cost benefits of a direct-connect topology while mitigating the operational challenges associated with this type of topology.

Similarly, cost-efficient topologies from direct-connect links also made their way to Google's production Jupiter fabrics, as highlighted in Poutievski et al. [30], leading to large CapEx and OpEx savings. Jupiter now uses optical circuit switches to directly connect larger aggregation blocks. More recent work from this year's ISCA [28] discusses how optical circuit switches directly attach TPU compute nodes to the interconnect in a TPU Superpod, enabling flexible topology changes for improved performance and availability.

IV. Closing Thoughts on Innovation

Finally we retrospect on the dynamics of innovation in large organizations. Organizations need to strike a balance between nurturing innovation and exploiting existing roadmaps. Some amount of separation between groups can be helpful to strike that balance, but too much separation often leads counterproductive conflicts and efforts that don't impact production. We are grateful to the early supporters of our efforts, including Bart Sano and Luiz Barroso, and the original Google-Madison site director, James Laudon, who all created the right organizational dynamics for innovation. While an interesting side project, this work was not our group's primary innovative contribution to the company. The "big bet" project, the Google-Madison lab[2], and our Platforms colleagues went on to seed the first Google TPU accelerator, developed low-latency host networking techniques that helped achieve the AI breakthrough of AlphaGo, developed novel packet processing techniques and frameworks that now power our most demanding use cases, pioneered cluster-scale in-memory storage systems that powers systems like BigQuery, and currently contributes to a range of infrastructure efforts in Google Cloud, data analytics, storage, and more.

References

Refer to original paper for references [1] - [26]

[27] L. Barroso, U. Hölzle, P. Ranganathan "The Datacenter as a Computer: Designing Warehouse Scale Machines," Morgan & Claypool Publishers, 2019

[28] Norman P. Jouppi, George Kurian, Sheng Li, Peter Ma, Rahul Nagarajan, Lifeng Nai, Nishant Patil, Suvinay Subramanian, Andy Swing, Brian Towles, Cliff Young, Xiang Zhou, Zongwei Zhou, David Patterson. "TPU v4: An Optically Reconfigurable Supercomputer for Machine Learning with Hardware Support for Embeddings". In ISCA '2023

[29] Arjun Singh, Joon Ong, Amit Agarwal, Glen Anderson, Ashby Armistead, Roy Bannon, Seb Boving, Gaurav Desai, Bob Felderman, Paulie Germano, Anand Kanagala, Hanying Liu, Jeff Provost, Jason Simmons, Eiichi Tanda, Jim Wanderer, Urs Hölzle, Stephen Stuart, and Amin Vahdat. 2015. "Jupiter Rising: A Decade of Clos Topologies and Centralized Control in Google's Datacenter Network". In SIGCOMM '15.

[30] Leon Poutievski, Omid Mashayekhi, Joon Ong, Arjun Singh, Mukarram Tariq, Rui Wang, Jianan Zhang, Virginia Beauregard, Patrick Conner, Steve Gribble, Rishi Kapoor, Stephen Kratzer, Nanfang Li, Hong Liu, Karthik Nagaraj, Jason Ornstein, Samir Sawhney, Ryohei Urata, Lorenzo Vicisano, Kevin Yasumura, Shidong Zhang, Junlan Zhou, and Amin Vahdat. 2022. "Jupiter Evolving: Transforming Google's Datacenter Network via Optical Circuit Switches and Software-Defined Networking". In SIGCOMM '22

[31] Abhishek Verma, Luis Pedrosa, Madhukar R. Korupolu, David Oppenheimer, Eric Tune, and John Wilkes. "Large-scale cluster management at Google with Borg". In Proceedings of the European Conference on Computer Systems (EuroSys) 2015

---

[1] The phrase "sifting and winnowing [by which alone the truth can be found]" is associated with the University of Wisconsin-Madison and is often mentioned in relation to the Wisconsin Idea that emerged in the early 20th century. The Wisconsin Idea is a guiding principle that emphasizes the university's commitment to extending its knowledge and resources to benefit society as a whole.

[2] The Google-Madison lab continues to thrive to this day and has grown to an engineering office that focuses primarily on infrastructure software, with some on-going hardware work on TPUs.