



Patents and High-Speed Datacenter Interconnects

ERIK R. ALTMAN
Thomas J. Watson Research Center

..... This issue of *IEEE Micro* deals with the increasingly important issue of novel architectures for high-speed datacenter interconnects. As Figure 1 shows, the bandwidth of such interconnects has been increasing at Moore's law rates—doubling every 18 months, which is faster than the 24-month doubling for server I/Os. Server I/Os in turn are growing differently than core counts, and these changes in relative system performance impact microarchitects in terms of things like cache hierarchy, and even how cheap it is to use remotely computed results, which in turn impacts how important it is for a processor to change contexts quickly. Interconnect design also impacts what computation is done on mobile CPUs versus CPUs in a cloud datacenter, with resulting microarchitectural impacts on each CPU type. These indirect impacts are reminiscent of the impacts of pervasive GPU presence on CPU design, outlined in the article “Redefining the Role of the CPU in the Era of CPU-GPU Integration,” by Arora, Nath, Mazumdar, Baden, and Tullsen.¹

I thank Guest Editors George Porter, Alex C. Snoeren, and George Papan from the University of California, San Diego, for their excellent efforts in recruiting and leading reviews for the large set of papers from which we selected the five articles in this issue. These five articles cover several key topics for high-speed datacenter interconnects. Two articles present results

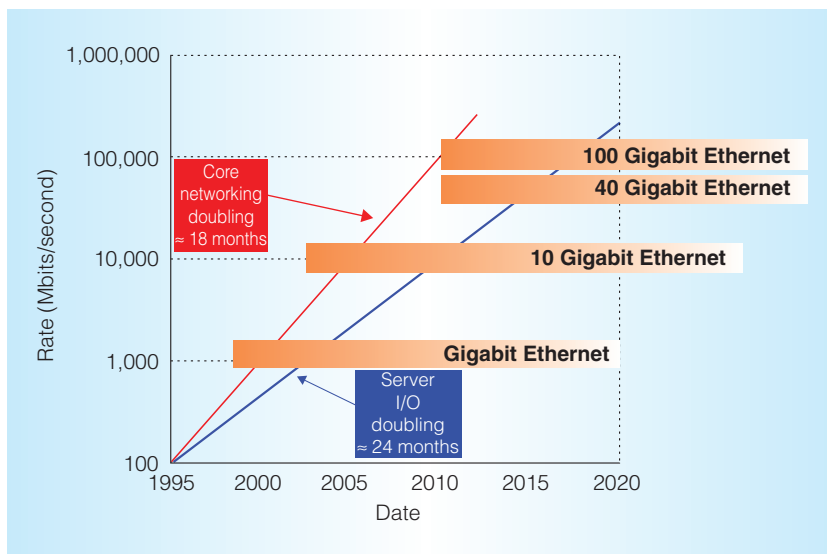


Figure 1. Growth rates of network bandwidth and server I/Os.² The bandwidth of high-speed datacenter interconnects has been doubling every 18 months, whereas the bandwidth of server I/Os has been doubling only every 24 months. (Copyright 2007, IEEE Higher Speed Study Group.)

on interconnection topologies, one discusses a network interface card (NIC) switching architecture to improve rack-level bandwidth, one outlines an optical switch that facilitates graph algorithm computations, and one discusses an FPGA-based network simulation system that can handle datacenter size simulations, as is needed to continue improving these systems.

This issue also begins a two-part series by Richard Stern, our Micro Law columnist, on business and software patents. As Figure 2 indicates, the number of

computing patents granted over the past 20 years has increased almost 12 times versus an overall patent increase of less than 3 times. Similarly, computing patents as a percentage of all patents have gone from 2.3 to 9.5 percent. As such, I think that the Micro Law column will be of interest to many in our *IEEE Micro* audience. The column focuses on the likely impacts of the recent U.S. Supreme court ruling in the case of *Alice Corporation Pty. Ltd. v. CLS Bank International*. The court unanimously held that, “Because Alice Corporation’s patent claims involving (1) a

method for exchanging financial obligations, (2) a computer system as a third-party intermediary, and (3) a computer-readable medium containing program code for performing the method of exchanging obligations are drawn to a patent-ineligible abstract idea under 35 U.S.C. Section 101, they are not patent eligible under Section 101.”³ The Micro Law column does an excellent job of outlining the history behind this ruling and the implications.

Whatever your background, I hope that this issue is interesting and informative—and as always, happy reading.

Erik R. Altman
Editor in Chief
IEEE Micro

References

1. M. Arora et al., “Redefining the Role of the CPU in the Era of CPU-GPU Integration,” *IEEE Micro*, vol. 32, no. 6, 2012, pp. 4-16.
2. H. Frazier, M. Nowell, and S. Trowbridge, “An Overview: The Next Generation of Ethernet,” IEEE 802.3 Higher Speed Study Group, 2007, slide 22; www.ieee802.org/3/hssg/public/nov07/HSSG_Tutorial_1107.zip.
3. “Alice Corporation Pty. Ltd. v. CLS Bank International,” *SCOTUS Blog*, 2014; www.scotusblog.com/case-files/cases/alice-corporation-pty-ltd-v-cls-bank-international.
4. “Patent Counts by Class by Year: January 1977–December 2013,” U.S. Patent and Trademark Office, Patent Technology Monitoring Team (PTMT), 2014; www.uspto.gov/web/offices/ac/ido/oeip/taf/cbcbcy.htm.

Erik R. Altman is the manager of the Dynamic Optimization Group at the Thomas J. Watson Research Center. Contact him at ealtman@us.ibm.com.

cn Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.

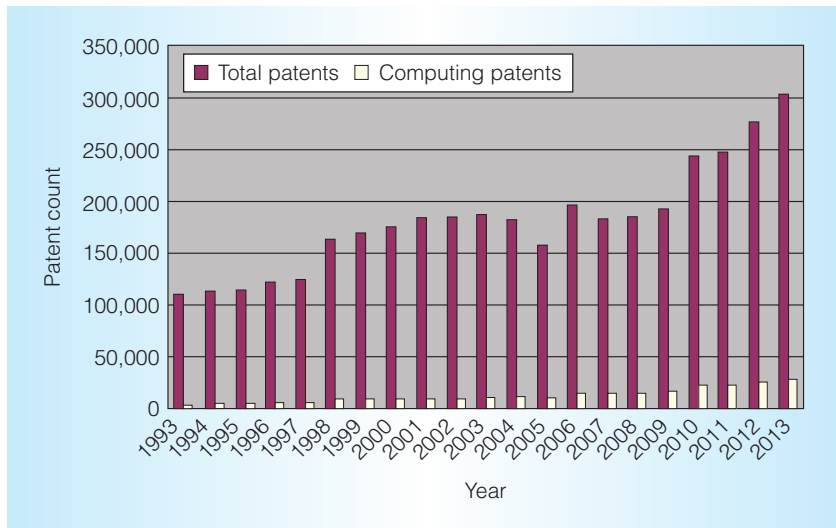


Figure 2. Patents granted by the U.S. Patent and Trademarks Office from 1993 through 2013, both overall patents (maroon) and computing patents (white).⁴ The number of computing patents issued over the past 20 years has increased 11.6 times, while the overall number of patents issued has increased less than 2.8 times.

IEEE micro Calls for Papers

IEEE Micro seeks general-interest submissions for publication in upcoming issues. These works should discuss the design, performance, or application of microcomputer and microprocessor systems. Of special interest are articles on performance evaluation and workload characterization. Summaries of work in progress and descriptions of recently completed works are most welcome, as are tutorials. *IEEE Micro* does not accept previously published material.

Visit our author center (www.computer.org/micro/author.htm) for word, figure, and reference limits. All submissions pass through peer review consistent with other professional-level technical publications, and editing for clarity, readability, and conciseness. Contact *IEEE Micro* at micro-ma@computer.org with any questions.

www.computer.org/micro/cfp

