Dark Silicon and Dangerous Predictions

Erik R. Altman
Thomas J. Watson Research Center

This issue features several articles exploring dark or dim silicon: ''dark silicon'' referring to the observation that power limitations may prevent all silicon on a chip from being used simultaneously, and hence require that some fraction of future chip area be unpowered or underpowered at all times; and ''dim silicon'' referring to the need to always keep some or all of the chip below peak power levels. We already see abundant precursors of dark silicon in existing systems, including smartphones and tablets that mostly sleep, sleeping processors in appliances and parked cars, and traditional computing systems from desktops to clouds when not at peak load.

The articles in this issue explore many key issues in future dark silicon systems, from a discussion of dark silicon trade-offs to near and subthreshold computing to computational sprinting and more. I thank Guest Editors Michael B. Taylor and Steve Swanson, who have done a fine job recruiting and editing this set of articles. I also note that the article authored by Taylor was handled under a completely separate review process, with editors and reviewers anonymous to him.

As evidence of how quickly the field is changing, it's interesting to note that the concept of dark silicon is so recent, it does not appear in the 2012 *International Technology Roadmap for Semiconductors* (http://www.itrs.net). Time will tell if the omission is astute prognostication or late acknowledgment of a key trend.

Table 1. Predictions from the 2003 *International Technology Roadmap for Semiconductors*.

Component	2004	2007	2010	2018
DRAM half-pitch (nm)	90	65	45	18
DRAM memory size (Gbits)	1G	2G	4G	32G
DRAM cost/bit (microcents)	2.7	0.96	0.34	0.02
Microprocessor physical gate length (nm)	37	25	18	7
Microprocessor speeds (GHz)	4.2	9.3	15	53

Table 1 depicts some key data from the 2003 *ITRS*. That roadmap did not mention flash and was clearly optimistic on processor frequency, although I would be happy to hear from anyone using a processor clocked faster than 15 GHz. On the other hand, the current DRAM cost per bit is approximately 0.08 microcents, approximately in line with the roadmap. DRAM half-pitch is currently around 23 nm on the way to about 15.9 nm in 2018—ahead of 2003 predictions.

Careful readers may also note that none of the articles in this issue cite patents, a situation that is common for *IEEE Micro* articles, and indeed I think for most publications in microarchitecture. Perhaps this is because patents are not peer-reviewed in the traditional sense, except by a patent examiner. The language of patents also appears more targeted to lawyers than to researchers and practitioners.

Nevertheless, patents of course play a large role in our field, and the law and rules surrounding them are ever changing, as outlined by Richard Stern in his Micro Law department column in our July/August 2013 issue (''Microsoft Tells Court That Without FRAND, Standard-Setting Would Be 'Blatant Antitrust Violation' ''). The main topic

Published by the IEEE Computer Society

Update on Software-Essential Patents

On 30 July 2013, the US Federal Trade Commission (FTC) made a prepared statement concerning ''Standard Essential Patent Disputes and Antitrust Law'' before the US Senate Subcommittee on Antitrust, Competition Policy, and Consumer Rights (see http://www.ftc.gov/os/testimony/113hearings/130730standardessentialpatents.pdf).

As background, many technology standards include patented components. A patent that is part of a standard gives its owner increased leverage in terms of licensing and royalty fees. As a result, such standard-essential patents (SEPs) have had limitations in law, regulation, and practice on the value of licensing fees allowed, as Richard Stern outlined in his Micro Law column in the July/August 2013 issue of *IEEE Micro*. More specifically, SEPs must be licensed to users on reasonable and nondiscriminatory (RAND) terms. As Stern noted, a SEP case is currently before the US Court of Appeals, with Microsoft arguing that Motorola/Google are not providing RAND terms for a patent that is part of the IEEE 802.11 Wi-Fi standard.

The FTC statement essentially argues in favor of the Microsoft position, that SEPs must be made available on RAND terms, and that Motorola/Google have failed to do so for IEEE 802.11 Wi-Fi patents. The FTC statement further encourages the International Trade Commission (ITC) to block imports of certain Motorola/Google products unless RAND terms are offered on SEP patents. This FTC advice presumably targets Samsung and other Android smartphones, although the testimony omits reference to particular products. The testimony also does not make clear how such ITC actions may relate to the court case Stern described in his July/August column.

Finally, a note of clarification: the International Trade Commission is formally the US ITC and is an agency of the US federal government operating under the US Department of Commerce; that is, the ITC is a US agency dealing with international trade, not a UN-like agency composed of many nations.

discussed in the column, standard-essential patents, continues to draw attention from lawmakers and regulators in Washington, D.C. The ''Update on Software-Essential Patents'' sidebar summarizes recent testimony from the US Federal Trade Commission to the US Senate. *IEEE Micro* takes no position on the merits of the various legal and political positions. However, we hope that Richard Stern's recent column and this update reflect how our microarchitecture-driven interests can shape and be shaped by legal and political forces.

As always, happy reading!

Erik R. Altman
Editor in Chief
IEEE Micro

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.

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