

Technical Perspective

Power Efficiency as the #1 Design Constraint

By Charles Moore

MOORE'S LAW,² AND the associated observations on scaling by Bob Dennard,¹ describe many of the key technical foundations that have given rise to the amazing growth of the modern semiconductor industry. But, taking a step back from these insightful assertions, we can see an even bigger picture emerge related to energy usage and power consumption. The earliest computers, such as the machine envisioned by Charles Babbage in 1822, were mechanical marvels. Although simultaneously amazing and underappreciated, they certainly set the stage for modern computing era.

As it turns out, these machines consumed very large amounts of power (for the time), and were quickly replaced by more power-efficient, electro-mechanical, relay-based systems. This shift enabled larger machines with more capability, but they too soon hit the practical "power wall" of the time. In 1946, the first vacuum tube-based computers were designed, and once again, these set a new standard in capability and power efficiency for the day, eventually replacing the relay-based systems. This pattern repeated with the invention of the discrete transistor, the integrated circuit, the bipolar-based integrated circuits, and the FET-based integrated circuits. The key point here is that these technology transitions were driven in equal parts by the new capability they brought to light *and* the improved power efficiency they offered.

That brings us to modern VLSI-based systems. Once again, our current technology of choice, CMOS-based integrated circuits in particular, has hit a modern day "power wall." There are, of course, lots of highly innovative process technology tweaks and circuit design tricks to extend the life of the CMOS-based era, but these are really only buying time until a fundamentally new technology option becomes practical. Unfortunately,

although there are several interesting contenders, it doesn't appear any of these will be practical within the next decade.

While the technologists and circuit designers continue to extend the life of CMOS, it is also time for computer architects to innovate on fundamentally more power-efficient algorithms and machine organizations. The following paper by Hameed et al. provides a deep dive into a specific application to highlight some of the most significant differences in power efficiency between a general-purpose processor and a special-purpose ASIC. By looking at a range of design options associated with improving the power efficiency of a general-purpose processor running H.264 HD video encode, they uncover an option that effectively uses the instruction orchestration aspect of the general-purpose processor to control the sequencing through a pipeline of customized special-purpose blocks. The new complexity associated with these customized logic blocks notwithstanding, they illustrate a power efficiency improvement of nearly three orders of magnitude.

This paper uncovers some of the most significant differences in power efficiency between a general-purpose processor and a special-purpose ASIC.

The power efficiency and performance advantages of special-purpose ASICs versus general-purpose processors is not new, nor is it surprising. In fact, the essence of these differences is reflected in the computer architects' mantra of "optimizing for the common case." In the 1968 seminar paper "On the Design of Display Processors" by T.H. Myer and Ivan Sutherland,³ this trade-off between generality and complexity is particularly well described with relevant examples from that time period. In particular, they observe that the appropriate solution is both application and technology dependent, and from that, they coin the phrase "The Wheel of Reincarnation" to illustrate these shifting optimizations.

But this paper goes a step beyond the general observation and quantitative analysis of a particular application. It also sets the stage for designing future machines that are prepared for higher-level hardware abstractions. This proposal implies some profound implications for application analysis, algorithm design, machine organization, and associated design methodologies. Combined, they may offer improvements in power efficiency, raw performance, and design productivity. This triple play is particularly significant at this point in time, as the industry must simultaneously work around the ongoing CMOS "power wall" while also investing to find that next technology to reset the power-efficiency bar. **C**

References

1. Dennard, R. et al. Design of ion-implanted MOSFETs with very small physical dimensions. *IEEE Journal of Solid State Circuits* SC-9, 5 (Oct. 1974).
2. Moore, G. Cramming more components onto integrated circuits. *Electronics Magazine*, Apr. 1965.
3. Myer, T.H., Sutherland, I.E. On the design of display processors. *Commun. ACM* 11, 6 (June 1968).

Charles Moore (chuck.moore@amd.com) is a Corporate Fellow at AMD, and the CTO for AMD's Technology Group.