

Life Cycle Aware Computing: Reusing Silicon Technology

John Y. Oliver, Cal Poly San Luis Obispo

Rajeevan Amirtharajah and Venkatesh Akella, University of California, Davis

Roland Geyer and Frederic T. Chong, University of California, Santa Barbara

Despite the high costs associated with processor manufacturing, the typical chip is used for only a fraction of its expected lifetime. Reusing processors would create a “food chain” of electronic devices that amortizes the energy required to build chips over several computing generations.

The past decade has seen unprecedented growth in the number of electronic devices available to consumers. Many of these devices, from computers to set-top boxes to cell phones, require sophisticated semiconductors such as CPUs and memory chips. The economic and environmental costs of producing these processors for new and continually upgraded devices are enormous.

Because the semiconductor manufacturing process uses highly purified silicon, the energy required is quite high—about 41 megajoules (MJ) for a dynamic random access memory (DRAM) die with a die size of 1.2 cm².¹ To illustrate the macroeconomic impact of this energy cost, Japan’s semiconductor industry is expected to consume 1.7 percent of the country’s electricity budget by 2015.² Approximately 600 kilograms of fossil fuels are needed to generate enough energy to create a 1-kilogram semiconductor.³ Furthermore, according to chip consortium Sematech, foundry energy consumption also continues to increase.⁴

In terms of environmental impact, 72 grams of toxic chemicals are used to create a 1.2 cm² DRAM die. The semiconductor industry manufactured 28.4 million cm² of such dies in 2000, which translates to 1.7 billion kilograms of hazardous material.² Due to the increasing number of semiconductor devices manufactured each year, semiconductor disposal costs are likewise increasing.

Despite these costs, the typical processor is used for only a fraction of its expected lifetime. While rapid tech-

nological advances are quickly making silicon obsolete, chips could be removed from recycled electronics and reused for less demanding computing tasks. A processor reuse strategy would create a “food chain” of computing devices that amortizes the energy required to build processors—particularly low-power, embedded processors—over several computing generations.

PROCESSOR LIFETIME ENERGY CONSUMPTION

The *lifetime energy consumption* of a processor or memory chip can be expressed as the sum of the

- manufacturing energy cost, including the creation of silicon wafers, the chemical and lithography processes, and chip assembly and packaging; and
- utilization energy cost.

A comparative analysis of these two components reveals that the energy required to manufacture a processor can dominate the energy consumed over the processor’s lifetime.

Manufacturing energy cost

Semiconductor manufacturing involves many steps, from crystal growth to dicing to packaging. Total energy cost can be expressed as $E_{\text{manufacturing}} = E_{\text{die}} + E_{\text{assembly}}$. E_{die} is the energy required to manufacture the die of the processor or memory chip and includes wafer growth, epitaxial layering, applying photo resists, etching,




Figure 1. Semiconductor yield and manufacturing energy costs over time. (a) Shrinking the processor increases yield, which (b) decreases manufacturing energy costs over subsequent generations.

implantation/diffusion, and managing these procedures. $E_{assembly}$ represents the cost to assemble the chip and includes wafer testing, dicing, bonding, encapsulation, and burn-in testing.

Based on this simple formula, the authors of a recent study¹ made several assumptions about the manufacturing energy cost for any CMOS-based semiconductor. First, they assumed that the energy required to manufacture a 1.2 cm² processor at any lithographical level is the same—thus, the energy costs of manufacturing a 1.2 cm² DRAM die and 1.2 cm² processor die are identical. Another assumption is that the manufacturing energy required is proportional to the semiconductor die area ($E_{die} = 1/yield \times area$), so that a 0.6 cm² processor requires half as much energy for die manufacture as a 1.2 cm² die, adjusted for yield. Finally, they assumed that the assembly energy cost is a constant 5.9 MJ, regardless of the die size. For a 1.2 cm² DRAM chip, $E_{manufacturing} = 41$ MJ, $E_{die} = 35.1$ MJ, and $E_{assembly} = 5.9$ MJ.

As part of their manufacturing energy analysis, the researchers employed the SUSPENS (Stanford University System Performance Simulator) yield model.⁵ According to this model, $yield = e^{D_0} \times area$, with the D_0 constant taken from the International *Technology Roadmap for Semiconductors*.⁶

Figure 1a shows the yield curves for four hypothetical processors over time. Shrinking the processor clearly increases yield. For example, the yield for a 200 mm² processor in 2006 is 60 percent; the same processor, shrunk using 2012 technology, has a yield over 80 percent. The manufacturing cost for subsequent generations of processors thus has the potential to decrease due to shrinking transistor geometry.

Figure 1b demonstrates the energy required to manufacture a processor with fixed functionality over time. The energy savings in subsequent years is due to shrinking transistor geometries and yield improvements. Processors with larger dies have a higher percentage of

energy savings because packaging costs are a smaller portion of the overall manufacturing cost. Also, in the extreme case, shrinking a processor might make it pad limited. The physical dimensions of a pad are unlikely to shrink far below 60 μm on a side.⁷

We believe that the projected figures shown in Figure 1b are on the conservative side. The data the researchers used is from a 4-inch wafer fab, and modern 12-inch wafers require more energy per unit area to process.⁴ In addition, many modern semiconductor processes have more layers than the process used in the study.

The amount of energy required to manufacture a processor die is clearly considerable. A 300 mm wafer uses 2 gigajoules of energy, which is roughly the amount contained within 200 gallons of gasoline. The good news is that the total manufacturing energy cost diminishes with every process shrink. Unfortunately, packaging and assembly costs are relatively fixed.

Utilization energy cost

A processor's utilization energy cost can be determined by simply multiplying its power consumption by the time it is operational. For example, the Intel XScale PX273 consumes 0.77 watts of power in full operation.⁸ Assuming that an XScale-based PDA is used two hours per day 365 days per year, the PX273 consumes just over 2 MJ of energy annually.

One factor that can impact a processor's power consumption is the manufacturing process technology. A benefit of shrinking transistor geometry is that circuits' switching capacitance decreases with each shrink. For low-end cell phones and other devices with relatively fixed performance, processor power consumption may benefit from process shrinks unless leakage current becomes problematic. Higher amounts of leakage make processor reuse a more attractive solution than upgrading to a new process technology, as the processors manufactured with older process technologies will have lower leakage current.




Figure 2. Potential benefits of processor reuse. (a) Upgrading a 1-W processor does not improve the lifetime energy consumption for at least 10 years, making processor reuse an attractive alternative. (b) For processors that use more power—in this case, 20 W—upgrading with newer, more efficient technology makes sense.

PROCESSOR REUSE

Figure 2a illustrates how processor reuse minimizes the lifetime energy consumption of a processor that uses 1 W of power. The two- and four-year upgrade curves increase every two and four years, depicting the high energy cost of manufacturing the processors. These results are based on the assumption that the processor has a die area of 1.2 cm², is operated three hours every day, and is dormant (but still leaking) when not in use. Processors with a 1-W rating or less clearly should not be upgraded with new processors to reduce their lifetime energy consumption. On the other hand, as Figure 2b shows, upgrading is a viable option to minimize lifetime energy consumption for a higher-power processor—in this case, a 20-W processor.

To minimize lifetime energy consumption, it makes sense to reuse a processor when it uses 100 kJ of energy per day or less. Assuming that upgrading occurs in three-year cycles and the device containing the processor is used three hours per day, this is roughly equivalent to the energy a 10-W processor consumes. For perspective, 100 kJ of energy is a bit less energy than is contained within a fully charged laptop battery, or about the same amount in 10 cell-phone batteries.

PROCESSOR FOOD CHAIN

Mobile device processors are typically used for only a fraction of their designed lifetime. “Computer chips can operate for 80,000 hours, and usually machines are thrown out after 20,000 hours,” observed *Guardian* columnist John Keeble. “However, at the moment, 60% of chips cannot be reused because of their specialized functions.”⁹

To facilitate reuse, researchers could standardize embedded processor footprints for a wide range of embedded devices. In addition, instead of reusing a processor in the same device, it could serve a next-generation device with lower performance requirements. Researchers also could apply power-savings techniques like voltage scaling, given the secondary device’s lower computational demand and corresponding operational frequency and voltage.

Example: ARM9

To illustrate how a food chain of electronic devices could reuse a processor, consider the ARM9 processor, which is featured in the Alpine Blackbird PMD-B100 and Sell GPS-350A automotive navigation systems. The ARM9 implementation in these systems runs at 266 MHz. Once the navigation system is recycled, the processor can be removed and placed into a mobile phone like the Sony Ericsson P800, which uses a similar ARM9 processor running at 156 MHz. When this phone is recycled, the processor can in turn be put into a Nintendo DS portable game system, which uses an ARM9 running at 77 MHz.

Table 1 compares the lifetime energy consumption of a processor reuse strategy with a strategy that uses new processors in this chain of devices. These results assume that the automotive navigation system is used one hour per day, the mobile phone three hours per day, and the Nintendo DS game system two hours per day, every day for three years, before being recycled.

Note that manufacturing energy constitutes a large portion of the processors’ lifetime energy consumption. In addition, the manufacturing energy cost of chips in 2009 and 2012 for the new-processor chain decreases only slightly. Some decrease is expected, as the die size shrinks in each generation, but the decrease is limited by the fixed amount of energy required to assemble the processors and the fact that pad size is unlikely to scale with technology.⁷ Also noteworthy is that reused processors have a higher utilization cost than new ones. The increase is small, but it could be important for severely power-constrained devices.

This study neglects the energy required to reclaim a processor, but processor reuse has other benefits that counterbalance this including reduced disposal costs and decreased toxic chemical use. Also, a processor recla-

mation infrastructure already exists, albeit in a black market fashion.¹⁰

Performance requirements

Figure 3 shows the BDTI_{mark} performance of a variety of electronic devices. The blue bars indicate devices that commonly use specialized hardware to accelerate processing and thus may have considerably higher requirements than indicated. The opportunities for processor reuse are evident: A processor used in a particular device should be capable of handling the processing required by all devices to the right of it in Figure 3. For example, the processor from a PDA could be reused in an automobile navigation system.

Over time, the range of performance requirements should continue to grow as the functionality of these devices expands. However, given the ever-present need for low-end processing, a food chain of applications will always exist in some form.

Battery-constrained devices

Because reused processors are manufactured with process technology that is potentially several years older than state of the art, reused processors have higher utilization energy requirements than new ones. Voltage scaling can mitigate this disadvantage. A reused processor that is higher up on the food chain will have a higher peak performance than what is required by a device that is lower on the food chain. Scaling back the frequency, and therefore the voltage of the reused processor, significantly reduces its energy requirements.

In addition, many mobile devices already have adequate battery life. For example, the Nintendo DS game system can run up to 10 hours on a single charge. If the system is used two hours per day, it would have to be recharged once every five days with a new processor but potentially once every four days with a reused processor.

REUSABLE PROCESSOR CHALLENGES

Despite its potential benefits, processor reuse poses both technical and economic obstacles.

Table 1. Lifetime energy consumption: processor reuse versus using new processors.

New processor every 3 years			Processor reused every 3 years	
Manufacturing energy cost (MJ)	Utilization energy cost (kJ)	Year	Manufacturing energy cost (MJ)	Utilization energy cost (kJ)
6.88	36.92	2006	6.88	36.92
0	36.92	2007	0	36.92
0	36.92	2008	0	36.92
6.40	28.87	2009	0	153.74
0	28.87	2010	0	153.74
0	28.87	2011	0	153.74
6.29	4.55	2012	0	50.59
0	4.55	2013	0	50.59
0	4.55	2014	0	50.59
19.57	211.02	Total	6.88	723.75
19.78 MJ		Lifetime	7.60 MJ	

Technical challenges

In order to facilitate processor reuse, it will be necessary to support some circuit flexibility on the die of a reusable processor. To ascertain how much circuit area overhead a reusable processor can tolerate, we compared the manufacturing and utilization energy costs for a strategy that uses new processors every three years with one that uses a single processor every three years for a total lifetime of nine years. Subtracting the energy for the latter strategy from that for the former, we then converted this energy differential to an amount of allowable “additional area” on a reusable processor—that is, we assumed this extra circuitry consumes the same




Figure 3. BDTI_{mark} performance of various electronic devices. A processor used in a particular device should be capable of handling the processing required by all devices to the right of it.




Figure 4. Chip area available for additional circuitry on three processor reuse chains while still maintaining lifetime energy efficiency.

amount of active energy per mm² as the processor core.

Figure 4 shows the additional circuit area that can support processor reuse while reducing the processor's lifetime energy consumption. This allowable area budget clearly depends on the processor's utilization. Processors used less frequently utilize less power and therefore have a higher allowable area budget.

The top line in Figure 4 represents a chain of three processors with capabilities similar to those of an ARM920T. The higher the processor's utilization, the less processor area that can be used for reuse support. For higher-power chips, such as the Intel XScale series illustrated by the bottom line, the reuse-support area decreases significantly. For reuse chains that involve devices with subsequently smaller computational requirements, the area available for reuse support is quite high due to low utilization energy. This is shown by the middle line, which is a reuse strategy based on an XScale in the first generation, ARM9 in the second generation, and ARM7 processor in the third generation. Overall, the additional area for supporting reuse is quite large: An XScale processor core is about 20 mm² in 130-nm technology.

Economic challenges

A major obstacle to processor reuse is that chipmakers would not profit from this strategy unless they become actively involved in salvaging and reselling operations. On the other hand, they would suffer financially only if third parties sold reused chips that competed with the manufacturer's new offerings. Conceptually, the easiest solution would be for chipmakers to charge a premium price for reusable processors that owners of the product containing the chip could recover when returning the product for recycling. Another option would be to credit the chipmaker when one of its processors is reused.

Free-market economic incentives, however, might be insufficient. Environmental protection is often within the purview of public policy. European Union directives to reduce hazardous waste in electronic devices, such as the Restriction of Hazardous Substances (RoHS 2002/95/EC),¹¹ have effectively led all major chipmakers to adopt plans such as moving to lead-free solder. More relevant to processor reuse, the Kyoto Protocol to the United Nations Framework Convention on Climate Change establishes a market economy for greenhouse gas emissions that creates an added financial incentive to reduce energy usage and create carbon-neutral products.¹²

Moore's law has led to a disposable-chip economy with increasingly severe economic and environmental costs. The energy required to manufacture low-power, embedded processors is so high that reusing them can save orders of magnitude of lifetime energy per chip. Processor reuse will require innovative techniques in reconfigurable computing and hardware-software codesign as well as governmental policies that encourage silicon reuse, but the potential benefits to society will be well worth the effort. ■

References

1. E.D. Williams, R.U. Ayres, and M. Heller, "The 1.7 Kilogram Microchip: Energy and Material Use in the Production of Semiconductor Devices," *Environmental Science and Technology*, vol. 36, no. 24, 2002, pp. 5504-5510.
2. R. Kuehr and E. Williams, eds., *Computers and the Environment: Understanding and Managing Their Impacts*, Kluwer Academic Publishers, 2003.
3. E.D. Williams, "Environmental Impacts of Microchip Manufacture," *Thin Solid Films*, vol. 461, no. 1, 2004, pp. 2-6.
4. "ISMI Study Finds Significant Cost Savings Potential in Fab Energy Reduction," Sematech news release, 22 Dec. 2005; www.sematech.org/corporate/news/releases/20051222a.htm.
5. H.B. Bakoglu, *Circuits, Interconnections, and Packaging for VLSI*, Addison-Wesley, 1990.
6. *International Technology Roadmap for Semiconductors: 2005 Edition—System Drivers*, ITRS, 2005; www.itrs.net/Links/2005ITRS/SysDrivers2005.pdf.
7. J. Courtney, N. Aldahhan, and M. Engloff, "The Probe-Centric Future of Test," presentation, 2006 Southwest Test Workshop; www.swtest.org/swtw_library/2006proc/PDF/S06_01_IMSI-SEMATECH.pdf.
8. L.T. Clark et al., "Standby Power Management for a 0.18 μm Microprocessor," *Proc. 2002 Int'l Symp. Low Power Electronics and Design*, ACM Press, 2002, pp. 7-12.
9. J. Keeble, "From Hackers to Knackers," *The Guardian*, supplement online, 21 May 1998; <http://online.guardian.co.uk>.
10. M. Pecht and S. Tiku, "Bogus: Electronic Manufacturing and Consumers Confront a Rising Tide of Counterfeit Electronics," *IEEE Spectrum*, vol. 43, no. 5, 2006, pp. 37-46.

- 
11. "Directive 2002/95/EC of the European Parliament and of the Council of 27 January 2003 on the Restriction of the Use of Certain Hazardous Substances in Electrical and Electronic Equipment," *Official J. European Union*, vol. 37, 2003, pp. 19-23; www.interwritelearning.com/rohs_compliance.pdf.
 12. "Kyoto Protocol to the United Nations Framework Convention on Climate Change," United Nations, 1998; <http://unfccc.int/resource/docs/convkp/kpeng.pdf>.

John Y. Oliver is an assistant professor in the Department of Computer Engineering at California Polytechnic State University, San Luis Obispo, California. His research interests include computer architecture, reliability in computing, and sustainable computing. Oliver received a PhD in computer engineering from the University of California, Davis. He is a member of the IEEE and the ACM. Contact him at jyoliver@calpoly.edu.

Rajeevan Amirtharajah is an assistant professor in the Department of Electrical and Computer Engineering at the University of California, Davis. His research interests include digital and mixed-signal circuit design, low-power signal processing, and system architectures. Amirtharajah received a PhD in electrical engineering and computer science from the Massachusetts Institute of Technology. He is a member of the IEEE and the American Association for the Advancement of Science. Contact him at amirtharajah@ece.ucdavis.edu.

Venkatesh Akella is a professor in the Department of Electrical and Computer Engineering at the University of California, Davis. His current research interests include FPGAs, computer architectures, and embedded systems with an emphasis on low power and reconfigurability. Akella received a PhD in computer science from the University of Utah. He is a member of the ACM. Contact him at akella@ucdavis.edu.

Roland Geyer is an assistant professor in the Donald Bren School of Environmental Science and Management at the University of California, Santa Barbara. His research focuses on the life cycle of manufactured goods, the environmental and economic potential of reuse and recycling activities, and the evolution of green business plans. Geyer received a PhD in engineering from the University of Surrey, Guildford, UK. Contact him at geyer@bren.ucsb.edu.

Frederic T. Chong is a professor in the Department of Computer Science as well as director of the Computer Engineering Program at the University of California, Santa Barbara. His research interests include next-generation embedded architectures, quantum computing architectures, and hardware support for system security. Chong received a PhD in electrical engineering and computer science from the Massachusetts Institute of Technology. He is a member of the ACM. Contact him at chong@cs.ucsb.edu.