

Experiences with Life-cycle Aware Computer Architecture*

JOHN OLIVER

California Polytechnic State University, San Luis Obispo, California, USA. E-mail: jyoliver@calpoly.edu

ROLAND GEYER, ALAN SAVAGE and FREDERIC T. CHONG

University of California, Santa Barbara, California, USA.

RAJEEVAN AMIRTHARAJAH and VENKATESH AKELLA

University of California, Davis, California, USA.

The dark side of Moore's Law is our society's insatiable need to constantly upgrade our computing devices. As a result, the typical processor is only used for a fraction of its expected lifetime, despite the immense cost to produce a processor. While the rapid advance of technology makes silicon obsolete in a few years, we propose that chips should be reused for less demanding computing tasks. This re-use strategy creates a food chain of computing devices which amortizes the energy required to build processors over several computing generations.

This paper is structured into two parts. First, we describe a proposed a processor re-use strategy, showing that processor re-use makes sense for low-power, embedded processors. These re-usable processors occupy a design space that requires us to implement flexible and reliable processors. The second part of this paper describes student efforts centered around re-usable processors at California Polytechnic State University, San Luis Obispo as well as the University of California, Santa Barbara.

Keywords: computer architecture; life-cycle analysis; sustainability, education

1. INTRODUCTION

DURING THE LAST decade, we have seen unprecedented growth in the numbers of electronic devices available to consumers. Many of these devices, from computers to set-top boxes to cell phones, all require highly sophisticated semiconductors such as CPUs and memory chips. The environmental impact of these semiconductors is high due to several reasons. First, the amount of toxic chemicals required to manufacture a semiconductor is staggering [1]. For a single DRAM die of 1.2 cm², 72 grams of chemicals are required. In the year 2000, the semiconductor industry manufactured a total of 28 million cm² of dies, requiring 1.7 billion kilograms of toxic chemicals. Also, due to the increasing number of semiconductor devices that are manufactured each year, the disposal costs of these semiconductors is likewise increasing.

Additionally, because the manufacturing process for semiconductors requires highly purified silicon, the energy required to manufacture semiconductors is very high. About 41 MJ of energy is required to manufacture a DRAM with a 1.2 cm² die [2]. The amount of energy required to manufacture semiconductors is high enough that 1.7% of Japan's national electricity budget is expected to be used by the semiconductor manu-

facturing industry by the year 2015 [1]. Approximately 600 kilograms of fossil fuels are required to generate the energy needed to manufacture a single kilogram of semiconductor [3]. Furthermore, according to Sematech, foundry energy consumption is still increasing [4]. Fortunately, new processors are typically more operationally power efficient per computation than old processors. We call the sum of the manufacturing energy cost and the energy consumed during operation of a processor the lifetime energy consumption of a semiconductor. It turns out that the energy required to manufacture a processor far outstrips the energy consumed during the processor's lifetime for most low-power embedded processors. The high cost of manufacturing a processor, coupled with the fact that many processors are used only a fraction of their lifetime [5], leads us to believe that processors should be re-used.

The rest of this paper is organized as follows. Section 2 looks at the lifetime energy consumption of various embedded processors. From this section, the reductions in lifetime energy consumption through re-using processors becomes clear. Next, Section 3 presents an example of processor re-use, where processors are removed from recycled electronics and used in future devices, thereby saving lifetime energy. Section 4 investigates the design space of re-usable processors. In particular, we look at how much inefficiency can be tolerated in a re-usable architecture and still

* Accepted 10 November 2009.

reduce the lifetime power consumption of the processor. Section 4 also investigates how to support the I/O demands of re-usable processors as well as reliability issues. We then cite some related work and conclude.

2. LIFETIME ENERGY CONSUMPTION OF PROCESSORS

The lifetime energy consumption of a processor or memory chip is comprised of two components. First is the energy required to manufacture the semiconductor, which includes the energy cost to create highly-pure silicon wafers, the chemicals needed to manufacture the processor, the lithography process and the assembly of the die with the packaging of the chip. The second component of the lifetime energy consumption of a processor is the energy consumed during use.

2.1 Manufacturing energy requirements

To manufacture a semiconductor, there are many steps required, from crystal growth to dicing to assembly of the packaging. For the purposes of this study, we can think of the manufacturing process as having two components. The first component is the energy required to manufacture the die of the processor or memory chip. This includes the cost of wafer growth, epitaxial growth, applying photo resists, etching, implantation/diffusion and the management of such procedures. The second component is the cost to assemble the chip which includes wafer testing, dicing, bonding, encapsulation and burn-in testing.

As stated in the introduction, for a DRAM chip with a die size of 1.2 cm^2 , 41 MJ of energy is required during the manufacturing process [2]. For this DRAM, approximately 35.1 MJ of energy is required for the die and 5.9 MJ of energy for the assembly and packaging. For the purposes of this paper, we make several assumptions about the energy required to manufacture any CMOS-based semiconductor, based on [2]. First, we assume that the energy required to manufacture 1.2 cm^2 of processor at any level of lithography is the same. We also assume that the amount of energy required to manufacture a DRAM is similar to the energy to make a processors of the same size. Additionally, we assume that the manufacturing energy required per unit die area is constant, so that a 0.6 cm^2 processor requires half as much energy for die manufacture as a 1.2 cm^2 die, adjusted for yield. Another assumption we make is that the assembly energy cost of a chip is constant, regardless of the size of the die. This assembly energy requirement is 5.9 MJ per processor.

We believe that this manufacturing energy consumption model is a conservative estimate of the manufacturing energy costs for a few reasons. The data used in this study is from a 4-inch wafer fab. Modern 12-inch wafers require more energy

per unit area to process [4]. Additionally, many modern semiconductor processes have more layers than the process used in [2]. Next, we will describe the power consumed by the processors while in operation over their lifetime.

2.2 Operational power model

To find the energy consumption of a processor over its lifetime, we can simply multiply the power consumption of the processor by the time it is operational. For instance, we know that the Intel XScale PX273 processor consumes 0.77 W of power when in full operation [6]. If we assume that an XScale processor is used in a PDA, and this PDA is used 2 hours per day, 365 days per year, this processor uses just over 2 MJ of energy per year.

One factor that can determine the power consumption of a processor is the process technology used to manufacture the processor. A benefit of shrinking transistor geometry is that the switching capacitance of the circuits decreases with each shrink. If a device has a fixed performance and is shrunk to a smaller geometry, we can expect a savings in the amount of operational energy required for that device. The exception is that if leakage current within the processor becomes a problem. Following the projected switching capacitance reductions in the ITRS roadmap [7], we can estimate the reduction in during-use energy consumption for subsequent generations of a chip with fixed amounts of functionality. Likewise, we can estimate the leakage current of a processor from the ITRS roadmap. For all leakage calculations, we assume a modest 40 degrees C, which is reasonable for low-power processors. This leads to modest amounts of leakage, as leakage is highly temperature dependent. Note that this is a conservative assumption for our re-use strategy—higher amounts of leakage would make upgrading a processor to a smaller feature size potentially less attractive, making re-use a more attractive solution.

2.3 Lifetime energy consumption

In Section 2.1, we showed how much energy is required to manufacture a processor, and in Section 2.2 we discussed much energy is required to operate a given processor. Here, we evaluate both the manufacturing and operational power of a processor over time with the goal of showing the impact of process technology on the overall lifetime energy consumption of a processor.

Using our power model, we generate the lifetime energy consumption curves of a 0.5 Watt (W) processor in Fig. 1 and of a 10 W processor in Fig. 2. In both of these figures, we have assumed that the processor is used 3 hours every day, and the rest of the time the processor is dormant. Alternatively, the 0.5 Watt processor could be a 1 Watt processor that is used 1.5 hours per day, every day. Additionally, this figure assumes that the processor has a die area of 1.2 cm^2 . Finally,




Fig. 1. To minimize the lifetime energy consumption, a processor that uses 0.5 W (or less) should never be replaced by a new processor if that processor is still viable.

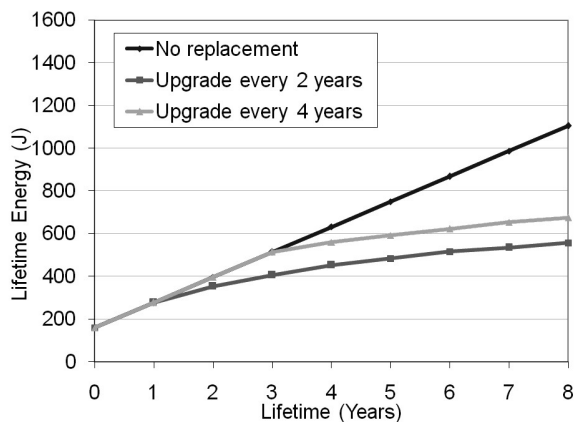


Fig. 2. Processors that use 10W should be upgraded with a processor built with newer, more efficient technology.

these graphs assume that the processor in the device can be upgraded with a processor with identical functionality, but scaled to finer geometries. The curve marked by diamonds assumes no upgrades, the curve marked by circles assumes processor upgrades every four years, and the curve marked by triangles assumes upgrades every two years.

From Fig. 1, we can see that upgrading a 0.5 W processor with the next generation processor does not improve the lifetime energy consumption of the processor for at least ten years. The two-year replacement curve (green) and the four-year replacement curve (yellow) jump up every two and four years respectively. This jump is due to the high energy cost of manufacturing the die. The results is that processors with a power consumption rating of 0.5 W or less should not be upgraded to a newer processor in order to reduce their lifetime energy consumption. This assumes that the processor is also not used more than 3 hours per day, 365 days per year. However, for processors that use more power, replacement can make sense, as shown in the case in Fig. 2 where the processor uses 10 W of power.

The point where it makes sense to replace versus

re-use a processor in order to reduce their lifetime energy consumption is when the processor uses on the order of 100 kJ of energy per day or less for a 1.2 cm² processor. For a 40 cm² processor, that point is about 50 kJ per day. Assuming that replacement is done on a three-year cycle basis, the device is used 3 hours per day, 100 kJ is roughly equivalent to a processor that consumes about 10 W of power. 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 of energy as 10 cell phone batteries.

3. PROCESSOR RE-USE

Because most mobile devices contain processors, and these processors are typically only used for a fraction of their lifetime, we propose that these processors be re-used. To facilitate re-use, we propose that standardized embedded processor footprints be used for a wide range of embedded devices. Additionally, instead of re-using a processor in the same electronic device, we propose that a processor from an electronic device be utilized in a future electronic device which has lower performance requirements. This allows us to make sure that the re-used processor is capable of handling the computational load of the secondary device. Furthermore, power savings techniques like voltage scaling could be applied to the future device, since the next device will have lower computational demand, and therefore operational frequency and voltage, than the previous device.

For instance, an ARM9 processor in an automotive navigation system can be removed from the system once it is recycled. Two automotive navigation systems that utilize ARM9 processors are the Alpine Blackbird PMD-B100 and the Sell GPS-350A. The ARM9 implementations in these systems run at 266 MHz and have 16 kB of cache. Once the navigation system is recycled, the processor can be removed from the navigation system and placed into a mobile phone. Several mobile phones utilize ARM9 processors, including Sony Ericsson, Siemens and Benq phones. For instance, the Sony Ericsson P800 uses an ARM9 processor running at 156 MHz. Once the phone is recycled, the processor can then be put into a Nintendo DS game system, which uses an ARM9 running at 77 MHz.

We call this chain of electronic devices a processor *food-chain*. This example food-chain is shown on the left of Fig. 3, where the mobile phone is 'down-stream' from the navigation system, and the Nintendo is 'down-stream' from the mobile phone.

To compare the lifetime energy cost of a processor re-use strategy against a strategy that uses new processors in each of the devices, we need both the manufacturing energy cost as well as the utilization energy of the processors. In order to find the operational energy required for these processors, we need to make an assumption about how

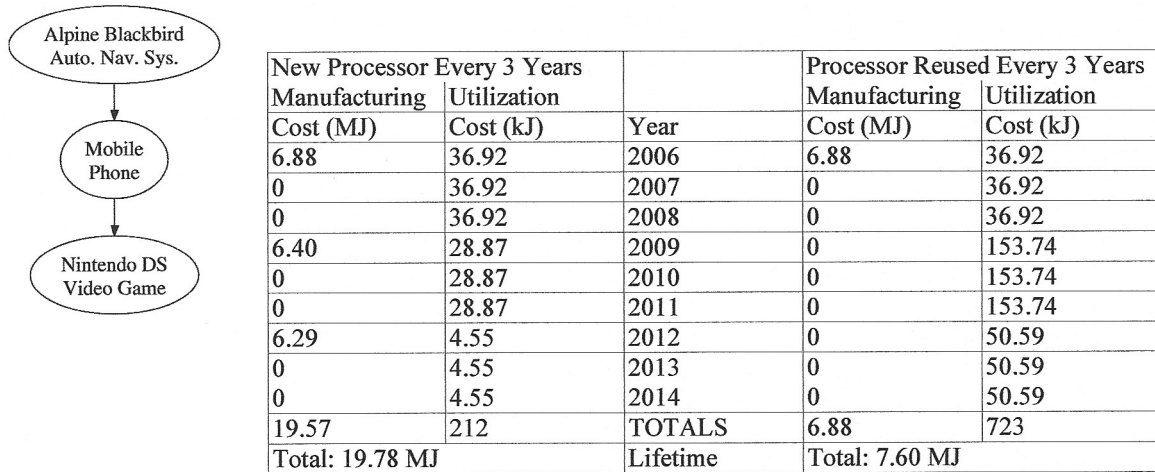


Fig. 3. On the left is an example processor food-chain based on an ARM9 processor. The Table shows the lifetime power consumption benefits of re-using this processor.

frequently these electronic devices are used. We assume that the automotive navigation system is used 1 hour per day, every day for 3 years, before being recycled. Likewise, the mobile phone's utilization is 3 hours per day, and the Nintendo DS game system is used 2 hours per day.

On the right of Fig. 3 is a table that shows the lifetime energy consumption of this food chain. The left side of the table is the case where new processors are used in each of the down-stream devices. The right side of the table in Fig. 3 shows the lifetime energy consumption of the same chain of devices that uses a re-use policy. There are several remarkable things about this table. First, we notice that the manufacturing energy of processors makes up a large portion of the lifetime energy consumption of the processors. Also, note that the manufacturing energy cost of the processors manufactured in 2009 and 2012 for the chain which uses new processors (left side of the table) does not decrease very much. Some decrease is expected, as the die size shrinks in each of future versions of the processor. However, the decrease is not large due to the fixed amount of energy required to assemble the processors. Another reason the manufacturing energy cost does not decrease as significantly as expected is due to the fact that pad size is unlikely to scale with technology [8]. Also of interest is that the operational power for the re-use strategy is higher than the operational power consumption of a strategy that utilizes new processors. This is true because new processors can take advantage of new process technologies, reducing their operational power. While the increase in operational power is small, this may be important for devices that are severely power constrained or for higher-powered processors. Note that we neglect the energy required to reclaim a processor. However, there are also other potential benefits of processor re-use, such as reduced disposal costs and a reduction in the number of highly-toxic chemicals required to

produce processors. It turns out that processor reclamation infrastructure already exists, albeit in a black-market fashion for counterfeiting chips [9].

3.1 Battery constrained devices

One disadvantage that re-usable processors have is higher operating energy requirements. This is due to the fact that they are manufactured in a process technology that is potentially several years older than state-of-the art.

One technique to mitigate this effect is to use voltage scaling. A processor that is re-used from higher on the food-chain will have more performance than what is required by a device that is lower on the food-chain. By scaling back the frequency, and therefore the voltage of the chip, we can significantly reduce the energy requirements of the re-used processor.

Additionally, many mobile devices may already have adequate battery life for many applications. For instance, the Nintendo DS portable game system can run up to 10 hours on a single charge. If a Nintendo DS with a new processor is only used 2 hours per day, it would only have to be recharged once every 5 days. Compare that to a Nintendo DS game system with a re-used processor that uses 15% more power. The user would then be forced to recharge the game system once every four days, instead of once every five.

4. DESIGN SPACE STUDY OF RE-USABLE PROCESSORS

In this section, we investigate the design space of re-usable processors. We look at how much circuit area overhead we can tolerate on a processor specifically designed for reuse. This is done by simply 'reversing' the lifetime energy model. We assume that a processor is re-used every 3 years, for a total lifetime of nine years.

We find the energy required for a strategy that

uses new processors and subtract from it the energy required for a strategy that implements processor re-use. This difference in lifetime energy consumption between a processor re-use strategy and a strategy that implements new processors in subsequent generations is then converted to ‘additional area’ that could be used on a re-usable processor following our manufacturing energy model. The additional area found is assumed to consume the same amount of active energy per mm^2 as the core of the processor.

Figure 4 shows the additional circuit area that we can spend supporting processor re-use while reducing the processor’s lifetime energy consumption. Essentially, this is an allowable area budget that we can use to support different requirements of re-usable processors. From Fig. 4, we can see that this budget depends on the power consumption of the processor and the utilization (in terms of hours per day) of the processor. Processors that are used less frequently utilize less power, and therefore have a higher allowable area budget. The top line in Fig. 4, marked by diamonds, is for a processor food-chain that consists of three processors with similar capabilities of an ARM920T processor. The higher the utilization of this processor, the smaller the area budget becomes. For higher power processors, such as the Intel Xscale PA255 processor, we can see that the amount of additional area is drastically reduced. For re-use chains that involve decreasing computational requirements in subsequent generations, we find that area available to support re-use is quite high. This is due to reductions in active power consumption of the subsequent devices, which is shown in Fig. 4 by the line marked by triangles. This line represents a re-use strategy where the first generation device has XScale performance requirements, the second generation has ARM9 performance requirements, and the third generation has ARM7 performance requirements. From Fig. 4, we can see that our area budget is quite large (although we would not want to use this entire budget if possible). For an Xscale processor that is used four hours per day,

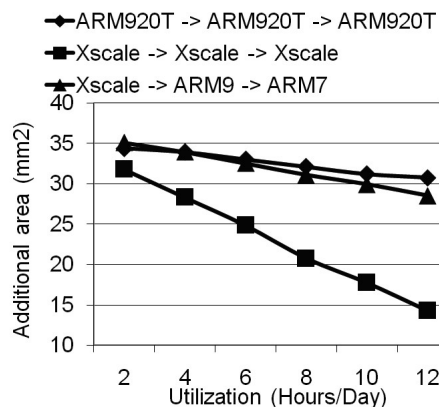


Fig. 4. The additional chip area available on three re-use chains while still maintaining lifetime energy efficiency.

we have sufficient area budget to put another one and a half Xscale processors!

5. STUDENT-BASED EFFORTS IN RE-USABLE PROCESSORS

In this section, we briefly describe the various student-related engineering efforts that are currently underway at California Polytechnic State University in San Luis Obispo (Calpoly). It should be noted, that some of the original research into the energy cost of semiconductor manufacturing was done the University of California, Santa Barbara and Davis. Generally, across all of these projects, it has been educational for students to have to design under a new design criteria: lifetime energy. This forces the students to find creative solutions that are not covered by any textbook.

5.1 Input/Output requirements for re-usable processors

One particular engineering effort that students at Calpoly are working on is integrating a common, flexible I/O circuit that can cover multiple I/O standards with a reduced hardware cost. The number of supported I/O protocols in embedded processors is numerous: I²C, RS232, PCI, USB, etc. We propose integrating I/O devices on-chip and muxing the on-chip I/O devices to provide flexibility in re-use. As long as the combination of I/O devices that we build on-chip does not exceed the area budget found in Section 4, we will still reduce the lifetime energy of a processor while supporting a wide array of I/O standards. The students on this project have developed a figure-of-merit which includes a lifecycle-aware portion by which to guide their design efforts and their development of an ‘integrated serial I/O driver’ is currently underway.

5.2 Wearout detection and wearout tolerant design

It is difficult to accurately estimate the amount of wearout of circuits at design-time, even with detailed reliability modeling. At Calpoly, there are several students who are analyzing the use of stability detection circuits that can be used to improve the lifetime of a processor. Our base design is similar to the stability detection circuit from Franco and McCluskey [21] and is small enough to be employed on widely. Using this circuit, if select structures being to slow down, we can detect this and either reduce the performance of the device, or swap in a spare structure. Similar ideas have been tried at other academic institutions [19], but we are evaluating them from the perspective of lifetime energy cost. Currently, students are investigating several different stability detection circuits and evaluating their suitability for use in a re-usable processor.

Another student initiative is to identify where the weakest links in a processor are located in terms of wearout. If we can identify those links,

potentially we can improve the mean-time to failure (MTTF) of the entire processor for a marginal increase in power or area, or slight degradation in performance. Our tool infrastructure takes circuit layout descriptions and simulates them to find the power density, as well as delay of the circuits. We feed the power density information into a spatial temperature simulator and then measure the MTTF based on some published equations [11]. This way, we have a simulator that we can use to evaluate trade-offs between power consumption, area, performance and MTTF.

5.3 Graceful degradation of performance in re-usable processors

While in Section 3 we stated that processors are typically only used for a fraction of their lifetime, wear-out is becoming an increasing concern as we move to finer geometries. In the worst case, we assume that a processor is designed for the lifetime of single task in our food chain and that we need to extend that lifetime by a factor of three. There are four different primary non-transient failure mechanisms that are of concern: electromigration, stress migration, time-dependent dielectric breakdown, and thermal cycling [10, 11] that we are modeling.

If we assume that the processor food-chain for re-usable processors is a down-stream food-chain, where subsequent devices have a smaller computational requirement than previous devices, we can make a further area savings to support reliability in re-usable processors. We note that, for instance, the functionality of the Blackfin DSP is a superset of the ARM9 processor and the PIC16 micro controller. Likewise, the ARM9's functionality is a superset of the PIC16 micro controller. Those structures that are shared between the Blackfin DSP with the ARM9 are shown in light and dark gray in Fig. 5. Those structures that are shared

between the Blackfin DSP and the ARM9 with the PIC 16 micro controller are shown in dark gray. The idea is to selectively grow or replicate those structures which are shared between the processors of the food-chain, in hopes of making a Blackfin-based re-usable processor more reliable while minimizing lifecycle energy costs. Currently, a student at Calpoly is designing the reusable hardware prototype on an FPGA.

5.4 Graceful degradation of performance in re-usable processors

While in Section 3 we stated that processors are typically only used for a fraction of their lifetime, wear-out is becoming an increasing concern as we move to finer geometries. In the worst case, we assume that a processor is designed for the lifetime of single task in our food chain and that we need to extend that lifetime by a factor of three. There are four different primary non-transient failure mechanisms that are of concern: electromigration, stress migration, time-dependent dielectric breakdown, and thermal cycling [10, 11] that we are modeling.

If we assume that the processor food-chain for re-usable processors is a down-stream food-chain, where subsequent devices have a smaller computational requirement than previous devices, we can make a further area savings to support reliability in re-usable processors. We note that, for instance, the functionality of the Blackfin DSP is a superset of the ARM9 processor and the PIC16 micro controller. Likewise, the ARM9's functionality is a superset of the PIC16 micro controller. Those structures that are shared between the Blackfin DSP with the ARM9 are shown in light and dark gray in Fig. 5. Those structures that are shared between the Blackfin DSP and the ARM9 with the PIC 16 micro controller are shown in dark gray. The idea is to selectively grow or replicate those

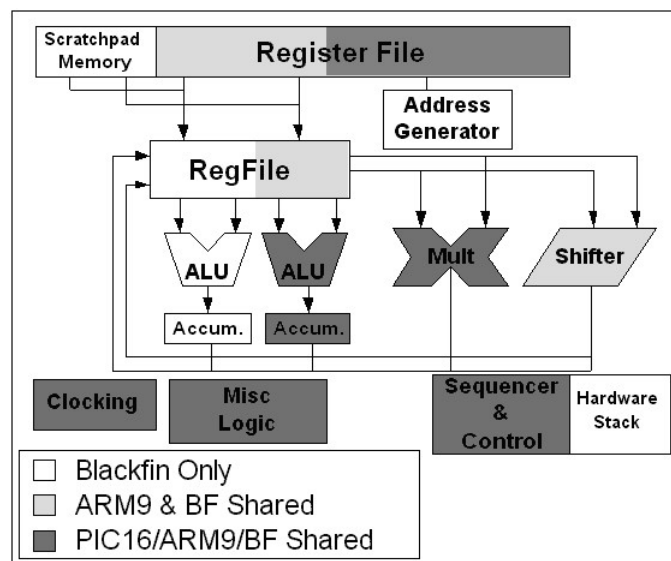


Fig. 5. Block diagram of shared micro-architectural structures for the Blackfin, ARM9 and PIC16 embedded processor cores.

structures which are shared between the processors of the food-chain, in hopes of making a Blackfin-based re-usable processor more reliable while minimizing lifecycle energy costs. Currently, a student at Calpoly is designing the reusable hardware prototype on an FPGA with two different CPU cores.

6. RELATED WORK

The impact of semiconductors on the environment have been studied before [12]. However, we believe that proposing processor re-use is novel in the literature.

Lately, the reliability of processors has been the subject of many papers in the literature. However, many of these papers deal with dynamic detection and/or correction of transient faults. The Razor paper explores dynamic error correction and detection while using aggressive voltage scaling to reduce power consumption [13]. Similarly, the Diva work uses a simple processor that verifies the correctness of the core processor's computation, only permitting correct results to commit [14]. The HotSpot toolset is a thermal modeling tool for processors and they show correlation between temperature and many failure modes within a processor [15]. There are many dynamic techniques that can be used to reduce the thermal profile of a processor [16, 17].

For a reliable computing substrate, we propose a method of processor duplication, which allows us to migrate processor activity from one processor to another. Activity migration is typically used on a micro-architectural level, rather than for individual cores, and has been studied before [18].

Perhaps most closely related to the reliability investigation in this paper is the RAMP paper from UIUC [10]. They investigate graceful degradation of a processor as the processor is subjected to wear-out. Many of their techniques take advantage of the dynamic issue of computations in a superscalar processor. For our study, the embedded domain is dominated by static-issue processors where many of their techniques do not apply.

7. FUTURE WORK AND CONCLUSIONS

The combination of pervasive electronics and Moore's law improvement in silicon has led to a disposable-chip economy of increasingly negative environmental impact. While substantial practical challenges remain, this paper motivated a technical strategy for processor re-use. In particular, we showed that the energy required to manufacture low-power, embedded processors is so high that processor re-use can save orders of magnitude of lifetime energy per processor. We identified processor wear-out as an important technical challenge for future process technologies and explored mitigation strategies and their effect on lifetime energy savings.

Calpoly students are tasked with taking circuits which were optimized for performance and re-designing them with the idea of minimizing their lifetime impact. While significant social, economic and technical road-blocks make re-usable computer processors an unattainable goal, this experience provides a foundation of thought for students and the sustainability of their designs.

REFERENCES

1. R. Kuehr and E. Williams, *Understanding and Managing their Impacts Computers and the Environment*, p. 1, Kluwer Academic Publishers, 2004.
2. 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*, 2002, **36**, pp. 4–10.
3. E. Williams, *Proceedings of symposium on semiconducting silicides: Science and future technology of the 8th IUMRS international conference on advanced materials*, 2004, p. 26.
4. Sematech, *ISMI study finds significant cost savings potential in fab energy reduction*. <http://ismi.sematech.org/corporate/news/releases/20051222a.htm>, 2002.
5. J. Keeble, From Hackers to Knackers, *The Guardian Supplement*, 1998, p. 13.
6. L. T. Clark, N. Deutscher, S. Demmons and F. Ricci, *Standby power management for a 0.18 μm microprocessor*, Proceedings of the 2002 international symposium on Low power electronics and design, 2002, p. 712.
7. ITRS, *International Technology Roadmap For Semiconductors*, System Drivers. Semiconductor Industry Association, 2005.
8. J. Courtney, N. Aldahhan and M. Engloff, *The probe-centric future of test*, Southwest Test Workshop, 2006.
9. M. Pecht and S. Tiku, Bogus: electronic manufacturing and consumers confront a rising tide of counterfeit electronics, *IEEE Spectrum*, 2006, **43**, pp. 37–46.
10. J. Srinivasan, S. V. Adve, P. Bose and J. A. Rivers, Lifetime reliability: Toward an architectural solution, *IEEE Micro*, 2005, **25**(3), p. 70–80.
11. Joint Electron Device Engineering Council, *Failure mechanisms and models for semiconductor devices*. www.jedec.org/download/search/jep122C.pdf, 2006.
12. S. Boyd, D. Dornfeld and N. Krishnan, *Life cycle inventory of a CMOS chip*, Proceedings of the 2006 IEEE International Symposium on Electronics and the Environment, 2006, p. 253–257.

13. D. Ernst, N. S. Kim, S. Das, S. Pant, R. Rao, T. Pham, C. Ziesler, D. Blaauw, T. Austin, K. Flautner and T. Mudge, *Razor: A low-power pipeline based on circuit-level timing speculation*, Proceedings of the 36th annual IEEE/ACM International Symposium on Microarchitecture, 2003, p. 7.
14. T. M. Austin, *Diva: a reliable substrate for deep submicron microarchitecture design*, Proceedings of the 32nd annual ACM/IEEE international symposium on Microarchitecture, 1999, p. 196–207.
15. K. Skadron, M. R. Stan, W. Huang, S. Velusamy, K. Sankaranarayanan and D. Tarjan, *Temperature-aware microarchitecture*, Proceedings of the 30th annual international symposium on Computer architecture, 2003, p. 213.
16. D. Brooks and M. Martonosi, *Dynamic thermal management for high-performance microprocessors*, Proceedings of the 7th International Symposium on High-Performance Computer Architecture, 2001, p. 171.
17. K. Skadron, T. Abdelzaher and M. R. Stan, *Control-theoretic techniques and thermal-rc modeling for accurate and localized dynamic thermal management*, Proceedings of the 8th International Symposium on High-Performance Computer Architecture, 2002, p. 17.
18. S. Heo, K. Barr and K. Asanovic, Reducing power density through activity migration, International Symposium on Low Power Electronics and Design, 2003, pp. 217–222.
19. J. Blome, S. Gupta, S. Feng, S. Mahlke and D. Bradley, *Online timing analysis for wearout detection*, The Second Workshop on Architectural Reliability (WAR), 2006.
20. J. Oliver, R. Amirtharajah, V. Akella, R. Geyer and F. T. Chong, Life Cycle Aware Computing: Reusing Silicon Technology. *IEEE Computer*, 2007, pp. 56–61.
21. P. Franco and E. J. McCluskey, *On-line Delay Testing of Digital Circuits*. Proceedings of the 12th IEEE VLSI Test Symposium, 1994.

John Oliver (Assistant Professor) joined the Electrical Engineering Department at California Polytechnic State University, San Luis Obispo in 2007. He completed the Ph.D. degree in Electrical and Computer Engineering at UC Davis in 2007. Dr. Oliver's field of research is computer architecture and teaches courses in digital systems design, circuits and computer architecture.

Rajeevan Amirtharajah received the S.B. and M.Eng. degrees in 1994, and the Ph.D. degree in 1999, all in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA. From 1999 to 2002, at High Speed Solutions Corp. (later Intel), Hudson, MA, he developed high performance memory buses. He is currently an associate professor at the University of California, Davis, where his research focuses on low power micro-architecture, circuit and interconnect design, energy scavenging, and signal processing for wireless sensor nodes. He received the National Science Foundation CAREER award in 2006.

Fred Chong is the Director of Computer Engineering and a Professor of Computer Science at UCSB. He also directs the Greenscale effort in Energy-Efficient Computing, which involves over 20 multi-disciplinary faculty. Chong received his Ph.D. from MIT in 1996 and was a faculty member and Chancellor's fellow at UC Davis from 1997–2005. He is a recipient of the NSF CAREER award and his research interests include emerging technologies for computing, multicore and embedded architectures, computer security, and sustainable computing.

Venkatesh Akella is a Professor of Electrical & Computer Engineering at University of California, Davis. He received his PhD in Computer Science from University of Utah. His current interests encompass embedded systems, computer architecture, field programmable logic and entrepreneurship.

Roland Geyer is Assistant Professor at the Bren School of Environmental Science and Management, University of California, Santa Barbara. He has a graduate degree in physics from the Technical University Berlin and a PhD in engineering from the University of Surrey, UK. In his research he uses the approaches and methods of industrial ecology, such as life cycle assessment and material flow analysis, to quantify the environmental performance and improvement potential of production and consumption systems.

Alan Savage received a BS in Computer Science from UCSB in 2007, and is currently pursuing my MS at UCSB. Alan's research interests are efficient computing in data centers, multi-core cache architectures, and value prediction for error tolerant applications.