March 2010

Final Report for CCC Cross-Layer Reliability Visioning Study: Executive Summary Version 0.5

The geometric rate of improvement of transistor size and integrated circuit performance known as Moore's Law has been an engine of growth for our economy, enabling new products and services, creating new value and wealth, increasing safety, and removing menial tasks from our daily lives. Affordable, highly integrated components have enabled both life-saving technologies and rich entertainment applications. Anti-lock brakes, insulin monitors, and GPS-enabled emergency response systems save lives. Cell phones, internet appliances, virtual worlds, realistic video games, and mp3 players enrich our lives and connect us together. Over the past 40 years of silicon scaling, the increasing capabilities of inexpensive computation have transformed our society through automation and ubiquitous communications.

Looking forward, **increasing unpredictability threatens our ability to continue scaling integrated circuits at Moore's Law rates.** As the transistors, wires, and other components that make up integrated circuits become smaller, they display both greater differences in behavior among devices designed to be identical and greater vulnerability to transient and permanent faults. Conventional design techniques expend energy to tolerate this unpredictability by adding safety margins to a circuit's operating voltage, clock frequency or charge stored per bit of data. Such approaches have been effective in the past, but their costs in energy and system performance are rapidly becoming unacceptable, particularly in an environment where power consumption is often the limiting factor on integrated circuit performance and energy efficiency is a national concern.

In the realms of memory and communication, engineers have a long history of success in tolerating unpredictable effects such as fabrication variability, transient upsets, and lifetime wear through strategic use of redundant information and multi-layer approaches that anticipate, accommodate, and suppress errors. Techniques to tolerate occasional and unpredictable deviations from intended functionality have reduced cost, reduced energy consumption, and increased performance while guaranteeing robust operation in the presence of noise. Unfortunately, mitigating logic errors is not as simple or as well researched as memory or communication systems. This lack of understanding has led to very expensive solutions. For example, triple-modular redundancy masks errors by triplicating computations in either time or area, albeit with a 200% increase in energy consumption.

If researchers are able to develop techniques that tolerate the growing unpredictability of silicon devices, Moore's Law scaling should continue until at least 2022. During this 12-year time period, transistors, which are the building blocks of electronic devices, will scale their dimensions (*feature sizes*) from 45nm to 4.5nm. This additional scaling will be possible only if we can mitigate unpredictability and noise effects in small feature-size devices as effectively as we have been able to compensate for these effects in memory—that is, without paying an increasing energy overhead. The challenge for the near future, then, is to achieve the density and full energy benefits of ideally-scaled smaller technologies with the reliability of our larger and older technologies. Over the longer term, techniques to tolerate unpredictable behavior at low-energy costs will also be necessary to make post-silicon technologies, such as quantum and molecular computation, feasible.

To better understand the reliability challenges facing future electronic systems, the limitations of current approaches, and the opportunities offered by techniques that distribute the responsibility for reliability across the entire system stack (See Figure 1), the Computing Community Consortium funded a study into cross-layer reliability, which was carried out from March-October of 2009. All told, eighty people participated in one or more of the study group's three meetings, representing academia, industry, and government.[1]

During the study, we formulated constituency groups on commercial and consumer electronics, infrastructure, life-critical, space/avionics and large-scale systems. These groups covered computing systems ranging from pacemakers and hand-held computers to satellites and supercomputers, as well as applications ranging from entertainment to protecting human life. Each of the groups had different perspectives on how scaling-induced power and reliability challenges impacted them and were willing to make different tradeoffs to address these challenges. All of the industries represented are feeling pain from these reliability problems and agree that the coming challenges will not be satisfied by current solutions.

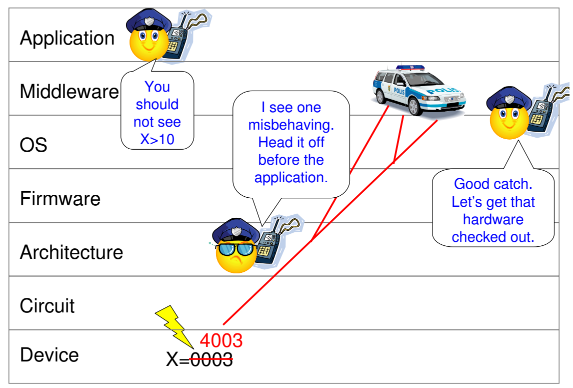

Figure 1: Cross-Layer Cooperation

Several corporations shared their perspective on how reliability affects them economically:

"Microprocessor and computer system reliability is a critical issue for Intel. As we look forward, we are deeply concerned about the techniques used to tolerate errors, device variations, and silicon aging and their impact on the cost, performance, and power consumption of our products. Meeting the reliability challenges of future products will require innovative new approaches in which the entire system contributes to overall reliability, and we strongly endorse research into these architectures and methods." – Justin Ratter, Vice-President and Chief Technology Officer, Intel Corporation.

"Reliability and fault tolerance are essential to the performance of many embedded systems. For example, automotive, industrial and medical applications may involve safety-critical functionality and harsh operating environments. As the scope and complexity of these applications continue to increase, new approaches are needed to design, validate and qualify highly reliable and resilient embedded systems." – Ken Hansen, Vice President and Chief Technology Officer, Freescale Semiconductor Inc.

"Achieving high reliability in the types of high end systems that IBM makes is a very high priority for us. We invest significant design and engineering resources into the hardware, firmware and software aspects of error detection, correction, and avoidance. As the potential for these errors increases due to technology scaling, additional investments must be made, and new ideas proposed to find solutions that are cost effective, practical and achievable within current design paradigms. This is not a problem that we expect to go away anytime soon, and solving it must have a high priority." – Carl J. Anderson, IBM Fellow.

All of the constituency groups reported significant reliability and efficiency challenges stemming from the fact that layers of current systems are designed and operated without key information that is available to other layers. Throughout the study, participants described situations where the lack of global information caused designers to make worst-case assumptions, resulting in systems that were over-designed, inefficient, and overly expensive. Looking forward, many of the participants identified cases where current design techniques would be unable to meet the needs of future

[1]See `http://www.relxlayer.org/Participants`.

systems, making a new approach to reliable system design necessary instead of merely desirable. **Common challenges that should be addressed include:**

- *Late-Bound Information:* Critical information about how a system or component will be used and/or the characteristics of the technology that will be used to manufacture it is often unavailable until late in the design process. Inflexible system designs cannot adapt to this information as it becomes available, resulting in over-design, limited ability to use the system in different contexts, and occasional need for complete re-design late in the design process.
- *Changing Operating Environments:* Error rates in electronic systems vary substantially depending on operating conditions (temperature, altitude, etc.) Systems that cannot sense and react to changes in their environment must always inefficiently assume worst-case conditions.
- *Application Variability:* Different applications (mp3 player vs. power plant control) have very different reliability requirements. Hardware that does not know an application's needs must provide complete reliability even if the application can tolerate errors.
- *Component Variation:* Electronic components vary from part to part and from manufacturer to manufacturer. Systems that cannot analyze and adapt to variation must assume worst-case behavior, leading to designs with poor performance and high power consumption.
- *Poor Understanding of Vulnerabilities:* Current design methodologies make it difficult to understand the real sources of weaknesses in a design, leading to over-design of components that are not critical to reliability.

The common causes of the challenges facing our constituency groups argue that we need to develop a better **scientific and engineering** understanding of information sharing and exploitation across the multi-layer system stack that supports computations. Specific components of this inquiry should include:

1. Design of hardware organizations that are prepared for repair.
2. Techniques to filter errors at multiple levels.
3. Formulation, analysis, and management of multilevel trade-offs that generalize the idea of hardware-software trade-offs, including the interfaces for cross-layer information sharing.
4. Theory and design patterns for efficient, light-weight error checking.
5. Theory and practical frameworks to express and reason about differential reliability, including both application needs and hardware and system organization to meet those needs.
6. Scalable solutions that can adapt to varying error rates and reliability demands.
7. Components that degrade gracefully and systems that are aware of their overall health.

Developing this understanding will directly impact many key **national missions**, including:

- **supercomputing**: Current supercomputers generate approximately 0.2 GigaFLOPs of computation per Watt of electrical power used and spend upwards of 20% of their time recovering from errors. In order to build the ExaFLOP supercomputers that will drive science, defense, and commerce in the 2020's while staying within the 20-Megawatt power budgets of major data centers, we will need to increase computational efficiency to greater than 50 GigaFLOPs/Watt, an improvement of more than 250 times, while simultaneously drastically reducing the fraction of time spent handling errors and extending mean-time-between-failures (MTBF) to months or years. To achieve this level of energy reduction, it will be necessary to mitigate a higher rate of device failure with lower energy overhead.
- **satellites**: Satellite technology provides cable television, supports in-theater warfighters, and supports science. Failures on these systems could cost millions from lost advertising revenue or loss of human life. On-board processing is necessary to optimize the limited communi-

cation bandwidth to the ground (as low as 9600 baud), but many payloads have 20-30 Watt power limits. As radiation-hardened technologies lag commercial technologies by several generations, they may not meet the increasing needs of satellite systems. We must find ways to use more advanced and energy-efficient technologies without sacrificing system reliability.
- **medical**: Reliable, ultra-low-power computing systems will enable a host of breakthroughs in medical technology, including personal genomics techniques that allow drugs to be tailored to an individual's biochemistry, sensing devices that help compensate for blindness, assistive technologies that allow the elderly to remain more independent, more life-like artificial limbs, and implantable devices that operate for years without failure or the need for recharging.
- **transportation**: Advanced safety features and drive-by-wire control demand the greater computational capabilities of advanced technologies, but also have even higher safety requirements to safeguard human life.
- **security**: As illustrated by a number of recent attacks that exploit transient hardware errors to compromise a system's security mechanisms, a computer system cannot be secure unless its data and computations are reliable. Cross-layer reliability techniques will help future computer systems provide the security required for electronic commerce, electronically-stored medical records, and military applications by providing a solid, reliable, foundation on which software security mechanisms can be built.

A **research program** in the nascent area of cross-layer reliable systems could have tremendous impact and influence over the development of this critical technology. Because of the cross-cutting nature of this area, it will be essential for this program to enable collaboration of researchers with a wide range of expertise. The program should focus on developing example systems and the standard, open platforms that will enable the subsequent engagement of a larger community of domain experts across academia, government, and industry. It should also support and encourage the development of tools to model and characterize cross-layer reliable systems.

Government leadership is essential. The work necessary to achieve cross-layer reliable systems crosses the entire computing system ecosystem from integrated circuits to software applications. Therefore, no one vendor or research laboratory will be able to effect change by themselves. Wide-scale cooperation across specialties and organizations is necessary to revolutionize computing systems in this manner, otherwise the community will be facing yet another stop-gap revision that will only postpone these problems for a few more years! Finally, changes are needed to Electrical Engineering, Computer Engineering, and Computer Science curricula so that young engineers and programmers are equipped to design and develop reliable computing systems.

There is also a strong **economic argument** for why this research needs to happen here in the United States. Over the last several decades, the US has been a primary contributor to the electronics revolution and has reaped substantial economic benefits from the products our inventors have developed. The research outlined above will allow US companies to sustain Moore's Law for at least another decade, will provide them with the technologies required to create innovative, reliable, products that will bolster our economy, and will create jobs in the US technology market, helping to fight the trend towards outsourcing the design and manufacture of electronics overseas.

Learn More `http://www.relxlayer.org`

Acknowledgment This material is based upon work supported by the National Science Foundation (NSF) under Grant 0637190 to the Computing Research Association (CRA) for the Computing Community Consortium (CCC). The opinions, findings, conclusions, and recommendations expressed in this material are those of the study leaders and do not necessarily reflect the views of the NSF, CRA, CCC, or the employers of the study leadership or participants.