

What to Do About the End of Moore's Law, Probably!

Krishna Palem
NTU-Rice Institute of Sustainable and Applied
Infodynamics, Nanyang Technological University
639798, Singapore
Department of CS & ECE, Rice University,
Houston, TX 77005, USA
kvp1@rice.edu

Avinash Lingamneni
Department of ECE, Rice University
Houston, TX 77005, USA
Wireless and Integration Systems, CSEM SA,
2000 Neuchatel, Switzerland
avinash.l@rice.edu

Categories and Subject Descriptors

B.8.0 [Hardware]: Performance and Reliability—General

General Terms

Algorithms, Design, Reliability

Keywords

Co-design, EDA, Energy-Accuracy Tradeoff, Moore's Law, Inexact Circuit Design, Probabilistic CMOS

1. RELIABLE COMPUTING IN THE BEGINNING

Computers process *bits* of information. A bit can take a value of 0 or 1, and computers process these bits through some physical mechanism. In the early days of electronic computers, this was done by electromechanical relays [28] which were soon replaced by vacuum tubes [6]. From the very beginning, these devices and the computers they were used to build were affected by concerns of reliability. For example, in a relatively recent interview with Presper Eckert [1] who co-designed ENIAC, widely believed to be the first electronic computer built, he notes: “we had a tube fail about every two days, and we could locate the problem within 15 minutes.”

The computer pioneer John von Neumann who had worked with Eckert and his colleague John Mauchly clearly understood the importance of reliability [38]. It is worth spending some time reviewing this relatively old history and issues from this period, both as a curiosity, but notably since there are important lessons to be learnt from his historically significant lectures delivered at Caltech [38]. Back in 1952, he notes, “*The subject matter, as the title suggests is error in logics, or in the physical implementation of logics—in automata synthesis. Error is viewed therefore, not as an extraneous and misdirected or a misdirecting accident, but as an essential part of the process under consideration . . .*” von Neumann is therefore concerned with errors that occur *intrinsically* in implementations of logics, and today, this would mean VLSI and ULSI circuits built out of CMOS transistors. Surprisingly, sixty years after von Neumann's lectures, we are again grappling with the same issue in the modern

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC 2012, June 3-7, 2012, San Francisco, California, USA.

Copyright 2012 ACM 978-1-4503-1199-1/12/06 ...\$10.00.

era, since many now claim that the laws of physics dictating the exponentially improving benefits of *Moore's Law* will end in the next 10 to 20 years. So, after six decades from the time von Neumann gave his lectures, we have returned to the same concerns again.

2. A WALK THROUGH HISTORY STARTING WITH VON NEUMANN'S SEMINAL LECTURES

More than half a century ago, as each new generation of technology emerged, the issue of failures and reliability invariably reared its ugly head. We believe that the awareness by working with engineers who built the early vacuum tube based computers motivated von Neumann and his collaborators to use *probability* to model error or failures, through abstract models of “hardware”. Since these models were based on Turing's now classical paper and were presented in the McCulloch-Pitts style, they had a cybernetic flavor, but nevertheless captured the essence of a state machine or automaton widely used to capture hardware behaviors today. We must remember that modern automata theory was nascent at this time and so, some of von Neumann's constructions might seem cumbersome. Nevertheless, we think the insights were incisive and conclusive. For example, the model in his lectures is essentially an abstract form of a modern computer, represented by a combinational logic component, and a communication component corresponding to wires or interconnect including a mechanism for encoding the state of a computation. In this model and with an element of error introduced into the elements, he considers ways for realizing, in modern terms, correct computational systems given that individual elements are vulnerable to failure.

Four years later, Moore and Shannon [26] in a sense reworked and extended the work presented in these lectures by introducing a model that is based on *switches*, which are the ubiquitous building blocks of computing systems even today. The particular model they used is based on Shannon's celebrated masters thesis, *A Symbolic Analysis of Relay and Switching Circuits*, from 1937. In this thesis, Shannon used the switch to abstractly represent an electro-mechanical relay to be used as a basis for building digital circuits. Once again, probabilities were used to model the correct or incorrect functioning of a switch. Their focus was on the important question of correcting errors introduced by potentially faulty switches and determining the cost of such correction in realizing digital circuits.

Since these early concerns, close to five decades passed with the spectacular success of the transistor, their integration at a very large scale leading to VLSI, the revolutionary invention of the microprocessor and the historically unprecedented march of ever-decreasing transistor feature sizes prophesied by Gordon Moore [27]. As physicist Gell-Mann notes in a recent interview [7] where he also mentions his role

with K. A. Bruckner in the work leading to von Neumann's celebrated lecture, concerns of reliability became insignificant with the remarkable reliability achieved with transistors replacing vacuum tubes as switches. However, by the 1990s, the call for approaches to help sustain Moore's law as CMOS transistor feature sizes approached the nanoscale dimensions, were becoming increasingly strident. Scholarly articles started appearing with daunting titles such as "*End of Moore's law: thermal (noise) death of integration in micro and nano electronics*" [16]. So, in keeping with what seems to be history's penchant for repeating itself, the need for realizing reliable computing architectures from unreliable switches resurfaced again close to five decades after von Neumann's lectures, this time in the modern context of CMOS transistor based switching devices [14].

3. CROSSING OVER TO THE DARK SIDE BY USING UNRELIABLE ELEMENTS UNRELIABLY

Around 2002, one of us asked the following question: what if we consider building unreliable circuits and computing blocks from unreliable elements, rather than striving to build reliable switches, circuits and computing hardware from potentially unreliable components? We could potentially have a richer domain of switches to draw upon and therefore be much less constrained, than having to strive for reliability all the time. Fifty years after von Neumann's lectures, this idea was shown to be viable using a variety of mathematical models [31] including a *random access machine* and those representing circuits as switches whose probabilities determine error [32]. Of course, why would anyone want to build unreliable circuits and computing elements from them, that do not compute correct answers? While this goal might seem counterintuitive, it turns out that there is an entirely different resource associated with the physical implementations of switches, that their probabilistic and erroneous behaviors are intimately tied to. This has to do with the *energy* consumed by the physical implementation of a switch. Notably, the work from [32] characterized, and as far as we can determine for the first time, the amount of energy savings in the physical implementation of switches as we vary the correctness or error.

Hand in hand with the concerns about being able to sustain Moore's law, by the late 1990s, there was a rapidly increasing concern, if not alarm, about the amount of energy consumed by computing systems. Therefore, implementing probabilistic switches designed to be erroneous as a basis for realizing energy savings and building computing systems from them *for doing useful work*, seemed very attractive. One of us became a "heretic" [12] as a result of starting an active project and building a group to work on this idea, starting in 2001-2002 time period. This effort received significant stimulation with support from DARPA through a seedling grant in 2002. A summary of our work over the next decade caricatured against the historical legacy of working with erroneous computing elements is the primary subject of this paper.

3.1 The Energy-error Relationship from a Thermodynamic Perspective

That error and the energy consumed by a switch are related is not entirely a surprising fact. von Neumann in fact came close to this connection when he remarks in his lecture that "Our present treatment of error is unsatisfactory and ad-hoc. It is the author's conviction, voiced over many years, that error should be treated by thermodynamical methods and be the subject of a thermodynamical theory, as information has been, by the work of L. Szilard and C. E. Shannon (Cf. 5.2). [38]" This seemingly passing comment is actually

central to the subject matter of this paper, where we are concerned with energy efficiencies, realized through using unreliable switches.

To understand its full import, a short digression first. We note that physical implementations of switches using electrical means, whether they are built using vacuum tubes or transistors, consume energy when they perform the switching function. Thus, they are governed by the laws of classical physics in general, and thermodynamics in particular. By the time von Neumann delivered his lectures, classical thermodynamics had a clear statistical foundation and interpretation following the seminal work of Maxwell, Boltzmann and Gibbs [2, 8]. In particular, Szilard's work [36] that von Neumann refers to is an important part of this development. One of the most important debates in classical physics which lasted over sixty years has to do with the celebrated second law of thermodynamics. Through a very clever construction, Szilard created a single object which has since come to be known as Szilard's engine, for trying to analyze and understand the validity of this law. Loosely speaking, Szilard's engine is a physical structure which we can think of as a cylinder delineated into two halves which are separated by a (weightless) trap door. For simplicity, we can think of a single molecule (of some gas) in motion in the cylinder. An external agent can, by raising and lowering the trapdoor, trap the molecule in either half of the cylinder. This agent also has the power to observe and *record* the half of the cylinder in which the molecule is trapped.

While it is not widely known, this construct had and continues to have a powerful influence on the way we reason about information, and its relationship to physical implementations, especially as they relate to issues of energy consumption. It is the first device that we know of which can be in one of two *states*, and which can be *switched* to induce a change of state by raising and lowering the trapdoor. It is therefore, the earliest instance of an abstract construction, as well as a physical implementation, of a bistable switch. In this sense, it is also the first known link between a physical object and a method for producing abstract information through switching, which in this instance, is recorded through the actions of the external agent. The record can be thought of as a precursor of our modern digital memory in technological terms. Thus, each act of raising and lowering the trapdoor represents a switching step and produces a bit of *information* which is recorded, encoding the *state* of the switch. With this interpretation, Szilard's construct can be viewed as an engine which, through a switching step, produces a single bit of information. Since the state of this device is determined by the location of the gas molecule, the laws of statistical thermodynamics naturally apply. Therefore, it provides a very natural and intrinsic statistical basis for relating the energy consumed to the information being produced or *computed* through switches.

Returning to von Neumann's comment, the "unsatisfactory" aspect, we believe, has to do with the fact that in the approach he described, the probability of error is modeled synthetically in a manner that is extrinsic to the implementation. The probability of a switch performing its activity correctly is simply a numerical value associated with it as a number which is not derived from the inherent, and therefore intrinsic, physical implementation of a switch. We believe this was intentional of course since it allows reasoning in purely abstract terms without being encumbered by the physical details as exemplified and exploited by Moore and Shannon subsequently. As a result, the relationship between the probabilistic error as it varies, and the associated thermodynamic or energy cost, was not a part of the development. An exception of course is the remarkable body of work leading to Landauer's [19] historic insight on the minimum amount of energy needed to perform the act of recording the bit of information by the agent, during a sin-

gle switching step of Szilard’s engine. However, to use our terminology, it is very important to observe that Landauer is concerned about producing a bit correctly and does not concern himself with the energy associated with probabilistic switches with varying probabilities of correctness. So, mathematical models of probabilistically correct switching which do not have a relationship with energy consumption is what we had on the one hand. On the other hand, through Landauer’s work, we have a connection between the energy cost as it relates to a switching step for producing one bit of information and recording it correctly without error.

3.2 Connecting Switches Back to the Physical Reality

What if were to go back to Szilard’s roots and rather than his focus on understanding the profound nature of the second law of thermodynamics, instead look at his engine as a technical and perhaps even as a technological construct. By doing this, we could extend his engine to a model of a switch which can be switching correctly with some probability of correctness p and relate this probability very naturally to the associated energy consumption. As a result, we now have a switching device which can produce a *bit* of information through some physical medium wherein, the probability of it being correct is the parameter p , and the greater its value, the greater the energy consumed. The first attribute of our switch is the probability of correctness p associated with each bit of information being produced by it, whereas the second attribute is the associated energy cost.

In 2003, this was done through a probabilistic switch that captures these two attributes simultaneously, and was the subject of our work referred to earlier [32]. In spirit and philosophy, it went against the direction of abstracting away the physical attributes through clean models as von Neumann and Moore-Shannon set out to do, and connected probabilistic error during switching back to the physics. In our experience, this formulation of a probabilistic energy aware switch has proven to be a very useful foundational construct to understand and reason about *potentially unreliable* and energy-efficient hardware.

However, in order to really use this idea, CMOS implementations had to be considered. By 2004, we could show that probabilistically correct CMOS switches, which have since been referred to as PCMOS (switches), did exhibit this trade-off (summarized in [15]. In fact, the trade-off between p and the associated energy was modeled mathematically and also measured physically. It turned out to be much more favorable than we had anticipated in the following sense. As the probability of a PCMOS switch being correct approached 1, the energy consumed grew extremely rapidly. In other words, a small decrease in the probability of correctness (below 1) can be traded for a significant drop in the energy consumed, illustrated in Figure 1 during one switching step involving an XOR gate. For example, the point A in that figure has a p of 0.988 with a corresponding energy consumption of 20 fJ while point B requires only 7.5 fJ with a p of 0.975 per switching step indicating a $\sim 3X$ reduction in energy consumption for a 2.3% decrease in the probability of correctness. Over the course of the next two years, we fabricated working devices and demonstrated that indeed, measured behaviors matching those predicted by the models and simulations could be realized [17].

4. THE LOGIC OF HARDWARE DESIGN AND A PROBABILISTIC BOOLEAN LOGIC

The behavior of probabilistic switches and the way they can be combined to design electrical circuits is based on Boole’s strikingly contemporary formulation of logic [3] going back to the nineteenth century! However, if we set out

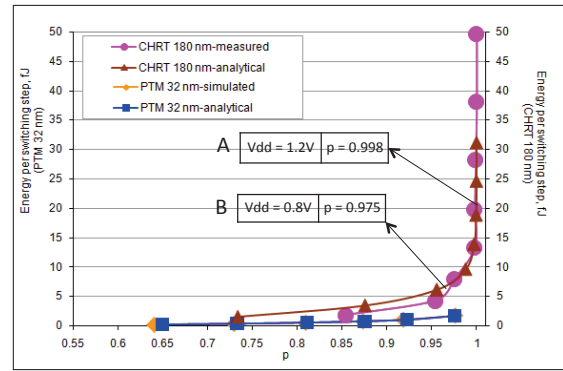


Figure 1: Measured and modeled energy-probability of correctness relationship for an XOR gate

Input			Truth Value
x	y	z	
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	1
1	0	0	0
1	0	1	1
1	1	0	1
1	1	1	1

(a)

Input			Probabilities	
x	y	z	Truth Value = 1	Truth Value = 0
0	0	0	¼	¾
0	0	1	¼	¾
0	1	0	¼	¾
0	1	1	¾	¼
1	0	0	¼	¾
1	0	1	1	0
1	1	0	1	0
1	1	1	1	0

(b)

Figure 2: (a) Conventional truth table for the carry logic function of a full adder $((x \wedge y) \vee (x \wedge z)) \vee (y \wedge z)$ built from reliable hardware; (b) A probabilistic Boolean truth table for $((x \wedge_1 y) \vee_1 (x \wedge_1 z)) \vee_1 (y \wedge_{3/4} z)$ for the full adder carry bit built from unreliable hardware [4].

to design hardware that is erroneous, Boolean logic has to be extended to admit incorrect outcomes. Specifically, we needed to develop a framework for modeling erroneous gates and furthermore, *extensions of Boole’s rules* for combining their outputs.

In our *probabilistic boolean logic* (PBL) [4] that we published in 2008, the operators have the probability parameter p tied to them inextricably as shown in Figure 2 (b). Thus, each probabilistic Boolean operator, say \wedge_p denoting a probabilistic AND, has an explicitly associated probability of correctness parameter p , which relates it simultaneously to its energy cost. Alternately, we can think of each of our operators \wedge_p , \vee_p and others as having a thermodynamic interpretation, in addition to capturing correctness, through p .

4.1 From Gates and Logic to Applications

The next obvious question was to see how best to use PCMOS hardware in mainstream computing. We considered applications that embodied probabilities naturally such as decision systems and pattern recognition based on bayesian and random neural networks respectively, as well as cryptographic applications. Using a system-on-a-chip type architecture, we could show that significant energy and speed gains were simultaneously possible through using PCMOS hardware [18]. Furthermore, it occurred to us that as we interact with information in our every day lives, we are, increasingly consuming it through our senses. So, the results of com-

puting can be erroneous, albeit *perceptually* acceptable, if not indistinguishable from those produced using correct and less efficient hardware which we did for embedded signal processing. As far as we can tell, this is the first foray of using hardware to design a system that embodies error, and significantly, there is no intention or effort to correct it.

5. FROM PROBABILISTIC TO INEXACT CIRCUITS AND COMPUTING

Thanks to the greatly successful effort that continues to sustain Moore’s law, transistors and VLSI circuits built from them are not very unreliable even today, probabilistically or otherwise. So, we were faced with the question as to how to use existing reliable VLSI technologies to be able to design circuits and architectures, and achieve gains in energy, and possibly in other physical attributes such as area and speed by exploiting this principle we have been developing.

For terminological clarity, let us refer to approaches that seek to garner resource gains by trading accuracy at the hardware level to belong to the *inexact design* class. We have shown the trajectory of inexact design in Figure 3 where we summarize the chronology of work starting with our departure in 2003 into the branch of *inexact* or unreliable yet useful computing. Traditionally, the term *reliable* has always been used to refer to component switches and circuits or systems built by composing them, when they have an associated probability of correctness of 1, or extremely close to one in practical settings; they are considered to be unreliable otherwise. As a result, realizing or synthesizing reliable systems from unreliable switches entails boosting the probability of correctness by compensatory error correction mechanisms. In contrast, as shown in Figure 3, since our work in 2003, the field has split from the legacy of von Neumann’s lecture by explicitly seeking to build *unreliable* yet useful computing systems from unreliable components, *without* compensating through error correction. This departure from traditional approaches to designing reliable systems had led to the emergence of the area of inexact design.

At this point, we will digress from inexact design to identify two very innovative ideas with great potential for utility and impact that appeared between 1995-2001 that aimed to realize *reliable* low-cost DSP primitives. Sometimes, we have found references in literature that indicate a close affinity of these ideas to our work. The first idea of *approximate signal processing* [13] was inspired by the approaches to designing resource-constrained (typically execution time) systems in the areas of Artificial Intelligence and Real-Time Systems. The second involved *algorithmic noise tolerance* [9] techniques which allow circuits to be unreliable but combining techniques from signal processing, correct the resulting errors. Another similar approach that appeared at the same time as the inception of our work is the RAZOR effort from University of Michigan [5]. We observe these approaches are fundamentally different from our work described so far and do not follow our split from the legacy embodied in von Neumann’s lectures as shown in Figure 3 since they eventually realize reliable computations.

Returning to inexact design, several research efforts have been launched by us as well as others to investigate techniques that would help realize inexact circuits starting with *mostly* reliable components of current CMOS technologies. Again, we refer the reader to Figure 3 where we have listed some of the highlights from this chronology along the inexact design branch. Most of the initial approaches to induce erroneous behavior in correctly functioning hardware involved varied forms of voltage overscaling. Voltage overscaled circuits were used as the basis for realizing cost-accuracy trade-offs for a wide variety of applications such as datapath circuits [24, 15] through Biased Voltage Overscaling (BiVOS), Motion Estimation, Discrete Cosine Transform [29], and




Figure 4: (a) Results of the Probabilistic Pruning technique on 64-bit Han-Carlson Adder; (b) Results of Probabilistic Logic Minimization technique on 16-bit array multiplier [23]

Image Processing. More recent and promising efforts applied voltage scaling at the granularity of processor modules. These approaches are popularly referred to as *stochastic computation* [11].

Canonically, voltage overscaling provides a fine-grained approach to enable energy-accuracy trade-offs, but have associated overheads due to a need for level-shifters, metastability resolutions circuits, and the routing of multiple voltage lines. The practical realizations only implement a subset of well characterized voltage levels owing to these overheads, and also due to the inevitable power supply fluctuations. While, these techniques do provide more dynamic control and have the potential for quadratic energy savings, the associated overheads are seldom amortized at fine-grain levels. Thus, more recent techniques have focused on the architecture as well as the logic layer for realizing circuits with *zero hardware overheads* while gleaming significant savings across the energy, delay and area dimensions simultaneously.

Continuing, in this context, we have proposed two novel techniques at the architectural and logic-layers which we have referred to as *probabilistic pruning* [20] and *probabilistic logic minimization* [21]. While the former is used to systematically *prune* or delete components and their associated wires along the paths of the circuit that have a lower significance or a lower probability of being active during circuit operation or both, the latter transforms logic functions to lower cost variants by manipulating the corresponding Boolean function through “bit-flips” (1 → 0 or vice-versa) for some of the input vector combinations that are not significant. Through these approaches, we have been able to demonstrate cumulative savings of a multiplicative factor of 8 in the energy-delay-area product (EDAP) for critical datapath elements such as adders and multipliers, as shown in Figure 4. The proposed pruning technique has been applied to a variety of standard 64-bit adder designs and a prototype chip has been fabricated using TSMC 180nm (low power) technology. A photograph of the 86-pin fabricated chip implementing our pruned adders along with its testing framework is shown in Figure 5 [22].

6. FROM INEXACT CIRCUITS TO DESIGN AUTOMATION

While holding great promise, the ability of the proposed inexact design techniques to influence the broader milieu of computing is limited as they are mostly based on ad-hoc hand designs and did not consider algorithmically well-characterized *automated* design methodologies. Also, existing design approaches were limited to particular layers of abstraction such as physical, architectural and algorithmic

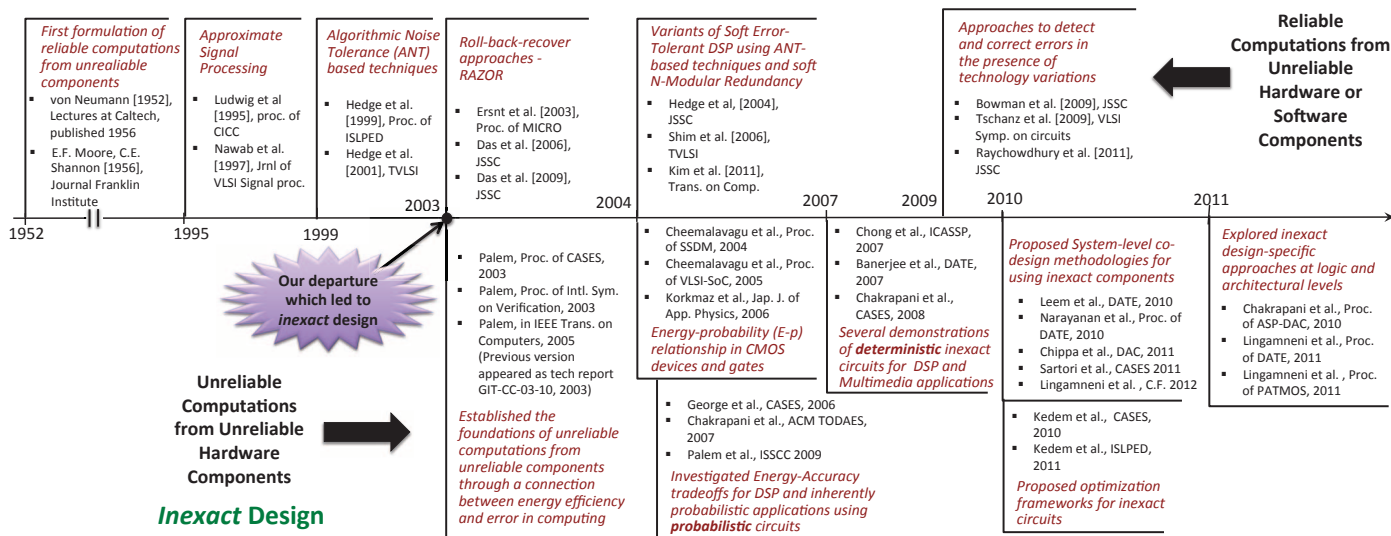


Figure 3: Timeline of important papers and innovations that shaped the domain of *inexact* circuit design




Figure 5: (a) A die photograph of the fabricated prototype chip; (b) The prototype chip integrated into the *icyboard* test platform [22]

or more broadly software. However, it is well-known that significant gains can be achieved by optimizing across the layers. This is in stark contrast to conventional computing systems and hardware design wherein the hardware is expected to operate correctly all the time. In this context, automatic algorithmic methods are used widely across the layers of abstraction to achieve gains, to great effect [25]. To respond to this need, we presented an algorithmically well-founded cross-layer co-design framework (CCF) for automatically designing *inexact* hardware in the form of datapath elements that could significantly reduce the time needed for design space exploration [22]. We believe that this is an important first step in injecting the twin ideas of EDA and co-design into the milieu of inexactness.

7. FUTURE DIRECTIONS AND REMARKS

When we were invited to submit a paper on the genesis of inexact design viewed against the backdrop of probabilistic methods in computing, we considered three possible approaches. The first was to present a survey of the field. An alternative was to discuss our own technical work in some depth. Upon consultation, we concluded that both of these alternatives could be easily done by compiling and perusing existing publications. Therefore, we decided to write this paper by not intending it to serve either of these purposes. However, what seemed to be missing and therefore of some value as a contribution to the scholarship of this emerging field, was in connecting it to its rich legacy spanning the last century. This seemed a worthwhile endeavor, especially given the powerful ideas that served as a basis. In doing

so, we were guided by Longfellow’s inspirational comments from his poem ‘*A Psalm of Life*’ where he eloquently and beautifully characterizes the legacy of great ideas and the influence of people behind them, in that their lives : “...remind us, We can make our lives sublime. And, departing, leave behind us Footprints on the sands of time.” Given the fast and furious pace at which we have been progressing in the CMOS world, we felt that to pause and reflect on the history behind topics in VLSI design and its automation wherein the “probability” is increasingly being used, would be a worthwhile effort. We hope to have done some justice to this goal within the restrictions of space that this paper has afforded us. A more detailed and complete version of this paper is to appear in a forthcoming special issue of the ACM Transactions on Embedded Computing Systems [30], and the reader is referred to this forthcoming paper for a more comprehensive overview and the many rich and exciting prospects for future work. Also, the deeper physical implications of CMOS devices and their switching has been analyzed comprehensively, and with great clarity, by Meindl [10]. The reader is referred to this paper for a deeper understanding of the issues underlying our discussion, with the context of reliable CMOS switches. Finally, probabilistic methods have been a topic of great import and utility in the software and algorithmic domains. Starting with Monte Carlo simulations [37] through the concepts of randomized algorithms [34, 35] to average case analysis [33], the fruitful use of probabilistic methods has deep roots. In contrast, the concepts and ideas that we have discussed in this paper represent a novel foray of probabilistic ideas and the frameworks of inexact computing that they have inspired, in the domain of hardware and its design.

8. ACKNOWLEDGMENTS

Parts of this work were supported in part by the US Defense Advanced Research Projects Agency (DARPA) under seedling contract number F30602-02-2-0124, the DARPA ACIP program under contract FA8650-04-C-7126 through a subcontract from USC-ISI, and an award from the Intel Corporation. We also wish to thank the Moore distinguished faculty fellow program at the California Institute of Technology, the Canon distinguished professorship program of the Nanyang Technological University (NTU) at Singapore, the NTU-Rice Institute for Sustainable and Applied Infodynam-

ics and CSEM SA at Switzerland, which enabled pursuing this work in part. We gratefully acknowledge the contributions of our colleagues who are co-authors in our work which we have surveyed through citations. And last but not least, we thank the organizers of the special session at DAC 2012, Christoph Kirsch and Vincent Mooney, for inviting us to share our perspective and present this paper.

9. REFERENCES

- [1] Alexander Randall 5th. A lost interview with ENIAC co-inventor J. Presper Eckert. Computer World, Retrieved 2011-04-25.
- [2] L. Boltzmann and S. Brush. *Lectures on Gas Theory, English translation by S.G. Brush*. Dover Publications, 1995.
- [3] G. Boole. The mathematical analysis of logic: Being an essay towards a calculus of deductive reasoning. 1847.
- [4] L. N. B. Chakrapani and K. V. Palem. A probabilistic boolean logic and its meaning. Technical Report TR08-05, Rice University, Department of Computer Science, Jun 2008.
- [5] D. Ernst et al. Razor: A low-power pipeline based on circuit-level timing speculation. *in proc. of MICRO*, pages 7–18, Oct. 2003.
- [6] J. A. Fleming. Intrument for converting alternating electric currents into continuous currents. US Patent 803684, Nov 1905.
- [7] M. Gell-Mann. Trying to make a reliable computer out of unreliable parts. <http://www.webofstories.com/play/10585?o=MS>.
- [8] J. Gibbs. *Elementary Principles in Statistical Mechanics*. Scribner, New York, 1902.
- [9] R. Hegde and N. R. Shanbhag. Energy-efficient signal processing via algorithmic noise-tolerance. *In Proc. Int. Symp. on Low Power Electronics and Design*, pages 30–35, 1999.
- [10] J. D. Meindl et al. Nanoelectronics in retrospect, prospect and principle. *in proc. of ISSCC*, pages 31–35, 2010.
- [11] John Sartori et al. Stochastic computing: Embracing errors in architecture and design of processors and applications. *in proc. of CASES*, 2011.
- [12] E. Jonietz. Probabilistic chips. *Technology Review, Published by MIT*, <http://www.technologyreview.com/energy/20246/>, 2008.
- [13] J.T. Ludwig et al. Low power filtering using approximate processing for DSP applications. *in proc. of CICC*, pages 185–188, 1995.
- [14] K. Nikolic et al. Architectures for reliable computing with unreliable nanodevices. *in the proc. of IEEE conf. on Nanotechnology*, pages 254–259, 2001.
- [15] K. V. Palem et al. Sustaining moore’s law in embedded computing through probabilistic and approximate design: retrospects and prospects. *In in proc. of CASES*, pages 1–10, 2009.
- [16] L. B. Kish. End of Moore’s law: Thermal (noise) death of integration in micro and nano electronics. *Physics Letters A*, 305:144–149, 2002.
- [17] P. Korkmaz. *Probabilistic CMOS (PCMOS) in the Nanoelectronics Regime*. PhD thesis, Georgia Institute of Technology, 2007.
- [18] L. N. B. Chakrapani et al. Probabilistic system-on-a-chip architectures. *in ACM Trans. on Design Automation of Elec. Sys*, 12(3):1–28, 2007.
- [19] R. Landauer. Irreversibility and heat generation in the computing process. *IBM J. Research and Development*, 3:183–191, July 1961.
- [20] A. Lingamneni et al. Energy parsimonious circuit design through probabilistic pruning. *in proc. of DATE*, pages 764–769, Mar 2011.
- [21] A. Lingamneni et al. Parsimonious circuit design for error-tolerant applications through probabilistic logic minimization. *in the proc. of the PATMOS*, pages 204–213, 2011.
- [22] A. Lingamneni et al. Algorithmic methodologies for ultra-efficient inexact architectures for sustaining technology scaling. *in the ACM International Conference on Computing Frontiers*, May 2012.
- [23] A. Lingamneni et al. Synthesizing parsimonious inexact circuits through probabilistic design techniques. *in the ACM Trans. on Embedded Computing Systems (spl. issue on Probabilistic Embedded Computing)*, 2012.
- [24] L.N.B. Chakrapani et al. Highly energy and performance efficient embedded computing through approximately correct arithmetic: A mathematical foundation and preliminary experimental validation. *In proc. of IEEE/ACM CASES*, pages 187–196, 2008.
- [25] G. D. Micheli. *Synthesis and Optimization of Digital Circuits*. McGraw-Hill, 1994.
- [26] E. Moore and C. Shannon. Reliable circuits using less reliable relays I. *Journal Franklin Institute*, 262:191–208, Sept 1956.
- [27] G. E. Moore. Cramming more components onto integrated circuits. *Electronics Magazine*, 38(8), 1965.
- [28] S. F. B. Morse. Improvement in the mode of communicating information by signals by the application of electro-magnetism. US Patent 1,647, June 1840.
- [29] N Banerjee et al. Process variation tolerant low power DCT architecture. *In Design, Automation and Test in Europe Conference*, Apr 2007.
- [30] K. Palem and A. Lingamneni. Computing unreliably from unreliable elements: Inexact computing in perspective. *in the ACM Trans. on Embedded Computing Systems (spl. issue on Probabilistic Embedded Computing)*, 2012.
- [31] K. V. Palem. Energy aware algorithm design via probabilistic computing: From algorithms and models to Moore’s law and novel (semiconductor) devices. *In proc. of CASES*, pages 113 – 116, 2003.
- [32] K. V. Palem. Energy aware computing through probabilistic switching: A study of limits. *IEEE Transactions on Computers*, 54(9):1123–1137, 2005; (abridged form appeared as – K.V. Palem, Energy Aware Computing through Randomized Switching, *Technical Report GIT-CC-03-16, Georgia Inst. of Technology*, May 2003).
- [33] R. M. Karp et al. Average case analysis of a heuristic for the assignment problem. *Mathematics of Operations Research*, 19(3):513–522, Aug 1994.
- [34] M. O. Rabin. Probabilistic automata. *Information and Control*, 6:230–245, 1963.
- [35] R. Solovay and V. Strassen. A fast monte-carlo test for primality. *SIAM Journal on Computing*, pages 84–85, 1977.
- [36] L. Szilard. Reduction in entropy of a thermodynamic system caused by the interference of intelligent beings. *Z. Physik*, 53:840–856, 1929.
- [37] S. Ulam, R. D. Richtmyer, and J. von Neumann. Statistical methods in neutron diffusion. *Los Alamos Scientific Laboratory report LAMS-551*, 1947.
- [38] J. Von Neumann. Probabilistic logics and the synthesis of reliable organisms from unreliable components. *In Automata Studies (C.E. Shannon and J. McCarthy eds.)*, Priceton Univ. Press, Princeton, N.J., 1956.