

Realizing Ultra-low Energy Application Specific SoC Architectures through Novel Probabilistic CMOS (PCMOS) Technology¹

Krishna V. Palem, Lakshmi N. Chakrapani, Bilge E. Akgul and Pinar Korkmaz

Center for Research on Embedded Systems and Technology

School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, U.S.A

Phone: +1-404-894-1295 Email: palem@ece.gatech.edu

1. Introduction

Process scaling beyond the sub-micron regime presents several hurdles to sustaining Moore's Law. Specifically, due to challenges posed by device scaling, effects such as deep sub-micron noise are expected to induce "probabilistic" behaviors in future CMOS designs [1,2]. As a promising approach to coping with these challenges, we describe a *probabilistic CMOS (PCMOS)* based computing methodology. In our approach, the devices tolerate noise, leading to inherently probabilistic devices that are guaranteed to compute correctly with a probability p , a design parameter; by design, they are expected to compute incorrectly with a probability $(1-p)$.

It was established earlier that ultra low-energy probabilistic switches can be realized by harnessing noise. In an earlier SSDM publication [3], the relationship between the energy consumed and the probability of correctness was characterized using analytical modeling and HSpice simulations. In this paper, our contributions are threefold:

1. We identify PCMOS characteristics that can be exploited at the application level for energy and performance benefits.
2. We demonstrate a methodology to use PCMOS devices in probabilistic system-on-a-chip (PSoC) architectures.
3. We show that in addition to harnessing noise as a resource, PSoCs also yield significant benefits simultaneously to the performance (measured in terms of the execution time) and the energy consumed (measured in terms of Joules) for the probabilistic cellular automata application [5] that we consider. The application and architecture-level savings are quantified using the product of the energy consumed (measured in Joules) and the performance (measured in number of clock cycles) denoted *energy x performance*: the PSoC based design achieves an impressive saving by a factor of more than 1900 in the context of a PSoC realizing probabilistic cellular automata.

The rest of the paper is organized as follows. Section 2 presents a brief overview of PCMOS technology with an emphasis on PCMOS characteristics that are critical for architectural benefits. Section 3 describes a canonical PSoC architecture which will be our primary vehicle to harness PCMOS in architectural designs that show application-level benefits. Section 4 describes the *probabilistic cellular automata* application and a PSoC for realizing this application. Section 5 presents the results and an analysis. Finally, we conclude the paper and present avenues for future work.

2. Overview of PCMOS and Its Characteristics

As described in detail in [3] and [4] coupling a conventional inverter with noise yields a probabilistic CMOS device which follows the energy-probability (E-p) relationship: The energy consumed per switching is exponentially related to the probability of correctness and is quadratically related to the noise RMS for a fixed probability of correctness. Through this relationship it is evident that a small change in the probability of correctness would yield significant energy benefits per switching due to the exponential relationship. In addition, to harness PCMOS devices to yield application level benefits, the following characteristics are critical and need to be exploited:

1. PCMOS devices are extremely (energy) efficient in producing random bits. When compared to a hardware based implementation of a random number generator, PCMOS devices are *2000* times (energy) efficient and with respect to a software based random number generator on a low-power (StrongARM SA1100) processor, PCMOS is almost *156000* times (energy) efficient.
2. PCMOS devices can be *tuned* a-priori to generate random bits with any probability p . In contrast, CMOS based or software based generation are capable of producing random bits only with a probability $p=0.5$ and therefore, require post processing for cases where $p \neq 0.5$ ---making them further inefficient.
3. However, the current implementation of PCMOS devices operate at a low frequency (of 1Mhz) necessitating replication of PCMOS devices for latency hiding.

Using these PCMOS characteristics listed above, we can harness PCMOS devices to form architectural building blocks to implement applications.

3. Probabilistic System on a Chip (PSoC) Architectures

To harness the benefits for PCMOS technology at the application level, we have developed a methodology for using PCMOS to realize ultra efficient embedded computing platforms. As shown in Figure 1, a canonical PSoC architecture consists of a host processor used to compute most of the control-intensive deterministic components of an application, whereas the co-processor realized using PCMOS will be used as an energy-performance accelerator and will execute the probabilistic part of probabilistic algorithms. In contrast, a conventional CMOS-based SoC is designed to be (functionally) identical to this PSoC by utilizing a hardware based implementation of a well known algorithm to generate random bits [6].

¹ This work was supported in part by DARPA Seedling Contract #F30602-02-2-0124.

The degree to which benefits can be derived through PCMOS technology in the context of an application is dependent on its “probabilistic content”---if an application has high probabilistic content, there will be more frequent activations of PCMOS-based components and less dependence on conventional CMOS components. Thus PSoCs yield high benefits to probabilistic algorithms and applications based on such algorithms. For concreteness, we will consider the probabilistic cellular automata algorithm [5].




Figure 1: A Canonical Probabilistic System on a Chip Architecture

4. Probabilistic Cellular Automata

Probabilistic cellular automata (PCA) are a class of cellular automata used to model stochastic processes and used in a wide range of applications like pattern and string matching, pattern generation and classification. Cellular automata consist of cells with local (typically nearest neighbor) communication. Each cell is associated with a state and a simple transition rule which specifies the next state of a state transition based on its current state and the states of its neighbors. In the probabilistic string classification algorithm [5], the state of each cell is either 0 or 1 , giving rise to 8 possible transition rules (each rule has two possible outcomes 0 or 1). In addition, the transition rules are probabilistic: For a transition rule τ_i ($0 \leq i \leq 7$) probability that the output state of the rule is 0 is denoted by $p_{i,0}$ and probability that the output state is 1 is denoted by $p_{i,1}$.

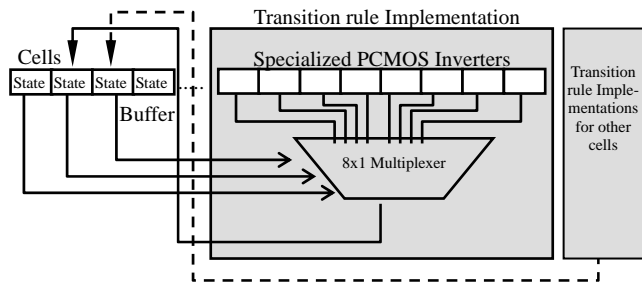


Figure 2: A PSoC implementation for PCA algorithm--Co-processor is shaded

The PSoC which implements this application is designed as shown in Figure 2: Each transition rule is implemented by a PCMOS inverter whose input is a 0 . The probability of correctness associated with the i^{th} inverter is $p_{i,1}$. The control-intensive part of choosing the transition rule (based on the state of a cell and the states of its neighbors) and updating the states is implemented using conventional CMOS devices. Since the rate at which the transition rules are evaluated exceeds the frequency of operation of the PCMOS devices, this structure is replicated many times. Since the probability associated with each of the transition rule is not 0.5 ($p \neq 0.5$), the conventional SoC equivalent architecture is implemented using a hardware based realization of the Park-Miller [6] algorithm coupled

with a comparator to yield the desired probability of random bits. The energy and performance modeling of the PSoC and SoC implementations is via HSpice simulations using a TSMC 0.25μ process.

5. PCMOS Benefits and Analysis

To study the gains of a PSoC implementation, we consider *energy x performance* (EPP) as the metric of interest. The gain of the PSoC over SoC is defined as the ratio of the EPP of the SoC design to the EPP of PSoC design: $\text{Gain} = \text{EPP}_{\text{SoC}} / \text{EPP}_{\text{PSoC}}$. The Gain of the PSoC design over a SoC design is $2900X$ in the context of a *core primitive probabilistic operation* of the PCA application (see Table 1) and it is $1900X$ in the context of the *overall execution* of the string classification application, which validates and even exceeds our claim about the utility of PCMOS devices [3,4].

Algorithm	Application	Primitive Operation	Gain over CMOS	Gain over Software
PCA	String Classification	Evaluating the probabilistic transition function	2.9×10^3	1.56×10^5

Table 1: EPP Gain of PCMOS over CMOS and over conventional software based implementation running on StrongArm SA-1100 processor to execute the primitive probabilistic operation of PCA.

6. Conclusion

PCMOS has been featured as a significant technological innovation in realizing ultra low-energy computing [7]. In this paper, we have shown that the device level benefits of PCMOS can be exploited at the application level through application specific probabilistic SoC architectures. Such implementations yield three orders of magnitude savings in terms of the energy-performance product (EPP) metric in the context of probabilistic cellular automata. For future work, we are exploring a larger suite of applications from the image/signal processing domain where PCMOS can be used to explore the trade-off between energy consumption and signal-to-noise-ratio (error rate) at the application level.

References

- [1] K. Shepard, “Design methodologies for noise in digital ICs,” *Design Automation Conference*, pp. 94-96, 1998.
- [2] L. Kish, “End of Moore’s law: thermal (noise) death of integration in micro and nano electronics,” *Physics Letters A*, vol. 305, pp. 144-149, Dec. 2002.
- [3] S.Cheemalavagu, P. Korkmaz and K. Palem, “Ultra low-energy computing via probabilistic algorithms and devices: CMOS device primitives and the energy-probability relationship,” *SSDM 2004*, pp. 402-403.
- [4] S.Cheemalavagu, P. Korkmaz, K. Palem, B. Akgul and L. Chakrapani, “A probabilistic CMOS switch and its realization by exploiting noise,” *to appear in IFIP-VLSI SoC*, October 2005.
- [5] H. Fuks, “Non-deterministic density classification with diffusive probabilistic cellular automata,” *Physical Review E, Statistical, Nonlinear, and Soft Matter Physics*, 2002.
- [6] S. Park and K. Miller, “Random number generators: good ones are hard to find,” *Communications of the ACM*, Oct. 1988.
- [7] O. Port, “Chips that thrive on uncertainty,” *Business Week*, February 4, 2005.