

A Probabilistic CMOS Switch and its Realization by Exploiting Noise

Suresh Cheemalavagu, Pinar Korkmaz, Krishna V. Palem, Bilge E. S. Akgul and Lakshmi N. Chakrapani
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, Georgia 30332–0250, USA

Abstract

By viewing noise as a resource rather than as an impediment, we demonstrate an entirely novel approach to ultra low-energy computing. The subject of this study is the probabilistic inverter, ubiquitous to the design of digital systems, whose behavior is rendered probabilistic by noise. Summarized through the concept of an energy-probability relationship for inverters based on AMI 0.5 μ m and TSMC 0.25 μ m processes, we quantitatively show that significant energy savings are possible when a probabilistic inverter is switched with probability $1/2 < p < 1$, and that these savings increase exponentially as p is lowered. We also quantitatively show that for a fixed p , increasing the noise RMS has the effect of increasing energy dissipation quadratically. Collectively, we refer to these two facts as the energy-probability laws governing probabilistic CMOS switches—these laws constitute the first contribution of this work. Furthermore, we also present a practical realization of a probabilistic inverter in a readily available TSMC 0.25 μ m technology. Finally, by using the probabilistic inverter as a building block, we provide early evidence that probabilistic switches can yield significant improvements to the energy \times performance metric at the application level, by a factor of more than 288, for a probabilistic neural network application.

1. Introduction and Background

As enormous gains are made in faster and smaller processors largely driven by Moore’s law, significant challenges have emerged, especially as CMOS devices approach the deep sub-micron “nano-scale”. Two of these challenges of particular significance are the impact of noise (see Shepard [21], Kish [6], Hernandez et. al [10, 11]), as well as lowering the energy consumption (see Mamun and Katti [8] and Meindl [9]). In the current approaches to overcome these twin challenges encountered in the semiconductor roadmap, noise is typically viewed as an impediment to scaling (see Kish [6], Natori and Sano [14]). *In this work and in a dramatic shift from known approaches, we show that viewed as a resource rather than as an impediment, noise can be a basis for achieving significant energy savings thus yielding an entirely new approach towards realizing ultra low-energy circuits and concomitant computing platforms.* Our ap-

proach, and the subject of this paper, is centered around the concept of a *probabilistic computing* switch or inverter, whose output is deliberately guaranteed to be correct only with a probability p , $1/2 < p < 1$ (where p is considered to be unity in the context of all conventional computing switches, in that the device is deemed to compute correctly).

In this paper, we provide a systematic and comprehensive characterization of a probabilistic CMOS switch, with an emphasis on its value as a building-block for applications with probabilistic workloads. In particular, we characterize the energy consumed per switching step and the associated probability of correctness p —referred to as the *energy-probability* (or *E-p* for short) *relationship*. Thus, by building on the mathematical foundations of probabilistic switching in the context of energy savings developed by Palem [15, 17] as well as a preliminary proof-of-concept within the context of CMOS based switches presented by Cheemalavagu, Korkmaz and Palem [1], we first show that the energy consumed in producing a single bit of information in Joules is exponentially related to p across CMOS technology generations—in our case, the AMI 0.5 μ m and TSMC 0.25 μ m processes. *First*, the energy per-switching step increases exponentially with p —the higher the probability with which the computed bit is guaranteed to be correct, the higher the energy consumed to compute it. This relationship allows us to trade a relatively small probability of error for significant savings in energy (and heat) dissipation of probabilistic applications. *Second*, we establish that for a fixed probability of correctness, energy is quadratically related to the amount of noise, quantified as the noise RMS value in this paper—the higher the noise amount exposed on the device, the higher the value of E . In Section 4.2, we characterize these two facts as the *energy-probability* law and the *energy-noise* law for probabilistic CMOS devices.

Conventional CMOS inverters (or switches) coupled with noise to induce the probabilistic behavior will be the subject of our study. The particular noise sources considered here include naturally occurring thermal noise sources as studied by Sano [20], as well as power-supply noise sources—as studied by Heydari and Pedram [4] and by Pant et. al [18]—increasingly common in Giga- and Tera-scale ICs based on deep sub-micron technologies. Given the comprehensive nature of this study, both scenarios wherein the thermal noise source is coupled to the

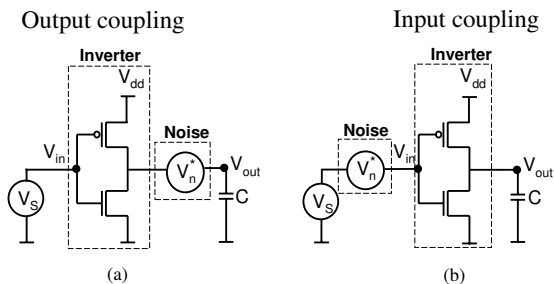


Figure 1. The two ways in which noise is coupled to a deterministic switch yielding its probabilistic variant: (a) noise is coupled to the output of the inverter, (b) noise is coupled to the input of the inverter.

input of the (inverter) switch as well as to the output have been characterized in this paper (see Figure 1) through the energy-probability and energy-noise laws.

With this as background, we provide a summary of related work in Section 2. In Section 3, we define the probabilistic switch and describe its CMOS implementation. In Section 4, we outline an analytical model that helps explain the empirical findings describing the two energy-probability laws. We quantify the application level benefits of probabilistic inverters or switches in terms of energy and performance in Section 6, and finally, we conclude the paper with remarks in Section 7.

2. Related Work

Low energy design has been an area of increasing interest especially since energy consumption has become a limiting factor in sustaining Moore’s law [5, 9, 13, 19]. Palem [15, 16, 17] showed for the first time that noise could be used for energy savings if it is harnessed as a source of randomness in realizing implementations of probabilistic algorithms. In all of this work [15, 16, 17], the energy characteristics are modeled in the framework of thermodynamics [7]. This new interpretation of noise as a utility was very different from the traditional interpretation where noise in the electrical domain is invariably viewed as an impediment to successful design [3, 6, 14, 24]. To leverage the results obtained in the framework of thermodynamics, and to prove the validity of this novel view of noise, the behavior of probabilistic devices realized out of CMOS technology was characterized by Cheemalavagu, Korkmaz and Palem [1] through the E - p relationship. In this paper, we extend this early characterization significantly by (i) providing a novel and explicit construction of a probabilistic inverter, (ii) validating it in AMI $0.5\mu\text{m}$ and TSMC $0.25\mu\text{m}$ technologies by characterizing the energy-probability laws governing its behavior and reconciling using an analytical model (from [1]), and (iii) demonstrating the value of these probabilistic inverters in realizing ultra-low energy implementation of a probabilistic neural network.

3. Probabilistic Inverter

In this section, we first give the definition of a probabilistic switch. Then, we describe a probabilistic switch realization using a CMOS inverter.

3.1. Definition of a Probabilistic Switch

A switch as a digital device with one input and one output. The switch can be either deterministic or probabilistic. The output of the switch is a function, say f , of its input. The act of switching is defined as the invocation of this function f , which corresponds to the event when the output is computed in a typical digital circuit. Switching takes a finite amount of time T . If we denote the output of the switch by $Y(t)$ and the input of the switch by $X(t)$ where t denotes a point in time, then for a deterministic switch $Y(t_2) = f(X(t_1))$, wherein $f : \{0, 1\} \rightarrow \{0, 1\}$ is a function of a single bit, t_2 is the time when the switching ends and t_1 is the time when the switching starts; thus, $T = (t_2 - t_1)$. For a probabilistic switch, on the other hand,

$$Y(t_2) = \begin{cases} f(X(t_1)) & \text{with probability } p \\ \overline{f(X(t_1))} & \text{with probability } (1-p) \end{cases}$$

where $\overline{f(X)}$ denotes the logical complement of the Boolean result of function $f(X)$.

3.2. CMOS Inverter Realization of a Probabilistic Switch

A CMOS inverter is a digital gate that realizes the inversion function with one input and one output. In the context of the switch definition above, for a deterministic inverter, $Y(t_2) = \overline{X(t_1)}$. For a probabilistic inverter, on the other hand,

$$Y(t_2) = \begin{cases} \overline{X(t_1)} & \text{with probability } p \\ X(t_1) & \text{with probability } (1-p) \end{cases}$$

To illustrate this concept, we show the input and output (waveforms) of a deterministic inverter in Figure 2(a) and Figure 2(b), and the corresponding output (waveform) of a probabilistic inverter (with $p = 0.87$) in Figure 2(c), respectively. Note that the input/output voltage levels for a deterministic inverter is higher (1V) than that for a probabilistic inverter (0.8V). This is because, the probabilistic behavior for the inverter is realized through varying two parameters (1) the noise amount coupled on the inverter characterized as its RMS value in Volts, and (2) its supply voltage V_{dd} , and in Figure 2(c), the supply voltage value of 0.8V corresponds to a probability value $p = 0.87$ with a noise RMS value of 0.4V. The details of the effects of the two parameters: the amount of noise, and the supply voltage, will constitute the two laws and are the subject of the following sections. Briefly, since the probability p results due to noise destabilizing the inverter as shown in Figure 1, the probability parameter p is decreased either by increasing the noise (RMS) magnitude, or by decreasing the operating supply voltage of the inverter, V_{dd} . As a result, incorrect switchings occur at the output of the inverter as shown in Figure 2(c).

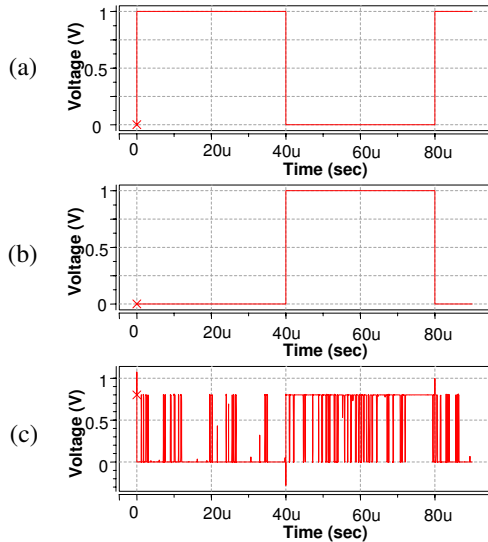


Figure 2. (a) Input signal of the inverter, (b) corresponding output signal of a deterministic inverter, and (c) corresponding output signal of a probabilistic inverter with $p = 0.87$.

4. Characterizing the E - p Relationship of a CMOS Inverter

Our experimental findings characterize the laws governing a probabilistic inverter caused by the effects of (i) noise coupling in the case of thermal noise, (ii) the power supply noise, and (iii) across technology generations. In Sections 4.1 to 4.4, we present our simulation framework and results that characterize the two laws.

4.1. Simulation Framework and Experimental Methodology

As pointed out in Section 1, an inverter is coupled with noise (Figure 1). We consider input-coupled noise as well as output-coupled noise, both of which are thermal, as well as power supply noise coupled to the supply node of the inverter. In modeling thermal noise and power supply noise, we follow the approach of Stein [24] and Pant [18], respectively, where the noise source is assumed to be a random process characterized by a Gaussian distribution with a standard deviation σ , referred to as root-mean-square (RMS) value.

In our simulations, noise is injected into an HSpice netlist in the form of a PWL (*piecewise linear*) voltage source. The data points of the PWL source are generated from a Gaussian distribution of random numbers generated by Matlab. The probability p is determined by sampling the output node, and it is computed as follows

$$p = \frac{\text{Number of samples with correct values}}{\text{Total number of samples}}$$

wherein the total number of samples is 14,000. Simulations were performed for AMI $0.5\mu\text{m}$ and TSMC $0.25\mu\text{m}$ processes. Table 1 summarizes the simulation parameters. As seen in the table, the two parameters being altered are V_{dd} and σ . C denotes the load capacitance corresponding to a fanout of four (which is a typical load present in digital circuits), leading to a load capacitance value of 60fF and 28fF for AMI $0.5\mu\text{m}$ and TSMC $0.25\mu\text{m}$, respectively.

	AMI $0.5\mu\text{m}$	TSMC $0.25\mu\text{m}$
V_{dd}	0.8V – 5V	0.5V – 2.5V
σ	0.2V – 0.8V	0.2V – 0.8V
C	60fF	28fF

Table 1. Simulation parameters




Figure 3. The simulation results when noise is input-coupled and output-coupled for TSMC $0.25\mu\text{m}$. Also shown is the analytical results for the same technology.

To reiterate, in our experimental framework, we vary the magnitude of the signal, namely, the supply voltage, V_{dd} , and the RMS value of the noise (shown in Table 1).

4.2. The E - p Relationship of a CMOS Inverter

As stated earlier (in Section 1 and illustrated in Figure 1), we consider two types of (thermal) noise coupling: the input-coupled noise and the output-coupled noise. This is because, the noise—which can appear in the form of interconnect noise, crosstalk, simultaneous switching noise, thermal noise, etc—could be aggregated at the output of a device as well as the input of a device as noted by Motchenbacher [12] and by Shepard [22]. The simulation results of both the input- and the output-coupled thermal noise cases are shown in Figure 3. The impact of both the output-coupled noise and the input-coupled noise cases follow each other very closely.

As validated in Section 4.3 below, the main conclusions that we drive from the observed trends in Figure 3 are as follows.

1. *The first law*—For any fixed technology generation (feature size), E consumed by a probabilistic inverter increases *exponentially* with p whenever the noise magnitude remains constant. For example, for a noise RMS value of 0.4V, the energy consumed by a probabilistic inverter (designed in TSMC $0.25\mu\text{m}$) rises from 20fJ to 32fJ, in going from a probability value of $p = 0.9$ to a (slightly) higher value of $p = 0.95$.
2. *The second law*—For any fixed technology generation (feature size), E consumed by a probabilistic inverter increases *quadratically* with noise magnitude (RMS), whenever p remains constant. For example, for $p = 0.9$, the energy consumed by a probabilistic inverter (designed in TSMC $0.25\mu\text{m}$) rises from 20fJ to 81fJ, by doubling the noise RMS value from $\sigma = 0.4\text{V}$ to $\sigma = 0.8\text{V}$.

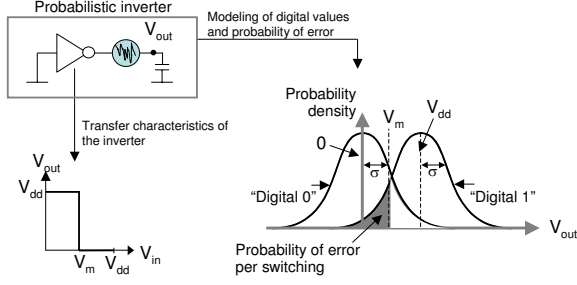


Figure 4. The digital values 0 and 1 represented by two Gaussian curves (shaded area equals to probability of error per switching, i.e., to $(1-p)$).

4.3. Explaining the Behavior of Probabilistic Inverter Analytically

Here, we sketch an analytical model to two energy-probability laws governing a probabilistic inverter for completeness. We will then use this analytical model to explain our HSpice simulation-based characterizations in Section 4.2 above.

Figure 4 illustrates the output voltage (V_{out}) of an inverter that is coupled with noise at its output (see Stein [24]). The probability of being correct can be computed as in Equation (1) and $V_m = V_{dd}/2$ (for a symmetric ideal inverter) yields Equation (2). Next, using $E = \frac{1}{2}CV_{dd}^2$ that gives the energy per switching step, the relationship of energy to probability can be characterized as shown in Equation (3), where C is the load capacitance of the inverter, σ is the RMS value of noise, and erf is the well known error function [25].

$$p = \frac{1}{2} + \frac{1}{4} \operatorname{erf}\left(\frac{V_m}{\sqrt{2}\sigma}\right) - \frac{1}{4} \operatorname{erf}\left(\frac{V_m - V_{dd}}{\sqrt{2}\sigma}\right) \quad (1)$$

$V_m = \frac{V_{dd}}{2}$ yields:

$$p = \frac{1}{2} + \frac{1}{2} \operatorname{erf}\left(\frac{V_{dd}}{2\sqrt{2}\sigma}\right) \quad (2)$$

$$E = 4C\sigma^2[\operatorname{erf}^{-1}(2p - 1)]^2 \quad (3)$$

As shown in Equations 2—3, p depends on the supply voltage, V_{dd} and standard deviation, σ referred as RMS value of noise. Moreover, E is exponentially related to p (first law) and quadratically to σ (second law). These expressions constitute the analytical model of a probabilistic CMOS inverter.

4.3.1. Comments on our analytical model As shown in Figure 3, the simulation and analytical results follow each other closely for the TSMC $0.25\mu m$ technologies, with a 7% average difference in E (for a fixed value of p).

The differences between the analytical and simulation results are due to the fact that our analytical model considers only the switching energy, whereas in the simulations, the total energy (which also includes the short circuit and leakage energies) is considered.

As a remark, both the analytical and simulation results show that given a fixed value of noise RMS, for the same




Figure 5. Coupling of the power supply noise to the CMOS (inverter) gate.

amount of change Δp to the probability, the corresponding energy saving ΔE would be higher as we approach $p = 1$. This is because E grows more rapidly with increasing p . For example, for $\Delta p = 0.1$, the energy savings would be $\Delta E = E_p - E_{(p-\Delta p)} = 45\text{fJ}$ for $p = 0.9$, whereas it would be $\Delta E = E_p - E_{(p-\Delta p)} = 21\text{fJ}$ for $p = 0.8$. This trade of between p and E would have implications to energy savings in the context of probabilistic applications and algorithms, that include a probabilistic inverter as a building block in hardware, in achieving the optimum point for energy savings, given a “quality of solution” determined with p .

4.4. Impact of Power Supply Noise

So far, we have considered the thermal noise as the noise source being coupled either to the input or to the output of the inverter. Again, following the stochastic modeling by Pant [18], the power supply noise is characterized by a Gaussian distribution with RMS value of a few hundred millivolts, derived by modeling a chip with a large number of power grids. Therefore, we induce a Gaussian noise source on the power supply node of the CMOS inverter. As illustrated in Figure 5, V_p^* is induced on the inverter in conjunction with the output-coupled (thermal) noise (V_n^*).

The corresponding results derived via simulations for the TSMC $0.25\mu m$ technology are shown in Figure 6. Again, for a fixed value of E , p decreases due to power supply noise. Hence, in the context of E - p relationship, power supply noise increases the effective RMS value of the thermal noise coupled to the inverter, or, equivalently, for a fixed RMS value, it has the effect of decreasing the value of p by 1.7% on the average. Thus, adding power supply noise affects the overall noise RMS and has the same effect as that shown in Figure 3 where the RMS value of thermal noise was altered explicitly. Similar trends are also seen if the thermal noise were to be input-coupled (as opposed to output-coupled). To conclude (Figure 6), both the first and second energy-probability laws are satisfied by reconciling them against the analytical model.

4.5. Impact of Technology Generations on the E - p Relationship

Considering technology scaling and comparing the AMI $0.5\mu m$ process based probabilistic inverter with the TSMC $0.25\mu m$ variant, as shown in Figure 7: to ob-




Figure 6. The simulation and analytical results for the power supply noise when thermal noise RMS value is fixed at 0.8V and the power supply noise RMS value is varied to be 0.3V and 0.1V.




Figure 7. The simulation and analytical results for AMI 0.5μm and TSMC 0.25μm inverters when the noise is output-coupled. Table on the right shows example points picked from the figure for a comparison reasons.

tain the same probability value p with a fixed RMS value of the noise, less energy is consumed as the feature size decreases. For example, for $p = 0.9$, the amount of switching energy consumed by the AMI 0.5μm inverter is 190.6fJ, whereas it is 69.7fJ in case of the TSMC 0.25μm inverter. This difference is mainly due the capacitance scaling between the two technologies. Thus both of the laws (see Section 4.2) are preserved across technology generations while the absolute values scale as anticipated by Moore’s Law.

5. Realizing a Probabilistic Inverter with Limited Available Noise

As mentioned in Section 1, noise is a paramount design challenge in sustaining Moore’s Law, and the resulting unstable or probabilistic behavior will be a natural feature of deep sub-micron and nano-scale technologies. Using such devices as a basis for realizing low energy and fast circuits for executing probabilistic algorithms is the basic motivation for us in studying such probabilistic devices. While this naturally induced probabilistic behavior is inevitable for future technologies, to study the impact of noise and the probabilistic devices in technologies currently available to us involves “approximation” of this behavior. To clarify this notion of an approximation, issue (as shown in Section 4.2) to produce a probabilistic inverter with $p = 0.7$ using a TSMC 0.25μm technology, we need a noise source with an RMS value of 0.8V (see Figure 3). However, the available noise source that we study in this paper has an RMS value of 12mV (the thermal noise

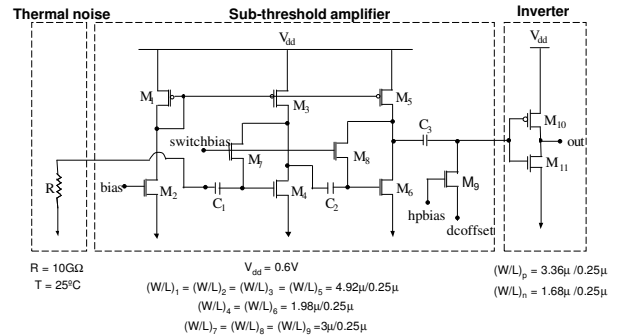


Figure 8. Amplifier circuitry.

of the 10GΩ resistor within the bandwidth of the system shown in Figure 8). Therefore, in order to achieve the approximate value of $p = 0.7$, from the first and the second laws, we also need an amplifier circuitry that raises the RMS value of noise input to the inverter. Figure 8 shows such an amplifier design that we verified using circuit simulations in HSpice also with TSMC 0.25μm technology.

Specifically, our design involves a transconductance amplifier, which is designed to operate in the *sub-threshold* region to minimize the energy consumption due to amplification. Thus we amplify the thermal noise across the resistor R (shown in Figure 8) which serves as a thermal noise source. In order to be able to tune the subthreshold amplifier, we included five bias signals in the design. Referring back to Figure 8, the signals denoted as dcoffset and hpbias control the DC offset at the output of the subthreshold amplifier. Also, the signal denoted as switchbias controls the DC operating points of the transistors M3–M6. Finally, the signal denoted as bias controls the bandwidth and the current consumption of the amplifier. The circuit simulations in HSpice have shown that the DC gain of the sub-threshold amplifier is 45.91dB, while the bandwidth (at the 3dB point) of the amplifier is 1MHz. The energy consumption of the amplifier is 221fJ per switching step, where switching occurs with a period of 1μs and corresponds to the 1MHz constraint bandwidth imposes. Using this integrated amplifier based probabilistic inverter design, we will now demonstrate significant energy as well as performance gains at the application level—our goal in innovating probabilistic switches.

6. Energy and Performance Savings for Probabilistic Applications

We quantify the impact of using a probabilistic inverter to solve the minimum vertex cover problem (which is well-known to be NP-complete), using a *randomized* (or, in our terminology, probabilistic) neural network algorithm following Gelenbe [2]. Such a network includes nodes (neurons) whose potentials are incremented or decremented, respectively, by the incoming positive or negative signals. Based on the magnitude of the signals, the nodes “fire” and this firing is determined by a Poisson process. Note that this algorithm differs from the standard neural network algorithm in the way the mechanisms controlling the firing

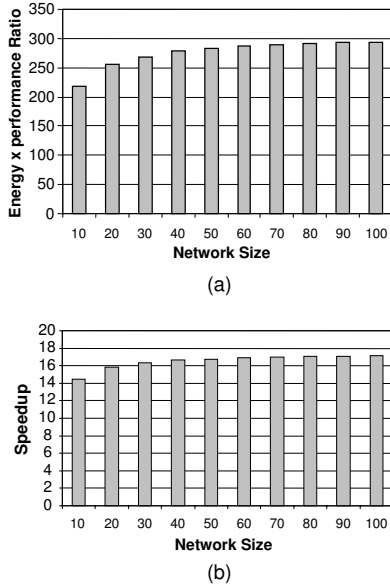


Figure 9. Improvements in (a) energy \times performance ratio and (b) only the performance ratio (speedup).

work: in the context of randomized neural network algorithm that we use, firing decisions are performed probabilistically from a Poisson distribution.

To realize a highly energy and performance (running time) efficient implementations of this algorithm using a probabilistic inverter, we have devised a custom architecture that consists of a set of probabilistic inverters that include the subthreshold amplifier shown in Figure 8 (and hence, the energy and time consumed by it), and a counter. This custom extension is accessed from a low-power embedded microprocessor, the StrongArm SA-1100 as if it were a (probabilistic) co-processor. We performed the energy and performance profiling of the whole application using the Trimaran infrastructure [26] for architecture level modeling, and Jouletrack [23] for energy estimation of the StrongArm SA-1100. Here, the energy is measured in Joules, and the performance, akin to delay in circuits, is measured in cycles. The energy consumption of the probabilistic co-processor is derived from the HSpice simulations. To determine the benefit of using the probabilistic inverter, we compared the energy and performance figures with those for the same application using StrongArm SA-1100—in this case, the pseudo-probabilistic source is implemented in software and thus, the entire application is executed using StrongArm SA-1100 in isolation.

Figure 9(a) shows the improvements achieved through the energy-performance (product) metric for different neural-network sizes. As seen in this figure, a network with 70 nodes for example realized using our probabilistic co-processor and StrongArm SA-1100, can achieve a striking factor of 288 improvement in the energy-performance product, compared to the case wherein the same application is executed on StrongArm SA-1100 only. Figure 9(b) shows the improvement to the performance ratio (speedup) alone, which is 17 for the same network size.

The gains from our approach are due to energy savings obtained by using a low-cost probabilistic inverter producing truly probabilistic bits, instead of using pseudo-randomness via software based technique run on a conventional microprocessor, as well as by being able to complete the execution of the application much faster and hence with a higher performance compared to any deterministic execution. Our approach could be a basis for an entirely novel approach to realizing highly efficient (energy \times performance) realizations of classes of applications: decision making problems using Bayesian networks (e.g. intensive care patient monitoring, remote sensing image retrieval, printer troubleshooting), classification problems using k-nearest neighbors algorithm (e.g., pattern recognition, spoken alphabet recognition), probabilistic routing, simulated annealing, and other applications that inherently admit a probabilistic component.

7. Remarks and Next Steps

In this paper, we have presented an extensive characterization of two basic laws governing the behavior of a probabilistic inverter, based on a study of input- and output-coupled thermal noise, as well as the effects of power-supply noise. These two laws provide a foundation for the design of ultra-low energy and fast realizations in CMOS for probabilistic algorithms. As preliminary evidences we demonstrated these benefits in the context of a probabilistic neural network solving the NP-complete vector cover problem in graphs.

As devices feature sizes scale down into the nano-regime, noise and the concomitant interference become significant, and are increasingly viewed as an impediment to sustaining Moore's law. Here, by using noise as a resource—through probability-energy laws, as far as we can determine, for the first time—we have provided a framework for using the probability of correctness p as an *explicit* design parameter. Furthermore, this explicit characterization encompasses a relationship to energy, which poses yet another serious impediment to CMOS technology scaling.

Essentially, we are able to show that while uncontrolled noise is an impediment to *deterministic designs*—and numerous efforts are underway to combat this issue—controlled and well-understood noise can be an asset in realizing probabilistic designs of value to classes of probabilistic algorithms. However, in currently and readily available technology, noise magnitudes and their behavior might not be sufficient to realizing any particular value of p that an application might need. Thus, if the benefits of probabilistic devices from an energy and performance perspective are to be realized using currently available technology—AMI $0.5\mu\text{m}$ and TSMC $0.25\mu\text{m}$ being our candidates for validating this idea—the noise attribute might have to be amplified. Thus, to facilitate immediate use of our approach, in this paper, we also provide a design and a detailed validation of an integrated structure that includes a (subthreshold) amplifier that is coupled to an inverter. This resulting composite structure is shown, in the context of probabilistic neural networks, to be signifi-

cantly energy and performance efficient—even after paying the penalty for noise amplification—compared to a conventional deterministic realization of these networks on a low energy microprocessor executing without the benefit of probabilistic switches. While using such a composite structure as a “proof” of concept for the utility of probabilistic CMOS and impressive energy savings of such devices, we expect more energy savings in the nano-scale regime where available noise is enough to elicit for probabilistic behavior and the amplifier can be dispensed with. In this context, out energy savings should be viewed as a conservative approach.

We are currently developing the techniques and methods from this paper along two dimensions. First, we are conducting a study to verify the validity of our energy-probability laws in VLSI device feature sizes at a finer-granularity than (coarse-grained) $0.25\mu\text{m}$ technology presented here. Our main reason for choosing this technology as a baseline is the ready availability to fabricate out probabilistic switches using this technology—a process that is underway now. Second, we are actively seeking and extending the suite of probabilistic applications which our probabilistic (CMOS) switching technology can be applied to yield ultra-low energy and high performance execution substrates.

Acknowledgment

This work is supported in part by DARPA under seedling contract #F30602-02-2-0124. We also wish to thank Venkatesh Srinivasan for his help with the amplifier design.

References

- [1] S. Cheemalavagu, P. Korkmaz, and K. V. Palem. Ultra low-energy computing via probabilistic algorithms and devices: CMOS device primitives and the energy-probability relationship. In *Proc. of The 2004 International Conference on Solid State Devices and Materials*, pages 402 – 403, Tokyo, Japan, Sept. 2004.
- [2] E. Gelenbe and F. Batty. Minimum graph covering with the random neural network model. In *Neural Networks: Advances and Applications*, volume 2, 1992.
- [3] R. Hegde and N. R. Shanbhag. Toward achieving energy efficiency in presence of deep submicron noise. *IEEE Trans. VLSI Syst.*, 8:379–391, Aug. 2000.
- [4] P. Heydari and M. Pedram. Analysis of jitter due to power-supply noise in phase-locked loops. In *Proc. of the IEEE Custom Integrated Circuits Conference*, pages 443 – 445, Orlando, Florida, May 2000.
- [5] N. Kim, T. Austin, D. Blaauw, T. Mudge, K. Flautner, J. Hu, M. Irwin, M. Kandemir, and V. Narayanan. Leakage current: Moore’s law meets static power. *IEEE Computer*, 36:65–77, 2003.
- [6] L. B. Kish. End of Moore’s law: thermal (noise) death of integration in micro and nano electronics. *Physics Letters A*, 305:144–149, Dec. 2002.
- [7] H. Leff and A. Rex. *Maxwell’s Demon 2: Entropy, Classical and Quantum Information, Computing*. Institute of Physics, 2003.
- [8] A. Mamun and R. Katti. A new parallel architecture for low power linear feedback shift registers. In *Proceedings of International Symposium on Circuits and Systems (ISCAS)*, pages 333–336, Vancouver, Canada, May 2004.
- [9] J. D. Meindl. Low power microelectronics: retrospect and prospect. *Proc. IEEE*, 83:619–635, Apr. 1995.
- [10] F. Mendoza-Hernandez, M. Linares-Aranda, and V. H. Champac-Vilela. An improved technique to increase noise-tolerance in dynamic digital circuits. In *Proceedings of International Symposium on Circuits and Systems (ISCAS)*, pages 489–492, Vancouver, Canada, May 2004.
- [11] F. Mendoza-Hernandez, M. Linares-Aranda, and V. H. Champac-Vilela. The noise immunity of dynamic digital circuits with technology scaling. In *Proceedings of International Symposium on Circuits and Systems (ISCAS)*, pages 493–496, Vancouver, Canada, May 2004.
- [12] C. D. Motchenbacher and J. A. Connelly. *Low-noise electronic system design*. John Wiley & Sons, Inc., New York, NY, 1993.
- [13] T. Mudge. Power: A first-class architectural design constraint. *IEEE Computer*, 34:52–58, 2001.
- [14] K. Natori and N. Sano. Scaling limit of digital circuits due to thermal noise. *Journal of Applied Physics*, 83:5019–5024, May 1998.
- [15] K. V. Palem. Energy aware algorithm design via probabilistic computing: from algorithms and models to Moores law and novel (semiconductor) devices. In *Proc. Intl. Conf. on Compilers, Architecture and Synthesis for Embedded Systems*, pages 113–117, San Jose, California, 2003.
- [16] K. V. Palem. Proof as experiment: Probabilistic algorithms from a thermodynamic perspective. In *Proc. Intl. Symposium on Verification (Theory and Practice)*, Taormina, Sicily, June 2003.
- [17] K. V. Palem. Energy aware computing through probabilistic switching: A study of limits. *IEEE Transactions on Computers*, to appear, September 2005.
- [18] S. Pant, D. Blaauw, V. Zolotov, S. Sundareswaran, and R. Panda. A stochastic approach to power grid analysis. In *Proc. Design Automation Conference*, pages 171 – 176, Anaheim, California, June 2004.
- [19] K. Roy, S. Mukhopadhyay, and H. Mahmoodi-Meimand. Leakage current mechanisms and leakage reduction techniques in deep-submicrometer cmos circuits. *Proc. of the IEEE*, 91:305–327, 2003.
- [20] N. Sano. Increasing importance of electronic thermal noise in sub-0.1mm Si-MOSFETs. *The IEICE Transactions on Electronics*, E83-C:1203–1211, Aug. 2000.
- [21] K. L. Shepard. Design methodologies for noise in digital integrated circuits. In *Proc. Design Automation Conference*, pages 94–99, San Francisco, California, June 1998.
- [22] K. L. Shepard and V. Narayanan. Conquering noise in deep-submicron digital ICs. *IEEE Design and Test of Computers*, 15:51–62, 1998.
- [23] A. Sinha and A. Chandrakasan. JouleTrack a web based tool for software energy profiling. *Proceedings of the 38th Conference on Design Automation*, pages 220–225, 2001.
- [24] K.-U. Stein. Noise-induced error rate as a limiting factor for energy per operation in digital ICs. *IEEE J. Solid-State Circuits*, 12:527–530, Oct. 1977.
- [25] A. J. Strecok. On the calculation of the inverse of the error function, mathematics of computation. *Mathematics of Computation*, 22:144–158, Jan. 1968.
- [26] Trimaran Consortium. Trimaran: An infrastructure for research in instruction-level parallelism.