

On the Impact of Manufacturing Process Variations on the Lifetime of Sensor Networks

Siddharth Garg
Carnegie Mellon University
Pittsburgh, PA, 15213
sgarg1@andrew.cmu.edu

Diana Marculescu
Carnegie Mellon University
Pittsburgh, PA, 15213
dianam@ece.cmu.edu

ABSTRACT

As an emerging technology, sensor networks provide the ability to accurately monitor the characteristics of wide geographical areas over long periods of time. The lifetime of individual nodes in a sensor network depends strongly on the leakage power that the nodes dissipate in the idle state, especially for low-throughput applications. With the introduction of advanced low power design techniques, such as sub-threshold voltage design styles, and the migration of fabrication processes to smaller technology generations, variability in leakage power dissipation of the sensor nodes will lead to increased variability in their lifetimes. In this paper, we analyze how this increased variability in the lifetime of individual sensor nodes affects the performance and lifetime of the network as a whole. We demonstrate how sensor network designers can use the proposed analysis framework to trade-off the cost of a sensor network deployment with the performance it offers. Our results indicate that up to 37% improvement in the critical lifetime of a sensor network (defined as the expected time at which the sensor network becomes disconnected) can be obtained over a baseline design with a 20% increase in the cost of the individual sensor nodes.

Categories and Subject Descriptors: B.8.2 [Hardware]: Performance and Reliability-*Performance Analysis and Design Aids*

General Terms: Design, Performance, Reliability

Keywords: Sensor networks, manufacturing process variations, leakage power variability, lifetime

1. INTRODUCTION

Wireless sensor networks are an emerging technology that combine the advantages of the small form-factor computing, communication and sensing resources that have been made possible by advances in micro-electro-mechanical systems (MEMS) and semiconductor technologies. A sensor network

typically consists of a large number of sensor *nodes*, each of which possess local, although limited, sensing, communication and computation resources. Sensor networks are thus able to sense the characteristics of the environment within which they are embedded over wide geographical areas and over large time scales. Sensor network applications typically have low real-time requirements; i.e., the rate at which they need to sample the characteristics of the surrounding environment is typically much smaller than the time required to sample, process and communicate the data. As a result, sensor nodes tend to spend a majority of their time in *idle* mode.

One of the primary design criteria for a sensor network is the ability to provide guaranteed quality-of-service over a time period which is in the order of months or even years. Furthermore, since the individual sensor network nodes need to be placed unobtrusively in the sensed environment, they must be powered by on-board batteries with limited battery capacity or equipped with energy harvesting systems. As a result, the energy efficiency of the sensor network nodes becomes a primary design concern. Since sensor nodes spend a significant portion of their life-time idling, sensor node implementations typically consist of components that can be put in one of several low-power states (also called *sleep* states) while the node is idling. Though the power dissipated in these *sleep* states can be orders of magnitude smaller than the power dissipated in the *active* state, its contribution to the overall energy budget of a sensor node cannot be neglected. This is especially true for sensor networks that are designed to sense rare or ephemeral events such as forest fires or the presence of chemical pollutants [4].

The power dissipation of a sensor node in *sleep* state is dominated by the leakage power of the components of the sensor node that cannot be turned off. These include segments of the on-board or on-chip memory that hold persistent data (the code segment for example) and the logic that is required to wake up the sleeping components when an event occurs or after a pre-specified time-out interval. As sensor network nodes move toward more aggressive technology generations, primarily to reduce the form-factor and cost of the sensor nodes (technology scaling has enabled a 29% reduction in cost per function of the manufactured die every year [1]), the leakage power dissipation of the digital components of the nodes will increase exponentially. In [7], the authors demonstrate that the move toward smaller feature sizes will only be feasible if explicit leakage power management features such as supply gating and adaptive body biasing are provided.

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.

CODES+ISSS'07, September 30–October 3, 2007, Salzburg, Austria.
Copyright 2007 ACM 978-1-59593-824-4/07/0009 ...\$5.00.

In this work, we consider the impact of another key concern in advanced CMOS technologies, manufacturing process variability, on sensor network design. Manufacturing process variations lead to variations in both the performance and power dissipation of the fabricated die. Specifically, in [2], the authors demonstrate that even at the 180 nm technology node, a 20X variation in leakage power is observed due to manufacturing process variations. Furthermore, sub-threshold processor designs, especially tailored for sensor network applications, show inherently high sensitivity to process variations due to their low supply voltages [6]. While there exists a significant body of prior research that proposes analytic techniques to determine the impact of process variations on the leakage power dissipation of a digital IC (this research can be directly used to compute the leakage power variability for sensor network *nodes*), our goal is to evaluate how inter-node leakage power variability affects the *characteristics of the sensor network as a whole*. Intuitively, leakage power variations will cause greater variability between the lifetimes of the individual nodes in a sensor network. However, how this increased variability in individual node lifetimes impacts the lifetime of the sensor network as a whole has been an unexplored topic so far. To this end, we consider two metrics that are closely related to the lifetime of a sensor network: (1) The evolution of the number of *blind spots* in the network as a function of time (blind spots are nodes that are disconnected from the network because all the other nodes within their transmission radius are dead) and (2) The *connectivity* of the sensor network, or more precisely the expected time at which the network graph becomes disconnected. Our results can be used by sensor network processor designers to determine how much effort they need to spend on reducing the impact of process variability on their designs and by sensor network designers to evaluate trade-offs between the deployment cost of a sensor network and its lifetime.

2. RELATED WORK

Since energy and power dissipation are first-class design concerns in sensor network design, there exists a significant body of research that aims to analyze the power dissipation in sensor networks, typically from a *node perspective*. In [11], the authors develop a detailed low-level simulation infrastructure to estimate the lifetime of a sensor network node. In [7], the authors consider the impact of technology scaling on the power dissipation of a sensor network processor, but do not consider the impact of leakage variability. In contrast to these works, we use high-level power models to determine the lifetime of sensor nodes and analytical results to relate the node lifetimes to the performance of the network as a whole, with a specific focus on the impact of leakage power variations. Finally, the power/performance impact of manufacturing process variations is a well studied topic in the computer-aided design (CAD) and digital circuits community [8]. We leverage this body of work to develop our leakage variability models.

Compared to the prior work in this area, in this paper we make the following novel contributions: First, we provide a *high-level modeling framework* to evaluate the impact of leakage power variations on sensor network deployments with a specific focus on *network* performance as opposed to *node* performance. To the best of our knowledge, existing work has not addressed the importance or effect of

manufacturing process variations on sensor network performance. Second, we propose a simple technique based on *leakage power thresholding* that designers can use to alleviate the impact of leakage power variability on their sensor network deployments, in return for an increased deployment cost. Finally, we demonstrate how designers can precisely quantify the cost-performance trade-off of a sensor network deployed using the proposed leakage thresholding technique.

3. PRELIMINARIES AND ASSUMPTIONS

We consider the case of low throughput sensor networks for surveillance or environmental monitoring applications. Furthermore, since the nodes of a low throughput sensor networks spend a majority of their time idling, we assume that the nodes can be placed in a low-power *sleep* state during the idle periods to save energy. When the nodes are performing sensing, computation or communication tasks they are assumed to be in *active* state.

We consider a sensor network consisting of N sensor nodes that are randomly located within the geographic area that needs to be sensed. This is typically the case for sensor networks deployed to cover large geographic areas where it may be too expensive to place the nodes in a regular topology [4]. Each node has a pre-specified transmission radius r_n and the nodes communicate via a multi-hop routing protocol [10]. Furthermore, each node can communicate to any other node within its transmission radius.

The nodes are assumed to possess limited on-board battery capacity. Once a node's energy dissipation exceeds its battery capacity, it is assumed to be *dead* and cannot perform any further sensing, computation or communication tasks. Furthermore, we assume that the only way a node can die is if its on-board capacity battery is exhausted; i.e. we ignore random failures of sensor nodes due to causes such as accidental destruction.

Finally, to model the impact of manufacturing process variations on leakage power variability, we consider only die-to-die (D2D) variations in gate length (L_g). While within-die (WID) variations in gate length and threshold voltage also contribute to leakage power variability, their effect is averaged out due to their uncorrelated nature and can be ignored with respect to the impact of D2D variability. We would like to note that, even though process variations also impact the operating frequency of the die, the clock frequency of chips fabricated for sensor network applications will most likely be provisioned for the worst case scenario. Therefore, we do not consider the impact of process variations on the operating frequency/performance of the sensor network nodes.

Before proceeding further, we briefly discuss the relevance of the two measures of network lifetime that we consider in this work. As mentioned before, *blind spots* are nodes in the network that still have remaining battery capacity, but cannot communicate with the rest of the network because all the nodes in their neighborhood, i.e., within the transmission radius of the node r_n , are dead. Since we assume that nodes communicate via multi-hop routing protocols to save energy, these "blind spot nodes" are rendered effectively useless since they cannot transmit their sensed data to the data aggregation centers for further processing. Therefore, at any given point of time, the number of *useful* nodes in the network depends not only on the nodes that are not already dead, but also on those that are not blind spots.




Figure 1: Two state Markov Chain used for application modeling. State a refers to the *active* state and state s refers to the *sleep* state.

However, the occurrence of blind spots does not fully capture the impact of the limited node lifetimes on the lifetime of the network. Specifically, it is also possible that the death of some nodes in the network causes it to become *disconnected*. As a result, the network gets partitioned into groups of nodes that can communicate within the groups, but cannot communicate across groups. This can have serious implications on the performance of a sensor network since it may mean that parts of the network that are otherwise functional cannot reach the data aggregation center. To take into account this phenomenon, we consider, as a second metric of sensor network lifetime, the *expected time* at which a sensor network becomes *disconnected*.

4. APPLICATION MODELING

Since we are primarily interested in analyzing the power dissipation characteristics of the sensor nodes, we model the application behavior as a series of events that transition the sensor node between *sleep* state and *active* state. Specifically, we model the sensor node as a two state Markov Chain with states a (representing *active* state) and s (representing *sleep* state). The transition probabilities between the states are given by $(p_{aa}, p_{as}, p_{sa}, p_{ss})$, where p_{ij} ($i, j \in \{a, s\}$) represents the probability of moving from state i to state j at a particular time step. We note that two recent works, [3] and [14], use similar models to characterize node behavior in sensor network applications. In [14], the authors use a two state Markov Chain to model a network of cameras deployed for security and surveillance applications. In their implementation, nodes time-out from *active* state into *sleep* state if they do not sense any motion in the captured video for a pre-determined amount of time. Nodes are woken up from *sleep* state either after pre-specified intervals or by neighboring nodes that detect motion in their captured frames. In [3], the authors model data dissemination in sensor networks using a three state Markov Chain for each sensor node. The transition probabilities between states are computed using collected application traces from live sensor network deployments.

Figure 1 shows the application model with the transition probabilities marked on the edges. To ensure that this model is representative of real applications, we need to ensure that:

$$p_{aa} \gg p_{as} \quad (1)$$

and

$$p_{ss} \gg p_{sa} \quad (2)$$

These conditions ensure that Markov Chain does not generate behavior patterns where the node toggles between *sleep*

and *active* states without spending sufficient time in either state. Besides being uncharacteristic of real applications, such a behavior would be inherently undesirable due to the additional power costs of transitioning between states. If T_a is a random variable that represents the time (measured in the number of discrete step of the Markov Chain) spent by the node in state a , we can compute the average time spent by the node in *active* state, $E(T_a)$ as [9]:

$$E(T_a) = \frac{p_{aa}}{1 - p_{aa}} \quad (3)$$

and, similarly, the average time spent in state s , $E(T_s)$ as [9]:

$$E(T_s) = \frac{p_{ss}}{1 - p_{ss}} \quad (4)$$

Finally, if we define the duty cycle, D , as the fraction of time the sensor node spends, on average, in *active* state, we can write:

$$D = \frac{E(T_a)}{E(T_a) + E(T_s)} \quad (5)$$

5. POWER AND VARIABILITY MODELING

Having modeled the application characteristics, we can now model the power and energy dissipation of each sensor node. Specifically, we assume that the average power that the sensor node dissipates in *active* state is P^a and the average power it dissipates in *sleep* state is P^s . These values can be derived from measurements from real hardware or from detailed power/energy simulators such as Powertossim [11]. The energy consumed by an application in a particular time interval can then be computed as the sum of the *active* and *sleep* state power values weighted by the fraction of time within the specified time interval spent in each of these states, respectively. Finally, a particular node is assumed *dead* at the time instant when its energy consumed becomes greater than the battery capacity of that node.

To model the impact of manufacturing process variation, we write the *active* state power of a node, P^a as the sum of its dynamic (P_d^a) and leakage components (P_l^a):

$$P^a = P_d^a + P_l^a \quad (6)$$

where the P_l^a component is subject to manufacturing process variations. Finally, the *sleep* state power is assumed to be a fraction of the leakage power in *active* mode, i.e.,

$$P^s = \alpha P_l^a \quad (7)$$

where α represents the fraction of the leakage power that still exists in the *sleep* state due to the modules that cannot be power gated. The effect of manufacturing process variability can now be modeled by writing the leakage power dissipation as a function of the process parameters that are subject to variability (as explained in Section 4, we only model D2D variations in the gate length L_g):

$$P_l^a = P_{l,0}^a e^{\gamma_1 \Delta L_g + \gamma_2 \Delta L_g^2} \quad (8)$$

where $P_{l,0}^a$ is the nominal value of the *active* state leakage power dissipation and ΔL_g is the deviation of gate length from its nominal value. ΔL_g is assumed to be a zero mean normal distribution with standard deviation σ_g . The parameters γ_1 and γ_2 are obtained from simulating the leakage current of an inverter in SPICE for a given technology library.




Figure 2: The node lifetime distribution $q_f(t)$ for a sensor node for values of $\sigma_g = (0\%, 2.5\%, 5\%, 7.5\%)$ of the nominal gate length L_g . The application is modeled using transition probabilities $p_{aa} = 0.99$ and $p_{ss} = 0.9995$. The time scale on the x-axis is normalized to the average lifetime of a sensor node without considering the impact of variability.

Finally, if T_f is a random variable that represents the lifetime of a sensor node, i.e., the time when it just runs out of battery capacity, we define $q_f(t)$ to be the probability that $T_f \leq t$, i.e., $P(T_f \leq t) = q_f(t)$. $q_f(t)$ is also called the *lifetime distribution* of a sensor node and can be obtained using Monte Carlo simulations of the Markov Chain representing the given application. In order to illustrate the impact of manufacturing process variability, in Figure 2, we plot $q_f(t)$ for increasing values of variability in gate length, L_g , by obtaining $q_f(t)$ for $\sigma_g = (0, 0.025L_g, 0.05L_g, 0.075L_g)$. Recall that σ_g is the standard deviation of variability in gate length; therefore, $\sigma_g = 0$ corresponds to the nominal case in which the impact of manufacturing process variations is ignored and $\sigma_g = 0.075L_g$ corresponds to the greatest variability in process parameters. From Figure 2, it is clear that manufacturing process variations can lead to significant variations in node lifetime compared to the nominal case, and that as the variability in process parameters increases due to technology scaling, the variations in the lifetimes of sensor nodes will only increase further. More precisely, the standard deviation of node lifetime, T_f , increases by 7.7%, 50.84% and 148.43% over the baseline without any process variations for σ_g equal to 2.5%, 5% and 7.5% of the gate length respectively.

6. LIFETIME METRICS

Having characterized the lifetime distributions of the individual sensor nodes in a sensor network deployment, we now consider the network lifetime metrics that we use to quantify the impact of manufacturing process variations. To do so, we build upon prior graph theoretic results that have been proved for random networks with specified node failure probabilities. We extend these results to model a scenario in which the node failure probabilities are time-variant and are derived from the node lifetime distribution, $q_f(t)$, described in the previous Section.

In [13], the authors show that the average number of useful node N_u (recall that a useful node is one that is still alive and has at least one active neighbor) in a network with N nodes that are distributed randomly within a unit square

(or a unit circle) with transmission radius r_n can be written as:

$$N_u = Npe^{-\pi r_n^2 p} \quad (9)$$

where each node is alive or dead with probability p and $1-p$ respectively. In our case, the probability of a node failing due to limited battery capacity is a time varying function as described the lifetime distribution $q_f(t)$. Therefore, the evolution of the average number of useful nodes in the network with time can be written as:

$$N_u(t) = N(1 - q_f(t))(e^{-\pi r_n^2 (1 - q_f(t))}) \quad (10)$$

Finally, in [5], the authors prove that the random network of N nodes with transmission radius r_n is fully connected *with high probability* (whp) if:

$$r_n \geq \sqrt{\frac{\log(N)}{N}} \quad (11)$$

This implies that as the number of useful nodes in the network decreases, the network will lose its probabilistic guarantees on connectivity at time t_c (which we call the *critical lifetime* of the sensor network) that satisfies the following equation:

$$r_n = \sqrt{\frac{\log(N_u(t_c))}{N_u(t_c)}} \quad (12)$$

which was derived by observing that a network starts losing connectivity when Equation 11 is just satisfied (i.e., by enforcing the equality in the Equation 11) and that at any given time, only the useful nodes in the network can participate in network formation. As it can be seen, for a given random network consisting of N nodes and transmission radius r_n , the *critical lifetime* of the network can be computed as a function of the lifetime distribution of the individual sensor nodes, $q_f(t)$.

7. COST-PERFORMANCE TRADE-OFFS

One technique that sensor network designers could use to alleviate the impact of process variations on their designs is to tune the behavior of each node based on its power dissipation; therefore, nodes that have high power dissipation may be tuned to spend more time in *sleep* state as opposed to nodes with lower leakage power. However, as mentioned in [4], manually programming a large number of sensor networks nodes individually is impossible. An automated solution would require on-chip test circuitry to measure the leakage power dissipation of each sensor node before appropriately tuning the parameters of the firmware running on it. More importantly, the on-chip power measurement circuits will add to the cost, power dissipation and form-factor of the sensor nodes, neither of which is desirable.

As an alternative solution, we assume a scenario in which sensor network designers could purchase sensor nodes from node manufacturers with specific guarantees on the leakage power dissipation of the sensor nodes; i.e., the node manufacturer guarantees that every sensor node has a leakage power lower than a specified value, called a *leakage threshold*. For the rest of this paper, we will assume that the leakage threshold is specified as a multiple, β_t , of the nominal leakage power of a sensor node. Therefore a leakage threshold of β_t implies that every sensor node is guaranteed have an *active* state leakage power lower than $\beta_t P_{l,0}^a$, where

$P_{l,0}^a$ is the nominal value of the *active* state leakage power (this automatically sets a threshold on the *sleep* state active power also). Clearly, as the value of β_t is reduced, a sensor network designer can be assured of better network lifetime and quality of service. However, designers will have to trade-off the improved network performance with a higher cost of deployment, since nodes with stricter guarantees on leakage power will cost more. To model this trade-off, we introduce a monotonically decreasing cost function of the leakage threshold, $c(\beta_t)$. $c(\beta_t)$ is assumed to be inversely proportional to the yield that a sensor node manufacturer would obtain with the specified leakage threshold, β_t . Therefore,

$$c(\beta_t) = \frac{1}{P(P_{l,0}^a \leq \beta_t P_{l,0}^a)} \quad (13)$$

This particular choice for the cost function is explained as follows: hypothetically, one way for a sensor network designer to obtain N sensor nodes with a leakage threshold of β_t is to purchase more than N nodes and select from these the N that fall within the specified leakage threshold (assuming, of course, that there are actually N such nodes that fall within the threshold). It can be shown that on average, a designer needs to purchase $c(\beta_t)N$ nodes to ensure that at least N of these nodes satisfy the leakage threshold β_t . We note that though the cost function will, in general, be decided by the sensor node manufacturers, we chose a cost function that makes intuitive sense for the sake of having meaningful results. Our methodology, however, does not hinge on a specific choice of cost function.

8. EXPERIMENTAL RESULTS

Since our model requires a number of parameters as input, we first indicate the parameters that we use uniformly for all our experiments. We assume that nominal *active* state leakage power ($P_{l,0}^a$) is 5% of the total *active* state power (P^a), which is typical of nodes fabricated in a 130 nm process. Furthermore, we assume that the transition from the *active* state to the *sleep* state leads to a 10X reduction in leakage power dissipation (i.e., $\alpha = 0.1$, as in Equation 7), which implies that the nominal *sleep* state power dissipation is only 0.5% of the total *active* state power dissipation. The standard deviation of gate length variation, σ_g , is assumed to be 5% of the nominal value of gate length [2, 12]. Finally, we assume that $N = 100$, i.e., the sensor network consists of 100 nodes, and that the transmission radius of the nodes, r_n , is always greater than the threshold required to ensure that the network will be fully connected *whp*.

We first characterize the evolution of the number of useful nodes in the sensor network, $N_u(t)$. To model the application characteristics, we use two sets of transition probabilities corresponding to *high* (10%) and *realistic* (1%) duty cycles respectively. The transition probability values for both these cases were obtained from measured data on a sensor network deployed for security and surveillance purposes [14]. Finally, the transmission radius, r_n , is set to be equal to its minimum value, i.e., $r_n = \sqrt{\frac{\log(N)}{N}}$. Figure 3 shows the evolution of the number of useful nodes in the network as a function of time (normalized to the average node lifetime of sensor nodes without any leakage power variability) for the *high* and *realistic* duty cycles respectively. Each graph plots three curves corresponding to network deployment costs of $c(\beta_t) = (1, 1.5, 2)$. From the graphs it is clear that applica-




Figure 3: Number of useful nodes in the network as a function of time and the cost of the network deployment. The time on the x-axis is normalized to the average lifetime of a sensor node with nominal gate length. The results are shown for (a) realistic duty cycle of 1% and (b) high duty cycle of 10%.

tions with low duty cycles are the most affected by variability in leakage power dissipation, while those with high duty cycles do not show much sensitivity to leakage power variations. For the realistic duty cycle case (Figure 3 (a)), a 50% increase in deployment cost increases the number of useful sensor nodes in the network by 56% compared to the time instant when the baseline design without leakage thresholding has 50 remaining useful nodes. Similarly, a 100% increase in deployment cost increases the number of useful sensors to 93% at the same time instant. The improvements for the high duty cycle case (Figure 3 (b)) are much lower- in this case, a 50% increase in deployment costs only affords a 27% increase in the number of remaining useful nodes at the time instant when the baseline design has 50 useful nodes left.

In the next set of experiments, we characterize the impact of leakage variability on the critical lifetime t_c of a sensor network. Since we have already ascertained that high duty cycle applications are not significantly impacted by variability, we only focus on the realistic duty cycle of 1%. Furthermore, to ensure that the network has some inherent fault-tolerance we overprovision the transmission radius, i.e., we set $r_n = (1 + \zeta)\sqrt{\frac{\log(N)}{N}}$ and experiment with values of $\zeta = (0.02, 0.1, 0.2)$. In Figure 4, we plot the increase in the critical lifetime of the sensor network for the three cases as a function of the deployment cost of the network. For $\zeta = 0.02$, we observe that a 20% increase in

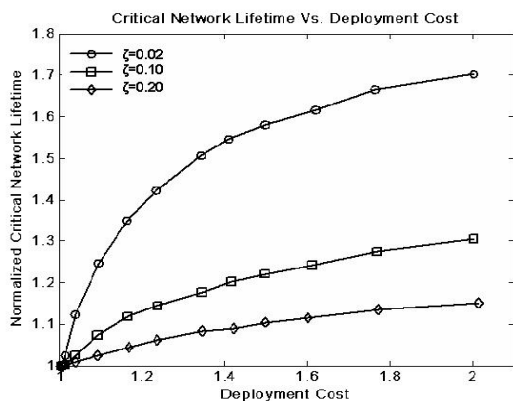


Figure 4: Critical network lifetime (t_c) as a function of the deployment cost $c(\beta_t)$ for values of transmit radius over provisioning, $\zeta = (0.02, 0.1, 0.2)$. The critical lifetimes are normalized to the critical lifetime at $c(\beta_t) = 1$.

deployment cost yields a 37% improvement in critical network lifetime. As the deployment cost increases, however, the percentage increase in critical lifetime levels off. Furthermore, it is evident that larger values of transmission radius provide a smaller percentage increase in critical network lifetime as the deployment cost is increased. This is not surprising since larger transmission radii provide greater fault tolerance which helps avoid the loss of connectivity in the network even when a number of nodes fail early due to higher leakage power dissipation than the nominal value. We note, however, that increasing the transmission radius of a sensor node can have a significant impact on the node power dissipation.

9. CONCLUSION

In this paper we provide a high-level modeling framework to analyze the impact of leakage power variability on the on the lifetime of a sensor network using two important metrics of network performance: the evolution of the number of useful nodes in the network with time and the expected time at which the network becomes disconnected. Our results indicate that variations in the leakage power dissipation of the individual sensor nodes can have a significant impact on both metrics of sensor network performance, especially for low throughput applications.

The proposed modeling framework can be used by sensor network processor designers to determine the how much design effort they need to expend to reduce the sensitivity of their designs to manufacturing process variations. Furthermore, the techniques proposed in this paper can also be used by sensor network designers to explore trade-offs between the cost of the network deployment and the performance of the deployed network.

10. ACKNOWLEDGEMENTS

The authors would like to thank Nick Zamora for helpful discussions and providing data for sensor network application characterization. This research was supported in part by Semiconductor Research Corporation contract 2005-HJ-1314.

11. REFERENCES

- [1] *International Technology Roadmap for Semiconductors 2005*, <http://public.itrs.net>.
- [2] S. Borkar, T. Karnik, S. Narendra, J. Tschanz, A. Keshavarzi, and V. De. Parameter variations and impact on circuits and microarchitecture. *Proceedings of DAC*, 2003.
- [3] D. Chakrabarti, J. Leskovec, C. Faloutsos, S. Madden, C. Guestrin, and M. Faloutsos. Information survival threshold in sensor and P2P networks. *Proceedings of 26th Annual IEEE ICC*, 2007.
- [4] P. Dutta, M. Grimmer, A. Arora, S. Bibyk, and D. Culler. Design of a wireless sensor network platform for detecting rare, random, and ephemeral events. *Proceedings of IPSN*, 2005.
- [5] P. Gupta and P.R. Kumar. Critical power for asymptotic connectivity in wireless networks. *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of WH Fleming*, 3(20):547–566, 1998.
- [6] S. Hanson, B. Zhai, D. Blaauw, D. Sylvester, A. Bryant, and X. Wang. Energy optimality and variability in subthreshold design. *Proceedings of ISLPED*, 2006.
- [7] M. Hempstead, G.Y. Wei, and D. Brooks. Architecture and circuit techniques for low-throughput, energy-constrained systems across technology generations. *Proceedings of CASES*, 2006.
- [8] S. Natarajan, M.A. Breuer, and S.K. Gupta. Process variations and their impact on circuit operation. *Proceedings of the IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems*, 1998.
- [9] A. Papoulis. *Probability, random variables, and stochastic processes*. McGraw-Hill New York, 1991.
- [10] R.C. Shah and J.M. Rabaey. Energy aware routing for low energy ad-hoc sensor networks. *Proceedings of WCNC*, 2002.
- [11] V. Shnayder, M. Hempstead, B. Chen, G.W. Allen, and M. Welsh. Simulating the power consumption of large-scale sensor network applications. *Proceedings of CASES*, 2004.
- [12] J.W. Tschanz, J.T. Kao, S.G. Narendra, R. Nair, D.A. Antoniadis, A.P. Chandrakasan, and V. De. Adaptive body bias for reducing impacts of die-to-die and within-die parameter variations on microprocessor frequency and leakage. *IEEE Journal Of Solid-State Circuits*, 37(11), 2002.
- [13] C.-W. Yi, P.-J. Wan, X.-Y. Li, and O. Frieder. Asymptotic distribution of the number of isolated nodes in wireless ad hoc networks with Bernoulli nodes. *IEEE Transactions on Communications*, 54(20):510–517, 2006.
- [14] N. Zamora, J-C. Kao, and R. Marculescu. Co-ordinated power management for wireless sensor network video systems. *Proceeding of DATE*, 2007.