Proceedings of the 2006 American Control Conference
Minneapolis, Minnesota, USA, June 14-16, 2006

Application of Stochastic Hybrid Systems in Power Management of Streaming Data

Jianghai Hu

***Abstract*— In this paper, we study the optimal power management problem for a pipeline of streaming data consisting of several components and buffers in between. The production rate of the source component is assumed to be random. We aim to find the optimal switching strategy and the optimal size of the buffers so that the expected average power consumption of the pipeline system is minimized. For the case of two components with one buffer in between, we model the system by a stochastic hybrid system, and derive analytically the solution to the optimal power management problem.**

I. INTRODUCTION

Data buffers are often used between two interacting components in a computer. For example, when a user watches a video clip using a PDA (personal digital assistant), the PDA first downloads sufficient amount of data into buffer memory so that the video can be played smoothly regardless of the variations of the network conditions. This process may consist of multiple components and buffers: The first component is the network card downloading the video data. The data can be stored in a local storage device, such as a microdrive or flash memory. The data then are retrieved by the processor to decode. The decoded video may be stored in a frame buffer before showing on the LCD display. These activities form a multiple-stage pipeline for the streaming data. In Fig. 1, we show schematically pipelines consisting of two components in (a) and three components in (b), with buffers inserted between successive components.

Conserving the energy of electronic systems is increasingly important as battery-powered portable systems become more popular. Many methods have been proposed for reducing the energy consumption of individual components, such as processors, wireless network interface cards, hard disks, and displays. These methods predict the periods when a component is idle or under-utilized. The component is turned off (sometimes called shut down, sleeping) or the performance is scaled down to reduce energy consumption. This is called dynamic power management [4]. Changing the power state requires additional energy; hence, a component cannot be turned off every time it is idle. Energy can be saved only if the hardware component can remain in the off state sufficiently long to compensate the energy overhead. This minimum time in the off state is called the component's break-even time [3].

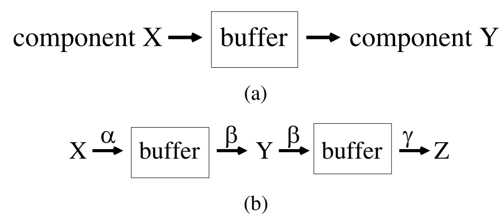

Fig. 1. (a) One buffer between a producer and a consumer. (b) Two buffers between three interacting components.

Adding buffer memory can create opportunities for power reduction for under-utilized components. For example, in Fig. 1(a), component X produces data for Y to consume. Since the production rate of X is typically higher than the consumption rate of Y, X can be turned off occasionally to save power, provided that there are enough data in the buffer for Y to retrieve. During any given time period, the power consumption of the whole pipeline system consists of the following parts: power consumed by X when it is turned on, by Y, and by the buffer memory. Thus an interesting problem is to determine the optimal size of the buffer and the optimal strategy for turning on and off the component X so that maximal power saving is achieved. This is the optimal power management problem for the pipeline.

In [6], [7], the optimal power management problem is studied for the one buffer case in Fig. 1(a). In particular, the optimal size of the buffer is derived. In [10], the two-buffer example in Fig. 1(b) is considered, and the optimal switching strategy is derived under the constraint that each component can turn on and off at most once in each cycle. A main drawback of these prior works is that they assume the production rates of the components are

J. Hu is with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA. `jianghai@purdue.edu`

1-4244-0210-7/06/$20.00 ©2006 IEEE

deterministic, while in real network environment, these rates may have random variations.

In this paper, we try to alleviate the above mentioned drawback by considering the simplest situation with random production rates, namely, we assume that the production rate of component X follows a continuous Markov process with two possible values. Under this assumption, we shall study the optimal power management problem for the pipeline by modeling it as a stochastic hybrid system. We will derive analytically the optimal buffer size for the one-buffer case, and discuss briefly how to extend the result to the two-buffer case.

The rest of the paper is organized as follows. In Section II, we present a general framework of stochastic hybrid systems. The model used in this paper will be a special case of this general framework. In Section III, we formulate the problem of optimal power management for a pipeline of streaming data, and model the system as a stochastic hybrid system. Thus the problem reduces to an optimal control problem for the stochastic hybrid system. In Section IV, we derive analytically the solution to the problem for the simple case of two components with one buffer in between. Finally, in Section V, we outline the conclusions and remark on some future directions.

II. Stochastic Hybrid Systems

In this section, we present a framework of stochastic hybrid systems originally proposed in [11] and subsequently extended in [5] by incorporating the concept of switching diffusions [8]. See [1], [2], [9], [13], [14] for other relevant works on stochastic hybrid systems.

In the extended framework, the state (X_t, m_t) of a stochastic hybrid system $\mathcal{H}$ consists of two parts: X_t is the continuous state and m_t is the discrete state (or mode), both of which vary with time t. The discrete state m_t takes values in a set S that is finite or countably infinite, while, depending on the value of $m_t = i$, the continuous state X_t takes values in a domain U_i in a Euclidean space $\mathbb{R}^n$. Thus the total state space of $\mathcal{H}$ is $U = \cup_{i \in S} U_i \times \{i\}$, a subset of $\mathbb{R}^n \times S$.

The dynamics of the state (X_t, m_t) is characterized by the following:

- **Continuous Dynamics.** During any time interval where m_t remains constant, the continuous state X_t evolves according to a SDE:
$$dX_t = \mu(X_t, m_t, u_t)\, dt + \sigma(X_t, m_t, u_t)\, dB_t, \quad (1)$$
where B_t is a d-dimensional Brownian motion; $u_t \in \mathbb{R}^p$ is the control input; and μ and σ are functions of proper dimensions and satisfy some smoothness conditions.
- **Discrete Dynamics.** The discrete state m_t jumps from its current mode i to another mode j when either of the following two events occurs. The first event, a *forced transition*, occurs whenever X_t hits a certain subset (typically a part of the boundary ∂U_i) of the domain U_i called the *guard* for mode j. The second event, a *spontaneous transition*, is generated by a (controlled) Markov process such that, given $X_t = x$, $m_t = i$ and $u_t = u$ at time t, the probability of m_t jumping spontaneously to mode j within time Δt is
$$P(m_{t+\Delta t} = j \,|\, m_t = i, X_t = x, u_t = u) = \begin{cases} \lambda_{ij}(x,u)\Delta t + o(\Delta t), & \text{if } i \neq j, \\ 1 + \lambda_{ii}(x,u)\Delta t + o(\Delta t), & \text{if } i = j. \end{cases} \quad (2)$$
Note that here $\lambda_{ij}(x,u) \geq 0$ for $i \neq j$, and $\lambda_{ii}(x,u) = -\sum_{j \in S, j \neq i} \lambda_{ij}(x,u)$.
- **Reset Conditions.** When the discrete state m_t jumps from mode i to mode j, the continuous state X_t will be restarted in the new domain U_j. The rules for deciding where X_t is restarted are called the *reset conditions*, which can be either random or deterministic.

Note that the continuous and discrete dynamics are coupled: the equation (1) governing the evolution of the continuous state varies with the discrete mode; and the transition probabilities of the discrete state jumps depend on the continuous state at the moments of the jumps. In addition, it is important to have both types of discrete transitions in the model: the forced transitions typically model those discrete mode jumps due to human intervention or hardware/safety constraints; while other jumps due to environmental changes or unplanned events such as component failures are best modeled by spontaneous transitions.

A typical solution (or *execution*) of a stochastic hybrid system is shown in Fig. 2a. In general, for an execution (X_t, m_t) of $\mathcal{H}$, $t \geq t_0$, there exists a partition of time $t_0 \leq t_1 \leq t_2 \leq \cdots$ such that in each subinterval $[t_i, t_{i+1})$, m_t is constant and X_t follows equation (1) for that particular m_t, and that (X_t, m_t) at the neighboring ends of successive subintervals satisfy the reset conditions. We assume that the guards and the reset conditions are chosen so that almost all executions are not zeno (a zeno execution has infinite number of transitions in a finite time period).

III. Problem Formulation

Consider the one-buffer example in Fig. 1(a). Assume that the consumption rate of component Y is a constant β. This is the rate at which Y retrieves data from

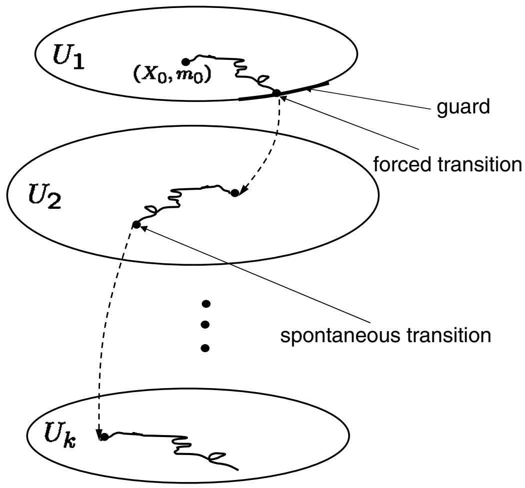

Fig. 2. An execution of a stochastic hybrid system

the buffer for its own use. The production rate α of component X, on the other hand, is assumed to be random with two possible values α_1 and α_2, and follows a continuous time Markov process with generator

$$\begin{bmatrix} -\lambda_1 & \lambda_1 \\ \lambda_2 & -\lambda_2 \end{bmatrix} \tag{3}$$

for some positive constants λ_1 and λ_2. In other words, if $\alpha = \alpha_i$ at the current time t, then within a short period of time Δt, α will change to α_j with probability $\lambda_{ij}\Delta t + o(\Delta t)$, provided $i \neq j$ and $i, j \in \{1, 2\}$. Furthermore, we assume that the production rate of X is zero if it is turned off.

A. Stochastic Hybrid System Model of the Pipeline of Streaming Data

From the above description, the pipeline can be modeled by a stochastic hybrid system $\mathcal{H}$. The continuous variable of $\mathcal{H}$ is the amount of data stored in the buffer, denoted by X_t; and the discrete mode m_t of $\mathcal{H}$ is the power state of component X, which has three possible values:

1) Mode 0 corresponds to when X is turned off;
2) Mode 1 corresponds to when X is turned on and has production rate α_1;
3) Mode 2 corresponds to when X is turned on and has production rate α_2.

In mode $m_t = m \in \{0, 1, 2\}$, the continuous variable X_t changes at the constant rate v_m defined by

$$v_0 = -\beta, \quad v_1 = \alpha_1 - \beta, \quad v_2 = \alpha_2 - \beta. \tag{4}$$

For example, if $m_t = 1$, then X is turned on and stores data into the buffer at the rate α_1, while Y retrieves data from the buffer at the rate β; thus the net increasing rate of data in the buffer is $\frac{dX_t}{dt} = v_1 = \alpha_1 - \beta$. Equation (4) specifies the continuous dynamics of $\mathcal{H}$. We assume that $\alpha_2 > \alpha_1 > \beta$. Thus $v_2 > v_1 > 0 > v_0$.

The discrete dynamics of $\mathcal{H}$ is as follows. The transitions between mode 1 and mode 2 are governed by the continuous time Markov chain with generator (3) described earlier, and belong to spontaneous transitions. The transitions between mode 0 and other modes 1 and 2 are forced transitions, and correspond to turning on/off component X. They are typically specified by the user as part of the overall power management strategy. To prevent buffer overflow and underflow, a transition from mode 0 to mode 1 or 2 is forced whenever $X_t = 0$, i.e., whenever the buffer becomes empty; while a transition from mode 1 or 2 to mode 0 is forced whenever X_t reaches the maximal capacity of the buffer.

The reset conditions of $\mathcal{H}$ are trivial. Namely, in the case a spontaneous or a forced transition occurs, X_t remains unchanged, as the amount of data in the buffer cannot experience instantaneous jumps.

B. Power Consumption of the Pipeline

In order to define the dynamic power management problem, we now analyze the power expenditure of the components and the buffer. The power consumed by each component consists of two parts: the dynamic power and the static power. The dynamic power includes the power to produce, process, and/or consume data, as well as the power for writing data into and retrieving data from the buffer memory. As the name suggests, the dynamic power depends on the data flow rate through the component. In contrast, the static power of the component is the part of the component's power used to sustain its normal operation, and is independent of the data rate. In this paper, we consider only the static powers, as it is shown in [6], [7], [12] that the dynamic powers do not affect the solution of the optimal buffer management problem defined later in this section.

More precisely, component X has two power states: on and off. When X is on, it consumes the (static) powers η; and when X is off, it consumes zero power. In addition, changing the power state of X may incur extra energy. For simplicity, we assume that the energy overhead for turning X from off to on is k_x, and the energy overhead for turning X from on to off is 0. In the hybrid system model, this implies that there is an extra energy η whenever there is forced transition from mode 0 to either mode 1 or mode 2. We also assume that there is no delay in changing the power states of X. Because the component Y is assumed to be always turned on, it consumes a constant power. Thus its power contribution can be ignored in finding the solutions for maximal power saving.

Remark 1: In practice, the power consumption of component X should depend on the rate at which it is transmitting data. Thus the powers of X in mode 1 and mode 2 should be different. In this paper, we assume that these two powers are identical for simplicity. The results in this paper, however, can be easily extended to the general case.

In addition to component powers, the buffer also consumes power that is proportional to it size. Denote by Q the size of the buffer. Then the power of the buffer is aQ for some constant a, where the constants a is the power per unit size of the buffer memory (W/MB).

To sum up, for a hybrid execution $z_t = (X_t, m_t)$ of the hybrid system $\mathcal{H}$, its average power consumption during a give time period $[0, t_f]$ is given by

$$\bar{P}(z_t; t_f, Q) = \frac{1}{t_f}\left[\eta S_1 + n_x k_x + aQt_f\right] = \frac{1}{t_f}\left[\eta S_1 + n_x k_x\right] + aQ, \tag{5}$$

where S_1 is the amount of time in $[0, t_f]$ that X is turned on, and n_x is the number of times X is turned from off to on within $[0, t_f]$. Note that the argument of $\bar{P}$ indicates its dependency on the time window size t_f and the buffer size Q.

C. Optimal Power Management Problem

The dynamic power management problem is then to find a switching strategy for turning on/off component X and the optimal buffer size Q to achieve the minimal expected average power. The switching strategy can be either deterministic or probabilistic. A deterministic switching strategy is specified by two subsets of the state space $[0, Q]$: one is the guard for turning component X on, i.e., the guard for the forced transitions from mode 0 to mode 1 or mode 2; and the other is the guard for turning component X off, i.e., the guard for the forced transitions from mode 1 or mode 2 to mode 0. A probabilistic switching strategy is specified by two functions $\lambda_{on}, \lambda_{off} : [0, Q] \to \mathbb{R}_+$. Specifically, if component X is off and $X_t = x$ at the current time t, the within a short period of time Δt, the probability of a force transition from mode 0 to either mode 1 or mode 2 is given by $\lambda_{on}(x)\Delta t + o(\Delta t)$. To prevent buffer underflow, such a force transition is necessary if $X_t = 0$. Thus $\lambda_{on}(x) \to \infty$ as $x \to 0$. On the other hand, if component X is on and $X_t = x$ at the current time t, the within a short period of time Δt, the probability of a force transition from mode 1 or mode 2 to mode 0 is given by $\lambda_{off}(x)\Delta t + o(\Delta t)$. To prevent buffer overflow, we need to have $\lambda_{off}(x) \to \infty$ as $x \to Q$. From the above description, forced transitions can be thought of as spontaneous transitions with infinite rates, and deterministic switching strategies can be thought of as special cases of probabilistic switching strategies.

In particular, if the look ahead time horizon t_f is large enough, the problem becomes:

Problem 1: Find the switching strategy and the buffer sizes Q so that the expected average power $\bar{P}_e = \lim_{t_f \to \infty} E[\bar{P}(z_t; t_f, Q)]$ is minimized.

Because of the infinite time horizon, the initial conditions of the hybrid execution z_t will not affect the value of $\bar{P}_e$. Thus Problem 1 becomes an optimal control problem for the stochastic hybrid system modeling the pipeline of streaming data.

IV. Solution of the One-Buffer Case

In this section, we try to find the solution to Problem 1 for the one-buffer case. The key to the solution is to evaluate $\bar{P}_e$ for the random process z_t.

For an arbitrary execution z_t of the stochastic hybrid system $\mathcal{H}$ modeling the pipeline, define a sequence of times $T_0, T_1, \ldots$, so that T_k is the k-th time that component X is turned from off to on. Denote by A_k the total energy consumed by component X and the buffer during the time period $[T_k, T_{k+1})$. Then

$$A_k = k_x + S_k\eta + aQ(T_{k+1} - T_k), \tag{6}$$

where S_k is the amount of time that X is turned on during the time interval $[T_k, T_{k+1})$. The desired quantity $\bar{P}_e$ can be rewritten as

$$\bar{P}_e = \lim_{t_f \to \infty} E[\bar{P}(z_t; t_f, Q)] = \lim_{k \to \infty} E\left[\frac{\sum_{k=0}^{\infty} A_k}{\sum_{k=0}^{\infty}(T_{k+1} - T_k)}\right].$$

Consider $\{T_{k+1} - T_k\}_{k=0}^{\infty}$ as a Markov regenerative process with rewards $\{A_k\}_{k=0}^{\infty}$, the standard results in stochastic processes theory [15] yield

$$\bar{P}_e = \frac{E(A_k)}{E(T_{k+1} - T_k)}$$

for arbitrary $k = 0, 1, \ldots$. Using (6), we have

$$\bar{P}_e = \frac{k_x + \eta E(S_k)}{E(T_{k+1} - T_k)} + aQ. \tag{7}$$

Note that the objective is to find the switching strategy and the buffer size Q to minimize $\bar{P}_e$. In (7), k_x, η, and aQ are constants independent of the switching strategy. $E(S_k)/E(T_{k+1} - T_k)$ is the proportion of time that X is turned on during $[T_k, T_{k+1})$. This proportion will be constant after the system reaches equilibrium state, namely, the random process z_t becomes stationary. Intuitively, in the equilibrium state, z_t is periodic probabilistically, and $E(S_k)/E(T_{k+1} - T_k)$ is equal to the ratio of the average increasing rate of X_t when X is on to the decreasing rate β when X is off. To sum up, in order to minimize $\bar{P}_e$, we need to make $E(T_{k+1} - T_k)$ as large as possible. Or equivalently, the optimal switching strategy is such that the forced switching from mode 0 to mode 1 or mode 2 is activated only if $X_t = 0$, and the forced switching from mode 1 or mode 2 to mode 0 is activated only if $X_t = Q$. Namely, component X is turned on only if the buffer is empty, and it is turned off only if the buffer is full.

Having determined the optimal switching strategy, we now study the optimal buffer size Q^*. To this purpose, we decompose the interval $[T_k, T_{k+1})$ into two stages. During the first stages with duration S_k, X is turned on, and X_t increases at the random rate of $\alpha_1 - \beta$ or $\alpha_2 - \beta$, depending on the production rate of X. During the second stages with duration S'_k, X is turned off, and X_t decreases at the constant rate β. Because of the choice of the optimal switching strategy, we have

$$S'_k = \frac{Q}{\beta},$$

which is deterministic. On the other hand, S_k is random, and its expectation in general does not admit a simple formula. However, since we consider the infinite time horizon case, the Markov chain modeling the evolution of the production rate α of X can be assumed to be in equilibrium state. The equilibrium probabilities of α are

$$P(\alpha = \alpha_1) = \frac{\lambda_2}{\lambda_1 + \lambda_2},$$
$$P(\alpha = \alpha_2) = \frac{\lambda_1}{\lambda_1 + \lambda_2}.$$

Thus the average production rate of X in equilibrium state is

$$\alpha_{av} = \frac{\lambda_2\alpha_1 + \lambda_1\alpha_2}{\lambda_1 + \lambda_2}.$$

If λ_1 and λ_2 are not exceedingly small, the Markov chain α will have relatively frequent transitions ("well mixed") during the first stage, and the expected duration S_k can be approximated by

$$E(S_k) \simeq \frac{Q}{\alpha_{av}} = \frac{Q(\lambda_1 + \lambda_2)}{\lambda_2\alpha_1 + \lambda_1\alpha_2}.$$

As a result,

$$\begin{aligned} E(T_{k+1} - T_k) &= E(S_k + S'_k) \\ &= Q\left(\frac{1}{\beta} + \frac{\lambda_1 + \lambda_2}{\lambda_2\alpha_1 + \lambda_1\alpha_2}\right). \end{aligned} \tag{8}$$

Plugging (8) into (7), we have

$$\bar{P}_e = aQ + \frac{b}{Q} + c, \tag{9}$$

where b and c are constants defined by

$$b = \frac{\beta k_x(\lambda_2\alpha_1 + \lambda_1\alpha_2)}{\lambda_2(\alpha_1 + \beta) + \lambda_1(\alpha_2 + \beta)}, \tag{10}$$

$$c = \frac{\beta\eta(\lambda_1 + \lambda_2)}{\lambda_2(\alpha_1 + \beta) + \lambda_1(\alpha_2 + \beta)}. \tag{11}$$

From (9), it is obvious that the optimal Q to minimize $\bar{P}_e$ is

$$Q^* = \sqrt{\frac{b}{a}} = \sqrt{\frac{\beta k_x(\lambda_2\alpha_1 + \lambda_1\alpha_2)/a}{\lambda_2(\alpha_1 + \beta) + \lambda_1(\alpha_2 + \beta)}}, \tag{12}$$

and the resulting optimal expected average power is

$$\begin{aligned} \bar{P}_e^* &= 2\sqrt{ab} + c \\ &= 2\sqrt{\frac{a\beta k_x(\lambda_2\alpha_1 + \lambda_1\alpha_2)}{\lambda_2(\alpha_1 + \beta) + \lambda_1(\alpha_2 + \beta)}} \\ &\quad + \frac{\beta\eta(\lambda_1 + \lambda_2)}{\lambda_2(\alpha_1 + \beta) + \lambda_1(\alpha_2 + \beta)}. \end{aligned} \tag{13}$$

We sum up the results in this section in the following proposition.

Proposition 1: Suppose that λ_1 and λ_2 are not exceedingly small with respect to Q so that the Markov chain α is well mixed between successive turning ons of component X. Then the optimal buffer size Q^* is given by (12), with the corresponding optimal expected average power $\bar{P}_e^*$ given by (13). The optimal switching strategy is such that component X is turned on only if the buffer is empty and turned off only the buffer is full.

V. Conclusion and Future Directions

We study the optimal power management problem for a pipeline of streaming data using the modeling tool of stochastic hybrid systems. For the simple case of two component with one buffer in between, we derived analytically the optimal switching strategy and the optimal buffer size so that maximal power saving can be achieved through dynamic power management.

Our result can be extended in several ways. For example, to more accurately model the randomness of the production rate α of component X, more possible values can be assigned to α. More importantly, it is interesting to study the problem for the general case of more components and buffers, such as the one in Fig. 1(b). It can be expected that, as the corresponding stochastic hybrid systems have more discrete modes and higher dimensional state space, the optimal switching strategy will be highly nontrivial, and determination of the optimal buffer sizes will be much more difficult. We intend to apply the technique proposed in [14] on the optimal control of general stochastic hybrid systems to study the resulting problems. In particular, the expected energy consumption is characterized by an ODE/PDE whose coefficients depend on the strategy functions λ_{on} and λ_{off}, and the adjoint method (a gradient-based method with PDE constraints) can be used to find the optimal solution. Research is currently carried out along these directions.

REFERENCES

[1] A. Abate, A. D. Ames, and S. Sastry. Stochastic approximation of hybrid systems. In *Proc. American Control Conf.*, pages 1557–1560, Portland, OR, Jun. 2005.

[2] S. Amin and S. Sastry. Reachability analysis for certain classes of stochastic hybrid systems. Preprint, 2005.

[3] L. Benini, A. Bogliolo, and G. D. Micheli. A Survey of Design Techniques for System-Level Dynamic Power Management. *IEEE Transactions on Very Large Scale Integration Systems*, 8(3):299–316, June 2000.

[4] L. Benini and G. D. Micheli. System-Level Power Optimization: Techniques and Tools. *ACM Transactions on Design Automation of Electronic Systems*, 5(2):115–192, April 2000.

[5] M. L. Bujorianu. Extended stochastic hybrid systems and their reachability problem. In R. Alur and G. J. Pappas, editors, *Hybrid Systems: Computation and Control, 7th Int. Workshop (HSCC 2004)*, volume 2993 of *Lecture Notes in Computer Science*, pages 234–249. Springer Verlag, 2004.

[6] L. Cai and Y.-H. Lu. Dynamic Power Management Using Data Buffers. In *Design Automation and Test in Europe*, pages 526–531, 2004.

[7] L. Cai and Y.-H. Lu. Energy Management Using Buffer Memory for Streaming Data. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 24(2):141–152, February 2005.

[8] M. K. Ghosh, A. Arapostathis, and S. I. Marcus. Optimal control of switching diffusions with applications to flexible manufacturing systems. *SIAM J. Control and Optimization*, 31:1183–1204, 1993.

[9] J. Hespanha. Polynomial stochastic hybrid systems. In M. Morari, L. Thiele, and F. Rossi, editors, *Hybrid Systems: Computation and Control, 8th Int. Workshop (HSCC 2005)*, volume 3414 of *Lecture Notes in Computer Science*, pages 322–338. Springer Verlag, 2005.

[10] J. Hu and Y.-H. Lu. Buffer management for power reduction using hybrid control. In *IEEE Conference on Decision and Control and European Control Conference*, 2005.

[11] J. Hu, J. Lygeros, and S. Sastry. Towards a theory of stochastic hybrid systems. In N. Lynch and B. H. Krogh, editors, *Hybrid Systems: Computation and Control, 3rd Int. Workshop (HSCC 2000)*, volume 1790 of *Lecture Notes in Computer Science*, pages 160–173. Springer Verlag, 2000.

[12] N. Pettis, L. Cai, and Y.-H. Lu. Dynamic Power Management for Streaming Data. In *International Symposium on Low Power Electronics and Design*, pages 62–65, 2004.

[13] S. Prajna, A. Jadbabaie, and G. J. Pappas. Stochastic safety verification using barrier certificates. In *Proc. 43rd IEEE Int. Conf. on Decision and Control*, pages 929–934, Atlantis, Paradise Island, Bahamas, Dec. 2004.

[14] R. Raffard, J. Hu, and C. Tomlin. Adjoint-based optimal control of the expected exit time for stochastic hybrid systems. In M. Morari, L. Thiele, and F. Rossi, editors, *Hybrid Systems: Computation and Control, 8th Int. Workshop (HSCC 2005)*, volume 3414 of *Lecture Notes in Computer Science*, pages 557–572. Springer Verlag, 2005.

[15] S. Ross. *Applied Probability Models with Optimization Applications*. Dover, 1970.