Optimal Power Modes Scheduling Using Hybrid Systems

Wei Zhang, Jianghai Hu and Yung-Hsiang Lu
School of Electrical and Computer Engineering
Purdue University, West Lafayette, IN 47907, USA.
{zhang70, jianghai, yunglu}@purdue.edu

***Abstract*— This paper studies a class of dynamic buffer management problem with one buffer inserted between two interacting components. Different from many previous studies, the component to be controlled is assumed to have multiple power modes, each of which corresponds to a different data processing rate. The overall system is modeled as a hybrid system and the buffer management problem is formulated as an optimal control problem. The cost function of the proposed problem depends on the switching cost and the size of the continuous state space, making its solutions much more challenging. By exploiting some particular features of the proposed problem, the best mode sequence and the optimal switching instants are characterized analytically using some variational approach. Simulation result shows that the proposed method can save at least 30% of energies compared with another heuristic scheme in several typical situations.**

I. Introduction

Energy conservation is a crucial issue for the design of portable electronic devices. Although many power management methods have been proposed [1], most of them can only reduce energies for individual components, such as processors, hard disks, cache, and memories. In contrast to these methods, the *dynamic buffer management (DBM)* technique considers the interactions among components and minimizes the power consumption of the overall system, including the components, the buffers and the switching costs.

Fig. 1. A simple example of the DBM problem

Fig. 1 illustrates a simple example of the DBM problem, where one buffer is inserted between two interacting components X and Y. Suppose that X produces data faster than Y consumes. Then if both of them are active, there will be some data accumulating in the buffer. Due to the data storage in the buffer, X can be turned off at some time to save energy. A larger buffer can save more energy for X but consumes more energy by itself. The goal of the DBM problem is to find the optimal buffer size and the best scheduling of the "on" and "off" modes of X, so that the power consumption of the overall system is minimized.

The optimal solution to the DBM problem is first derived in [2] for one buffer between two components. The result is further extended in [3][4] where two buffers are inserted between three streamlined components. A major limitation of these previous studies is that they all assume that the components to be controlled have only two power modes, "on" and "off". However, in practice, many components can work in more than two power modes, such as variable speed processors [5][6] and multi-speed disks [7]. For such a component, instead of completely turning it off, we can properly design a switching strategy, namely the scheduling of different power modes of the component, to further reduce the overall power consumptions.

This paper studies a more general DBM problem, where the component to be controlled has multiple power modes. Since different power modes correspond to different data accumulation/depletion rate in the buffer, the overall system is modeled as a piecewise-constant hybrid system, or more accurately, a multi-rate automata [8]. The DBM problem is thus formulated as an optimal control problem of the underlying hybrid system. Despite the richness of the literature in optimal control of hybrid systems [9][10][11][12], previous results cannot be directly applied to our problem as it has the following distinct features: (1) transitions among discrete modes depend on the evolution of the continuous state; whereas most previous studies ignore such dependence; (2) the switching (mode) sequence is a decision variable that cannot be assumed fixed as in [10][11]; (3) the switching cost ignored in most previous papers is an important part of our cost function; (4) The buffer size that determines the range of the continuous states is variable, indicating that both the optimal control and the optimal size of the continuous state space are to be designed at the same time. Few existing results have addressed all of the above issues. Despite these difficulties, this paper derives the optimal solution to the proposed problem analytically, under some practically reasonable conditions, using a variational approach. It represents an important application of the optimal control theory of hybrid systems in the low power design of embedded systems.

The rest of this paper is organized as follows. In Section II, the dynamic buffer management problem is introduced and formulated as an optimal control problem of a piecewise constant hybrid system. The optimal solution is derived analytically in Section III using some variational approach. A numerical simulation is performed in Section IV to demonstrate the effectiveness of the proposed method. Finally, some conclusion remarks and possible future research directions are discussed in Section V.

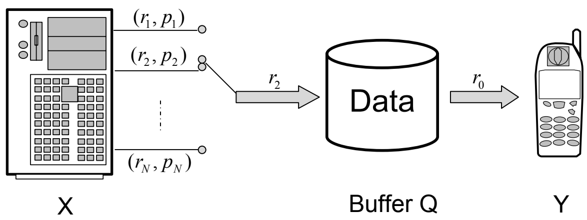

Fig. 2. System Configuration

II. Problem Formulation

A. System Description

Consider two interacting components X and Y as shown in Fig. 2, where X produces data for Y to consume. Suppose Y is always "on" and consumes data at a constant speed r_0. On the other hand, assume that X has N different operation modes where in mode i, $i = 1, 2, \ldots, N$, it produces data at a constant speed r_i and consumes energy at the rate of p_i. Without loss of generality, assume $r_1 < r_2 < \cdots < r_N$. Usually, a lower data processing rate corresponds to a lower power consumption; thus we require $p_1 < p_2 < \cdots < p_N$. Denote by I and J the sets of indices whose corresponding data rates are greater and smaller than r_0, respectively, i.e.,

$$\begin{aligned} & I = \{i \,|\, r_i > r_0, i = 1, \ldots, N\}, \\ \text{and} \quad & J = \{j \,|\, r_j < r_0, j = 1, \ldots, N\}. \end{aligned}$$

Assume that both I and J are nonempty, i.e., $r_N > r_0 > r_1$. A mode σ is called an *ascending mode* if $\sigma \in I$ and a *descending mode* otherwise. To ensure smooth operation, a buffer B with capacity Q is inserted between X and Y. See Fig. 2 for the configuration of the overall system.

Many real-world applications can be described by the above system. One simple example is the data-copying process, where a device Y copies data from a hard drive X. The hard drive has two power modes "on" and "off". If the speed of the hard drive is faster than the speed of Y, then X can be turned off during some time intervals to save energy. In this case, the system memory, which serves as the buffer B in our model, is also needed to temporarily store the data from X for later delivery. As another example, consider the video playing process. Let X be the Intel Xscale processor [13] that can operate on multiple voltages corresponding to different speeds r_i's and powers p_i's; let Y be a video card that needs data from X at a constant speed, say 30frame/sec. To ensure smooth operation, the system memory is needed as a buffer to store the data that has been decoded by X but yet to be displayed by Y. Therefore, the abstract system as shown in Fig. 2 represents a class of practical systems. Minimizing the power consumption of such a system is a meaningful and important research problem.

If X is switched to a mode with a speed higher than r_0, data will accumulate in B. When B has enough data for Y to consume, one can switch X to a lower power mode to save its energy. In this way, the buffer reduces the power consumption of X. On the other hand, it introduces additional energy consumption, namely, the buffer energy and the energy to switch between different modes. Assume that the buffer power is proportional to the buffer capacity Q and denoted by $p_m Q$, where p_m is a constant. Suppose that switching between different modes costs the same amount of energy k_s. Under these assumptions, the goal of this paper is to find an optimal switching strategy, in the sense that X switches to certain modes at certain times so that the data rate as demanded by component Y is guaranteed, and that the overall system consisting of X, Y and the buffer B, consumes the least amount of average power.

B. Hybrid System Model

The above system can be modeled as a hybrid system H. The discrete state space of the hybrid system consists of N modes: $S = \{1, 2, \ldots, N\}$, representing all the operation modes of X. The continuous state $q(t)$ is defined as the amount of data stored in the buffer B, and is thus required to take values in the interval $[0, Q]$. The evolution of $q(t)$ is determined by the speed difference between the two components, i.e., $\dot{q}(t) = r_i - r_0$ for mode i. As a physical constraint, there should be no buffer underflow or overflow. Thus we require that whenever $q(t)$ hits the boundary of its domain, namely, $q(t) = 0$ or Q, the system must transit to another mode that can bring $q(t)$ back to the inside of $[0, Q]$. Except for this, there are no other transition rules or guard conditions. The reset map of the system is trivial, i.e., there is no jump in $q(t)$ at the transition instant.

Given a time period $[0, t_f]$, the behavior of the above system can be uniquely determined by the switching strategy $\sigma : [0, t_f] \to S$, which determines the active mode of the system over $t \in [0, t_f]$. The overall trajectory $z(t) = (q(t), \sigma(t))$ of the hybrid system consists of the trajectories of both the continuous state $q(t)$ and the discrete state $\sigma(t)$. For a given initial value $q(0)$, the system is governed by the following differential equation:

$$\frac{dq(t)}{dt} = r_{\sigma(t)} - r_0, \qquad \forall t \in [0, t_f]. \tag{1}$$

We assume that there is a partition of $[0, t_f]$, $t_0 = 0 \leq t_1 \leq \ldots \leq t_n = t_f$, for some $n \geq 0$, so that $\sigma(t) \equiv \sigma_i \in S$ is constant in each subinterval $[t_{i-1}, t_i)$, $i = 1, \ldots, n$. The sequence $(\sigma_1, \ldots, \sigma_n)$ is called the *switching sequence* and $(t_0, \ldots, t_{n-1})$ is called the *switching instants*[1].

C. Problem Statements

Let $z(t) = (q(t), \sigma(t))$ be a hybrid trajectory over the time interval $[0, t_f]$. Suppose that $(\sigma_1, \ldots, \sigma_n)$ is the switching sequence associated with $\sigma(t)$. Thus n is the number of switchings during $[0, t_f]$. Since $\sigma(t) \in S$, $p_{\sigma(t)}$ is the instantaneous power of X at time t. The energy associated with $z(t)$ can be written as

$$E_\sigma = \int_0^{t_f} p_{\sigma(t)} dt + n k_s + p_m Q \cdot t_f.$$

The three terms on the right hand side of the above equation represent the *running energy*, namely, the energy

[1]The system is turned on at $t = 0$. Hence, we assume that there is always a switching at $t = 0$. We ignore the switching, if any, at $t = t_f$ for all trajectories.

consumed by X, the *switching energy*, and the *buffer energy*, respectively. The average power of the system during $[0, t_f]$ is

$$\bar{P}(\sigma, Q) = \frac{E_\sigma}{t_f} = \frac{1}{t_f}\left(\int_0^{t_f} p_{\sigma(t)} dt + nk_s\right) + p_m Q. \quad (2)$$

In this paper, we study the power consumption of the whole process of transferring a certain amount of data from X to Y. It is thus required that the system must start with an empty buffer at $t = 0$ and end up with an empty buffer at $t = t_f$ when Y have received all the data produced by X. This yields two boundary conditions for the continuous state, namely, $q(0) = 0$ and $q(t_f) = 0$. Hence, minimizing the average power can be formulated as the following optimal control problem.

Problem 1: Find a switching strategy $\sigma(t)$ over $[0, t_f]$ and a proper buffer size Q that

$$\text{Minimize} \quad \bar{P}(\sigma, Q) = \frac{1}{t_f}\left(\int_0^{t_f} p_{\sigma(t)} dt + nk_s\right) + p_m Q$$

$$\text{Subject to} \quad \max_{t\in[0,t_f]} q(t) \le Q, \text{ and } \min_{t\in[0,t_f]} q(t) \ge 0, \quad (3)$$

$$\frac{dq(t)}{dt} = r_{\sigma(t)} - r_0, \text{ with } q(0) = q(t_f) = 0. \quad (4)$$

D. Problem Simplification

Problem 1 can be greatly simplified using some particular features of the hybrid system H. Note that since $\sigma(t)$ is piecewise constant, so are $r_{\sigma(t)}$ and $p_{\sigma(t)}$. Therefore, the integral term in (2) reduces to a sum. Let $(\sigma_1, \ldots, \sigma_n)$ and $(t_0, \ldots, t_{n-1})$ be the switching sequence and switching instants of $z(t)$, respectively. Then (2) is equivalent to

$$\bar{P}(\sigma, Q) = \frac{1}{t_f}\left(\sum_{i=1}^{n} p_{\sigma_i}\Delta t_i + nk_s\right) + p_m Q,$$

where $\Delta t_i = t_i - t_{i-1}$, for $i = 1, \ldots, n$. Since the solution $q(t)$ of equation (4) is piecewise linear, the constraint in (3) is equivalent to

$$0 \le q(t_m) = \sum_{i=1}^{m}(r_{\sigma_i} - r_0)\Delta t_i \le Q, \ m = 1, \ldots, n-1,$$

$$\text{and} \quad \sum_{i=1}^{n}(r_{\sigma_i} - r_0)\cdot\Delta t_i = 0.$$

In other words, to guarantee that the entire trajectory stays inside $[0, Q]$ during the interval $[0, t_f]$, it is sufficient to require that $q(t)$ lies in $[0, Q]$ at every switching instant, due to the piecewise linearity of $q(t)$.

In many real-world applications, t_f is very large and we are interested in periodic switching strategies that can be easily implemented in computers. Hence, we adopt the following assumption to further simplify our problem.

Assumption 1: Assume that $z(t)$ is periodic and t_f is infinity.

In other words, in this paper we only focus on periodic solutions with infinite time horizon to Problem 1. For convenience, we redefine $(\sigma_1, \ldots, \sigma_n)$ and $(t_1, \ldots, t_{n-1})$ as the switching sequence and switching instants during the first period of $z(t)$.

A hybrid trajectory $z(t) = (q(t), \sigma(t))$ over $[0, \infty)$ is called *periodic* with period T if $q(t + T) = q(t)$ and $\sigma(t+T) = \sigma(t)$ for all $t \in [0, \infty)$. For such $z(t)$ the average power is equal to the average power during the first period $[0, T]$. Hence, the average power to be minimized reduces to:

$$\bar{P}(\sigma, Q, T) = \frac{1}{T}\left(\sum_{i=1}^{n} p_{\sigma_i}\Delta t_i + nk_s\right) + p_m Q. \quad (5)$$

Note that by limiting ourselves to the class of periodic solutions in finding the optimal ones, we introduce a new decision variable: the period T. Also, since we require $q(0) = 0$, we must have $q(T) = 0$, i.e., at the end of each period, continuous state q must come back to zero. Therefore, Problem 1 reduces to the following problem.

Problem 2: Find a proper buffer size Q and a periodical switching strategy $\sigma(t)$ with a period T that

$$\text{Minimize} \quad \bar{P}(\sigma, Q, T) = \frac{1}{T}\left(\sum_{i=1}^{n} p_{\sigma_i}\Delta t_i + nk_s\right) + p_m Q$$

$$\text{Subject to} \quad 0 \le \sum_{i=1}^{m}(r_{\sigma_i} - r_0)\cdot\Delta t_i \le Q,$$

$$\text{for } m = 1, \ldots, n-1, \text{ and } \sum_{i=1}^{n}(r_{\sigma_i} - r_0)\cdot\Delta t_i = 0. \quad (6)$$

III. Optimal Periodic Solutions

A. Necessary Conditions

A solution to Problem 2 consists of three parts, the hybrid trajectory $z(t) = (q(t), \sigma(t))$, the buffer size Q and the period T. If $((q(t), \sigma(t)), Q, T)$ is an optimal solution, then we have

$$\min_{t\in[0,T]} q(t) = 0, \quad \text{and} \quad \max_{t\in[0,T]} q(t) = Q. \quad (7)$$

The first equality in (7) is due to the constraints that $q(t) \ge 0$ and $q(0) = 0$. The second equality is also straightforward. To see this, suppose that $\max_{t\in[0,T]} q(t) < Q$. Let

$$\bar{Q} = \max_{t\in[0,T]} q(t).$$

Then $((q(t), \sigma(t)), \bar{Q}, T)$ satisfies all the constraints in (3) and (4), but consumes a less average power.

Lemma 1 (Tightness Condition): If $((q(t), \sigma(t)), Q, T)$ is an optimal solution to Problem 2, $q(t)$ must touch the boundary of its domain, i.e., $q(t)$ must satisfy equation (7).

The necessary condition given in the following lemma is the key result of this paper, and can be used to derive the optimal solutions to Problem 2.

Lemma 2 (Fundamental Lemma): Let $(t_0, \ldots, t_{n-1})$ be the switching instants corresponding to the first period of a trajectory $z(t)$. If $(z(t), Q, T)$ is an optimal solution, then $q(t_i) = 0$ or Q, for $i = 0, 1, \ldots, n-1$. In other words, the optimal solution only switches when the continuous state $q(t)$ hits the boundary of its domain $[0, Q]$.

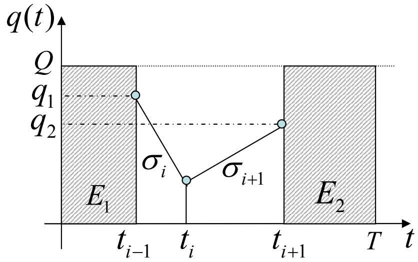

(a) Switching at Interior Point

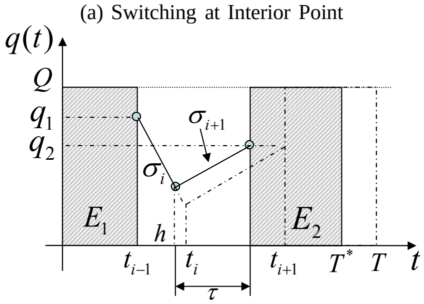

(b) Variation on t_i

Fig. 3. Scheme of Variation

Proof: The basic idea of the proof is as follows. First we assume that the i^{th} switching of an arbitrary trajectory $z(t)$ occurs at t_i for which $q(t_i)$ is an interior point of $[0, Q]$. Next we define a new trajectory $z^*(t; h)$ as shown in Fig. 3-(b), which is obtained from $z(t)$ by perturbing t_i to h and properly translating the other parts of $z(t)$ to the left or right so that the parts of the trajectory $z^*(t; h)$ before the $(i-1)^{th}$ switching and after the $(i+1)^{th}$ switching are kept the same as $z(t)$. In this way, by varying the variable h in the neighborhood around t_i, we can obtain different trajectories corresponding to different average powers $\bar{P}^*(h)$. Finally, we show that $\frac{d\bar{P}^*(h)}{dh}$ has a constant sign for all h near t_i, which indicates that for some h sufficiently close to t_i, $z^*(t; h)$ is a better solution than $z(t)$.

Following the above idea, let $(z(t), Q, T)$ be a solution to Problem 2 with switching sequence $(\sigma_1, \dots, \sigma_n)$ and switching instants $(t_0, \dots, t_{n-1})$. Suppose that $z(t)$ has a switching at some interior point of $[0, Q]$, i.e., $0 < q(t_i) < Q$ for some i. Define

$$E_1 = \sum_{j=1}^{i-1} p_{\sigma_j}(t_j - t_{j-1}),$$

$$\text{and} \quad E_2 = \sum_{j=i+2}^{n} p_{\sigma_j}(t_j - t_{j-1}).$$

As shown in Fig. 3-(a), E_1 and E_2 represent the running energy during the time interval $[0, t_{i-1})$ and $[t_{i+1}, t_n)$, respectively. Any variation on t_i will affect the trajectory after t_i. We want to study a particular class of variations on t_i that changes the times spent in mode σ_i and σ_{i+1} in such a way that the shape of the rest of the trajectory is kept the same. To this end, let $q_1 = q(t_{i-1})$, $q_2 = q(t_{i+1})$ and h be a parameter representing the new switching instant from mode σ_i to mode σ_{i+1}. Define the new switching strategy as a function of the variational parameter h as

$$\sigma^*(t; h) = \begin{cases} \sigma(t) & 0 \le t < h \\ \sigma_{i+1} & h \le t < h + \tau \\ \sigma(t + t_{i+1} - t_i - \tau) & h + \tau \le t \le T^* \end{cases}$$

where

$$\tau = \frac{q_2 - q_1 - (h - t_{i-1})(r_{\sigma_i} - r_0)}{r_{\sigma_{i+1}} - r_0},$$

and the new period T^* is

$$T^* = T + (h - t_{i-1} + \tau) - (t_{i+1} - t_{i-1}).$$

The corresponding new continuous state $q^*(t; h)$ can be obtained by:

$$\frac{dq^*(t; h)}{dt} = r_{\sigma^*(t;h)} - r_0.$$

Intuitively speaking, $(q^*(t; h), \sigma^*(t; h))$ is obtained from $(q(t), \sigma(t))$ by changing t_i to h and, at the time when the new continuous state evolves to q_2, concatenating the rest of the original trajectory to the new one. Refer to Fig. 3-(b) for a graphical illustration. In this way, the energy consumptions during the first $i-1$ modes and the last $n-i-1$ modes are kept the same and Q is also unchanged. Thus the average power associated with $((q^*(t; h), \sigma^*(t; h)), Q, T^*)$ can be written as

$$\begin{aligned} \bar{P}^*(h) = &\frac{1}{T^*}\big[E_1 + E_2 + nk_s + p_{\sigma_i}(h - t_{i-1}) \\ &+ p_{\sigma_{i+1}}\tau\big] + p_m Q. \end{aligned}$$

Taking the derivative of $\bar{P}^*(h)$ with respect to h, we have

$$\begin{aligned} \frac{d\bar{P}^*(h)}{dh} = &\frac{1}{(T^*)^2}\Bigg[\left(p_{\sigma_i} - p_{\sigma_{i+1}}\frac{r_{\sigma_i} - r_0}{r_{\sigma_{i+1}} - r_0}\right) \\ &\cdot \left(T - (t_{i+1} - t_{i-1}) + \frac{q_2 - q_1}{r_{\sigma_{i+1}} - r_0}\right) - \left(1 - \frac{r_{\sigma_i} - r_0}{r_{\sigma_{i+1}} - r_0}\right) \\ &\cdot \left(E_1 + E_2 + nk_s + p_{\sigma_{i+1}}\frac{q_2 - q_1}{r_{\sigma_{i+1}} - r_0}\right).\Bigg] \end{aligned} \tag{8}$$

Note that the h-related terms in the numerator have been canceled out. It is clear that the sign of $\frac{\bar{P}^*(h)}{dh}$ does not depend on h, which indicates that $\bar{P}^*(h)$ is monotone with respect to h. Thus the minimum power $\bar{P}^*(h)$ can only occur at the boundary of its feasible domain of h. Since we assume that $q(t_i) \in (0, Q)$, there exist $\epsilon_1 > 0, \epsilon_2 > 0$ such that $\forall h \in [t_i - \epsilon_1, t_i + \epsilon_2]$, the trajectory $q^*(t; h) \in [0, Q]$ for any $t \in [0, T^*]$. Therefore, $h = t_i$ cannot be optimal. In other words, if $q(t_i) \in (0, Q)$, for either $h = t_i - \epsilon_1$ or $h = t_i + \epsilon_2$, the new solution $(q^*(t; h), \sigma^*(t; h), Q, T^*)$ satisfies the constraints in (6) but corresponds to a less average power. These two extreme cases are plotted in Fig. 4. Thus we can conclude that the optimal solution can only switch when $q(t)$ is 0 or Q. ■

Lemma 3: There exists an optimal solution to Problem 2 that contains exactly two switchings ($n = 2$) in each period.

Proof: Suppose that $z(t) = (q(t), \sigma(t))$ is an optimal solution to Problem 2 with an average power $\bar{P}$ and period

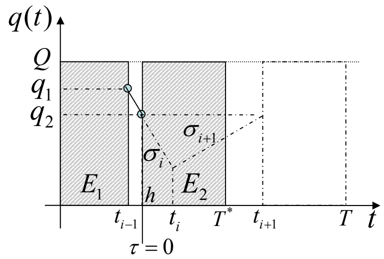

(a) Degenerate to no switch

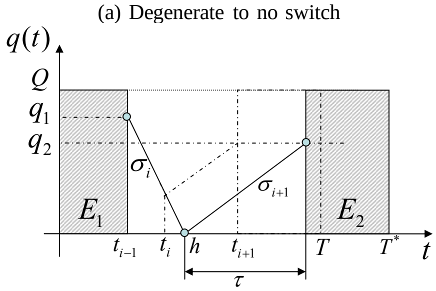

(b) Switch at the boundary

Fig. 4. Two Extreme Cases of Variations on t_i

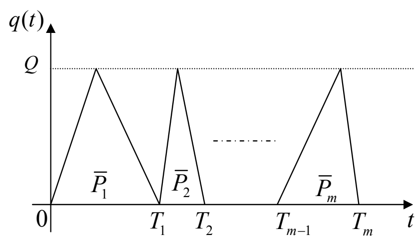

Fig. 5. Example for Lemma 3

T_m. Since $z(t)$ can not switch at any interior point of $[0, Q]$ (Lemma 2), it must be bouncing up and down between Q and 0 as shown in Fig. 5. Let T_i, $i = 0, \ldots, m$, be the times for which $q(T_i) = 0$. If $m = 1$, there is nothing to prove since $z(t)$ itself is optimal with exactly two switchings in each period. Thus we can assume $m > 1$. Denote by $\{z(t)\}_{[i]}$, $i = 1, \ldots, m$, the part of $z(t)$ within the interval $[T_{i-1}, T_i)$ and by $\bar{P}_i$ the average power of $\{z(t)\}_{[i]}$, $i = 1, \ldots, m$. Let $i^* = \arg\min_i \bar{P}_i$. Then it is obvious that $\bar{P} \geq \bar{P}_{i^*}$. Define $\hat{z}(t)$ as the periodic extension of $\{z(t)\}_{[i^*]}$ with average power $\hat{P}$. Then $\hat{P} = \bar{P}_{i^*} \leq \bar{P}$. Thus $\hat{z}(t)$ must also be an optimal solution, as we assume $z(t)$ is optimal. By definition, $\hat{z}(t)$ has exactly two switchings in each period, i.e., $n = 2$. Hence, we can conclude that there always exists an optimal periodic solution to Problem 2 with exactly two switchings in each period. ∎

B. Optimal Solutions

Lemma 3 indicates that we can obtain an optimal solution by only checking all the possible 2-mode switching sequence, namely, the switching sequences with $n = 2$. For each such switching sequence, since $q(t)$ can only switch when $q(t) = 0$ or Q, the average power over a period becomes a function that only depends on Q. Therefore, we can optimize the average power with respect to Q for each 2-mode switching sequence, and compare the optimal powers obtained for different sequences to get an optimal solution to Problem 2. If X has N modes, there are at most $N(N-1)/2$ sequences to compare. Thus if the optimal buffer size Q for each given 2-mode switching sequence can be derived analytically, the total computation cost of finding the optimal solution will be trivial. The following theorem gives an analytical expression of the optimal Q for each given 2-mode switching sequence and characterizes analytically an optimal solution to Problem 2.

Theorem 1: An optimal solution $((q(t), \sigma(t)), Q, T)$ to Problem 2 is given by

$$Q = \sqrt{\frac{k_s(r_{\sigma_1} - r_0)(r_0 - r_{\sigma_2})}{p_m(r_{\sigma_1} - r_{\sigma_2})}}$$
$$T = \frac{Q}{r_{\sigma_1} - r_0} + \frac{Q}{r_0 - r_{\sigma_2}}$$
$$\sigma(t) = \begin{cases} \sigma_1 & 0 \leqslant t < Q/(r_{\sigma_1} - r_0) \\ \sigma_2 & Q/(r_0 - r_{\sigma_2}) \leqslant t < T \end{cases}$$
$$q(t) = \int_0^t r_{\sigma(\tau) - r_0} d\tau,$$

where (σ_1, σ_2) is defined by

$$(\sigma_1, \sigma_2) = \arg\min_{(i,j) \in I \times J} \left(\frac{(r_i - r_0)p_j + (r_0 - r_j)p_i}{r_i - r_j} \right.$$
$$\left. + 2\sqrt{\frac{p_m k_s (r_i - r_0)(r_0 - r_j)}{r_i - r_j}} \right).$$

Proof: By Lemma 3, there exists an optimal solution that contains two switchings in each period. Suppose that the first and the second mode are i and j, respectively. Considering the constraint in (6), we must have $i \in I$ and $j \in J$. For a fix pair (i, j), the period T can be computed as

$$T = \frac{Q}{r_i - r_0} + \frac{Q}{r_0 - r_j},$$

and the average power over one period is

$$\bar{P} = \frac{1}{T}\left(\frac{p_i Q}{r_i - r_0} + \frac{p_j Q}{r_0 - r_j} + k_s \right) + p_m Q. \quad (9)$$

Taking the derivative of (9) with respect to Q and setting it to zero, we obtain the optimal buffer size in terms of i and j as:

$$Q = \sqrt{\frac{k_s(r_i - r_0)(r_0 - r_j)}{p_m(r_i - r_j)}}. \quad (10)$$

Substitute (10) back to (9), we have

$$\bar{P} = \frac{(r_i - r_0)p_j + (r_0 - r_j)p_i}{r_i - r_j}$$
$$+ 2\sqrt{\frac{p_m k_s (r_i - r_0)(r_0 - r_j)}{r_i - r_j}}. \quad (11)$$

The optimal $\bar{P}$ can be obtained by minimizing (11) with respect to (σ_1, σ_2). Hence,

$$(\sigma_1, \sigma_2) = \arg\min_{(i,j)\in I\times J} \left(\frac{(r_i - r_0)p_j + (r_0 - r_j)p_i}{r_i - r_j} + 2\sqrt{\frac{p_m k_s (r_i - r_0)(r_0 - r_j)}{r_i - r_j}} \right). \quad (12)$$

■

IV. Simulation

Our theoretical results can be applied in many real-world applications, such as the power management problem of a multiple-speed disk [7] and the dynamic voltage scheduling (DVS) problem of a variable speed processor [5]. In this section, we use a DVS example to illustrate the effectiveness of our result.

TABLE I

PROCESSOR PARAMETER

Mode	1	2	3	4	5
Voltage (v)	0	0.9	1.1	2.5	3.3
Speed (Kbps)	0	3000	3667	8333	11000
Power (w)	0	0.15	0.22	1.16	2.01

The power of a processor is approximately proportional to the square of its supply voltage; and DVS tries to reduce the system power consumption by appropriately arranging the supply voltages for a given task [6]. In our simulation, the interacting components X and Y are a processor and a video, respectively. Suppose that the video consists of 45000 frames; each is of size 300Kb. As a strict requirement, the video data must be processed by X at a constant speed 25 frame/sec, or equivalently, 7500Kb/sec. We assume that the processor has five power modes and the voltage, speed, and power in each mode are given in Table I. In addition, we suppose that the power parameter of the buffer is $p_m = 8.24 * 10^{-8}$ W/B and the switching cost k_s varies from 1mJ to 100mJ. One way to reduce the power consumption is to use the highest speed to finish a frame and turn off the processor for the rest of the frame period. We refer this method as Scheme 1, while Scheme 2 represents the proposed method which schedules the voltages according to the formulas in Theorem 1. In Figure 6, for each scheme we plot the corresponding energy consumption of processing the whole video as a function of the switching costs. The result is normalized with respect to the energy cost of Scheme 1 for each switching cost. It is clear that the proposed method can save at least 30% of energy compared with the heuristic method and more energy can be saved as the switching cost increases. In practice, the switching cost is application dependent [5]. The proposed method can guarantee a minimum energy consumption for any switching cost.

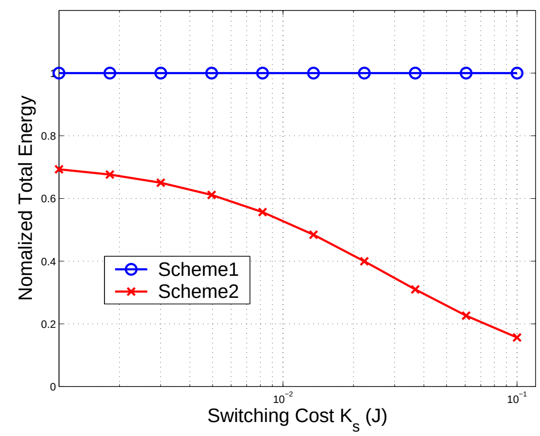

Fig. 6. Simulation Result

V. Conclusion

This paper studies the optimal power management problem of two interacting components with a buffer. We assume that the component to be controlled has multiple power modes and model the overall system as a hybrid system. The power management problem is formulated as an optimal control problem and the optimal buffer size and the optimal switching strategy are derived analytically. Future research will focus on the case where the data consumption rate r_0 is time varying or random.

References

[1] L. Benini and G. D. Micheli. System level power optimization: Techiniques and tools. *ACM Transactions on Design and Automation of Electronic Systems*, 5(2):115–192, April 2000.

[2] 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*, pages 141–152, February 2005.

[3] J. Hu and Y.-H. Lu. Buffer management for power reduction using hybrid control. In *Proc. IEEE Int. Conf. Decision and Control*, pages 6997–7002, Seville, Spain, 2005.

[4] J. Ridenour, J. Hu, and Y.-H. Lu. Low power buffer management using hybrid control. In *Proc. American Control Conference*, 2006.

[5] T. D. Burd and R. W. Brodersen. Design issues for dynamic voltage scaling. In *Proceedings of the international symposium on Low power electronics and design*, pages 9–14. ACM Press, 2000.

[6] T. Ishihara and H. Yasuura. Voltage scheduling problem for dynamically variable voltage processors. In *Proceedings of the international symposium on Low power electronics and design*, pages 197 – 202. ACM Press, 1998.

[7] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir, and H. Franke. Drpm: Dynamic speed control for power management in server class disks. In *Proceedings of the International Symposium on Computer Architecture (ISCA)*, pages 169–179, June 2003.

[8] R. Alur, T. Henzinger, G. Lafferriere, and G. J. Pappas. Discrete abstractions of hybrid systems. *Proceedings of the IEEE*, 88(2):971–984, 2000.

[9] S. Hedlund and A. Rantzer. Optimal control of hybrid system. In *Proceedings of the IEEE Conference on Decision and Control*, volume 4, pages 3972–3977, Phoenix, AZ, December 1999.

[10] M. Egerstedt, Y. Wardi, and F. Delmotte. Optimal control of switching times in switched dynamical systems. *IEEE Transactions on Automatic Control*, 51(1):110–115, 2006.

[11] X. Xu and P.J. Antsaklis. Optimal control of switched systems based on parameterization of the switching instants. *IEEE Transactions on Automatic Control*, 49(1):2–16, 2004.

[12] S. C. Bengea and R. A. DeCarlo. Optimal control of switching systems. *Automatica*, 41(1):11–27, 2005.

[13] Y.-H. Lu, L. Benini, and G. D. Micheli. Dynamic frequency scaling with buffer insertion for mixed workloads. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, pages 1284–1305, November 2002.