

Optimal Power-Down Strategies

John Augustine*
jea@ics.uci.edu

School of Information and Computer Science,
Univ. of California at Irvine, Irvine, CA 92697

Sandy Irani*
irani@ics.uci.edu

School of Information and Computer Science,
Univ. of California at Irvine, Irvine, CA 92697

Chaitanya Swamy†
swamy@cs.cornell.edu
Dept. of Computer Science,
Cornell University, Ithaca, NY 14853

Abstract

We consider the problem of selecting threshold times to transition a device to low-power sleep states during an idle period. The two-state case in which there is a single active and a single sleep state is a continuous version of the ski-rental problem. We consider a generalized version in which there is more than one sleep state, each with its own power consumption rate and transition costs. We give an algorithm that, given a system, produces a deterministic strategy whose competitive ratio is arbitrarily close to optimal. We also give an algorithm to produce the optimal online strategy given a system and a probability distribution that generates the length of the idle period. We also give a simple algorithm that achieves a competitive ratio of $3 + 2\sqrt{2} \approx 5.828$ for any system.

1. Introduction

Suppose you are about to go skiing for the first time in your life. Naturally, you ask yourself whether to rent skis or to buy them. Renting skis costs, say, \$30, whereas buying skis costs \$300. If you knew how many times you would go skiing in the future (ignoring complicating factors such as inflation, and changing models of skis), then your choice would be clear. If you knew you would go at least 10 times, you would be financially better off by buying skis right from the beginning, whereas if you knew you would go less than 10 times, you would be better off renting skis every time.

Alas, the future is unclear, and you must make a decision nonetheless.

Although the *Ski-Rental* problem is very simple abstraction, this basic paradigm arises in many applications in computer systems. In these situations, there is a system that can reside in either a low-cost or a high-cost state. Occasionally, it is forced to be in the high-cost state (usually to perform some task). A period between any two such points in time is called an *idle period*.

The system pays a per time unit cost to reside in the high-cost state. Alternatively, it can transition to the low-cost state at a fixed one-time cost. If the idle period is long, it is more advantageous to transition to the low cost state immediately. If the idle period is short, it is better to stay in the high-cost state. An online algorithm which does not know the length of the idle period must balance these two possibilities.

This problem has been studied in the context of shared memory multiprocessors in which a thread is waiting for a locked piece of data and must decide whether to spin or block [7, 8]. Researchers investigating the interface between IP networks and connection-oriented networks have discovered this same underlying problem in deciding whether to keep a connection open between bursts of packets that must be sent along the connection [9]. The problem also arises in cache coherency in deciding whether to update or invalidate data that has been changed in a processor's local cache [4, 2].

An important application of the ski-rental problem is in minimizing the power consumed by devices that can transition to a low power *sleep* state when idle. The sleep state consumes less power; however, one incurs a fixed start-up cost in making the transition to the high-power *active* state in order to begin work when a new job arrives. At the architectural level, the technique of eliminating power to a func-

* Research supported partially by NSF grant CCR-0105498 and by ONR Award N00014-00-1-0617.

† Research supported partially by NSF grant CCR-9912422.

tional component is called clock/power gating. At a higher level, the powered-down component might be a disk drive or even the whole system (e.g., a laptop that hibernates). The embedded systems community has invested a great deal of effort into devising policies governing the selection of power states during idle periods (termed *Dynamic Power Management* in their literature). See [3] for a survey. These techniques have been critical to maximizing battery use in mobile systems. While power is already a first-class parameter in system design, it will become increasingly important in the future since battery capacities are increasing at a much slower rate than power requirements.

Most of the previous work on this problem has been concerned with two-state systems which have an active state and a single sleep state. This paper focuses on finding power-down thresholds for systems that have more than one low-power state. An example of such a system is the Advanced Configuration and Power Interface (ACPI) included in the BIOS on most newer computers which has five power states, including hibernation and three levels of standby [1].

2. Previous Work and New Results

For the two-state problem, an online algorithm consists of a single threshold T after which time the algorithm will transition from the active to the sleep state. The input to the problem is the length of the idle period and the cost of an algorithm is the total amount of energy it consumes over a single idle period. Typically, an online algorithm is evaluated in terms of its competitive ratio — the ratio of the cost of the online algorithm to the cost of the optimal offline algorithm, maximized over all inputs. When randomized algorithms are considered, we look at the ratio of the expected cost of the online algorithm to the cost of the offline algorithm. Previous work has also addressed this problem when the idle period is generated by a known probability distribution. In this case, the online algorithm will choose a threshold which minimizes its expected cost. We call such algorithms *probability-based* algorithms.

The best deterministic online algorithm will stay in the high power state until the total energy spent is equal to the cost to power up from the low power state. It is known that this algorithm achieves the optimal competitive ratio of two. If the idle period is generated by a known probability distribution, then the algorithm that minimizes its expected cost is always within a factor of $e/(e-1)$ of optimal. Furthermore, this bound is tight since there is a distribution over the idle period lengths which will force any online algorithm to consume an expected $e/(e-1)$ times more energy than the optimal offline algorithm [7].

Note that in the context of power-down systems, it may not be the case that the power usage in the sleep state is zero or even that the start-up cost in the active state is zero.

In these cases, both the online and the offline algorithm will have an identical additional cost. Thus, the ratio of the online to the offline cost will decrease and the optimal competitive ratio will be strictly less than two. These additional costs do not change the optimal online or offline strategy in either the deterministic or the probability-based case. The optimal competitive ratio that can be achieved for such systems can easily be determined as a function of all the parameters of the system.

The problem with multiple sleep states is formalized as follows: the sequence of states is denoted $\langle s_0, \dots, s_k \rangle$. There is also a vector of power-consumption rates $\mathcal{K} = \langle \kappa_0, \dots, \kappa_k \rangle$, where κ_i is the power consumption rate of the system in state s_i . We assume as a convention that the states are ordered so that $\kappa_i > \kappa_j$ for $0 \leq i < j \leq k$. Thus, the system must transition to s_0 (the *active* state) at the end of the idle period. There is an associated transition cost $d_{i,j}$ to move from state s_i to s_j . A *system* is described by a pair (\mathcal{K}, d) . Note that there can be costs to move from high-power states to low-power states and vice versa. However, the only power-up costs that are of interest are the costs to transition from a particular state s_i to the active state s_0 since the only reason to transition to a higher power state is when a new task arrives. A *schedule* or *strategy* $\mathcal{A} = (\mathcal{S}_{\mathcal{A}}, \mathcal{T}_{\mathcal{A}})$ consists of a sequence of $n_{\mathcal{A}}$ states $\mathcal{S}_{\mathcal{A}}$ that is a subsequence of S and a sequence of transition times $\mathcal{T}_{\mathcal{A}}$. Where obvious, we will omit the subscript \mathcal{A} . We require that $\mathcal{S}(0) = s_0$ and $\mathcal{T}(0) = 0$. $\mathcal{A}(t)$ will denote the cost of the schedule produced by strategy \mathcal{A} for an idle period of length t .

The power-down problem for systems with multiple states has been examined for the special case where the cost to power-down is zero and the algorithm only pays to move from low power states to higher power states [5]. Note that this also includes the case where the transition costs are additive ($d_{i,j} + d_{j,k} = d_{i,k}$ for $i < j < k$) since the costs to power down can then be folded into the costs to power up. Generalizations of the algorithms for the two-state case are given along with upper bounds on the competitive ratio of these algorithms for multi-state systems with additive transition costs. Namely, it is shown that the generalized deterministic algorithm is 2-competitive and the probability-based algorithm is $e/(e-1)$ -competitive.

There are two important directions left open in this work. The first is based on the observation that systems, in general, do not have additive transition costs. In many scenarios, additional energy is spent in transitioning to lower power states. Furthermore, there is overhead in stopping at intermediate states, i.e. the transition costs are not additive. (See [3] for an example.) The second point is that these upper bounds are typically not optimal. For multi-state systems, the optimal competitive ratio that can be achieved will, in

general, be a complicated function of all the parameters of the system (the power consumption rates as well as transition costs). While it may not be feasible to express the optimal competitive ratio as a function of all these parameters, a system designer would, in general, like to design a power-down strategy that obtains the best possible competitive ratio given the constraints of his or her particular system.

This paper establishes the following results.

- We give an algorithm that takes as input a description of a system (\mathcal{K}, d) , and an error parameter ϵ , and produces a power-down strategy $\mathcal{A} = (\mathcal{S}_{\mathcal{A}}, \mathcal{T}_{\mathcal{A}})$ whose competitive ratio is within ϵ of the best competitive ratio that can be achieved for that system. The algorithm runs in time $O(k^2 \log k \log(1/\epsilon))$, where $k + 1$ is the number of states in the system, and also outputs the competitive ratio of \mathcal{A} . The algorithm works via a decision procedure which determines for a system and a constant ρ if there is a ρ -competitive strategy for that system. We use this decision procedure to also obtain lower bounds for specific systems. In particular, we show a lower bound of 2.45 for the best competitive ratio that can be achieved for a particular system. This is the first lower bound known that is greater than 2.
- Experimental results show that there are significant performance gains to be made by estimating the distribution governing the length of an idle period based on recent history and using this estimate to drive a probability-based strategy [6]. We give an algorithm that takes as input a description of a system and a probability distribution generating the idle period length and produces the optimal power-down strategy. Naturally, the running time of the algorithm will depend on the representation of the distribution. In practice, this is most likely to be a histogram. Our algorithm runs in time $O(k^2(\log k + B))$ where B is the number of bins in the histogram and $k + 1$ is the number of states. One outcome of the proof is that it also establishes the optimality of the strategy given in [5] for *additive* systems.
- We give a deterministic strategy that achieves a competitive ratio of $3 + 2\sqrt{2} \approx 5.8284$ for all systems. This result gives a bound on the competitive ratio achieved by the optimal strategies generated by our algorithms. Note that $3 + 2\sqrt{2}$ also serves as a bound on the ratio of the expected costs of the online and offline when the input is probabilistically generated.

For the remainder of this paper, we use the terms *schedule* or *strategy* to refer to the choices of states and threshold times for powering down. The term *algorithm* will re-

fer to a procedure that produces a schedule or strategy based on a particular system.

3. Preliminaries

First we will establish that we can assume without loss of generality that the power-up transition costs are zero. If this is not the case for some system (\mathcal{K}, d) , we can define a new system such that for any $i < j$, the cost to transition from i to j is $d_{i,j} + d_{j,0} - d_{i,0}$ and the cost to go from j to i is 0. Any set of actions in the original system will incur the same cost in the new system. For the remainder of this paper, we will assume that $d_{i,0} = 0$ for all i .

Let $D(i)$ denote $d_{0,i}$. Then $OPT(t) = \min_i (D(i) + \kappa_i t)$. Let $S(t)$ denote the state which attains the minimum — the optimal state. The optimal strategy is to transition to state $S(t)$ at time 0, and stay there through time t . We assume that the optimal strategy will actually “use” every state, i.e., $range(S(t)) = \{s_0, \dots, s_k\}$. None of the online strategies we present will make use of a state that is not used by the offline strategy for some t .




Figure 1. Energy consumed by the optimal strategy as a function of idle period length.

Note that $S(t)$ is non-decreasing with t . As the idle period length gets longer, it becomes more worthwhile to pay the extra cost to transition to a lower power state. Let b_i denote the first time instant at which state i becomes the optimal state, so $b(0) = 0$ and $D(i - 1) + \kappa_{i-1} b_i = D(i) + \kappa_i b_i \Rightarrow b_i = \frac{D(i) - D(i-1)}{\kappa_{i-1} - \kappa_i}$. We have $b(0) < b(1) < \dots < b(k)$, so we can write $OPT(t) = D(i) + \kappa_i t$ for $t \in [b_i, b_{i+1}]$. We compare our online strategy with $OPT(t)$ and want to get a strategy \mathcal{A} which minimizes the competitive ratio, $c_{\mathcal{A}} = \sup_t \frac{\mathcal{A}(t)}{OPT(t)}$ where $\mathcal{A}(t)$ denotes the total power consumption of \mathcal{A} by time t . Figure 1 shows the total energy consumed by OPT as a function of the length of the idle period. There is a line for each state. The y -intercept is the transition cost to move to that state from the active

state and the slope is the power consumption rate. The energy consumed by the optimal strategy is the lower envelope of these lines since it will pick the single state which minimizes the cost for a given idle period length.

4. A $(3 + 2\sqrt{2})$ -competitive strategy

First we establish that we can assume that for all $i < j$, $d_{i,j} < d_{0,j}$. Recall that we are really using $d_{i,j}$ to denote $d_{i,j} + d_{j,0} - d_{i,0}$ and $d_{0,j}$ to denote $d_{0,j} + d_{j,0}$. Thus, the assumption that $d_{i,j} < d_{0,j}$ really amounts to assuming that $d_{i,j} < d_{i,0} + d_{0,j}$. If this were not the case, we could just transition from state s_i to state s_j by first going to s_0 and then down to s_j .

Let us for the moment assume that $D(i) \geq \gamma D(i - 1)$ for all $i = 1, \dots, k$ and for some $\gamma > 1$. This is a non-trivial assumption that we will have to handle later. Consider the strategy, \mathcal{A} , which always stays in state $S(t)$, the same state as OPT , at every time t . This is the strategy proposed in [5] and shown to be 2-competitive for additive systems. Note that this strategy is the same as the 2-competitive balance strategy for the two-state case.

For $t \in [b_i, b_{i+1}]$ the online cost is, $\mathcal{A}(t) = \sum_{j=0}^{i-1} (\kappa_j(b_{j+1} - b_j) + d_{j,j+1}) + \kappa_i(t - b_i)$, and the ratio $\frac{\mathcal{A}(t)}{OPT(t)}$ will be maximized at $t = b_i$. We have,

$$\begin{aligned} \mathcal{A}(b_i) &= \sum_{j=0}^{i-1} (\kappa_j(b_{j+1} - b_j) + d_{j,j+1}) \\ &\leq \sum_{j=0}^{i-1} \kappa_j(b_{j+1} - b_j) + \sum_{j=1}^i D(j) \\ &\leq OPT(b_i) + D(i) \sum_{j=1}^i \gamma^{-(i-j)} \\ &\leq \left(1 + \frac{\gamma}{\gamma-1}\right) OPT(b_i) \\ &= \frac{2\gamma-1}{\gamma-1} \cdot OPT(b_i). \end{aligned} \quad (1)$$

This holds for any t implying a competitive ratio of $\frac{2\gamma-1}{\gamma-1}$.

Now suppose the assumption $D(i) \geq \gamma D(i - 1)$ does not hold. We consider a new *offline* strategy OPT' that only uses a subset of states S' for which the property does hold, and is a γ -approximation of OPT , i.e., $OPT'(t) \leq \gamma \cdot OPT(t)$. We now view our problem as specified by just the states in S' , and execute strategy \mathcal{A} as specified above, emulating OPT' instead of OPT . We get that $\mathcal{A}'(t) \leq \frac{2\gamma-1}{\gamma-1} OPT'(t) \leq \frac{\gamma(2\gamma-1)}{\gamma-1} OPT(t)$. Setting $\gamma = 1 + \frac{1}{\sqrt{2}}$, we get a competitive ratio of $3 + 2\sqrt{2} \approx 5.8284$.

We determine OPT' as follows. Let $S' = \{s_k\}$ initially. Consider the states in S in reverse order. Let s_i be the last state added to S' . We find the largest j , $0 \leq j < i$ s.t.

$D(j) \leq D(i)/\gamma$. We add s_j to S' and continue until no such j exists. Note that $s_0 \in S'$ since $D(0) = 0$. OPT' will execute the optimal offline strategy assuming that only the states in S' are available. Consider i, j s.t. $s_i, s_j \in S'$ and no s_ℓ is in S' for $i < \ell < j$. We have $OPT'(t) = OPT(t)$ for $t \in [b_i, b_{i+1})$ and $t \in [b_j, b_{j+1})$. For ℓ s.t. $i < \ell < j$ and time $t \in [b_\ell, b_{\ell+1})$. $OPT'(t) = \min(D(i) + \kappa_i t, D(j) + \kappa_j t)$ and $OPT(t) = D(\ell) + \kappa_\ell t > OPT'(t)/\gamma$ since $D(\ell) > D(j)/\gamma$, $\kappa_\ell > \kappa_j$. So OPT' is a γ -approximation to OPT .

Theorem 1 *There is a $(3 + 2\sqrt{2})$ -competitive strategy for any system.*

5. A Near Optimal Deterministic Algorithm

In this section, we turn our attention to obtaining a near optimal schedule for a particular system. More precisely, given a system (\mathcal{K}, d) with state sequence S for which the optimal online schedule has competitive ratio ρ^* , we give an algorithm that returns a $(\rho^* + \epsilon)$ -competitive online schedule in time $O(k^2 \log k \log(1/\epsilon))$. The algorithm is based on a decision procedure which determines whether a ρ -competitive schedule exists for a given value of ρ . Theorem 1 establishes an upper bound of $3 + 2\sqrt{2}$ on the optimal competitive ratio, so we perform a bisection search in the range $[1, 3 + 2\sqrt{2}]$ to find the smallest ρ such that there exists a ρ -competitive schedule. We also output the resulting schedule.

The following lemma shows that the online strategy must eventually get to a sufficiently low-power state. Lemma 3 allows us to limit our concern to just the transition points in any online schedule. It is straightforward to show the following.

Lemma 2 *If $\mathcal{A} = (S, T)$ is a ρ -competitive strategy and s_ℓ is the last state in S , then $\kappa_\ell \leq \rho \cdot \kappa_k$.*

Proof: For the sake of contradiction, assume that $\kappa_\ell > \rho \cdot \kappa_k$. For \mathcal{A} to be ρ -competitive, it will have to lie entirely within the convex $\rho \cdot OPT$. However, in $\rho \cdot OPT$, the last line has a slope $\rho \cdot \kappa_k$, which will meet the last line in \mathcal{A} because it has larger slope of κ_ℓ , after which \mathcal{A} will start exceeding $\rho \cdot \kappa_k$. This is a contradiction. ■

Lemma 3 *If a schedule \mathcal{A} has finite competitive ratio, then the time t at which $\frac{\mathcal{A}(t)}{OPT(t)}$ is first maximized is a transition point in the strategy \mathcal{A} .*

Proof: For the sake of contradiction, assume that t is in the interval (t_1, t_2) , where t_1 and t_2 are two consecutive elements of $\mathcal{T}_\mathcal{A}$. If in case $t > \mathcal{T}(n_\mathcal{A} - 1)$, i.e., it occurs when the system is in sleep state according to \mathcal{A} , then assume $t_2 = \infty$. There are two possible cases that we analyze




Figure 2. Energy consumed by the online and optimal strategy as a function of idle period length. The solid line is $\rho \cdot OPT(t)$. The dashed line is the online cost. t is the first transition time that is not eager. t' shows the transformed strategy which now has an eager transition.

here. In the first case, t is not one of the transition points of OPT . Let ϵ be some infinitesimally small positive constant such that $[t - \epsilon, t + \epsilon]$ does not contain any transition points either in \mathcal{A} or OPT . Define $r(\tau) = \mathcal{A}(\tau)/OPT(\tau)$ for any time τ . We know that $r(t - \epsilon) < r(t)$. Since both the cost of \mathcal{A} and OPT are linear in the interval $[t - \epsilon, t + \epsilon]$, $r(t) < r(t + \epsilon)$, which is a contradiction. Now in the second case, t is a transition point in OPT . Hence OPT will start a new line in its cost curve starting at t , whose slope will be lower than the line before t and hence $r(t) < r(t + \epsilon)$. This, again, is a contradiction. ■

We now explore ways to restrict the space of schedules we need to consider in searching for a ρ -competitive schedule. For a strategy $\mathcal{A} = (S, T)$, we say that a transition at time $t \in T$ is ρ -eager (or just eager if ρ is clear from the context) if $\mathcal{A}(t) = \rho \cdot OPT(t)$. We say that \mathcal{A} is a ρ -eager strategy if $\mathcal{A}(t) = \rho \cdot OPT(t)$ for every $t \in T$. Note that by Lemmas 2 and 3, a ρ -eager strategy that ends at state s such that $\kappa_s \leq \rho \cdot \kappa_k$ is ρ -competitive.

Lemma 4 *If $\mathcal{A} = (S, T)$ is a ρ -competitive strategy, then there exists an eager strategy $\mathcal{A}' = (S, T')$ that is also ρ -competitive.*

Proof Sketch : Figure 2 shows a schematic of the proof. The jumps in the online cost (the dashed line) are transition costs. The solid line is $\rho \cdot OPT(t)$. The figure shows a transition time t at which the online cost is less than $\rho \cdot OPT(t)$. The idea is that we can slide such a transition time earlier until it hits the function $\rho \cdot OPT(t)$. Since the slope of the line before time t is steeper than after time t , that is, since we transition to state with lower power consumption rate at time t , transitioning earlier only decreases the cost. We do

this for each transition point from left to right to get an eager strategy. ■

Lemma 4 immediately gives an algorithm that is exponential in k , the number of states, and determines whether a ρ -competitive algorithm exists for the system. This algorithm enumerates all subsequences of states, and determines the ρ -eager strategy for that subsequence. Determining the ρ -eager strategy is done by finding the eager transition to each state based on the eager transitions to all the previous states in the sequence as described in the proof of Lemma 4. A ρ -competitive strategy for the system exists if and only if one of these ρ -eager strategies is ρ -competitive. The remainder of this section presents a way to remove the exponential dependence on k .

Let $S = \langle s_0, s_1, \dots, s_k \rangle$ be a sequence of states that form a system. Define $S_{s_i \rightarrow s_j}$ to be the subsequence $\langle s_i, \dots, s_j \rangle$, where s_i and s_j are elements of S such that $i < j$. Let Ψ_s be the set of subsequences of $S_{s_0 \rightarrow s}$ that include s_0 and s such that for each $\psi \in \Psi_s$, one can find transition times for the state sequence ψ so that in the resulting schedule, each transition up to and including the transition to state s is a ρ -eager transition. For a state $q \in \psi$, we will use $t_{\psi, q}$ to denote this ρ -eager transition time of q for the sequence ψ . (Note that ψ uniquely determines the transition times $t_{\psi, q}$.)

We define the *earliest* transition time $E(s, \rho)$ of state s for the given system as $E(s, \rho) = \min_{\psi \in \Psi_s} t_{\psi, s}$, that is, $E(s, \rho)$ is the earliest time at which any online strategy can transition to state s while remaining ρ -eager over all its transitions up to (and including) the transition to state s . Observe that if there is ρ -competitive strategy that uses state s , then by Lemma 4, there is such a ρ -early strategy, so $\Psi_s \neq \phi$ and $E(s, \rho)$ is well defined. We call a transition

to state s ρ -early if it happens at time $E(s, \rho)$ (or simply early). A strategy that consists entirely of early transitions is called an ρ -early strategy.

Lemma 5 *If there is a ρ -competitive strategy $\mathcal{A} = (S, T)$, then there is an eager and early ρ -competitive strategy.*

Proof: Let s be the last state in S . Consider the sequence $\psi \in \Psi_s$ such that $t_{\psi, s} = E(s, \rho)$ and the strategy π that uses only the states in ψ , transitioning to state $q \in \psi$ at time $t_{\psi, q}$, i.e., $\pi = (\psi, \{t_{\psi, q}\}_{q \in \psi})$. We will argue that π is an eager and early strategy, and since $\kappa_s \leq \rho \cdot \kappa_k$, π is ρ -competitive. Note that by definition, all transitions of π are ρ -eager, and the transition to state s is ρ -early. We have to show that the remaining transitions of π are also ρ -early.

Suppose not. Consider the latest transition that is not ρ -early. Suppose this happens for state r ($\neq s$), so $t_1 = t_{\psi, r} > E(r, \rho)$. Let r' be the state just after r in sequence ψ , let $T = t_{\psi, r'} = E(r', \rho)$. Let $\psi' \in \Psi_r$ be the sequence for which $t_{\psi', r} = E(r, \rho) = t_2$. Consider the hybrid strategy π' that uses the states in ψ' followed by the states in ψ that appear after r , with the transition times being $t_{\psi', q}$ for $q \in \psi'$ and $t_{\psi, q}$ for $q \in \psi_{r' \rightarrow s}$. For every time $t \geq t_1$, we have $\pi'(t) = \pi(t) - (\pi(t_1) - \pi'(t_1))$. We will show that $\pi'(t_1) < \pi(t_1)$ which implies, in particular, that $\pi'(T) < \pi(T)$. So in π' the transition to r' is no longer ρ -eager. Arguing as in Lemma 5 this means that we can shift the transition to r' to get an eager transition at an earlier time T' . But this contradicts the definition of $E(r', \rho)$ since we now have a sequence $\psi'' = \langle \psi, r' \rangle \in \Psi_{r'}$ with ρ -eager transition times such that $T' = t_{\psi'', r'} < T = E(r', \rho)$.

We now prove the claim. The points $(t_2, \pi'(t_2))$ and $(t_1, \pi(t_1))$ both lie on the $\rho \cdot OPT(t)$ curve. The slope of this curve at time $t = t_1$ is at least the slope κ_r of curve π at time t_1 since $\pi(t) \leq \rho \cdot OPT(t)$ for $t > t_1$. The slope of $\rho \cdot OPT(t)$ decreases with increasing t , so its slope is at least κ_r for all times $t \in [t_2, t_1]$. Hence, $\pi(t_1) = \rho \cdot OPT(t_1) \geq \rho \cdot OPT(t_2) + \kappa_r(t_1 - t_2) = \pi'(t_2) + \kappa_r(t_1 - t_2) = \pi'(t_1)$. The idea is illustrated in Figure 3. ■

From Lemma 5 we can deduce that we only need to consider a specific early and eager schedule, the one that is determined by the $E(\cdot, \rho)$ values, to determine if a ρ -competitive strategy exists. We now define procedure EXISTS that takes a system and a constant $\rho \in [1, 3 + 2\sqrt{2}]$. The output is a YES if a ρ -competitive strategy exists for the system. The algorithm can also be modified to give us the ρ -competitive strategy. If no such strategy exists, then the output is NO. We employ a dynamic programming approach to calculate $E(s_i, \rho)$, for $0 < i \leq k$. We always start with the high power state and hence $E(s_0, \rho) = 0$. Suppose we have computed $E(s_j, \rho)$ for all $j = 0, \dots, i-1$. Let t_j be the earliest time at which the system ρ -eagerly transitions from s_j to s_i given that the transition to s_j is




Figure 3. The solid line is $\rho \cdot OPT$. The dashed line is the schedule π' from Lemma 5 and the dashed/dotted line is π . The point labeled p_2 is $(t_2, \pi'(t_2))$ and p_1 is $(t_1, \pi(t_1))$. The idea is to show that at time t_1 , π' has a lower cost than π .

ρ -eager and occurs at time $E(s_j, \rho)$. If such a transition is not possible, then we assign $t_j = \infty$. We can compute t_j in $O(\log k)$ time by binary searching through the list of line segments in $\rho \cdot OPT$. Then, $E(s_i, \rho) = \min_{j < i} t_j$. Determining each $E(s_j, \rho)$ requires examining j different possibilities, so finding all the early transition times takes time $O(k^2 \log k)$. By Lemma 2, we know that if $E(s_i, \rho)$ is finite for some state s_i where $\kappa_i \leq \rho \cdot \kappa_k$, we know that a ρ -competitive strategy exists. We use the above procedure to search for $\rho \in [1, 3 + 2\sqrt{2}]$ within an error ϵ of ρ^* . The total time taken is $O(k^2 \log k \log(1/\epsilon))$. One can quickly elicit the schedule by starting from state k and retracing the states that minimized the earliest transition time.

6. A Probability-based Algorithm

Karlin *et al.* study the two-state case when the length of the idle period is generated by a known probability distribution p [7]. (Although they examined the problem in the context of the spin-block problem, their problem is identical to our two-state case.) They observed that the expected cost of the online strategy that makes the transition to the sleep state at time T is

$$\int_0^T p(t)(\kappa_0 t) dt + \int_T^\infty (\kappa_0 T + \kappa_1(t - T) + \beta) dt,$$

where κ_0 is the power consumption rate in the active state, κ_1 is the power consumption rate in the sleep state and β is the transition cost between the two states. The online strategy then should select the transition time T that minimizes this cost.

The multi-state case presents two distinct challenges. The first is to determine the optimal sequence of states

through which an online strategy should transition throughout the course of the idle period. Then once this sequence has been determined, the optimal transition times need to be determined. Our proof proceeds by establishing that the only transition times that need to be considered are the optimal transition times for two-states systems. Suppose, for example, that we are considering a sequence of state transitions in which state s_i is followed by state s_j . Let $T_{i,j}$ denote the optimal transition time from state s_i to s_j if these were the only two states in the system (that is, if s_i were the active state and s_j were the only sleep state). Note that $T_{i,j}$ can be determined by the expression above. We establish that regardless of the rest of the sequence, the optimal transition point from state s_i to s_j is $T_{i,j}$. We call the $T_{i,j}$'s the *pairwise-optimal* transition times.

Lemmas 6 and 7 establish that the pairwise-optimal transition times happen in the right order. That is for $i < k < j$, $T_{i,k} \leq T_{k,j}$. If this is not the case, then any subsequence that has s_i followed by s_k followed by s_j can not possibly be the best sequence of states. Note that the $T_{i,j}$'s may not necessarily be unique. In general, we will select the earliest transition time that minimizes the cost.

Lemma 8 then shows that as long as the pairwise-optimal transition times are in the right order, they give the globally optimal set of transition times for that subsequence. Our algorithm then uses this fact to find the optimal sequence of states by dynamic programming. Note that it is not necessary to exhaustively consider all possible subsequences.

6.1. Optimal Transition Times

Consider a particular subsequence of $l + 1$ states s_{a_0}, \dots, s_{a_l} . In order to avoid the double subscripts throughout this subsection, we will re-name our subsequence. q_0, q_1, \dots, q_l . Since the strategy must start in state s_0 , we can assume that $q_0 = s_0$. For $i < j$, define $\beta_{i,j}$ to be the cost to transition from state q_i to state q_j so that is $\beta_{i,j} = d_{a_i, a_j}$. Furthermore, we will refer to the power consumption rate of state q_i as α_i , that is, $\kappa_{a_i} = \alpha_i$.

We will consider the strategy that transitions through the states in the subsequence q_0, q_1, \dots, q_l . Suppose that we use transition time T_i to transition from state q_{i-1} to state q_i . It will be convenient for notation to define $T_{l+1} = \infty$ and $T_0 = 0$. The cost of the strategy that uses these transition times is:

$$\begin{aligned} \text{cost}(T_1, \dots, T_l) &= \sum_{j=1}^{l+1} \int_{T_{j-1}}^{T_j} p(t) \alpha_{j-1} (t - T_{j-1}) dt \\ &+ \sum_{j=1}^l \int_{T_j}^{\infty} p(t) [\alpha_{j-1} (T_j - T_{j-1}) + \beta_{j-1,j}] dt. \quad (2) \end{aligned}$$

The goal is to pick the T_1, \dots, T_l so as to minimize the above cost. This is the optimal cost for the subsequence q_0, \dots, q_l .

For each $i \in \{1, \dots, l\}$, let $\gamma_i = \frac{\alpha_{i-1} - \alpha_i}{\beta_{i-1,i}}$.

Lemma 6 *Suppose that there is an $i < j$ such that $\gamma_i < \gamma_j$, then there is a strict subsequence of q_0, \dots, q_l whose optimal cost is no greater than the optimal cost for q_0, \dots, q_l .*

Proof : Consider the first j such that $\gamma_{j-1} < \gamma_j$. Let $(\bar{t}_1, \dots, \bar{t}_{j-1}, \bar{t}_j, \dots, \bar{t}_l)$ be the sequence of thresholds that minimizes the cost of this sequence of states. Define the following quantities:

$$\begin{aligned} F_{j-1,j} &= \text{cost}(\bar{t}_1, \dots, \bar{t}_{j-1}, \bar{t}_j, \dots, \bar{t}_l) \\ F_{j-1,j-1} &= \text{cost}(\bar{t}_1, \dots, \bar{t}_{j-1}, \bar{t}_{j-1}, \dots, \bar{t}_l) \\ F_{j,j} &= \text{cost}(\bar{t}_1, \dots, \bar{t}_j, \bar{t}_j, \dots, \bar{t}_l) \end{aligned}$$

We will show that $F_{j-1,j}$ is greater than or equal to a weighted average of $F_{j-1,j-1}$ and $F_{j,j}$ which means that it must be greater than or equal to at least one of these values. This means that the strategy that transitions from state q_{j-2} to state q_{j-1} and then immediately transitions to state q_j at either time \bar{t}_{j-1} or \bar{t}_j is at least as good as the original strategy. Since $\beta_{j-1,j+1} \leq \beta_{j-1,j} + \beta_{j,j+1}$, skipping state j altogether can only improve the strategy.

Using the definition for the cost, we observe that

$$\begin{aligned} F_{j,j} - F_{j-1,j} &= \int_{\bar{t}_{j-1}}^{\bar{t}_j} p(t) (t - \bar{t}_{j-1}) (\alpha_{j-2} - \alpha_{j-1}) dt \\ &+ \int_{\bar{t}_j}^{\infty} p(t) (\bar{t}_j - \bar{t}_{j-1}) (\alpha_{j-2} - \alpha_{j-1}) dt \\ &- \int_{\bar{t}_{j-1}}^{\bar{t}_j} \beta_{j-2,j-1} p(t) dt. \end{aligned}$$

This is the expected difference in the power consumed by strategies $F_{j,j}$ and $F_{j-1,j}$. Dividing by $(\alpha_{j-2} - \alpha_{j-1})$, This becomes

$$\begin{aligned} \frac{F_{j,j} - F_{j-1,j}}{\alpha_{j-2} - \alpha_{j-1}} &= \int_{\bar{t}_{j-1}}^{\bar{t}_j} p(t) (t - \bar{t}_{j-1}) dt \\ &+ \int_{\bar{t}_j}^{\infty} p(t) (\bar{t}_j - \bar{t}_{j-1}) dt \\ &- \int_{\bar{t}_{j-1}}^{\bar{t}_j} \frac{1}{\gamma_{j-1}} p(t) dt. \quad (3) \end{aligned}$$

Again, using the definition for cost, we get that

$$\begin{aligned} F_{j-1,j} - F_{j-1,j-1} &= \int_{\bar{t}_{j-1}}^{\bar{t}_j} p(t) (t - \bar{t}_{j-1}) (\alpha_{j-1} - \alpha_j) dt \\ &+ \int_{\bar{t}_j}^{\infty} p(t) (\bar{t}_j - \bar{t}_{j-1}) (\alpha_{j-1} - \alpha_j) dt \\ &- \int_{\bar{t}_{j-1}}^{\bar{t}_j} \beta_{j-1,j} p(t) dt. \end{aligned}$$

Dividing by $(\alpha_{j-1} - \alpha_j)$, this becomes

$$\begin{aligned} \frac{F_{j-1,j} - F_{j-1,j-1}}{\alpha_{j-1} - \alpha_j} &= \int_{\bar{t}_{j-1}}^{\bar{t}_j} p(t)(t - \bar{t}_{j-1})dt \\ &+ \int_{\bar{t}_j}^{\infty} p(t)(\bar{t}_j - \bar{t}_{j-1})dt \\ &- \int_{\bar{t}_{j-1}}^{\bar{t}_j} \frac{1}{\gamma_j} p(t)dt. \end{aligned} \quad (4)$$

Comparing, equations (3) and (4), the expressions are almost identical except for the γ in the last term. Since $\gamma_{j-1} < \gamma_j$ and $\int_{\bar{t}_{j-1}}^{\bar{t}_j} p(t)dt \geq 0$, We have that

$$\frac{F_{j,j} - F_{j-1,j}}{\alpha_{j-2} - \alpha_{j-1}} \leq \frac{F_{j-1,j} - F_{j-1,j-1}}{\alpha_{j-1} - \alpha_j}.$$

Let $\omega_1 = 1/(\alpha_{j-2} - \alpha_{j-1})$ and $\omega_2 = 1/(\alpha_{j-1} - \alpha_j)$. Note that both ω_1 and ω_2 are at least 0. Rearranging, we get that

$$\left(\frac{\omega_1}{\omega_1 + \omega_2} \right) F_{j,j} + \left(\frac{\omega_2}{\omega_1 + \omega_2} \right) F_{j-1,j-1} \leq F_{j-1,j}. \quad \blacksquare$$

Now suppose that we consider only the two-state system consisting of state q_{i-1} and state q_i . We will let τ_i denote the optimal threshold time if these are the only two states in the system. We have that τ_i is the time T that minimizes

$$\int_0^T p(t)\alpha_{i-1}tdt + \int_T^\infty p(t)[\alpha_{i-1}T + \alpha_i(t - T) + \beta_{i-1,i}]dt.$$

Note that the value of T that results in the minimum above may not be unique. In this case, we take τ to be the smallest value which achieves the minimum. Also note that by subtracting the term $\int_0^\infty p(t)\alpha_i t dt$ (which is independent of T) and dividing by $\beta_{i-1,i}$ in the above definition, it can be seen that $\tau_i = \arg \min_T f(\gamma_i, T)$ where

$$f(\gamma, T) = \int_0^T p(t)\gamma t dt + \int_T^\infty p(t)[\gamma T + 1]dt.$$

Note that for a two-state system whose active state and sleep states has power consumption rates of γ and 0 respectively and whose transition cost is 1, $f(\gamma, T)$ denotes the expected power consumed by an online strategy that transitions to the sleep state at time T . We will show that for a particular subsequence of states, if we minimize the cost over all choices for the thresholds, the resulting thresholds are those obtained by the pair-wise optimization above. First, however, we must establish that the τ_i values have the correct ordering.

Lemma 7 *If $\gamma_i > \gamma_{i+1}$, then $\tau_i \leq \tau_{i+1}$.*

Proof : Intuitively, γ_i is the ratio of the additional power cost of being in state q_i instead of state q_{i-1} over the transition costs between the two states. It stands to reason that the larger this cost, the sooner one would want to transition from state q_{i-1} to state q_i .

We will formalize this argument using a proof by contradiction. Suppose that we have $\tau_i > \tau_{i+1}$ and $\gamma_i > \gamma_{i+1}$. The proof will make use of the definition of $f(\gamma, T)$ given above. τ_i is the smallest value for T which attains the minimum of $f(\gamma_i, T)$. Since $\tau_{i+1} < \tau_i$, we know that $f(\gamma_i, \tau_{i+1}) > f(\gamma_i, \tau_i)$. By the definition of τ_i , we have that $f(\gamma_{i+1}, \tau_i) \geq f(\gamma_{i+1}, \tau_{i+1})$. Thus, it should be the case that

$$f(\gamma_{i+1}, \tau_i) - f(\gamma_{i+1}, \tau_{i+1}) \geq 0 > f(\gamma_i, \tau_i) - f(\gamma_i, \tau_{i+1}). \quad (5)$$

Using the definition of $f(\gamma, T)$ above, for any $T_1 < T_2$,

$$\begin{aligned} f(\gamma, T_2) - f(\gamma, T_1) &= \gamma \left[\int_{T_1}^{T_2} p(t)(t - T_1)dt \right. \\ &+ \left. \int_{T_2}^{\infty} p(t)(T_2 - T_1)dt \right] \\ &- \int_{T_1}^{T_2} p(t)dt. \end{aligned}$$

The quantity inside the square braces above is non-negative. This implies that the quantity $f(\gamma, T_2) - f(\gamma, T_1)$ is non-decreasing in γ . This, however, contradicts Inequality 5 and the fact that $\gamma_i > \gamma_{i+1}$. \blacksquare

Finally, we prove the main lemma which states that the transition times are simultaneously optimized at the pairwise-optimal transition points.

Lemma 8 *For a given subsequence of states q_0, \dots, q_l , if $\tau_{i-1} < \tau_i$ for all $i \in \{1, \dots, l\}$, then the minimum total cost is achieved for $\text{cost}(\tau_1, \dots, \tau_l)$.*

Proof : The basic idea is that we can interpret $\text{cost}(T_1, \dots, T_l) - \int_0^\infty p(t)\alpha_l t dt$ as the sum of the power consumed in l two-state systems, where the i^{th} system, $i = 1, \dots, l$, has power consumption rates of $(\alpha_{i-1} - \alpha_i)$ and 0 respectively in the active and sleep states, and transition cost $\beta_{i-1,i}$. Note that $\int_0^\infty p(t)\alpha_l t dt$ is a constant, independent of the choice of T_i 's. After rescaling, one can write this expression as a linear combination of the $f(\gamma_i, T_i)$ terms. Since τ_i minimizes $f(\gamma_i, T)$, and the τ_i values have the right ordering, this implies that $\text{cost}(T_1, \dots, T_l)$ is minimized by setting $T_i = \tau_i$ for $i = 1, \dots, l$.

We will establish below that we can rewrite (2) as fol-

lows:

$$\begin{aligned} \text{cost}(T_1, \dots, T_l) &= \int_0^\infty p(t) \alpha_l t dt \\ &+ \sum_{i=1}^l \left[\int_0^{T_i} p(t) (\alpha_{i-1} - \alpha_i) t dt \right. \\ &\left. + \int_{T_i}^\infty p(t) \left((\alpha_{i-1} - \alpha_i) T_i + \beta_{i-1,i} \right) dt \right]. \quad (6) \end{aligned}$$

So by rescaling, we get that

$$\text{cost}(T_1, \dots, T_l) - \int_0^\infty p(t) \alpha_l t dt = \sum_{i=1}^l \beta_{i-1,i} f(\gamma_i, T_i).$$

We want to choose $T_1 \leq \dots \leq T_l$ to minimize this expression. Since $\tau_1 \leq \dots \leq \tau_l$ and each $\tau_i = \arg \min_T f(\gamma_i, T)$ it follows that the minimum is attained by setting $T_i = \tau_i$ for each i .

To complete the proof we show the equivalence of (2) and (6). It suffices to show that (2) and (6) integrate the same expression over each interval $[T_{i-1}, T_i], i = 1, \dots, l + 1$. The integrand in (2) (in this interval) is

$$p(t) \left[\alpha_{i-1} (t - T_{i-1}) + \sum_{j=1}^{i-1} (\alpha_{j-1} (T_j - T_{j-1}) + \beta_{j-1,j}) \right],$$

and the integrand in (6) is

$$\begin{aligned} p(t) \left[\sum_{j=1}^{i-1} (p(t) (\alpha_{j-1} - \alpha_j) T_j + \beta_{j-1,j}) + \right. \\ \left. \left(\sum_{j=i}^l (\alpha_{j-1} - \alpha_j) + \alpha_l \right) t \right]. \quad (7) \end{aligned}$$

The summations over the j indices in (7) telescope to show that the two expressions are identical. ■

6.2. The Optimal State Sequence

We now present a simple polynomial time algorithm to obtain the optimal state sequence for a given system. First, for each pair $(i, j), 0 \leq i < j \leq k$, let $T_{i,j}$ denote the optimal transition point if s_i and s_j were the only two states in the system. The time complexity of determining a single $T_{i,j}$ depends on the representation of the probability distribution. In practice, this is most likely to be estimated by a finite histogram with B bins starting at time 0 and sampled at a uniform discrete interval of δ . It follows that bin i corresponds to time δi . It is not difficult to generalize this for variable sized bins. We will also assume that all transition times occur at some δi . The height of bin i is $H(i)$ and this implies that the probability that the idle time t equals δi is given by $\frac{H(i)}{\sum_i H(i)}$. In what follows below, $ACC[i]$ is

$\int_0^{i\delta} p(t) dt$ and $ACCT[i] = \int_0^{i\delta} t p(t) dt$. The $T_{i,j}$ values can be evaluated in $O(k^2 B)$ time as follows:

```

ACC[0] ← H[0]
ACCT[0] ← 0
for k = 1 to B do
  ACC[k] ← ACC[k - 1] + H[k]
  ACCT[k] ← ACCT[k - 1] + H[k] × k · δ
end for
TOTAL ← ACC[B]
TOTALT ← ACCT[B]
for all (i, j) pairs such that 0 ≤ i < j ≤ k do
  min ← ∞, argmin ← -1
  for l = 0 to B - 1 do
    val = κi · ACCT[l] + (κilδ - κjlδ +
      βi,j)(TOTAL - ACC[l]) + κj(TOTALT -
      ACCT[l])
    if val < min then
      min ← val
      argmin ← l
    end if
  end for
  Ti,j ← argmin · δ
end for

```

Once the $T_{i,j}$'s are found, we sweep through them in non-decreasing order, keeping a running tab of the best sub-schedules that we can achieve ending in each state s_i at each point in time. When we encounter a $T_{i,j}$, we check to see if transitioning from s_i to s_j can improve the current best sub-schedule ending in s_j , and if it does, update our data structure to reflect it.

A given strategy divides time into intervals where each interval is the period of time spent in a particular state. The expected cost for a strategy given in Equation 2 is obtained by summing over the expected cost incurred in each interval. The cost for each interval is divided into two parts which results in two separate summations in Equation 2. We define the function Q for the first term which is

$$Q(i, t_i, j, t_j) = \int_{t_i}^{t_j} p(t) \kappa_i (t - t_i) dt.$$

This is the expected cost in the interval $[t_i, t_j]$ for those idle periods whose length is also in the interval $[t_i, t_j]$. The second part is the expected cost in the interval $[t_i, t_j]$ for those idle periods that are longer than t_j . Define

$$R(i, t_i, j, t_j) = \int_{t_j}^\infty p(t) [\alpha_i (t_j - t_i) + \beta_{i,j}] dt.$$

Note that $Q(i, \delta l_i, j, \delta l_j)$ and $R(i, \delta l_i, j, \delta l_j)$ can both be evaluated in constant time given $ACC[l_i], ACC[l_j], ACCT[l_i]$ and $ACCT[l_j]$ defined above.

At each transition $T_{i,j}$, we check to see if the current best schedule that ends in state s_j can be improved by transitioning to j from the current best schedule that ends in state s_i .

For this purpose, we maintain two arrays of size $k+1$: $t[i]$ is the time at which the current best schedule that ends at state s_i transitions to s_i and $h[i]$ is the cost at t_i of that schedule. Initially, $h[0] \leftarrow 0$ and all other $h[i] \leftarrow \infty$. $t[i]$, for all i can be initialized to 0. We provide the pseudocode for processing at each transition point $T_{i,j}$.

Current Status: $T_{i,j}$ is the transition point that is being processed
 {The cost up to time $T_{i,j}$ if transitioning from i to j at $T_{i,j}$ }
 $h1 \leftarrow h[i] + Q(i, t[i], j, T_{i,j}) + R(i, t[i], j, T_{i,j})$
 {The cost up to time $T_{i,j}$ if not transitioning to j at $T_{i,j}$ }
 $h2 \leftarrow h[j] + Q(j, t[j], j, T_{i,j}) + R(j, t[j], j, T_{i,j})$
if $h1 < h2$ **then**
 $h[j] \leftarrow h1$
 $t[j] \leftarrow T_{i,j}$
end if

It is easy to see that each transition point takes a constant amount of processing. The sorting takes an overhead of $O(k^2 \log k)$. The initial preprocessing to calculate the transition points takes $O(k^2 B)$. Hence, the total running time is $O(k^2(\log k + B))$.

References

- [1] <http://www.microsoft.com/windows2000/techenthusiast/features/standby1127.asp>.
- [2] C. Anderson and A. Karlin. Two adaptive hybrid cache coherency protocols. In *Proceedings of the Second International Symposium on High-Performance Computer Architecture*, pages 303–313, 1996.
- [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 (TVLSI) Systems*, 8(3):299–316, 2000.
- [4] S. J. Eggers and R. H. Katz. Evaluating the performance of four snooping cache coherency protocols. In *Proceedings of the 16th annual international symposium on Computer architecture*, pages 2–15. ACM Press, 1989.
- [5] S. Irani, R. Gupta, and S. Shukla. Competitive analysis of dynamic power management strategies for systems with multiple power savings states. In *IEEE Conference on Design, Automation and Test in Europe*, 2002.
- [6] S. Irani, S. Shukla, and R. Gupta. Online strategies for dynamic power management in systems with multiple power saving states. *Trans. on Embedded Computing Sys.*, 2003. Special Issue on Power Aware Embedded Computing.
- [7] A. Karlin, M. Manasse, L. McGeoch, and S. Owicki. Randomized competitive algorithms for non-uniform problems. In *ACM-SIAM Symposium on Discrete Algorithms*, pages 301–309, 1990.
- [8] A. R. Karlin, K. Li, M. S. Manasse, and S. Owicki. Empirical studies of competitive spinning for a shared-memory multiprocessor. In *Proceedings of the thirteenth ACM symposium on Operating systems principles*, pages 41–55. ACM Press, 1991.
- [9] S. Keshav, C. Lund, S. Phillips, N. Reingold, and H. Saran. An empirical evaluation of virtual circuit holding time policies in ip-over-atm networks. *IEEE Journal on Selected Areas in Communications*, 13(8):1371–1382, 1995.