

Stochastic Modeling of a Power-Managed System: Construction and Optimization

Qinru Qiu, Qing Wu and Massoud Pedram
Department of Electrical Engineering-Systems
University of Southern California
Los Angeles, CA 90089

Abstract -- *The goal of a dynamic power management policy is to reduce the power consumption of an electronic system by putting system components into different states, each representing certain performance and power consumption level. The policy determines the type and timing of these transitions based on the system history, workload and performance constraints. In this paper, we propose a new abstract model of a power-managed electronic system. We formulate the problem of system-level power management as a controlled optimization problem based on the theories of continuous-time Markov decision processes and stochastic networks. This problem is solved exactly using linear programming or heuristically using “policy iteration”. Our method is compared with existing heuristic methods for different workload statistics. Experimental results show that power management method based on Markov decision process outperforms heuristic methods by as much as 44% in terms of power dissipation savings for a given level of system performance.*

I. INTRODUCTION

With the rapid progress in the semiconductor technology, the chip density and operation frequency have increased, making the power consumption in battery-operated portable devices a major concern. High power consumption reduces the battery service life. The goal of low-power design for battery-powered devices is thus to extend the battery service life while meeting performance requirements. Reducing power dissipation is a design goal even for non-portable devices since excessive power dissipation results in increased packaging and cooling costs as well as potential reliability problems. The focus of this paper is however on portable electronic systems.

Portable electronic devices tend to be much more complex than a single VLSI chip. They contain many components, ranging from digital and analog to electro-mechanical and electro-chemical. Much of the power dissipation in a portable electronic device comes from non-digital components. Dynamic power management – which refers to selective shut-off or slow-down of system components that are idle or underutilized – has proven to be a particularly effective technique for reducing power dissipation in such systems. Incorporating a dynamic power management scheme in the design of an already-complex system is a difficult process that may require many design iterations and careful debugging and validation.

To simplify the design and validation of complex power-managed systems, a number of standardization attempts have been initiated. Best known among them is the *Advanced Configuration and Power Interface (ACPI)* [6] that specifies an abstract and flexible interface between the power-managed hardware components (VLSI chips, hard disk drivers, display drivers, modems, etc.) and the *power manager* (the system component that controls the turn-on and turn-off of the system components). The functional areas covered by the ACPI specification are:

- System power management – ACPI defines mechanisms for putting the computer as a whole in and out of system sleeping states. It also provides a general mechanism for any device to wake the computer.
- Device power management – ACPI tables describe motherboard devices, their power states, the power planes the devices are connected to, and controls for putting devices into different power states. This enables the OS to put devices into low-power states based on application usage.
- Processor power management – While the OS is idle but not sleeping, it will use commands described by ACPI to put processors in low-power states.

ACPI does not however specify the power management policy. It is the objective of the proposed research to provide a framework and supporting tools for constructing optimal power management policies based on modeling the power-managed system as a continuous-time Markov decision process.

The problem of finding a power management scheme (or policy) that minimizes power dissipation under performance constraints is of great interest to system designers. A simple power management system includes four components: Service Provider (SP), Service Requestor (SR), Service Queue (SQ) and Power Manager (PM). Figure 1 shows the information/command flow in a power-managed system. The SR generates service requests for the SP. The SQ buffers the service requests. The SP provides service to the requests in a top-down manner. The PM monitors the states of the SR, SQ and SP and issue state-transition commands to the SP. A simple and well-known heuristic policy is the “time-out” policy, which is widely used in today’s portable computers. In the “time-out” policy, the SP is shut down after it has been idle for a certain amount of time. The predictive system shutdown approach in [7][8] tries to achieve better power-delay trade-off by predicting the “on” and “off” time of the component. This prediction approach uses a regression equation based on the component’s previous “on” and “off” times to estimate the next “turn-on” time, such that the SP can be turned on just before the request comes. Therefore, the system performance can be improved. This method is only applicable to the special cases where the requests are highly correlated.




Figure 1 A power-managed system.

In general, heuristic policies cannot achieve the best power-delay trade-off for the system, cannot deal with complex components that have more than two (on and off) operating modes such as defined in ACPI, and cannot deal with a complex system with multiple interactive components.

A power management approach based on Markov decision process has been proposed in [8]. The system is modeled as a **discrete-time** Markov decision process by combining the stochastic models of each component: Once the model and its parameters are determined, an optimal power management policy for achieving the best power-delay trade-off in the system is generated. This approach offers significant improvements over previous power management techniques in terms of its theoretical framework for modeling and optimizing the system. There are however some shortcomings. Firstly, because the system is modeled in the discrete-time domain, some assumptions about the system components may not hold for real applications, such as the assumption that each event comes at the beginning of a time slice, the assumption that the transition of SQ is independent of the transition of SP, etc. Secondly, the state transition probability of the system model cannot be obtained accurately. For example, the discrete-time model cannot distinguish the busy state and the idle state because the transition between these two states are instantaneous. However the transition probability of SP when it is in these two states are different. Moreover, the power management program needs to send control signals to the components in every time-slice, which results in heavy signal traffic and heavy load on the system resources (therefore more power).

In this work, we overcome these shortcomings by introducing a new system model based on **continuous-time** Markov decision processes. More precisely:

1. The new model is based on the continuous-time Markov decision processes, which is more suitable for modeling real systems.
2. The resulting power management policy is asynchronous which is more appropriate for implementation as part of the operating system.
3. The new model explicitly distinguishes the busy state and the idle state of SP so that the system characterization becomes more accurate.
4. The new model considers the correlation between the state of the SQ and the state of the SP, which is the real-life scenarios.

5. The model for the service queue consists of a normal queue and a high priority queue. This is important since some service requests are "urgent" and need immediate response from the server.
6. The service requester model is capable of capturing complex workload characteristics.
7. The overall system model is constructed exactly and efficiently from those of the component models. We use an analytical base approach to calculate the generator matrix for the joint process of SP-SQ and a tensor sum based calculation to calculate the generator matrix of the joint process of SP-SQ and SR.
8. Both (exact) linear programming and (heuristic) policy iteration algorithms are used to solve the policy optimization problem.

A preliminary work of this topic can be found in [10] and [11]. This paper is organized as follows, Section II gives theoretical background of continuous-time controllable Markov process, Sections III and IV describes the models for the components and the system. Section V describes the solution technique for the optimal policy. Sections VI and VII present the experimental results and conclusions.

II. BACKGROUND

This section provides theoretical background on continuous-time Markov decision processes. Readers who are familiar with this topic can skip this part.

We first give the notation that will be used throughout the paper:

$P_{i \rightarrow j}(t)$: transition probability from state i (directly or indirectly) to state j during time 0 to t

$p_i(t)$: probability of that the system is in state i at time t

$X(t)$: value of the stochastic process X at time t

S, T : state space and parameter space of a stochastic process

G : the *generator matrix* of a continuous-time Markov process

A_i : set of available actions when a system is in state i

$a_i(t)$: the action that the system takes when it is in state i at time t , $a_i(t) \in A_i$

$p_i^{a_i}(t)$: the probability that action $a_i(t)$ is taken when the system is in state i at time t

$\mathbf{p}_i^{A_i}(t)$: the vector of $p_i^{a_i}(t)$, for all $a_i \in A_i$

$\boldsymbol{\pi}$: the power management policy

$\sigma_{i,j}$: transition rate from state i to state j

$\sigma_{i,j}^{a_i(t)}$: transition rate from state i to state j at time t when action $a_i(t)$ is taken

$\sigma_{i,j}^{\mathbf{p}_i^{A_i}(t)}$: transition rate from state i to state j at time t when actions are taken with probability $\mathbf{p}_i^{A_i}(t)$

$r_{i,i}$: reward rate (per unit time) of the system when it is in state i

$r_{i,j}$: transition reward of the system during the time when it makes a transition from state i to state j

r_i : earning rate of the system during the time it is in state i

$r_i^{a_i(t)}$: reward rate of the system when it is in state i and action $a_i(t)$ is taken at time t

$r_i^{\mathbf{p}_i^{A_i}(t)}$: reward rate of the system when it is in state i and actions are taken with probability $\mathbf{p}_i^{A_i}(t)$ at time t

$v_i(t)$: the total expected reward of the system from time 0 to time t with initial state i

$v_i^\pi(t)$: the total expected reward of the system from time 0 to time t with initial state i and policy π

$v_{i,avg}^\pi$: the limiting average reward of the system with the initial state i and policy π , $v_{i,avg}^\pi = \lim_{t \rightarrow \infty} \frac{1}{t} v_i^\pi(t)$

χ_{s_i, s_j} : the inverse of the average switching time of SP from state s_i to state s_j

τ_{r_i, r_j} : the inverse of the average switching time of SR from state r_i to r_j

$\lambda_l (\lambda_h)$: SR request generating rate of low (high) priority request

$\mu_l (\mu_h)$: SP service rate of low (high) priority request

A. Stochastic process

Definition 2.1 A *stochastic process* is a family of random variables $\{X(t), t \geq 0\}$, one for each t . t denotes the time parameter. For a specific t , $X(t)$ is a random variable with distribution $F(x, t) = P[X(t) \leq x]$. The values assumed by the process are called the *states*, and the set of possible values is called the *state space*.

Definition 2.2 A stochastic process $X(t)$ is called a *Markov process* if for any set of time instances $t_0 < t_1 < \dots < t_n < t$ its conditional distribution has the property:

$$P[X(t) \leq x \mid X(t_n) = x_n, X(t_{n-1}) = x_{n-1}, \dots, X(t_0) = x_0] = P[X(t) \leq x \mid X(t_n) = x_n]$$

where $t_0, t_1, \dots, t_n, t \in \mathbf{T}$ and $x_0, x_1, \dots, x_n \in \mathbf{S}$. When \mathbf{T} is a continuous space and \mathbf{S} is a discrete space, the Markov process is called the *continuous-time Markov process*.

Given a continuous-time Markov process with n states, its *generator matrix* \mathbf{G} is defined as an $n \times n$ matrix as shown in Eqn. (2.1). An entry $\sigma_{i,j}$ in \mathbf{G} is called the *transition rate* from state i to state j . All entries are defined in Eqn. (2.2) and Eqn. (2.3). Eqn. (2.4) gives the relationship between $\sigma_{i,i}$ and $\sigma_{i,j}$.

$$\mathbf{G} = \begin{bmatrix} -\sigma_{0,0} & \sigma_{0,1} & \sigma_{0,2} & \cdots \\ \sigma_{1,0} & -\sigma_{1,1} & \sigma_{1,2} & \cdots \\ \sigma_{2,0} & \sigma_{2,1} & -\sigma_{2,2} & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} \quad (2.1)$$

$$\sigma_{i,i} = \lim_{t \rightarrow 0} \frac{1 - p_{i \Rightarrow i}(t)}{t} = -p'_{i \Rightarrow i}(0) \quad i=1, 2, \dots, n \quad (2.2)$$

$$\sigma_{i,j} = \lim_{t \rightarrow 0} \frac{p_{i \Rightarrow j}(t)}{t} = p'_{i \Rightarrow j}(0) \quad i, j = 1, 2, \dots, n; i \neq j \quad (2.3)$$

$$\sum_{j \neq i} \sigma_{i,j} = \sigma_{i,i} \quad i, j = 1, 2, \dots, n; i \neq j \quad (2.4)$$

where $p'_{i \Rightarrow j}(t)$ is the derivative of $p_{i \Rightarrow j}(t)$. Obviously, $p_{i \Rightarrow j}(t) \geq 0$ and $\sum_{j \in \mathbf{S}} p_{i \Rightarrow j}(t) = 1$.

The generator matrix in the continuous-time Markov process is the analogue of the transition probability matrix in the discrete-time Markov process. We can calculate the limiting distribution (steady) state probabilities of the continuous-time Markov process from its generator matrix. Theorem 2.1 shows the relation between this matrix and the limiting distribution probabilities [7]. Before stating the theorem, we give some definitions.

Definition 2.3 Let T_{ij} be the first time instance at which that the Markov process visits state j starting from state i . A state i is called *recurrent* if $P(T_{ii} < \infty) = 1$. A state i is called *transient* if $P(T_{ii} < \infty) < 1$. A recurrent state is said to be *positive recurrent* if $E(T_{ii}) < \infty$, where $E(T_{ii})$ is the expectation of T_{ii} .

Definition 2.4 A recurrent state i is said to be *periodic* with period d if $d > 1$ is the greatest common divisor of all t_n , which is the n th time instance at which the Markov process returns state i if it starts from state i in time 0. If there is no such d , the state is called *aperiodic*.

Definition 2.5 State j is said to be *accessible* from state i if j can be reached from i within finite time, which is denoted as $i \rightarrow j$. If $i \rightarrow j$ and $j \rightarrow i$, they are said to be *communicate*, which is denoted as $i \leftrightarrow j$. The set of all states of a Markov process that communicate with each other forms a *communicating class*. If the set of all states of a stochastic process \mathbf{X} form a single communicating class, then \mathbf{X} is *irreducible*.

Definition 2.6 A Markov chain is called *ergodic* if it is irreducible, positive recurrent and aperiodic.

Theorem 2.1 ([12])

1. If the Markov process is irreducible, then the limiting distribution $\lim_{t \rightarrow \infty} p_i(t) = p_i$, $i \in \mathbf{S}$, exists and is independent of the initial conditions of the process, The limits $\{p_n \mid n \in \mathbf{S}\}$ are such that they either vanish identically (i.e., $p_i = 0$ for all $i \in \mathbf{S}$) or are all positive and form a probability distribution (i.e., $p_i > 0$ for all $i \in \mathbf{S}$, $\sum_{i \in \mathbf{S}} p_i = 1$).
2. The limiting distribution $\{p_i, i \in \mathbf{S}\}$ of an irreducible positive recurrent Markov process is given by the unique solution of the equation: $\mathbf{pG} = 0$ and $\sum_{j \in \mathbf{S}} p_j = 1$, where $\mathbf{p} = (p_0, p_1, p_2, \dots)$.

B. Continuous-time Markov decision processes

We now give a brief introduction to *continuous-time Markov decision processes*. For the discussions in the rest of this paper, we will omit the term ‘‘continuous-time’’ for more brevity description. Unless otherwise stated, all processes are assumed to be continuous-time.

Firstly, we describe a Markov process with *reward*. Assume the system earns a reward at rate $r_{i,i}$ (per unit time) during the time that it occupies state i . When it makes a transition from state i to state j ($i \neq j$), it receives a reward of $r_{i,j}$. Note that $r_{i,i}$ and $r_{i,j}$ have different dimensions. It is not necessary that the system earns according to both reward rates and transition rewards, but these definitions give us generality. We define the ‘‘earning rate’’ of state i as:

$$r_i = r_{i,i} + \sum_{j \neq i} \sigma_{i,j} r_{i,j} \quad (2.5)$$

Let $v_i(t)$ be the expected total reward that the system will earn during a time period of t if it starts in state i . The total expected reward during a time period of $t+dt$, that is $v_i(t+dt)$, can be written as:

$$v_i(t+dt) = (1 - \sum_{j \neq i} \sigma_{i,j} dt) [r_{i,i} dt + v_i(t)] + \sum_{j \neq i} \sigma_{i,j} dt [r_{i,j} + v_j(t)] \quad (2.6)$$

Equation (2.6) may be interpreted as follows. During the time interval dt the system may remain in state i or make a transition to some other state j . If it remains in state i for a time dt , it will earn a rate $r_{i,i} dt$ plus the expected reward that it will earn in the remaining t units of time, $v_i(t)$. The probability that it remains in state i for a time dt is $(1 - \sum_{j \neq i} \sigma_{i,j} dt)$. On the other hand, the system may make a transition to some state $j \neq i$ during the time interval

dt with probability $\sigma_{i,j} dt$. In this case the system would receive the reward $r_{i,j}$ plus the expected reward to be made if it starts in state j with time t remaining, $v_j(t)$. The product of probability and reward must then be summed over all states $j \neq i$ to obtain the total contribution to the expected values.

With $dt \rightarrow 0$ and using the definition of earning rate r_i , we have:

$$\frac{d}{dt}v_i(t) = r_i + \sum_{j=1}^n \sigma_{i,j} v_j(t) \quad i = 1, 2, \dots, n \quad (2.7)$$

where n is the total number of states of the process. Eqn. (2.7) gives a set of linear, constant coefficient differential equations that relate the total reward in time t from a starting state i to the values of r_i and $\sigma_{i,j}$.

Secondly, a *controllable Markov process* is a Markov process whose state transition rates can be controlled by controlling commands (defined as *actions*). When the system is in state i , an action a_i is chosen from a finite set A_i which includes all possible actions in state i . We denote this state-action relation by $\langle i, a_i \rangle$. If the chosen action changes as the time changes, we denote the action as a time-dependent variable $a_i(t)$. Hence the state-action pair is written as $\langle i, a_i(t) \rangle$.

Definition 2.7 A policy π is the set of state-action pairs for all the states of a controllable Markov process.

A policy can be either deterministic or randomized. If the policy is deterministic, then when the system is in state i at time t , an action $a_i(t)$ is chosen with probability 1. We denote a deterministic policy as: $\pi = \{ \langle i, a_i(t) \rangle \mid a_i(t) \in A_i, 1 \leq i \leq n \}$. If the policy is randomized, then when the system is in state i at time t , an action a_i is chosen with probability $p_i^{a_i}(t)$, such that $\sum_{a_i \in A_i} p_i^{a_i}(t) = 1$. We denote a randomized policy as: $\pi = \{ \langle i, p_i^{A_i}(t) \rangle \mid 1 \leq i \leq n \}$, where

$p_i^{A_i}(t)$ is a vector of $p_i^{a_i}(t)$, for all $a_i \in A_i$. Notice that the deterministic policy is a special case of randomized policy with one of the $p_i^{a_i}(t)$ equals 1.

In a controllable Markov process, the state transition rates $\sigma_{i,j}$ have different values when different actions are taken. In a deterministic policy, we denote it as $\sigma_{i,j}^{a_i(t)}$. In a randomized policy, we denote it as $\sigma_{i,j}^{p_i^{A_i}(t)}$, and $\sigma_{i,j}^{p_i^{A_i}(t)} = \sum_{a_i \in A_i} \sigma_{i,j}^{a_i} p_i^{a_i}(t)$. As a result, the generator matrix of a controllable Markov process is a parameterized matrix (action is the parameter). A *Markov decision process* is a controllable Markov process with rewards. In a Markov decision process, since $\sigma_{i,j}$ is action-dependent, the reward rate r_i becomes also action-dependent and will thus be denoted as $r_i^{a_i(t)}$ for deterministic policy and $r_i^{p_i^{A_i}(t)}$ for randomized policy, $a_i(t) \in A_i$. The expected total reward $v_i(t)$ depends on the chosen action in each state, i.e., it becomes policy-dependent and will be denoted as $v_i^\pi(t)$. An example of Markov decision process is given in Example 3.1.




Figure 2 Example of Controllable Markov Process.

Example 3.1 Consider a controllable Markov process which consists of only two states: i and j . When the system is in state i , the set of available actions is denoted as A_i ; we assume that $A_i = \{a, b\}$. The transition rate from state i to state j equals $\sigma_{i,j}^a$ when action a is taken and $\sigma_{i,j}^b$ when action b is taken. Similarly, we define $A_j = \{x, y\}$, $\sigma_{j,i}^x, \sigma_{j,i}^y$. If we always take action a when the system is in state i and always take action x when the system is in

state j , we can write the resulting policy as: $\boldsymbol{\pi} = \{ \langle i, a \rangle, \langle j, x \rangle \}$. The system generator matrix using policy $\boldsymbol{\pi}$ can be written as $\mathbf{G} = \begin{bmatrix} -\sigma_{i,j}^a & \sigma_{i,j}^a \\ \sigma_{j,i}^x & -\sigma_{j,i}^x \end{bmatrix}$. Now consider a randomized policy $\boldsymbol{\pi}' = \{ \langle i, (0.5, 0.5) \rangle, \langle j, (0.4, 0.6) \rangle \}$. This means that: when the system is in state i , action a is taken with probability 0.5 and action b is taken with probability 0.5; when the system is in state j , action x is taken with probability 0.4 and action y is taken with probability 0.6. Therefore $\sigma_{i,j}^{p_i^{A_i}} = 0.5\sigma_{i,j}^a + 0.5\sigma_{i,j}^b$ and $\sigma_{j,i}^{p_j^{A_j}} = 0.4\sigma_{j,i}^x + 0.6\sigma_{j,i}^y$. The system generator matrix is constructed based on these two values. □

For the remainder of this paper, we will only use the notation for a deterministic policy.

Let $p_{i \Rightarrow j}^{\boldsymbol{\pi}}(t)$ denote the probability of being in state j at time t when the initial state at time 0 is i and the state transition rates are determined by policy $\boldsymbol{\pi}$. Similarly, let $r_j^{a_j(t)}$ denote the earning rate in state j at time t and action a_j is taken. The total expected reward that the process can earn for a time period of t using policy $\boldsymbol{\pi}$, can be written as ([14]):

$$v_i^{\boldsymbol{\pi}}(t) = \int_0^t \sum_{j=1}^n p_{i \Rightarrow j}^{\boldsymbol{\pi}}(\tau) r_j^{a_j(\tau)} d\tau \quad (2.8)$$

Given two policies $\boldsymbol{\pi}_1$ and $\boldsymbol{\pi}_2$, if we can find a time ξ , such that $v_i^{\boldsymbol{\pi}_1}(t) \geq v_i^{\boldsymbol{\pi}_2}(t)$, for all $t > \xi$, $i = 1, 2, \dots, n$, then we say that policy $\boldsymbol{\pi}_1$ is superior to $\boldsymbol{\pi}_2$, denoted as $\boldsymbol{\pi}_1 \geq \boldsymbol{\pi}_2$. A policy $\boldsymbol{\pi}$ is *optimal*, if it is superior to all other policies for the Markov decision process.

Let $v_i^{\boldsymbol{\pi}} = \lim_{t \rightarrow \infty} v_i^{\boldsymbol{\pi}}(t)$. The goal of a Markov decision process is to find the optimal policy that maximizes $v_i^{\boldsymbol{\pi}}$ for all i . However, in practice, we cannot use $v_i^{\boldsymbol{\pi}}$ directly for calculating the optimal policy since $v_i^{\boldsymbol{\pi}}(t)$ approaches infinity when t approaches infinity. A commonly used alternative quantity is the *limiting average reward*:

$$v_{i,avg}^{\boldsymbol{\pi}} = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \sum_{j=1}^n p_{i \Rightarrow j}^{\boldsymbol{\pi}}(\tau) r_j^{a_j(\tau)} d\tau, \quad (2.9)$$

Obviously, maximizing the limiting average reward for any fixed t is the same as maximizing $v_i^{\boldsymbol{\pi}}$.

Definition 2.8 A policy $\boldsymbol{\pi}$ is *stationary* if the action is only a function of the state and independent of time, that is: $\boldsymbol{\pi} = \{ \langle i, a_i \rangle \mid a_i \in A_i, 1 \leq i \leq n \}$ for a deterministic policy, or $\boldsymbol{\pi} = \{ \langle i, p_i^{A_i} \rangle \mid 1 \leq i \leq n \}$ for a randomized policy.

Definition 2.9 A policy $\boldsymbol{\pi}$ is *piecewise-stationary* if for any τ , interval $[0, \tau)$ can be divided into a finite number of intervals $[0, t_1), [t_1, t_2), \dots, [t_{m-1}, \tau)$ such that inside each interval, the policy is stationary.

Theorem 2.2 [14] There exists a stationary policy that maximizes $v_{i,avg}^{\boldsymbol{\pi}}$ over the class of piecewise-stationary policies.

Based on this theorem, we conclude that we will not lose generality if our search for the optimum policy is restricted to the set of stationary policies. Therefore, actions and policies that we will discuss from now on are all time-independent and we will reduce the notation $a_j(t)$ and $p_i^{A_i}(t)$ to a_i and $p_i^{A_i}$.

The goal of a Markov decision process is to find a policy that maximizes the expected reward. In our case, we want to find a policy that minimizes our cost function (i.e. power dissipation). These two problems become equivalent if we use the negative of cost as the reward. In the remainder of the paper, we will use the term *cost*

instead of *reward* and use $c_{i,i}$ and $c_{i,j}$ instead of $r_{i,i}$ and $r_{i,j}$, $c_{i,avg}$ instead of $v_{i,avg}$. In our work, we actually used a constraint Markov decision process to model the power-managed system. In the constraint Markov decision process, for each state, there is an object cost c_{obj} and several constraint costs c_{con} . The goal of the constraint Markov decision process is to find a policy that minimizes the expected value of objective cost while meeting the given constraint. That is:

$$\text{Minimize}_{\pi}(c_{obj_{i,avg}}^{\pi}), \text{ s.t. } c_{con_{i,avg}}^{\pi} \leq \text{Constraint}$$

The value of $C_{obj_{i,avg}}^{\pi}$ and $C_{con_{i,avg}}^{\pi}$ can be calculated based on c_{obj} and c_{con} using (2.9). We will introduce how we define and calculate the objective cost and the constraint cost in our system model in Section V.

III. COMPONENT MODELING

In this section, we describe the mathematical models of the components in a power-managed system.

A. Assumptions

The power-managed system consists of four components: a server that processes requests with different power modes (SP), a generator that generates the service request (SR), a queue which stores the requests that cannot be immediately serviced upon arrival (SQ), and a power manager (PM) that issues mode-switching commands to the SP. The SR is independent to the rest of the system. Requests generated by the SR can be divided into two categories: *low-priority requests* and *high-priority requests*, which are generated independently of each other.

Both the request arrival events and the request service completion events are stochastic processes and follow the Poisson distribution. When we state that the request arrival event is a Poisson process, it means that during time $(0, t]$, the number of the events follows a Poisson distribution with mean λt . Consequently, the request inter-arrival time follows an exponential distribution with mean $1/\lambda$. A Poisson distribution with a mean value of 5 ($t=1$) and an exponential distribution with a mean value of $1/5$ are illustrated in Figure 3. Notice a request that arrives when the SQ is full will be rejected. We also assume that the time that is needed for the SP to switch from one state to another follows an exponential distribution. In reality, the switching time for the SP is usually a small fixed value. We know that if the expected value of an exponentially distributed random variable is a then its variance is a^2 . If a is very small, then the variance will be negligible. Therefore, if the switching time of the SP is much shorter than the service time or input generation time, then it can be modeled by an exponential distribution without introducing much error.



Figure 3 Examples of Poisson distribution and exponential distribution.

In the remainder of this paper, we will use upper case bold letters (*e.g.*, \mathbf{M}) to denote matrices, lowercase bold letters (*e.g.*, \mathbf{v}) to denote vectors, italicized Arial-Font letters (*e.g.*, \mathcal{S}) to denote sets, uppercase italicized letters (*e.g.*, S) to denote scalar constants and lower case italicized letters (*e.g.*, x) to denote scalar variables.

B. Model of the Service Provider

The **Service Provider** (SP) is modeled as a stationary, continuous-time Markov decision process with state (operation mode) set $\mathbf{S}=\{s_i \text{ s.t. } i=1, 2, \dots, S\}$, action set \mathbf{A} , and parameterized generator matrix $\mathbf{G}_{SP}(a)$, $a \in \mathbf{A}$. It can be described by a quadruple $(\boldsymbol{\chi}, \mu(s), pow(s), ene(s_i, s_j))$ where: (i) $\boldsymbol{\chi}$ is an $S \times S$ matrix; (ii) $\mu_l(s)$ and $\mu_h(s)$ are functions, $\mu_l, \mu_h: \mathbf{S} \rightarrow \mathbf{R}$; (iii) $pow(s)$ is a function, $pow: \mathbf{S} \rightarrow \mathbf{R}$; (iv) $ene(s_i, s_j)$ is a function, $ene: \mathbf{S} \times \mathbf{S} \rightarrow \mathbf{R}$.

We call $\boldsymbol{\chi}$, the switching speed matrix of the SP. The (i,j) th entry of $\boldsymbol{\chi}$ is denoted by χ_{s_i, s_j} and represents the switching speed from state s_i to state s_j . The average switching time from state s_i to state s_j is then $1/\chi_{s_i, s_j}$. We set χ_{s_i, s_i} to be ∞ , because the switch from state s_i to itself is instantaneous.

The entries of the parameterized generator matrix $\mathbf{G}_{SP}(a)$ can be calculated as:

$$\sigma_{s_i, s_j}(a) = \delta(s_j, a) \cdot \chi_{s_i, s_j}, \quad s_i \neq s_j; \quad (3.1)$$

$$\sigma_{s_i, s_i}(a) = - \sum_{s_j \neq s_i} \sigma_{s_i, s_j}(a) \quad (3.2)$$

$$\text{where } \delta(s, a) = \begin{cases} 1 & \text{if } s \text{ is the destination state of action } a \\ 0 & \text{otherwise} \end{cases} \quad (3.3)$$

The *service rates* $\mu_l(s)$ and $\mu_h(s)$ represent the service speed of the SP for low-priority requests and high-priority requests in state s , respectively. Therefore, $1/\mu_l(s)$ ($1/\mu_h(s)$) gives the average time which is needed by the SP to complete the service for a low (high) priority request when it is in state s .

A *power consumption* $pow(s)$ is associated with each state $s \in \mathbf{S}$. It represents the power consumption of the SP during the time it occupies state s . The cost rate $c_{s,s}$ of state s is equal to $pow(s)$.

A *switching energy* $ene(s_i, s_j)$ is associated with each state pair (s_i, s_j) , $s_i, s_j \in \mathbf{S}$, $s_i \neq s_j$. It represents the energy needed for the SP to switch from state s_i to state s_j . The cost c_{s_i, s_j} is equal to $ene(s_i, s_j)$.

From Eqn. (2.5), we know that the expected power consumption (cost rate) of the SP when it is in state s and action a_s is chosen, can be calculated as:

$$c_s = pow(s) + \sum_{s' \neq s} \sigma_{s, s'}(a_s) ene(s, s').$$

In reality, states (i.e. working modes) of the SP can be divided into three groups: busy, idle, and power-down. In busy states, the SP is fully powered and working on the first request in the SQ. We assume that each request service is atomic so that the SP cannot switch to any other state when it is working on some request. In other words, the switch from the busy state to another states is not controllable. It only occurs when the SP finishes one service. For each busy state, there exists a corresponding idle state. In the idle states, the SP is fully powered, but it is not working on any request. An idle state is the only state that connects to its corresponding busy state. When the SP finishes a service, it will automatically switch from the busy state to its corresponding idle state. When SP wants to switch from some other state to a busy state, it will first switch to the corresponding idle state then goes to the busy state. Notice that the idle states are not physical states of SP. When the SP is in an idle state, it is in the same power mode as when it is in the corresponding busy state. We only use the idle state to convenient our modeling. In power-down states, the SP is partially or completely shut down, i.e., it is not operational. The SP may have multiple power-down states (e.g. standby, soft off, hard off).

Not all actions in \mathbf{A} are valid in all SP states. Constraints on a valid action can be stated as follows:

1. When the SP is in a busy state, its transition is not controllable, $\mathbf{A}_{busy} = \Phi$.
2. The action cannot make a transition from a power-down state to a busy state.
3. The action cannot make a transition from an idle state to a busy state other than its corresponding busy state.

Definition 3.1 Power down state s_1 is more *vigilant* than inactive state s_2 if the SP in state s_1 wakes up (switches to an active state) faster than the same SP does in state s_2 .

Different busy states may be used to model a component working under different supply voltages. We associate different power and delay (service rate) values to each of these states to model the server performance under different supply voltages. Therefore, our policy optimization approach (cf. Section V) can also find the best policy for dynamic voltage scaling as it finds the optimal policy for power management.

Example 3.1 Consider a SP with six states, $\mathcal{S}=\{busy_1, busy_2, idle_1, idle_2, wait, sleep\}$. When the SP is in state $busy_1$, it services the requests at a low speed. Assume that the average time needed for each service (for both low-priority requests and high-priority requests) is 5 ms. Therefore, $\mu_l(busy_1)$ and $\mu_h(busy_1)$ are 0.2. Also assume that in state $busy_2$, the SP services the request at a higher speed, e.g., $\mu_l(busy_2)=\mu_h(busy_2)=0.4$. $\mu_l(idle_1)$, $\mu_h(idle_1)$, $\mu_l(idle_2)$, $\mu_h(idle_2)$, $\mu_l(wait)$, $\mu_h(wait)$, $\mu_l(sleep)$ and $\mu_h(sleep)$ are all 0. Let the command set be defined as $\mathcal{A}=\{go_busy_1, go_busy_2, go_idle_1, go_idle_2, go_wait, go_sleep\}$. Notice that not all four commands are valid (or available) in all states. The switching speed matrix χ is given by:

$$\chi = \begin{bmatrix} \infty & 0 & 0.2 & 0 & 0 & 0 \\ 0 & \infty & 0 & 0.4 & 0 & 0 \\ \infty & 0 & \infty & 10 & 1 & 0.5 \\ 0 & \infty & 10 & \infty & 1 & 0.5 \\ 0 & 0 & 0.454 & 0.454 & \infty & 1.5 \\ 0 & 0 & 0.166 & 0.166 & 1.5 & \infty \end{bmatrix}$$

By default, the order of states in rows and columns of this matrix are the same as the left-to-right order of states in \mathcal{S} . $\chi_{s_i, s_i} = \infty$ means that the SP can transfer from state s_i to s_j immediately. $\chi_{s_i, s_j} = 0$ means that the SP can never transfer from state s_i to s_j . In this example, the transfer from any state to itself needs no time. The SP can transfer from *busy* state to *idle* state with the transition rate equal to the service rate because it autonomously goes to the *idle* state immediately after it finishes a request. The SP cannot switch between the *busy* state and *wait* state (or *sleep* state) directly (it must go through the *idle* state), therefore the corresponding entries in the matrix are 0.

The power consumption is: $pow(busy_1)=2.3W$, $pow(busy_2)=6.5W$, $pow(idle_1)=2.3W$, $pow(idle_2)=6.5W$, $pow(wait)=0.8W$, $pow(sleep)=0.1W$.

The switching energy $ene(s_i, s_j)$ matrix is:

$$ene(s_i, s_j) = \begin{bmatrix} 0 & \infty & 0 & \infty & \infty & \infty \\ \infty & 0 & \infty & 0 & \infty & \infty \\ 0 & \infty & 0 & 0.3mJ & 1mJ & 2mJ \\ \infty & 0 & 0.3mJ & 0 & 1mJ & 2mJ \\ \infty & \infty & 4.4mJ & 4.4mJ & 0 & 0.66mJ \\ \infty & \infty & 30mJ & 30mJ & 9mJ & 0 \end{bmatrix}$$

$ene(s_i, s_j) = \infty$ means that the SP cannot switch between the corresponding states. Note that the energy cost of autonomous state change (busy to idle) is zero.

A graphical illustration of the SP is shown in Figure 4. The transition rates associated with the edges have not been shown in the figure. They can be extracted from $\mathbf{G}_{SP}(a)$ for specific actions. The self-transitions of each state are not shown in the figure. \square

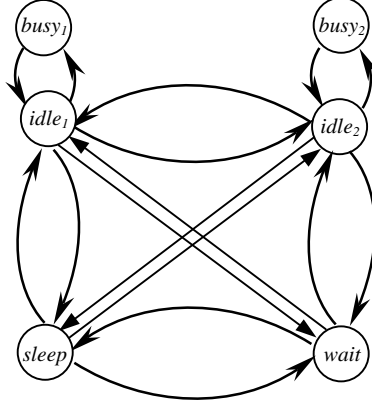


Figure 4 Markov process model of the SP.

C. Model of the Service Requester

The **Service Requester** (SR) is modeled as a stationary, continuous-time Markov process with state set $R=\{r_i \text{ s.t. } i=0, 1, 2, \dots, R\}$ and generator matrix \mathbf{G}_{SR} . It can be characterized by a pair $(\boldsymbol{\tau}, \lambda(r))$, where: (i) $\boldsymbol{\tau}$ is an $R \times R$ matrix, (ii) $\lambda_l(r)$ and $\lambda_h(r)$ are functions $\lambda: R \rightarrow \mathbf{R}$.

We call $\boldsymbol{\tau}$ the switching speed matrix of the SR. The (i,j) th entry of $\boldsymbol{\tau}$ is denoted as τ_{r_i, r_j} . We assume that the time needed for the SR to switch from one operation state to another is a random variable with exponential distribution. The average switch time from state r_i to state r_j is given by $1/\tau_{r_i, r_j}$. We set τ_{r_i, r_i} to be ∞ because the switch from state r_i to r_i is instantaneous. The SR model is a continuous-time Markov process with the generator matrix \mathbf{G}_{SR} . The value of σ_{r_i, r_j} (the transition rate from state r_i to state r_j) can be calculated as:

$$\sigma_{r_i, r_j} = \tau_{r_i, r_j}, r_i \neq r_j; \quad \sigma_{r_i, r_i} = - \sum_{r_j \neq r_i} \sigma_{r_i, r_j} \quad (3.4)$$

The *request rates* $\lambda_l(r)$ and $\lambda_h(r)$ are associated with state $r \in R$. When the SR is in state r , the generation of the low-priority requests follows the Poisson process with mean value $1/\lambda_l(r)$, whereas the generation of the high-priority requests follows the Poisson process with mean value $1/\lambda_h(r)$.

Example 3.2 Consider a SR with two states, r_1 and r_2 . When it is in state r_1 , it generates a low priority request every 30 ms in the average, a high priority request every 50 ms in the average. When it is in state r_2 , it generates a low priority request every 60 ms, a high priority request every 90 ms. So the request rates in each state is defined as follows: $\lambda_l(r_1)=1/30$, $\lambda_h(r_1)=1/50$, $\lambda_l(r_2)=1/60$, $\lambda_h(r_2)=1/90$. The switch matrix $\boldsymbol{\tau}$ is a 2×2 matrix, with one entry for each state pair. For instance:

$$\boldsymbol{\tau} = \begin{bmatrix} \infty & 1/200 \\ 1/400 & \infty \end{bmatrix}.$$

It means that when SR is in state r_0 the expected time that it will switch to state r_1 is 200 ms, when SR is in state r_1 the expected time that it will switch to state r_0 is 400 ms. Therefore, the generator matrix of SR is:

$$\mathbf{G}_{\text{SR}} = \begin{bmatrix} -1/200 & 1/200 \\ 1/400 & -1/400 \end{bmatrix}.$$

Figure 5 gives the illustration of the Markov process model of this SR. The self-transition of each states are not shown. □




Figure 5 Markov process model of the service requester.

D. Model of the Service Queue

A **Single Service Queue (SSQ)** is modeled as a stationary, continuous-time Markov process, with state set $Q_{SSQ}=\{q_i, i=0, 1, 2, \dots, Q\}$ and the generator matrix $G_{SSQ}(s, r)$, where Q is the maximum length of the queue, s is the state of SP, r is the state of SR.

The number of waiting requests in the queue decides the state of the SSQ. If there are i requests waiting in the queue, then the queue is in the state of q_i . The entries of the parameterized generator matrix $G_{SSQ}(s, r)$ can be calculated as:

$$\sigma_{q_i, q_{i+1}} = \lambda(r), \sigma_{q_{i+1}, q_i} = \mu(s), \sigma_{q_i, q_i} = - \sum_{q_j \neq q_i} \sigma_{q_i, q_j}, 0 \leq i \leq Q-1$$

$$\sigma_{q_i, q_j} = 0, \text{ for other state pairs}$$

$\lambda(r)$ and $\mu(s)$ are the request input rate and service rate of the queue. The shortcoming of using SSQ as the stochastic model of the service queue is that we can assign only one delay constraint (i.e. the constraint on the average waiting time of the requests) during the policy optimization. In a power-managed system, the SP, under control of the PM, may thus not service an incoming request immediately in order to achieve better power-delay trade-off. However, there may exist high-priority requests that need immediate service by the SP. In this case, if we use a loose delay constraint (which means the power management policy may not service the request immediately) the consequent long latency is however not acceptable for high-priority requests. If instead use a tight delay constraint to ensure that the high-priority requests are serviced immediately, then there will be undesirable power dissipation related to unnecessarily tight delay constraint on low-priority requests.

We henceforth model the service queue as a combination of two SSQs: one (denoted as HSQ) for the high-priority requests and the other (denoted as LSQ) for the low-priority requests. The relationships between these two queues are:

1. Two different delay constraints are assigned to HSQ and LSQ separately such that the requests in HSQ have smaller waiting time than those in LSQ.
2. The SP will not start serving the requests in LSQ until it finishes all the requests in HSQ. Therefore, we define the service rate of LSQ as a function that relate to both the state of SP and the state HSQ: $\mu'_l(s, hq_i)$, where s is the SP state, hq_i is the HSQ state. If $i=0$ then $\mu'_l(s, hq_0)=\mu(s)$, otherwise, $\mu'_l(s, hq_i)=0$.

Although we have introduced two queues in our stochastic model of the service queue, we are actually modeling a single priority queue. The SQ model can be used to model the commonly used priority queue in an operating system where two different priorities are assigned to tasks and high-priority tasks, when they come, are inserted into the front of the queue. Moreover, obviously, the SQ model can be extended to model a queue of requests that have more than two priority levels.

The formal definition of the SQ model is as follows.

*The **Service Queue (SQ)** is modeled as a stationary, continuous-time Markov process, which is the combination of two SSQs: LSQ and HSQ. The state set of the SQ is given by $Q = Q_{LSQ} \times Q_{HSQ}$ and the generator matrix is given by*

$G_{SQ}(s, r) = G_{LSQ}(s, r) \oplus G_{HSQ}(s, r, hq)$, where s is the state of SP, r is the state of SR state, and the “ \oplus ” operation is the tensor sum defined in Definition 4.1.

The number of waiting requests in the HSQ and LSQ decides the state of the SQ. If there are i requests waiting in the LSQ and j requests waiting in the LSQ, then the queue is in the state of (lq_i, hq_j) . The entries of both the parameterized generator matrix $G_{HSQ}(s, r)$ and $G_{LSQ}(s, r, hq)$ is calculated in the same way as $G_{SSQ}(s, r)$ except that they use $\lambda_l(r)$, $\mu'_l(s, hq)$ and $\lambda_h(r)$, $\mu_h(s)$ as input rate and service rate respectively.

Example 3.3 Consider a HSQ and a LSQ. Assume that the maximum length of HSQ is 1, the maximum length of LSQ is 2. The generator matrix of HSQ and LSQ are:

$$G_{HSQ}(s, r) = \begin{bmatrix} -\lambda_h(r) & \lambda_h(r) \\ \mu_h(s) & -\mu_h(s) \end{bmatrix}, G_{LSQ}(s, r, hq) = \begin{bmatrix} -\lambda_l(r) & \lambda_l(r) & 0 \\ \mu'_l(s, hq) & -\mu'_l(s, hq) - \lambda_l(r) & \lambda_l(r) \\ 0 & \mu'_l(s, hq) & -\mu'_l(s, hq) \end{bmatrix}.$$

The generator matrix of SQ is:

$$G_{SQ}(s, r) = \begin{bmatrix} \sigma_{0,0} & \lambda_h(r) & \lambda_l(r) & 0 & 0 & 0 \\ \mu_h(s) & \sigma_{1,1} & 0 & \lambda_l(r) & 0 & 0 \\ \mu_l(s) & 0 & \sigma_{2,2} & \lambda_h(r) & \lambda_l(r) & 0 \\ 0 & 0 & \mu_h(s) & \sigma_{3,3} & 0 & \lambda_l(r) \\ 0 & 0 & \mu_l(s) & 0 & \sigma_{4,4} & \lambda_h(r) \\ 0 & 0 & 0 & 0 & \mu_h(s) & \sigma_{5,5} \end{bmatrix}$$

To save place, we did not write out the value of diagonal entries. Figure 5 shows an example of HSQ, LSQ and SQ.

□




Figure 6 Markov process model of the service queue.

Notice that request starvation may occur when high priority requests keep coming while there are some low priority requests waiting. However, this problem can be solved by the OS. If a low priority request has been

waiting in the queue for too long, the OS will change its priority to high so that it will get serviced. Most of the existing operating system codes have a similar mechanism to prevent starvation. As a result, the power manager need not be concerned with this problem.

IV. SYSTEM MODELING

We first show how to construct the model of the entire system by combining the component models. Next we explain how the power-managed system model is applied to practical applications.

A. Model of the Power-Managed System

The **Power-Managed System (SYS)** can be modeled as a controllable continuous-time Markov process which is the composition of the models of the SP, the SR and the SQ. The state set is given by: $X = \mathcal{S} \times \mathcal{Q} \times \mathcal{R} - \{\text{invalid states where SP is busy and SQ is empty}\}$. An action set of all possible actions which is the same as \mathcal{A} in the SP model. A parameterized generator matrix $\mathbf{G}_{\text{SYS}}(a)$ gives the state transition rates under action a . The SYS state can be represented as $(s, r, (lq, hq))$, where $s \in \mathcal{S}$, $r \in \mathcal{R}$, $lq \in \mathcal{Q}_{\text{LSQ}}$ and $hq \in \mathcal{Q}_{\text{HSQ}}$.

Similar to the case of the SP model, not all actions are valid for any system state. The action constraints (which are described in Section III.A) for the SP model still apply to the model of SYS. In addition, we add the following constraints related to the SYS model.

- (1) When both LSQ and HSQ are full and the SP is in an inactive state, the SP cannot make a transition to another inactive state which is less vigilant (c.f. Definition 4.1) than the current one. This constraint is reasonable because the SP must go to a fully functional state as soon as possible in this situation.
- (2) When both LSQ and HSQ are full and the SP is in an idle state, the SP cannot make a transition to a power-down state or another idle state whose corresponding busy state has a slower service rate. This constraint is also reasonable because when SP and SQ are in the above states, it means that the service speed is not enough to match the incoming speed of the requests. Therefore, we need to increase the service rate, not decrease it.

Proposition 4.1 If the SYS satisfies the above two constraints, there is only one ergodic chain in the system and the states outside the ergodic chain are all transient states.

B. Calculating the generator matrix

We next introduce an efficient method for calculating the generator matrix $\mathbf{G}_{\text{SYS}}(a)$ of the system from the generator matrices of the system components: $\mathbf{G}_{\text{SP}}(a)$, \mathbf{G}_{SR} , and $\mathbf{G}_{\text{SQ}}(s, r)$.

First, we show how to calculate the generator matrix of a joint process of two independent continuous-time Markov processes. Proposition 4.2 gives a method to obtain the joint transition rate of two independent continuous-time Markov processes. Proposition 4.3 gives the method for generating the generator matrix of the joint system using matrix operations.

Proposition 4.2 Given two independent stochastic processes X and Y , let $\sigma_{(x,y),(x',y')}$ denote the transition rate of the joint process from the joint state (x,y) to joint state (x',y') , where x and $x' \in$ state space of X , y and $y' \in$ state space of Y . Let $\sigma_{x,x'}$ denote the transition rate of process X from state x to state x' and $\sigma_{y,y'}$ denote the transition rate of process Y from state y to state y' . Then $\sigma_{(x,y),(x,y')} = \sigma_{y,y'}$, $\sigma_{(x,y),(x',y)} = \sigma_{x,x'}$, $\sigma_{(x,y),(x',y')} = 0$, $\sigma_{(x,y),(x,y)} = \sigma_{x,x} + \sigma_{y,y}$.

Let matrices \mathbf{A} and \mathbf{B} be defined as:

$$\mathbf{A} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \quad \text{and} \quad \mathbf{B} = \begin{bmatrix} b_{11} & b_{12} & b_{13} \\ b_{21} & b_{22} & b_{23} \\ b_{31} & b_{32} & b_{33} \end{bmatrix}$$

Definition 4.1 The tensor product $C=A\otimes B$ is given by $C = \begin{bmatrix} a_{11}\mathbf{B} & a_{12}\mathbf{B} \\ a_{21}\mathbf{B} & a_{22}\mathbf{B} \end{bmatrix}$. The tensor sum $C=A\oplus B$ is given by: $C = A \otimes \mathbf{I}_{n_2} + \mathbf{I}_{n_1} \otimes B$, where n_1 is the order of A , n_2 is the order of B , \mathbf{I}_{n_i} is the identity matrix of order n_i .

Proposition 4.3 Given two independent continuous-time Markov processes with generator matrices A and B , the generator matrix of the joint process is given by $A\oplus B$.

We have mentioned in Section III.A that the SR is independent from the rest of the system. Therefore, $\mathbf{G}_{\text{SYS}}(a)$ can be calculated as:

$$\mathbf{G}_{\text{SYS}}(a) = \mathbf{G}_{\text{SP-SQ}}(a, r) \oplus \mathbf{G}_{\text{SR}} \quad (4.1)$$

where $\mathbf{G}_{\text{SP-SQ}}(a, r)$ is the generator matrix of the joint process of SP and SQ. Notice that $\mathbf{G}_{\text{SYS}}(a)$ generator matrix is also a parameterized matrix of action a .

The transition of SP from a busy state to its corresponding idle state is correlated with the transition of SQ from state (lq_i, hq_i) to state (lq_i, hq_{i-1}) or the transition of SQ from state $(lq_i, 0)$ to state $(lq_{i-1}, 0)$. This is because whenever the SP makes a transition from a busy state to an idle state (finishes the service for a request), the SQ makes a transition from state (lq, hq) to state (lq', hq') where $lq+hq=lq'+hq'+1$. The other transitions of SP and SP are independent. Therefore, we can calculate each entry of $\mathbf{G}_{\text{SP-SQ}}$ as this:

Let $\sigma_{x,x'}$ denote the transition rate for the transition from state $x=(s, (lq, hq))$ to state $x'=(s', (lq', hq'))$.

If s is a busy state, s' is the idle state corresponding to s and $(lq+hq)-(lq'+hq')=1$, then $\sigma_{x,x'}$ equals $\sigma_{(lq, hq), (lq', hq')}$, which is the transition rate of SQ from state (lq, hq) to state (lq', hq') .

If s is a busy state and s' is the corresponding idle state, but $(lq+hq) - (lq'+hq') \neq 1$, then $\sigma_{x,x'}$ equals 0.

If s is a busy state and $(lq+hq)-(lq'+hq')=1$, but s' is not the corresponding idle state of s , then $\sigma_{x,x'}$ equals 0.

For all the other $\sigma_{x,x'}$, we can use the value of the corresponding entry of matrix $\mathbf{G}_{\text{SP}}(a) \oplus \mathbf{G}_{\text{SQ}}(s, r)$. Notice that, after the operation, parameter s in $\mathbf{G}_{\text{SQ}}(s, r)$ has been removed by substituting the real state of the SP. Therefore, the generator matrix of $\mathbf{G}_{\text{SP-SQ}}$ is parameterized matrix which depends on variables a and r .

The values of diagonal entries of $\mathbf{G}_{\text{SP-SQ}}$ need to be recalculated using (2.4).

Example 4.1 Consider the SP and SQ model given in Example 3.1 and 3.3. The generator matrix of their joint process can be calculated like this:

$$\sigma_{(busy_1, (0,1)), (idle_1, (0,0))} = \sigma_{(0,1), (0,0)}(busy_1) = \mu_h(busy_1),$$

$$\sigma_{(busy_1, (1,0)), (idle_1, (0,0))} = \sigma_{(busy_1, (2,0)), (idle_1, (1,0))} = \mu'_l(busy_1, 0) = \mu_l(busy_1),$$

$$\sigma_{(busy_1, (1,1)), (idle_1, (0,0))} = \sigma_{(busy_1, (2,1)), (idle_1, (1,1))} = \mu'_l(busy_1, 1) = 0,$$

$$\sigma_{(busy_1, (0,0)), (idle_1, (0,1))} = \sigma_{(busy_1, (0,0)), (idle_1, (1,0))} = \sigma_{(busy_1, (0,0)), (idle_1, (1,1))} = \dots = \sigma_{(busy_1, (2,1)), (1,0)} = 0,$$

$$\sigma_{(busy_1, q), (s', q')} = 0, \text{ where } s' \neq idle_1, q, q' \in Q.$$

The value of $\sigma_{(busy_2, q), (s', q')}$ can be calculated in the same way as $\sigma_{(busy_1, q), (s', q')}$. The values of the other $\sigma_{x,x'}$, equals the value of the corresponding entry of matrix $\mathbf{G}_{\text{SP}}(a) \oplus \mathbf{G}_{\text{SQ}}(s, r)$.

□

C. Application issues

Using the SYS model, a power-managed system in real application can work in the following way. When the SP of the system changes state, it sends an interrupt signal SWITCH_DONE to the PM. The PM then reads the states of all components in the power-managed system (hence obtains the joint system state) and issues a command according to the chosen policy. The SP receives the command and immediately starts to switch to a state that is dictated by the command. Notice that the command may ask the SP to switch to its current state; therefore the SP state does not change. We assume that after the SP finishes a service, it stays in the idle state for a period of time long enough for it to accept the command from the PM and to switch to another state.

V. SOLUTION TECHNIQUES

A. Problem formulation

The power management problem is to find the optimal policy (set of state-action pairs) for the PM such that the average system power dissipation is minimized subject to the performance constraints. The performance of a system is usually measured by the average delay of each request. We have the following theorem:

Theorem 5.0 In a power-managed system, if the loss rate of request is small enough, then $D = Q \cdot \lambda$, where D is the average request delay, Q is the average number of waiting requests in the queue, and λ is the average incoming request speed. Furthermore, during any time period of length T , $E_T(d) = E_T(q) \cdot T / X$, where $E_T(d)$ and $E_T(q)$ denote the average request delay and the average number of waiting requests in the queue during time T , and X is the number of incoming requests of this system during time period T .

Therefore, we define the objective cost of the constraint Markov decision problem $c_obj_x^{a_x} = c_pow(x, a_x)$, $c_pow(x, a_x)$ denotes the power consumption of the system when it is in state x and action a_x is used. In this work, we only consider the power consumption of the SP. We will explain how to calculate the objective cost in difference solution approaches later in this section. We define two constraint cost $c_1_con_x^{a_x} = c_lsq(x)$, $c_2_con_x^{a_x} = c_hsq(x)$, $c_lsq(x)$ and $c_hsq(x)$ denotes the number of waiting requests in low priority queue and high priority queue respectively when the system is in state x . Notice that, although the constraint cost for each state is policy independent, since the state probability for each state is policy dependent, thus the total expected constraint cost is policy dependent.

The problem can be formally described as:

$$\begin{aligned} \text{Find a policy } \boldsymbol{\pi} \text{ to minimize: } & \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \sum_{x' \in X} p_{x \Rightarrow x'}^{\boldsymbol{\pi}}(\tau) \cdot c_pow(x', a^{\boldsymbol{\pi}}) d\tau, \\ \text{s.t. } & \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \sum_{x' \in X} p_{x \Rightarrow x'}^{\boldsymbol{\pi}}(\tau) \cdot c_lsq(x') d\tau \leq D_L \\ \text{and } & \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \sum_{x' \in X} p_{x \Rightarrow x'}^{\boldsymbol{\pi}}(\tau) \cdot c_hsq(x') d\tau \leq D_H, \forall x \in X \end{aligned}$$

where $p_{x \Rightarrow x'}^{\boldsymbol{\pi}}(\tau)$ is the state transition (direct or indirect) probability from state x to x' in a time period of τ under policy $\boldsymbol{\pi}$. D_H and D_L are performance constraints for high priority queue and low priority queue.

B. The linear programming approach

If the delay constraint for the power management system is an *active constraint* [19], generally, the optimal power management policy will be a randomized policy [15]. The randomized optimal policy can be obtained by solving a linear programming problem. In this section, we first introduce how to solve the unconstrained continuous-time Markov decision process using a linear programming approach, and then show how to apply this linear programming approach to our constraint power management problem.

First we introduce the definition of some variables, which are used to measure the Markov process. Let t_{ij} denotes the time that the system needs to switch from state i to state j . $Q_{ij}(t)$ denotes the probability of the event that the next observed state is state j and that state j is observed no later than time $T+t$, given that state i is observed at time T . We know that $Q_{ij}(t)$ is the probability that t_{ij} is the smallest among all t_{il} and t_{ij} is less than t , where l is the possible next state of state i . We use $t_{ij}^{a_i}$ and $Q_{ij}^{a_i}(t)$ to denote the value of t_{ij} and $Q_{ij}(t)$ when action a_i is taken.

Proposition 5.1 If t_{ij} follows exponential distribution $f(t_{ij}) = \sigma_{ij} \cdot e^{-\sigma_{ij}t}$, $j \in J$, (J is the set of all possible next states of state i), then $Q_{ij}(t) = \sigma_{ij} (1 - e^{-\sum_{l \in J} \sigma_{il}t}) / \sum_{l \in J} \sigma_{il}$.

Let $\tau_i^{a_i}$ denotes the expectation of the time that the system will be in state i if action a_i is chosen in this state, then $\tau_i^{a_i} = \sum_{j \in J} \int_0^\infty t dQ_{ij}^{a_i}(t)$, J is the set of all possible next state of i . It is easy to show that, in a Markov process, $\tau_i^{a_i} = 1 / \sum_{j \in J} \sigma_{ij}^{a_i}$. Let $p_{ij}^{a_i}$ denotes the probability that the next system state is j if the system is currently in state i and action k is taken, the $p_{ij}^{a_i} = Q_{ij}^{a_i}(\infty)$. In a Markov process we have $p_{ij}^{a_i} = \sigma_{ij}^{a_i} / \sum_{l \in J} \sigma_{il}^{a_i}$. Let $\gamma_i^{a_i}$ denotes the expected cost of the system during the time it stays in state i and action a_i is taken, then

$$\gamma_i^{a_i} = c_{ii} \tau_i^{a_i} + \sum_{j \in J} c_{ij} p_{ij}^{a_i} \quad (5.1)$$

Let $x_i^{a_i}$ denotes the frequency that the next state of the system will be i and action a_i will be taken if we take a random observation of the system. Then $x_i^{a_i} \tau_i^{a_i}$ is the probability that the system is in state i and action a_i is taken in a random observation, which is also called *state-action probability*. We know from the definition that $p_i^{a_i} = x_i^{a_i} \tau_i^{a_i} / \sum_{a_i \in A_i} x_i^{a_i} \tau_i^{a_i}$, $p_i^\pi = \sum_{a_i \in A_i} x_i^{a_i} \tau_i^{a_i}$ and $\sum_{i \in S} \sum_{a_i \in A_i} x_i^{a_i} \tau_i^{a_i} = 1$. We also know that if a Markov process is stationary, the input rate to each state needs to be equal to the output rate from that state [15]. That means $\sum_{a_i} x_i^{a_i} - \sum_j \sum_{a_i} x_j^{a_i} p_{ji}^{a_i} = 0$.

For a given policy, if we know that the resulting Markov process is irreducible, Theorem 5.1 shows that we can use the variable x and γ to calculate the limiting average cost.

Theorem 5.1 If a Markov process is irreducible, then $\sum_{i \in S} \sum_{a_i \in A_i} x_i^{a_i} \gamma_i^{a_i} = C_{i,avg}^\pi \quad \forall i \in S$.

To prove this, we need the following lemmas:

Lemma 5.1 Given a Markov process, $\sum_{i \in S} \sum_{a_i \in A_i} x_i^{a_i} \gamma_i^{a_i} = \sum_{i \in S} p_i^\pi c_i^{A_i}$.

Lemma 5.2 If the Markov process is irreducible, there exists a $C_{avg}^\pi, C_{i,avg}^\pi = C_{avg}^\pi \quad \forall i \in S$. Furthermore,

$$\sum_{i \in S} p_i^\pi c_i^{A_i} = C_{avg}^\pi.$$

Using these lemmas, the statement in Theorem 5.1 is very obvious.

With the above-mentioned variables and their characteristics, we can write the linear programming as [15]:

$$\text{LP1:} \quad \text{Minimize}_{\{x_i^{a_i}\}} \left(\sum_i \sum_{a_i} x_i^{a_i} \gamma_i^{a_i} \right), \quad a_i \in \mathcal{A}_i \quad (5.2)$$

$$\text{subject to } \sum_{a_i} x_i^{a_i} - \sum_j \sum_{a_j} x_j^{a_j} p_{ji}^{a_i} = 0 \quad i \in \mathcal{S} \quad (5.3)$$

$$\sum_i \sum_{a_i} x_i^{a_i} \tau_i^{a_i} = 1 \quad (5.4)$$

$$x_i^{a_i} \geq 0 \quad \text{all } i, a_i \quad (5.5)$$

Since the action in each state is unknown in the Markov decision process, in general cases, we cannot guarantee that for every possible policy the resulting Markov process is irreducible. For those resulting Markov processes, expression (5.2) may not equal to their average cost because Theorem 5.1 is no longer applicable. However we have the following theorem:

Lemma 5.3 [15] Let $\{x_i^{a_i}\}$ be a basic feasible solution to LP1. Then set $E = \{i, \sum_{a_i} x_i^{a_i} > 0\}$ and $F = \{(i, a_i) : x_i^{a_i} > 0\}$ identify a unique ergodic chain, and expression (5.2) is the cost rate for this chain.

Theorem 5.2 [15] LP1 is feasible and bounded. The positive variables in its optimal solution identify the states of an ergodic chain whose cost rate is minimized over all chains.

In our power management system, we set constraints on the action sets so that the Markov process model of the resulting power management system contains only one ergodic chain. Therefore, expression (5.2) gives the cost rate for the entire system and the solution of LP1 identifies an optimal policy.

In our case, we add the delay constraint to the linear program, and formulate it as:

$$\text{LP2:} \quad \text{Minimize}_{\{x_i^{a_i}\}} \left(\sum_i \sum_{a_i} x_i^{a_i} c_{-pow_i^{a_i}} \right) \quad (5.6)$$

$$\text{subject to } \sum_{a_i} x_i^{a_i} - \sum_j \sum_k x_j^{a_j} p_{ji}^{a_i} = 0 \quad i \in \mathcal{X} \quad (5.7)$$

$$\sum_i \sum_{a_i} x_i^{a_i} \tau_i^{a_i} = 1 \quad (5.8)$$

$$x_i^{a_i} \geq 0 \quad \text{all } i, a_i \quad (5.9)$$

$$\sum_i \sum_{a_i} x_i^{a_i} c_{-hsq_i^{a_i}} < D_H \quad (5.10)$$

$$\sum_i \sum_{a_i} x_i^{a_i} c_{-lsq_i^{a_i}} < D_L \quad (5.11)$$

In LP2, $c_{-pow_i^{a_i}}$, $c_{-hsq_i^{a_i}}$, $c_{-lsq_i^{a_i}}$ denote the power consumption of the system, the high priority request waiting cost and the low priority request waiting cost during the time it stays in state i and action a_i is taken. They can be calculated based on (5.1): $c_{-pow_i^{a_i}} = pow_i \cdot \tau_i^{a_i} + \sum_j ene_{ij} \cdot p_{ij}^{a_i}$, $c_{-hsq_i^{a_i}} = hq_i \cdot \tau_i^{a_i}$, $c_{-lsq_i^{a_i}} = lq_i \cdot \tau_i^{a_i}$. D_H and D_L are performance constraints for high priority queue and low priority queue. If constraint (5.10) and (5.11) are implied in the constraints (5.7), (5.8) and (5.9), (i.e. they are inactive constraint), the resulting optimal policy must be a deterministic policy. Otherwise, the resulting optimal policy will be a randomized policy.

C. The nonlinear programming approach

Because a randomized policy may be undesirable in some applications, we are interested in finding the optimal deterministic policy instead of the optimal randomized policy. In a deterministic policy, for each state i , there is one and only one $x_i^{a_i}$, $a_i \in A_i$, which is non-zero. Therefore, we can formulate the problem of searching for an optimal deterministic policy into a nonlinear program NLP1, which has a polynomial objective function and a set of linear constraints. In NLP1 λ is an arbitrary large number.

$$\begin{aligned}
 \text{NLP1:} \quad & \text{Minimize}_{\{x_i^{a_i}\}} (\lambda \sum_i \sum_{a_i \neq l, a_i, l \in A_i} x_i^{a_i} \cdot x_i^l + \sum_i \sum_{a_i} x_i^{a_i} c_{-pow_i^{a_i}}) \\
 & \text{subject to } \sum_{a_i} x_i^{a_i} - \sum_j \sum_{a_i} x_j^{a_i} p_{ji}^{a_i} = 0 \quad i \in X \\
 & \sum_i \sum_{a_i} x_i^{a_i} \tau_i^{a_i} = 1 \\
 & x_i^{a_i} \geq 0 \quad \text{all } i, a_i \\
 & \sum_i \sum_{a_i} x_i^{a_i} c_{-hsq_i^{a_i}} < D_H, \quad \sum_i \sum_{a_i} x_i^{a_i} c_{-lsq_i^{a_i}} < D_L
 \end{aligned}$$

There are specific algorithms to solve this kind of nonlinear programming problem. A classical algorithm is *feasible direction algorithm* which can be found in [16].

D. The branch-bound approach

Another way to find the optimal deterministic policy is to use a branch and bound approach. In a power management system, making different decisions in one system state has significant effect on the power and performance of the system. For example, given two policies, the first one chooses action “Go_busy” when the server is idle and there are 1 request waiting, however, the second one chooses action “Go_sleeping” when the system is in the same state. No matter how the actions are chosen in the other state, the best performance of the system using the second policy will be worse than that of the system using the first policy, because there is always at least one request waiting in the queue. This observation is the root cause of the efficiency of the branch and bound algorithm.

```

I: set of states which have deterministic policy
S: set of states which have non-deterministic policy
K: set of actions for state I, i ∈ I
I = ∅, S = {all states}
Search_Deterministic_Policy(I, S, K)
{
  if (S ≠ ∅) {
    pop state i from S, push i into I;
    for each action k ∈ A_i {
      set action of state a to be k, push k in K.
      pow = power_of_randomized_policy_of_parital_decision_problem()
      if (pow < min_pow) Search_Deterministic_Policy(I, S, K);
      pop k from K;
    }
  }
  else {
    if the policy given by K satisfies the delay constraints and has less pow than
    minpow {
      record K, and update minpow;
    }
  }
}

```

Figure 7 Algorithm of searching for an optimal deterministic policy using branch and bound.

The optimal deterministic power-management policy decision tree is a level X full tree, where X is the number of states in the controllable Markov process. Each node in level x has A_x child nodes, where A_x is number of available action in state x . Therefore for each leaf in the decision tree there is a corresponding deterministic power-management policy. We are going to search for the power optimal performance constraint power-management policy based on the decision tree using a branch and bound algorithm.

For each internal node in the decision tree, there is a partial decision problem with the same constraints and objective as the original problem. We call it partial decision problem because its variables are a subset of that of the original problem. We can find an optimal randomized policy for this partial decision problem using a linear programming approach. The power consumption of this policy is the lower bound of the policies in this branch. Therefore, we define the branch operation as: finding the optimal randomized policy for the partial decision problem; We define the prune operator as: comparing the power consumption of that randomized policy with that of the best deterministic policy we already have. If it consumes less power then we continue with this branch, otherwise, we prune this branch. If we cannot find a policy that satisfies the performance constraint for the partial decision problem, then we also prune this branch. The algorithm is given in Figure 7.

E. The policy iteration approach

The policy iteration algorithm is widely used in searching for an optimal policy without constraint [17]. It starts from a set of randomly selected actions for each state that is called the “initial policy”. It then calculates the expected cost of the system under this policy, which is called the “reference cost”. Using this reference cost, by some calculation we can find a new policy that has lower cost than the initial policy. The process is iterated until the resulting policy cannot be improved further. Because of the mathematical nature of the policy optimization problem, this simple greedy algorithm yields the provably optimal solution. Experiments show that the “policy iteration” technique is a fast technique. For a system with 23 states and assuming that three commands are available for each state, the “policy iteration” algorithm can find out the optimal solution within 4 iterations. However, the policy iteration approach cannot directly be applied to the problem of constrained optimization. We make some modification to this approach.

We first consider the power-managed system with only single priority SQ. We define a joint cost as a weighted summation of the power and delay costs as:

$$joint_cost_x^{a_x} = c_pow_x^{a_x} + w \cdot c_sq_x \quad (5.12)$$

Where $c_pow_x^{a_x}$ denotes the average power consumption when the system is in state x and action a_x is chosen, c_sq_x is the number of waiting request in sq when the system is in state x . Let x be denoted by (s, q_i) , where $s \in \mathbf{S}$, $q_i \in \mathbf{Q}$. Using (2.5) the power cost can be calculated as: $c_pow_x^{a_x} = pow(s) + \sum_{s' \in \mathbf{S}, s' \neq s} \sigma_{s,s'}(a_x) ene(s, s')$. The

delay cost is: $c_sq_x = i$. The constraint policy optimization procedure is given in Figure 8. We start from a randomly selected weight w . Then we find a policy which has optimal $joint_cost$ using policy iteration algorithm. If the delay of this policy is approximately same as the given constraint, then report this policy, otherwise, we increase or decrease the weight and continue searching using the new weight value.

This modified policy iteration algorithm is a heuristic algorithm. For certain performance constraint, it may not find the policy consumes the minimum power. However, we can prove later that its output policy is the one that consumes minimum power among a class of *convex policies* which satisfies the performance constraint. We first gives the definition of an *effective policy* and a *convex policy*:

Definition 5.1 An effective policy is a policy that consumes the minimum power comparing to other policies which have the same or better performance.

Definition 5.2 If we sorted the policy based on the increasing order of their power consumption and denote the power and delay of a policy p_i as pow_i and $delay_i$. A policy p_i is a convex policy if it is an effective policy and $(delay_i - delay_j) / (pow_j - pow_i) < (delay_l - delay_j) / (pow_l - pow_i)$, $\forall j, j > i$ and $\forall l, l < i$.

Comparing to the policies that consumes less power, a convex policy has the better power-delay trade off (smaller ratio of increasing delay to the decreasing power). Figure 9 gives an illustration of convex policy and effective policy. In that figure, policy a, b, c, d, e are all effective policies. However, policy c is not a convex policy.

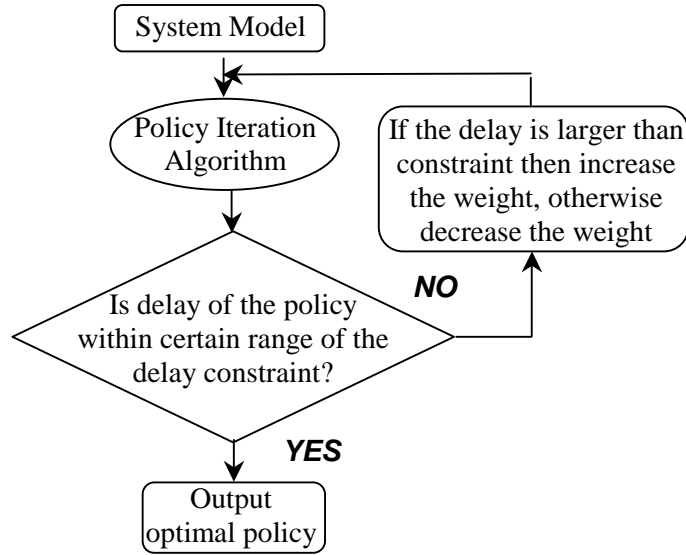


Figure 8 Policy optimization workflow.




Figure 9 Example of effective, inferior and superior policy.

Theorem 5.3 The output of the modified policy iteration algorithm is a policy that consumes the minimum power among the class of convex policies which satisfies the performance constraint.

To prove this theorem we will need the following lemma.

Lemma 5.4 Given a set of points S on the 2D plan. Let a convex point to be a point that is a vertex of the minimum convex hull of these points. $\forall(x', y')$, if and only if (x', y') is a convex point of S , $\exists w, x'+w \cdot y' \leq x+w \cdot y$, $(x, y) \in S$.

We will not provide a strict mathematical proof for the theorem and lemma; however, we will show that this theorem is true geometrically in Figure 10. Because (x', y') is a convex point, we can find a line l which pass this point and all the other point in S are on the same side of line l . Let line k denote another line $z = x$ and $y = 0$. Let a denote an arbitrary point on line k . From a we can draw a line l' which is parallel to line l . We can always construct a 3D plan H based on the two lines k and l' . It is easy to see that the function of H has the form $z = x + w \cdot y$, and the value $z' = x' + w \cdot y'$ is smaller than or equal to $z = x + w \cdot y$, where $(x, y) \in S$. For all these non-convex points, we cannot construct such a plan.

Given a set of policies and their power consumption and delay value pow and $delay$. We can associate each policy with a pointer on 2D plan by letting $x=delay$, $y=pow$. It is easy to see that the points associated with convex

policies are the convex points. From Lemma 5.4 we know that if a policy has the minimum joint cost then it is a convex policy; if a policy is a convex policy then we can use our algorithm to find it by adjusting the value of w .




Figure 10 Illustration of Lemma 5.4.

In a power management system model with multiple priority queue, the similar cost function is used: $cost_x^{a_x} = c_pow_x^{a_x} + w_1 \cdot c_lsq_x + w_2 \cdot c_hsq_x$. Similarly, by adjusting both the weight coefficient w_1, w_2 , we can find the convex policy.

The policy iteration algorithm can be used as a fast algorithm to do on-line policy optimization when system parameters change at runtime because it finds the optimal policy based on the previous optimal policy. It is more efficient than the branch and bound algorithm as well as the NLP approach which starts solving from the very beginning.

VI. EXPERIMENTAL RESULTS

Experiments have been designed to examine the performance of our system model and optimization method.

In the first part of the experiment, we will present the comparison of our power management policy for a simple DPM system with an SSQ with the heuristic policies including greedy policy, time-out policy, and N-policy.

6.1 Experimental setups using the basic DPM model

In this experimental setup, the service requestor (SR) is implemented as an input trace file that stores the time instances when the disk drive read/write requests arrive. The service queue (SQ) and service provider (SP) are implemented in the event-driven simulator. The power manager (PM) is also implemented in the simulator that controls the state transition of the SP based on the policy it reads from user-specified input file. Q , the capacity of the SQ, is set to 20.

We will use two abstract models of the SP for the Fujitsu disk drive [24] in the experiments, which are shown in Figure 11.

When the SP is in the “sleep” and “standby” state, it is inactive and therefore no request can be serviced. When the SP is in the “idle” state, the SP is active; however, it is not working on any requests, which means that the SQ is empty. The SP makes a transition automatically from “idle” to “busy” whenever a request arrives for service. Similarly, the SP makes a transition automatically from “busy” to “idle” state whenever it finishes the service for a request. The transitions between the “sleep,” “standby,” and “idle” are controlled by the PM, i.e., the DPM policy.




Figure 11 Three-state and four-state SP models.

The average power consumption and service time (time to finish a read/write operation) for each state of the SP is shown in Table 1. The average energy consumption values for the SP to make transitions among the various states are shown in Table 2. The average transition times for the SP to make transitions among the various states are shown in Table 3. The data is obtained from [24] and [21].

Table 1 Average power consumption values and service times of the SP.

State	“sleep”	“standby”	“idle”	“busy”
Ave. power (W)	0.13	0.35	0.95	2.15
Ave. service time (s)	0	0	0	0.008


Table 2 Average energy consumption values for state transition of the SP.

Ave. energy (J)	“sleep”	“standby”	“idle”	“busy”
“sleep”	0	5.1	7.0	-
“standby”	0.006	0	2	-
“idle”	0.067	0.001	0	0
“busy”	-	-	0	0


Table 3 Average transition times for state transition of the SP.

Ave. transition time (s)	“sleep”	“standby”	“idle”	“busy”
“sleep”	0	0.6	1.6	-
“standby”	0.3	0	1.2	-
“idle”	0.67	0.4	0	0
“busy”	-	-	0	0


We have used five different distributions for the request inter-arrival time. They are:




Inter-arrival time (s)
(a) Exponential distribution




Inter-arrival time (s)
(b) Combination of exponential and Pareto distributions



Inter-arrival time (s)
(c) Combination of two exponential distributions



Inter-arrival time (s)
(d) Uniform distribution



Inter-arrival time (s)
(e) Normal distribution

Figure 12 Five distributions of request inter-arrival time.

1. Exponential distribution. This is the key assumption made by the continuous-timed Markov decision process (CTMDP) approach. For the CTMDP policy to be optimal, the request inter-arrival time must follow exponential distribution. In many research works related to hard drive performance, it is found that the request inter-arrival time follows the exponential distribution. Figure 12 (a) gives the illustration of the density function of exponential distribution.
2. Combination of the exponential and Pareto distribution. This combination is tried because of the observation made in [21]. The Pareto distribution has a density function: $f(t)=1-at^{-b}$. Figure 12 (b) gives the illustration of the density function of exponential and Pareto distribution. Compared with exponential distribution, it models an input sequence that shows more bursts.
3. Combination of two exponential distributions. We are using this distribution to approximate distribution 2. Figure 12 (c) gives the illustration of the density function for the distribution.
4. Uniform distribution whose density function is illustrated in Figure 12 (d).
5. Normal distribution whose density function is illustrated in Figure 12 (e).

We have used two different distributions for the state transition time of the SP. They are:

1. Exponential distribution. This is another assumption made by the CTMDP approach. For the CTMDP policy to be optimal, the SP transition time must follow exponential distribution.
2. Uniform distribution. This distribution is proposed and used by the authors of [21] in their work.

The mean values we used for these distributions are shown in Table 3. The deviation for the uniform distribution is set to 0.1s, which is again observed in [21].

The policies we have used in this work are as follows:

1) *The “time-out” policy*

The “time-out” policy has a single parameter T_{out} . This policy can only be used for the three-state system model shown in Figure 11 (a). Under the “time-out” policy, the PM switches the SP from “idle” to “sleep” after the SP has been in the idle state and the SQ has been empty for a time period of T_{out} ; the PM switches the SP from “sleep” to “idle” immediately after a request has arrived.

In our experimental setup, we have used different values of T_{out} to study the performance and power characteristics of the “time-out” policy.

2) *The N-Policy*

The N-policy has a single parameter N , which must be smaller than the capacity of the SQ, Q . Similar to the “time-out” policy; this policy is only applicable to the three-state system model. Under the N-policy, the PM switches the SP from “idle” to “sleep” immediately after the SQ is empty; the PM switches SP from “sleep” to “idle” only after there are more than N requests waiting in the SQ.

In our experimental setup, we have used different values of N to study the performance and power trade-offs of the N-policy.

3) *The “always on” and “greedy” policies*

The “always on” policy can be regarded as a special case of either the “time-out” policy when T_{out} is infinity, or the N-policy when N is zero. As indicated by its name, the “always on” policy never turns off the SP.

The “greedy” policy can be regarded as a special case of either the “time-out” policy when T_{out} is zero, or the N-policy when N is one.

4) *The CTMDP policy*

The CTMDP policy we use as described in Chapter 3. This policy is applied to both the three-state model and the four-state model in Figure 11. Given any performance constraint, the power-optimal CTMDP policy can be calculated by solving a linear program.

In our experiments, a series of CTMDP policies were generated for different performance constraints. The CTMDP policies were applied using both the three-state SP model and the four-state SP model. Then the simulator and real traces were used to evaluate this policy in real applications.

5) *The 3CTMDP-Poll Policy*

Because the CTMDP policy is a randomized policy, at times it may not turn off the SP even when there is no request in the SQ. If our stochastic model exactly represents the system behavior, then this policy is optimal. However, in practice, because the stochastic model is not accurate enough, the CTMDP policy may cause unnecessary energy dissipation by not turning off the SP. For example, the real requests pattern on the SP may be quite different from what has been assumed in theory, and the SP idle time may be much longer than we expected based on the assumption of exponential input inter-arrival time. In this case keeping the SP on while it is idle wastes much power. We need to put some mechanism in place that will prevent such a case in real applications.

Based on the original description of the CTMDP policy, we have designed a modified CTMDP policy by adding a *polling state*. The functionality of the polling state is very simple. After adding this state, if the CTMDP policy allows the SQ to stay on when the SQ is empty, the policy will re-evaluate this decision after some random-length period of time. For example, if the SQ is empty and the PM has made a decision (with probability of 0.1) of letting the SQ stay on, then after 2s, if there is no change in the SQ, we enter the polling state, and the PM has to re-evaluate its decision. At this time, the probability for it to still let SQ remain on is again 0.1. So as the time goes on, the total probability of the SQ remaining on reduces in a geometric manner. In this way, we can make sure that the SP will not be idle for too long, resulting in less wasteful energy dissipation. We name this modified policy the **3CTMDP with Polling** (3CTMDP-Poll) policy and implement it in our experiments based on 3CTMDP policy.

The simulation results are presented in the following tables. We use the average number (#) of waiting requests in the SQ as the performance metric and the average power consumption (W) as the power metric.

Table 4 Simulated performance and power values for the N-policy.

N	Transition time distribution	Exponential					Uniform				
	Inter-arrival time distribution	Exp.	Exp. & Par.	Exp. & Exp.	Uni.	Nor.	Exp.	Exp. & Par.	Exp. & Exp.	Uni.	Nor.
0	Perf.	0.011	0.003	0.007	0.011	0.110	0.011	0.003	0.007	0.011	0.110
	Power (W)	0.963	0.955	0.959	0.963	1.082	0.963	0.955	0.959	0.963	1.082
1	Perf.	1.737	0.453	1.110	1.715	1.776	1.016	0.264	0.657	0.986	0.934
	Power (W)	2.658	0.809	1.800	2.868	2.926	2.453	0.738	1.629	2.743	2.636
2	Perf.	2.016	0.693	1.352	2.035	2.110	1.397	0.510	0.969	1.414	1.362
	Power (W)	2.166	0.674	1.462	2.314	2.370	2.054	0.632	1.368	2.240	2.215
3	Perf.	2.381	0.982	1.709	2.391	2.486	1.823	0.822	1.321	1.865	1.815
	Power (W)	1.849	0.583	1.258	1.930	2.003	1.772	0.558	1.185	1.902	1.928
4	Perf.	2.777	1.335	2.100	2.820	2.904	2.256	1.187	1.720	2.332	2.280
	Power (W)	1.606	0.520	1.099	1.677	1.755	1.561	0.503	1.051	1.659	1.717
5	Perf.	3.183	1.760	2.502	3.237	3.347	2.712	1.579	2.142	2.809	2.754
	Power (W)	1.420	0.470	0.981	1.470	1.576	1.403	0.461	0.949	1.479	1.554
8	Perf.	4.577	3.031	3.872	4.646	4.718	4.117	2.884	3.536	4.254	4.201
	Power (W)	1.094	0.382	0.749	1.117	1.222	1.082	0.378	0.742	1.121	1.239
10	Perf.	5.494	3.917	4.818	5.571	5.686	5.079	3.818	4.420	5.230	5.174
	Power (W)	0.944	0.343	0.658	0.960	1.085	0.947	0.343	0.653	0.975	1.108
13	Perf.	6.915	5.340	6.252	6.998	7.111	6.539	4.932	5.900	6.714	6.657
	Power (W)	0.802	0.305	0.562	0.807	0.939	0.806	0.305	0.562	0.821	0.969
15	Perf.	7.880	6.283	7.137	7.973	8.081	7.532	6.257	6.761	7.688	7.637
	Power (W)	0.734	0.284	0.516	0.735	0.875	0.736	0.287	0.519	0.750	0.902

Table 5 Simulated performance and power values for the “time-out” policy.

T_{out} (s)	Transition time distribution	Exponential					Uniform				
	Inter-arrival time distribution	Exp.	Exp. & Par.	Exp. & Exp.	Uni.	Nor.	Exp.	Exp. & Par.	Exp. & Exp.	Uni.	Nor.
0.1	Perf.	1.606	0.421	1.054	1.607	1.651	0.943	0.247	0.613	0.910	0.740
	Power (W)	2.539	0.778	1.740	2.756	2.806	1.409	0.776	0.946	1.561	1.896
0.3	Perf.	1.408	0.370	0.910	1.375	1.386	0.796	0.212	0.525	0.746	0.301
	Power (W)	2.332	0.724	1.595	2.526	2.521	1.544	0.767	1.035	1.745	2.200
0.5	Perf.	1.186	0.309	0.794	1.137	0.973	0.659	0.179	0.445	0.522	0.127
	Power (W)	2.127	0.665	1.463	2.269	2.074	1.675	0.757	1.115	1.948	2.312
0.7	Perf.	0.971	0.265	0.667	0.892	0.503	0.542	0.149	0.374	0.187	0.111
	Power (W)	1.919	0.618	1.346	1.983	1.525	1.789	0.748	1.188	2.194	2.322
0.9	Perf.	0.803	0.218	0.545	0.586	0.187	0.445	0.121	0.305	0.011	0.110
	Power (W)	1.755	0.573	1.228	1.647	1.168	1.887	0.739	1.252	2.302	2.322
1.2	Perf.	0.596	0.167	0.406	0.175	0.112	0.317	0.088	0.223	0.011	0.110
	Power (W)	1.532	0.518	1.089	1.157	1.083	2.008	0.728	1.331	2.302	2.322
1.5	Perf.	0.418	0.127	0.336	0.011	0.110	0.219	0.073	0.192	0.011	0.110
	Power (W)	1.362	0.479	1.009	0.963	1.082	2.100	0.729	1.380	2.302	2.322
1.7	Perf.	0.326	0.107	0.275	0.011	0.110	0.174	0.069	0.186	0.011	0.110
	Power (W)	1.273	0.460	0.963	0.963	1.082	2.145	0.734	1.404	2.302	2.322
1.9	Perf.	0.263	0.095	0.263	0.011	0.110	0.138	0.066	0.181	0.011	0.110
	Power (W)	1.209	0.452	0.952	0.963	1.082	2.181	0.739	1.428	2.302	2.322
2.5	Perf.	0.119	0.082	0.247	0.011	0.110	0.067	0.058	0.168	0.011	0.110
	Power (W)	1.070	0.445	0.951	0.963	1.082	2.249	0.754	1.495	2.302	2.322
5	Perf.	0.013	0.056	0.180	0.011	0.110	0.012	0.042	0.122	0.011	0.110
	Power (W)	0.966	0.442	0.955	0.963	1.082	2.301	0.811	1.726	2.302	2.322
10	Perf.	0.011	0.038	0.090	0.011	0.110	0.011	0.028	0.065	0.011	0.110
	Power (W)	0.963	0.457	0.954	0.963	1.082	2.302	0.897	2.012	2.302	2.322
20	Perf.	0.011	0.027	0.034	0.011	0.110	0.011	0.019	0.022	0.011	0.110
	Power (W)	0.963	0.490	0.960	0.963	1.082	2.302	1.014	2.229	2.302	2.322
40	Perf.	0.011	0.018	0.008	0.011	0.110	0.011	0.013	0.008	0.011	0.110
	Power (W)	0.963	0.540	0.959	0.963	1.082	2.302	1.170	2.298	2.302	2.322
80	Perf.	0.011	0.013	0.007	0.011	0.110	0.011	0.009	0.007	0.011	0.110
	Power (W)	0.963	0.609	0.959	0.963	1.082	2.302	1.364	2.301	2.302	2.322

Table 6 Simulated performance and power values for the CTMDP policy using a three-state SP model.

Perf. Cons.	Transition time distribution	Exponential					Uniform				
	Inter- arrival time distribution	Exp.	Exp. & Par.	Exp. & Exp.	Uni.	Nor.	Exp.	Exp. & Par.	Exp. & Exp.	Uni.	Nor.
0.01	Perf.	0.017	0.030	0.011	0.019	0.114	0.017	0.004	0.009	0.018	0.115
	Power (W)	0.963	0.954	0.958	0.963	1.082	0.963	0.955	0.959	0.963	1.082
0.02	Perf.	0.020	0.121	0.013	0.026	0.128	0.021	0.009	0.016	0.024	0.120
	Power (W)	0.963	0.950	0.958	0.963	1.083	0.963	0.954	0.958	0.963	1.082
0.03	Perf.	0.034	0.164	0.025	0.043	0.142	0.026	0.030	0.026	0.030	0.136
	Power (W)	0.963	0.951	0.957	0.962	1.082	0.963	0.952	0.957	0.963	1.082
0.04	Perf.	0.041	0.093	0.030	0.044	0.140	0.041	0.009	0.018	0.040	0.141
	Power (W)	0.963	0.953	0.956	0.963	1.081	0.963	0.952	0.958	0.963	1.082
0.05	Perf.	0.056	0.175	0.047	0.058	0.165	0.039	0.016	0.028	0.054	0.150
	Power (W)	0.963	0.952	0.955	0.963	1.082	0.963	0.953	0.958	0.963	1.082
0.1	Perf.	0.094	0.342	0.073	0.111	0.234	0.075	0.036	0.062	0.101	0.184
	Power (W)	0.962	0.947	0.956	0.963	1.081	0.962	0.951	0.956	0.963	1.082
0.15	Perf.	0.141	0.612	0.134	0.156	0.296	0.105	0.085	0.084	0.158	0.252
	Power (W)	0.960	0.939	0.952	0.961	1.080	0.962	0.947	0.955	0.962	1.082
0.2	Perf.	0.199	0.702	0.154	0.217	0.324	0.144	0.080	0.139	0.174	0.269
	Power (W)	0.959	0.941	0.951	0.960	1.080	0.962	0.946	0.951	0.962	1.081
0.25	Perf.	0.235	0.896	0.185	0.246	0.418	0.178	0.154	0.183	0.242	0.331
	Power (W)	0.959	0.934	0.948	0.961	1.082	0.961	0.937	0.947	0.962	1.081
0.5	Perf.	0.476	1.811	0.459	0.550	0.720	0.369	0.301	0.286	0.482	0.556
	Power (W)	0.955	0.911	0.936	0.955	1.080	0.958	0.917	0.943	0.959	1.080
1	Perf.	0.943	3.729	0.830	1.046	1.301	0.723	0.597	0.679	0.983	1.013
	Power (W)	0.950	0.868	0.908	0.949	1.072	0.955	0.878	0.918	0.954	1.077
2	Perf.	1.977	7.682	1.711	2.184	2.430	1.569	1.203	1.365	1.965	1.970
	Power (W)	0.937	0.764	0.864	0.939	1.061	0.943	0.800	0.877	0.947	1.074
3	Perf.	3.129	12.035	2.680	3.266	3.501	2.530	1.835	2.283	2.963	2.933
	Power (W)	0.928	0.663	0.803	0.932	1.051	0.933	0.721	0.816	0.938	1.069
4	Perf.	4.181	15.670	3.641	4.392	4.649	3.624	2.765	3.133	4.065	4.032
	Power (W)	0.901	0.576	0.741	0.915	1.039	0.921	0.602	0.763	0.928	1.064
5	Perf.	5.405	20.563	4.702	5.512	5.753	4.924	3.852	4.327	5.198	5.136
	Power (W)	0.892	0.422	0.671	0.894	1.045	0.901	0.427	0.673	0.920	1.058

Table 7 Simulated performance and power values for the CTMDP-Poll policy using a three-state SP model.

<i>Perf.</i> <i>Cons.</i>	Transition time distribution	Exponential					Uniform				
	Inter-arrival time distribution	Exp.	Exp. & Par.	Exp. & Exp.	Uni.	Nor.	Exp.	Exp. & Par.	Exp. & Exp.	Uni.	Nor.
0.01	Perf.	0.020	0.015	0.019	0.014	0.018	0.016	0.020	0.020	0.019	0.115
	Power (W)	0.963	0.936	0.963	0.963	0.963	0.963	0.933	0.958	0.963	1.082
0.02	Perf.	0.035	0.042	0.034	0.031	0.026	0.024	0.045	0.028	0.028	0.122
	Power (W)	0.962	0.899	0.963	0.963	0.963	0.963	0.907	0.957	0.963	1.082
0.03	Perf.	0.031	0.078	0.049	0.043	0.049	0.027	0.074	0.059	0.033	0.127
	Power (W)	0.963	0.878	0.962	0.963	0.963	0.963	0.870	0.955	0.963	1.082
0.04	Perf.	0.062	0.092	0.048	0.052	0.057	0.045	0.085	0.069	0.045	0.144
	Power (W)	0.963	0.861	0.963	0.963	0.962	0.963	0.854	0.953	0.963	1.082
0.05	Perf.	0.067	0.136	0.064	0.068	0.067	0.059	0.136	0.062	0.057	0.145
	Power (W)	0.962	0.835	0.963	0.962	0.964	0.963	0.831	0.954	0.963	1.082
0.1	Perf.	0.124	0.237	0.131	0.119	0.109	0.093	0.240	0.138	0.099	0.195
	Power (W)	0.960	0.748	0.960	0.960	0.962	0.962	0.755	0.947	0.962	1.082
0.15	Perf.	0.185	0.315	0.185	0.209	0.196	0.159	0.319	0.244	0.150	0.222
	Power (W)	0.959	0.699	0.960	0.960	0.960	0.961	0.702	0.939	0.962	1.082
0.2	Perf.	0.261	0.470	0.256	0.255	0.238	0.182	0.459	0.289	0.192	0.279
	Power (W)	0.959	0.653	0.960	0.959	0.960	0.961	0.641	0.933	0.962	1.081
0.25	Perf.	0.302	0.509	0.302	0.321	0.309	0.225	0.523	0.359	0.251	0.334
	Power (W)	0.960	0.615	0.960	0.960	0.958	0.961	0.625	0.926	0.961	1.081
0.5	Perf.	0.609	0.866	0.615	0.625	0.600	0.498	0.725	0.643	0.481	0.538
	Power (W)	0.954	0.528	0.954	0.953	0.957	0.956	0.531	0.904	0.959	1.080
1	Perf.	1.238	1.359	1.195	1.192	1.241	0.957	1.224	1.248	0.949	0.981
	Power (W)	0.951	0.451	0.950	0.949	0.949	0.952	0.456	0.853	0.955	1.077
2	Perf.	2.391	2.289	2.333	2.377	2.407	1.878	2.005	2.283	1.927	1.955
	Power (W)	0.931	0.390	0.931	0.930	0.935	0.939	0.394	0.766	0.947	1.073
3	Perf.	3.455	2.963	3.470	3.529	3.484	2.904	2.632	3.044	2.971	2.976
	Power (W)	0.915	0.362	0.922	0.929	0.922	0.927	0.361	0.710	0.937	1.068
4	Perf.	4.528	3.411	4.514	4.501	4.485	3.994	3.128	3.824	4.039	4.042
	Power (W)	0.906	0.341	0.910	0.900	0.901	0.914	0.344	0.660	0.930	1.063
5	Perf.	5.535	4.191	5.527	5.505	5.529	5.062	3.826	4.539	5.195	5.161
	Power (W)	0.894	0.332	0.898	0.890	0.891	0.901	0.332	0.630	0.920	1.059

Table 8 Simulated performance and power values for a CTMDP policy using a four-state SP model.

Perf. Cons.	Transition time distribution	Exponential					Uniform				
	Inter- arrival time distribution	Exp.	Exp. & Par.	Exp. & Exp.	Uni.	Nor.	Exp.	Exp. & Par.	Exp. & Exp.	Uni.	Nor.
0.01	Perf.	0.027	0.010	0.022	0.032	0.127	0.029	0.014	0.025	0.028	0.128
	Power (W)	0.962	0.953	0.957	0.962	1.081	0.962	0.952	0.957	0.962	1.081
0.02	Perf.	0.027	0.010	0.023	0.031	0.133	0.028	0.014	0.022	0.026	0.129
	Power (W)	0.963	0.952	0.957	0.963	1.081	0.962	0.952	0.957	0.963	1.082
0.03	Perf.	0.042	0.017	0.034	0.044	0.151	0.036	0.020	0.024	0.039	0.140
	Power (W)	0.962	0.952	0.955	0.962	1.082	0.962	0.950	0.957	0.962	1.081
0.04	Perf.	0.049	0.025	0.038	0.062	0.168	0.042	0.044	0.041	0.050	0.149
	Power (W)	0.962	0.948	0.955	0.962	1.081	0.962	0.943	0.954	0.962	1.081
0.05	Perf.	0.063	0.039	0.052	0.071	0.177	0.056	0.038	0.046	0.064	0.159
	Power (W)	0.961	0.946	0.954	0.961	1.081	0.961	0.944	0.954	0.961	1.081
0.1	Perf.	0.124	0.056	0.091	0.138	0.252	0.119	0.048	0.090	0.113	0.213
	Power (W)	0.959	0.943	0.947	0.960	1.081	0.958	0.944	0.949	0.959	1.079
0.15	Perf.	0.186	0.095	0.148	0.206	0.330	0.168	0.083	0.133	0.174	0.269
	Power (W)	0.957	0.929	0.944	0.959	1.079	0.956	0.930	0.941	0.957	1.077
0.2	Perf.	0.246	0.134	0.196	0.261	0.406	0.222	0.122	0.172	0.222	0.337
	Power (W)	0.955	0.922	0.934	0.956	1.079	0.954	0.916	0.937	0.955	1.075
0.25	Perf.	0.301	0.179	0.225	0.347	0.472	0.268	0.156	0.225	0.278	0.380
	Power (W)	0.953	0.906	0.933	0.956	1.076	0.951	0.909	0.928	0.953	1.074
0.5	Perf.	0.608	0.351	0.482	0.666	0.831	0.544	0.288	0.457	0.570	0.667
	Power (W)	0.943	0.858	0.899	0.947	1.073	0.939	0.865	0.896	0.941	1.066
1	Perf.	1.203	0.668	0.980	1.299	1.473	1.100	0.664	0.882	1.148	1.226
	Power (W)	0.922	0.758	0.843	0.932	1.063	0.914	0.754	0.834	0.922	1.049
2	Perf.	2.407	1.494	1.966	19.999	2.570	2.226	1.419	1.834	2.271	2.339
	Power (W)	0.868	0.485	0.686	0.130	1.032	0.850	0.480	0.668	0.871	1.008
3	Perf.	3.335	2.341	2.890	3.440	3.560	3.220	2.341	2.775	3.291	3.370
	Power (W)	0.737	0.452	0.601	0.754	0.902	0.730	0.449	0.593	0.742	0.887
4	Perf.	4.332	3.447	3.845	4.429	4.572	4.216	3.375	3.774	4.292	4.380
	Power (W)	0.668	0.433	0.553	0.675	0.831	0.659	0.430	0.548	0.666	0.817
5	Perf.	5.336	4.251	4.811	5.416	5.557	5.212	4.144	4.720	5.298	5.383
	Power (W)	0.621	0.420	0.524	0.624	0.781	0.611	0.418	0.517	0.617	0.772

The comparisons of performance-power trade-off curves for the DPM policies for all the combinations of SP state transition distribution (TD) and request inter-arrival time distribution (RD) are shown in the following figures. In the legends of the figures, “3CTMDP” means CTMDP policy using a three-state SP model and “4CTMDP” means CTMDP policy using a four-state SP model. The X-axis gives the performance value, which is represented

by the average number of waiting requests in queue. In all the figures, the X-axis is given in logarithmic scale. The Y-axis gives the system power consumption.

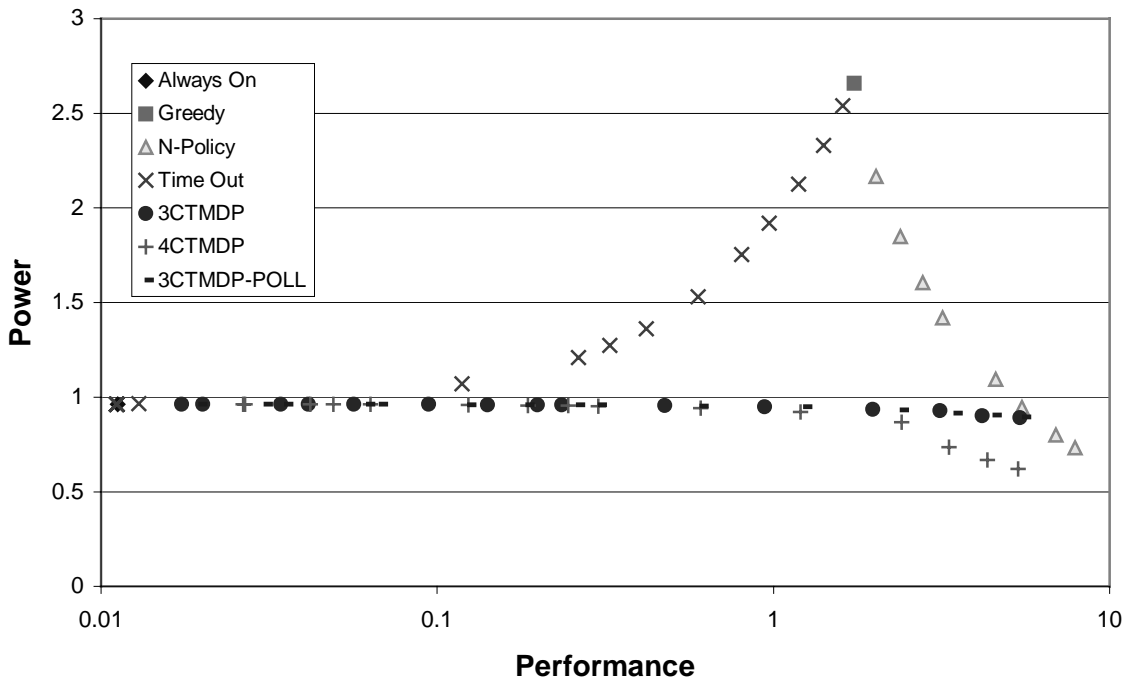


Figure 13 Performance-power trade-off curves for Exp. TD and Exp. RD.

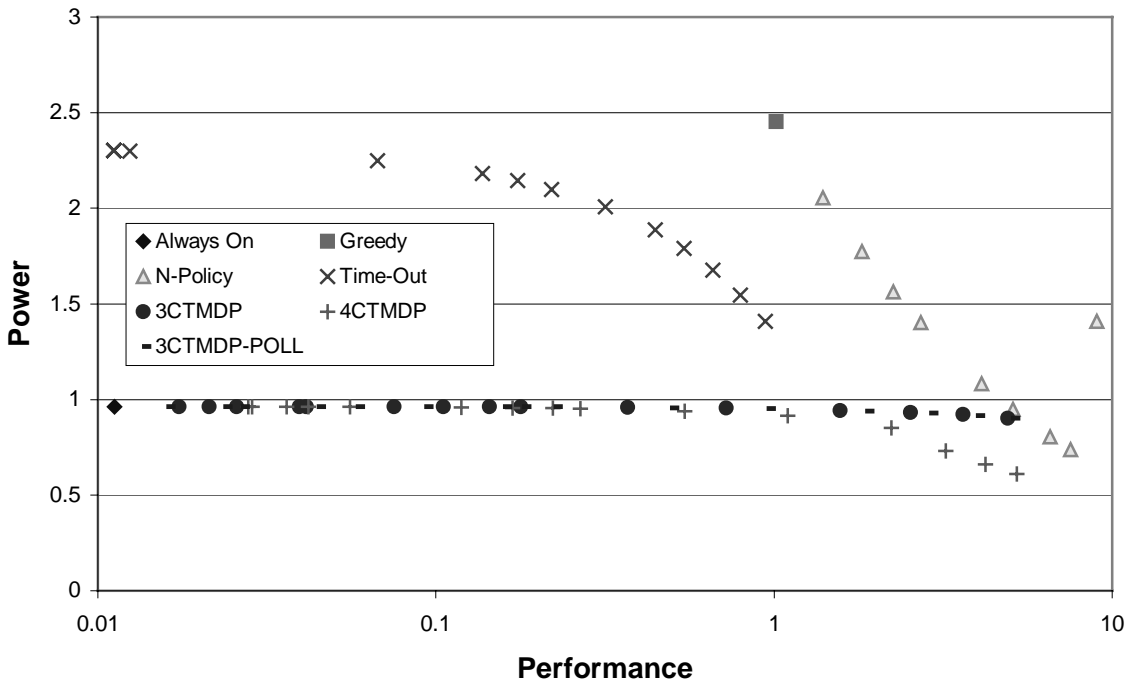


Figure 14 Performance-power trade-off curves for Uni. TD and Exp. RD.

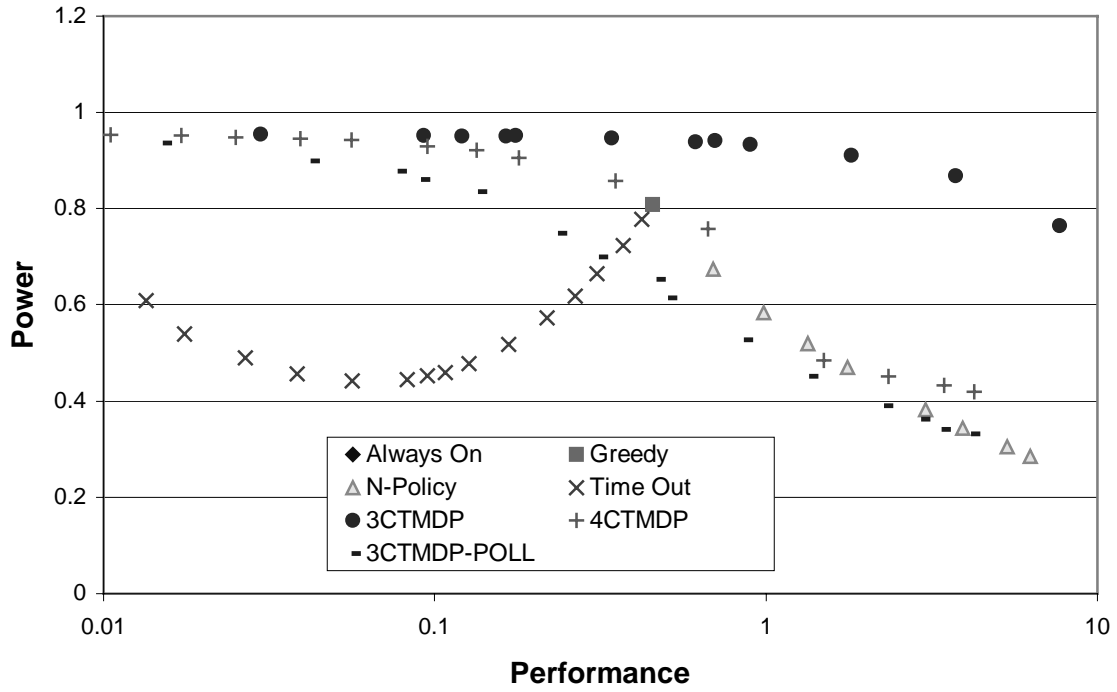


Figure 15 Performance-power trade-off curves for Exp. TD and Exp. & Par. RD.

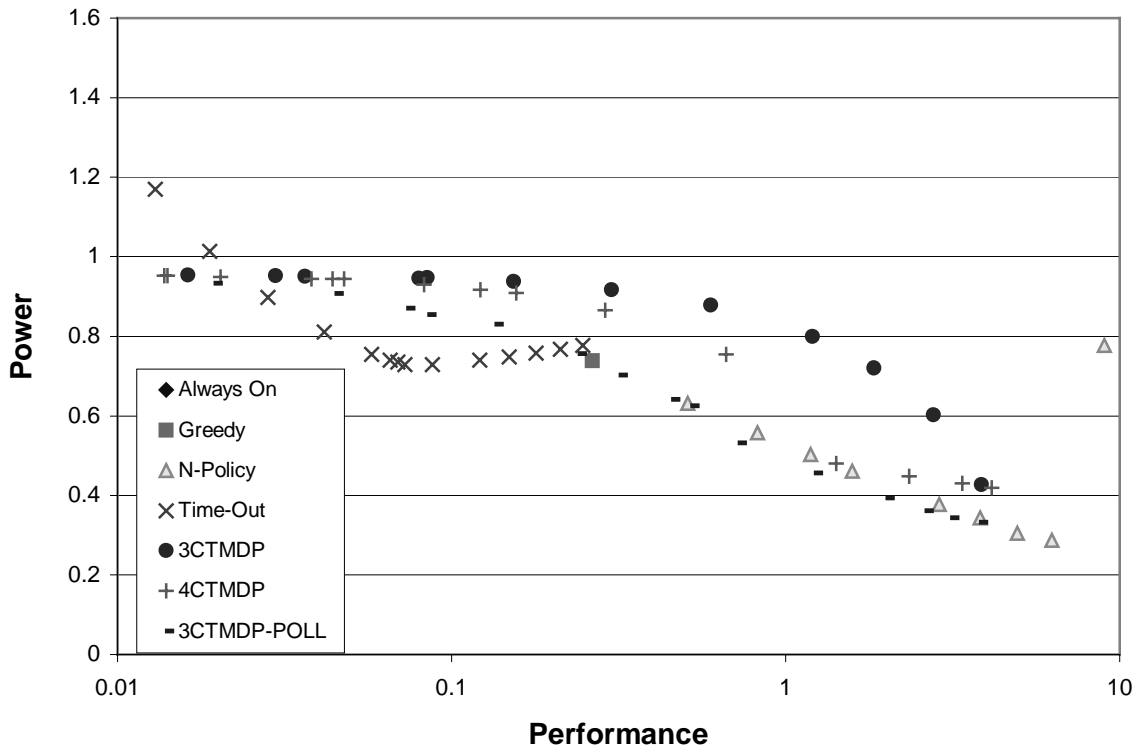


Figure 16 Performance-power trade-off curves for Uni. TD and Exp. & Par. RD.




Figure 17 Performance-power trade-off curves for Exp. TD and Exp. & Exp. RD.




Figure 18 Performance-power trade-off curves for Uni. TD and Exp. & Exp. RD.




Figure 19 Performance-power trade-off curves for Exp. TD and Uni. RD.




Figure 20 Performance-power trade-off curves for Uni. TD and Uni. RD.




Figure 21 Performance-power trade-off curves for Exp. TD and Nor. RD.




Figure 22 Performance-power trade-off curves for Uni. TD and Nor. RD.

From above figures we can see that,

1. In most situations, the stochastic policies out perform the heuristic policies.
2. The stochastic policies provide a good power-delay trade-off. We can always trade performance to reduce power. However, the heuristic policies such as the time-out policy cannot provide a valid power-delay trade-off. In all experiments, decreasing the time out threshold increases the average latency for the request. However, decreasing the time out threshold does not necessarily decrease the average power consumption. In Figure 15, 16, 17, 19, and 21, with the decrease of the time out threshold and the increase of the average request latency, the system power consumption first decreases then increase. In these experiments, there is an optimal time-out threshold, under which both the average request latency and system power consumption are minimized. However, there is no formal way to find out this optimal threshold.
3. Three-states CTMDP policy is not efficient with input sequence with Exp. & Pareto inter-arrival time because our stochastic model does not give as accurate representation of the real system. However, the 3CTMDP-poll policy solves the above problem. The 3CTMDP-poll policy works almost as well as the four-state CTMDP.
4. Four-state CTMDP is robust in different TD, RD distributions. This means that if the device can provide us with more operation modes of different performance and power consumption, the CTMDP policy can explore these modes to make itself an always-robust DPM policy independent of the distribution of request inter-arrival time or the distribution of SP transition time.

6.2 Experimental setup for DPM models with priority queue

In this setup, we call our algorithm POMP (Policy Optimization by Markov Decision Process).

A. Comparison of POMP with heuristic policy

In this experiment, we will present the comparison of POMP policy and the heuristic policies including greedy policy, time-out policy and show that our policy consumes less power than the heuristic policies.

The system model used in this part includes:

A SP model that is the same as in Example 3.1 except that it has only one busy state, $busy_1$, and one idle state, $idle_1$.

A SR model with two states r_1 and r_2 , $G_{SR}(r_1, r_2)=1/200$, $G_{SR}(r_2, r_1)=1/400$, $\lambda_l(r_1)=1/30$, $\lambda_h(r_1)=1/50$, $\lambda_l(r_2)=1/60$, $\lambda_h(r_2)=1/90$.

A SQ model with a LSQ of length 3 and a HSQ of length 2.

We compare the POMP policy with Greedy and Time-Out policies.

Two different traces of requests are used for simulation:

1. Requests are generated to exactly follow the SR model.
2. Requests are generated to follow the SR state transition rate between r_1 and r_2 . However, in state r_1 , the inter-arrival time of low-priority requests follows the uniform distribution (comparing with the exponential distribution in trace 1) with mean value $1/\lambda_l(r_1)=30$, the inter-arrival time of high-priority requests follows the uniform distribution with mean value $1/\lambda_h(r_1)=50$. In state r_2 , the inter-arrival time of low-priority requests follows the uniform distribution with mean value $1/\lambda_l(r_2)=60$, the inter-arrival time of high-priority requests follows the uniform distribution with mean value $1/\lambda_h(r_2)=90$.

In both cases, we simulated the heuristic policies to get the average delay of the low priority and high priority requests. We then use these delay values as the delay constraint and searched for a randomized POMP policy

using linear programming. Finally, We simulated the POMP policies and compare the power and delay value with that of heuristic policies.

Table 9 Comparison with heuristic policy with input trace 1.

Heuristic Policy				LP-based POMP policy	
Name of policy	P (w)	D_l	D_h	P	ΔP (%)
Timeout policy $t_{out}=20$	1.665	0.399	0.143	0.942	43.43
Timeout policy $t_{out}=40$	2.080	0.232	0.118	1.156	44.42%
Timeout policy $t_{out}=60$	2.244	0.183	0.104	1.865	16.89
Greedy policy	1.303	0.151	0.255	1.067	18.11

Table 10 Comparison with heuristic policy with input trace 2.

Heuristic Policy				LP-based POMP Policy					
Name of Policy	P(w)	D_l	D_h	P(w)	ΔP (%)	D_l	ΔD_l (%)	D_h	ΔD_h (%)
Timeout policy $t_{out}=20$	1.515	0.439	0.121	1.121	26.0	0.406	7.51	0.119	1.7
Timeout policy $t_{out}=40$	2.071	0.243	0.097	1.674	19.2	0.227	6.6	0.094	3.1
Timeout policy $t_{out}=60$	2.242	0.151	0.085	2.142	4.5	0.149	1.3	0.083	2.4
Greedy policy	1.655	0.227	0.137	1.522	8.0	0.169	25.6	0.107	21.9

Table 9 and Table 10 present the comparison results by simulation. The column P, D_l and D_h give for the power, low priority requests delay, high priority request delay for each power management policy respectively. The column ΔP , ΔD_l and ΔD_h give the POMP policy's improvement in power, low priority requests delay and high priority request delay. The values are simulated using a commercial stochastic activity network analysis tool: UltraSAN [20]. The conclusions we have from these experiments are,

If the inter-arrival time of the input requests follows exponential distribution, the POMP policy saves more than 30% power on average over heuristic methods while keeping the same delay for both the high priority and low priority requests.

If the inter-arrival time of the input requests follows uniform distribution, the POMP policy cannot exactly meet the delay constraint. It saves about 14% power on average, at the same time it reduces the delay of low priority requests by about 10% and reduces the delay of high priority requests by about 7%.

Further reduces can be achieved using POMP method if we increase the delay constraint. In other words, the POMP method can make different power-delay trade-offs by changing the delay constraints. Little can be done by the heuristic method.

The POMP method can do optimal management for a complex system, while the heuristic methods can only do simple management such as turn-on and turn-off.

The POMP method can adjust the optimal policy when workload characteristics changes, while the greedy and timeout methods are not adjustable to workload change.

We did not compare DPM method with predictive method, because in our experiments the inter-arrival time of each requests are assumed to be independent. Therefore, the predictive method is not applicable.

B) Comparison of our policy with N-policy

In this section we will show that the POMP method is powerful in finding power and delay trade-offs. When the server has only two states: *active* and *sleeping*, it can easily be shown that the N-policy gives the minimum power compared to other stationary policies with the same performance constraint. Our experiments show that, however, for a system with more than two server states, the N-policy does not give the optimal power-delay tradeoff. In the experimental results, we will present the comparison of our policy and the N-policy and show that our policy is more efficient in finding power and performance trade-off than the N-policy.

In this experiment, the system model includes:

A SP model that is the same as that in the previous experiment.

A SR model with only one state r , $\lambda(r) = 1/20$, the inter-arrival time of each requests follows exponential distribution.

A SQ model with a SSQ of length 5

Table 11 gives the result of comparison of our policy with N-policy, where N varies from 0 to 5. The result shows that there is always a randomized policy, which has the exact delay as the N-policy and consumes less power than N-policy. The deterministic policy consumes more power than a randomized policy, however, it is still more efficient than N-policy. Furthermore, a N-policy in this experiment only gives 5 different policies. However, a randomized policy can find power optimal policy under any delay constraint, which is achievable. Therefore, our algorithm is more flexible and effective.

Table 11 Comparison of our policy with N-policy.

N-policy			Randomized-policy				Deterministic-policy			
N	Power	Delay	Power	$\Delta P(\%)$	Delay	$\Delta D(\%)$	Power	$\Delta P(\%)$	Delay	$\Delta D(\%)$
5	0.860	2.346	0.754	12.3	2.304	1.8	0.754	12.3	2.304	1.8
4	0.944	1.933	0.802	15.0	1.933	0	0.8107	14.1	1.869	3.3
3	1.045	1.483	0.861	38.7	1.483	0	0.869	16.8	1.425	3.9
2	1.216	1.027	0.952	20.9	1.027	0	0.966	20.6	1.005	2.1
1	1.623	0.608	1.160	28.5	0.607	0.2	1.213	25.3	0.606	0.3
0	2.3	0.332	2.3	0	0.332	0	2.3	0	0.32	0

VII. CONCLUSIONS

We have proposed a new system model and method for dynamic power management at the system-level. The problem of system-level power management was formulated as an optimal policy selection problem based on the theories of continuous-time Markov decision process, and stochastic network. Comparing with previous work, we introduced new and more complete model of the system components, as well as the model of the whole system. The proposed mathematical framework captures the characteristics of the real applications more accurately compared to the models proposed by previous researchers. This is mainly because we solve the problem in continuous-time domain while previous authors solve the problem in discrete-time domain. We have also proposed CTMDP-Poll techniques, which is better than the DMTDP-based technique and the vanilla CTMDP-based techniques in all cases. Furthermore, we have used a priority queue model, which is more general and adaptable to real applications.

In the examples and the experimental results included in this paper, we have assumed that the system contains a single SP. Therefore, all the requests in the SQ are targeted for that SP (which means that they have identical type). However, please note that this does not mean that all the requests pose the same amount of work load to the SP. The difference in the request work loads is captured by the exponentially distributed service time in the model. More generally, our mathematical framework can handle systems with more than one SP and therefore, service requests of different types can be handled. In this case, we have to create multiple SQ's one for each type (or class) of service requests. There will also be a compatibility relation describing what type of requests can be serviced by which SPs. This adds complexity to the problem but is quite manageable within our proposed mathematical framework.

A shortcoming of our CTMDP-based technique is that it is very difficult to use the model when attempting to represent complex systems that consist of multiple closely interacting Service Providers and must cope with complicated synchronization schemes. In this case, we need to use the modeling techniques based on the theory of generalized stochastic Petri Nets (GSPN). Finally, in real applications, the inter-arrival time of service requests may not follow an exponential distribution. Although our experimental results demonstrate that, in practice, the our method remain effective for a wide range of distributions, it is still desirable from a theoretical point of view to develop a new mathematical framework where arbitrary distributions can be easily handled. Again this problem can be solved by using the "stage method" (series-parallel connection of exponentially distributed sources) based on the GSPN framework. However, GSPN is not in the scope of this manuscript.

REFERENCES

- [1] A. Chandrakasan, R. Brodersen, *Low Power Digital CMOS Design*, Kluwer Academic Publishers, July 1995.
- [2] M. Horowitz, T. Indermaur, and R. Gonzalez, "Low-Power Digital Design", IEEE Symposium on Low Power Electronics, pp.8-11, 1994.
- [3] A. Chandrakasan, V. Gutnik, and T. Xanthopoulos, "Data Driven Signal Processing: An Approach for Energy Efficient Computing", 1996 International Symposium on Low Power Electronics and Design", pp.347-352, Aug. 1996.
- [4] J. Rabaey and M. Pedram, *Low Power Design Methodologies*, Kluwer Academic Publishers, 1996
- [5] L. Benini and G. De Micheli, *Dynamic Power Management: Design Techniques and CAD Tools*, Kluwer Academic Publishers, 1997.
- [6] Intel, Microsoft and Toshiba, "Advanced Configuration and Power Interface specification", URL: <http://www.intel.com/ial/powermgm/specs.html>, 1996.
- [7] M. Srivastava, A. Chandrakasan. R. Brodersen, "Predictive system shutdown and other architectural techniques for energy efficient programmable computation," *IEEE Transactions on VLSI Systems*, Vol. 4, No. 1, pp.42-55, 1996.
- [8] C.-H. Hwang and A. Wu, "A Predictive System Shutdown Method for Energy Saving of Event-Driven Computation," *Proceedings of the International Conference on Computer Aided Design*, pp. 28-32, Nov. 1997.
- [9] G. A. Paleologo, L. Benini, et.al, "Policy Optimization for Dynamic Power Management", *Proceedings of the Design Automation Conference*, pp.182-187, Jun. 1998.
- [10] Q. Qiu, M. Pedram, "Dynamic Power Management Based on Continuous-Time Markov Decision Processes", *Proceedings of the Design Automation Conference*, pp. 555-561, Jun. 1999.
- [11] Q. Qiu, Q. Wu, M. Pedram, "Stochastic Modeling of a Power-Managed System: Construction and Optimization", *Proceedings of the International Symposium on Low Power Electronics and Design*, 1999.
- [12] U. Narayan Bhat, "*Elements Of Applied Stochastic Processes*", John Wiley & Sons, Inc. 1984.
- [13] A. Chandrakasan, V. Gutnik, and T. Xanthopoulos, "Data Driven Signal Processing: An Approach for Energy Efficient Computing", *Proceedings of International Symposium on Low Power Electronics and Design*, pp. 347-352, Aug. 1996.
- [14] B. Miller, "Finite State Continuous Time Markov Decision Processes With an Infinite Planning Horizon". *J. Of Mathematical Analysis and Applications*, No. 22, pp. 552-569, 1968.
- [15] E. V. Denardo, "On Linear Programming in a Markov Decision Problem", *Management Science*, Vol. 16, No. 5, pp. 281-288, January, 1970.

- [16] J. F. Shapiro, *Mathematical Programming: Structures and Algorithms*, John Wiley & Sons, Inc, 1979.
- [17] R. A. Howard, *Dynamic Programming and Markov Processes*, Wiley, New York, 1960.
- [18] D. P. Heyman, M. J. Sobel, *Stochastic Models in Operations Research*, McGraw-Hill Book Company, 1982.
- [19] L. E. Scales, *Introduction to Non-Linear Optimization*, Springer-Verlag New York Inc, 1985.
- [20] UIUC, Performability Engineering Research Group, "The UltraSAN Software", <http://www.crhc.uiuc.edu/UltraSAN/>, 1997.
- [21] T. Simunic, L. Benini, G. De Micheli: "Dynamic Power Management of Portable Systems", MOBICOM, 2000.
- [22] T. Simunic, L. Benini, G. De Micheli: "Dynamic Power Management of Laptop Hard Disk", DATE, 2000.
- [23] Y. Lu, E. Chung, T. Simunic, L. Benini, G. De Micheli: "Quantitative Comparison of Power Management Algorithms", DATE, 2000.
- [24] "MHE2064AT, MHE2043AT, MHF2043AT, MHF2021AT Disk Drive Product Manual", Fujitsu Limited, 1998.

APPENDIX

Proof of Proposition 4.1

Let s^* denote the state of SP in which it has the highest service rate, q^* denote the SQ state that both LSQ and HSQ are full. From the definition of SR we also know that the states $\{(s^*, r, q^*), \forall r \in \mathcal{R}\}$ communicate with each other. The two constraints ensures that for any state $x \in \mathcal{X}$, $x \rightarrow (s^*, r, q^*), \forall r \in \mathcal{R}$. Therefore, there cannot exist two separate communicating class in the system. According to the definition of ergodic chain, we know that there is only one ergodic chain in the Markov process. For any state $x \in \mathcal{X}$, if $x \leftarrow (s^*, r, q^*)$, then $x \leftrightarrow (s^*, r, q^*)$ which means that x belongs to the ergodic chain. If x does not belong to the ergodic chain, then $x \nleftrightarrow (s^*, r, q^*)$. Therefore, $P(T_{xx} < \infty) < 1$ and x is a transient state.

Proof of Proposition 4.2

Let $o(t)$ denote the class of functions that $\lim_{t \rightarrow \infty} o(t)/t = 0$. From the definition of transition rate (2.2) and (2.3), we know that:

$$\lim_{t \rightarrow 0} p_{x \Rightarrow x'}(t) = \lim_{t \rightarrow 0} (\sigma_{x,x'} t + o(t))$$

$$\lim_{t \rightarrow 0} p_{y \Rightarrow y'}(t) = \lim_{t \rightarrow 0} (\sigma_{y,y'} t + o(t))$$

$$\lim_{t \rightarrow 0} p_{x \Rightarrow x}(t) = \lim_{t \rightarrow 0} (1 - \sum_{x' \neq x} \sigma_{x,x'} t + o(t))$$

$$\lim_{t \rightarrow 0} p_{y \Rightarrow y}(t) = \lim_{t \rightarrow 0} (1 - \sum_{y' \neq y} \sigma_{y,y'} t + o(t))$$

From (2.4) we know that:

$$\sigma_{x,x} = \sum_{x' \neq x} \sigma_{x,x'}, \quad \sigma_{y,y} = \sum_{y' \neq y} \sigma_{y,y'}$$

Since X_1, X_2 are independent, we have the following equations:

$$\begin{aligned} \lim_{t \rightarrow 0} p_{(x,y) \Rightarrow (x,y')}(t) &= \lim_{t \rightarrow 0} (p_{x \Rightarrow x}(t) p_{y \Rightarrow y'}(t)) \\ &= \lim_{t \rightarrow 0} (\sigma_{y,y'} t + \sum_{x' \neq x} \sigma_{y,y'} \sigma_{x,x'} t^2 + o(t)) \\ &= \lim_{t \rightarrow 0} (\sigma_{y,y'} t + o(t)) \end{aligned} \tag{P1.1}$$

$$\begin{aligned} \lim_{t \rightarrow 0} p_{(x,y) \Rightarrow (x',y)}(t) &= \lim_{t \rightarrow 0} (p_{x \Rightarrow x'}(t) p_{y \Rightarrow y}(t)) \\ &= \lim_{t \rightarrow 0} (\sigma_{x,x'} t + \sum_{y' \neq y} \sigma_{y,y'} \sigma_{x,x'} t^2 + o(t)) \\ &= \lim_{t \rightarrow 0} (\sigma_{x,x'} t + o(t)) \end{aligned} \tag{P1.2}$$

$$\begin{aligned} \lim_{t \rightarrow 0} p_{(x,y) \Rightarrow (x',y')}(t) &= \lim_{t \rightarrow 0} (p_{x \Rightarrow x'}(t) p_{y \Rightarrow y'}(t)) \\ &= \lim_{t \rightarrow 0} (\sigma_{x,x'} \sigma_{y,y'} t^2 + o(t)) = \lim_{t \rightarrow 0} (o(t)) \end{aligned} \tag{P1.3}$$

Therefore,

$$\sigma_{(x,y),(x,y')} = \lim_{t \rightarrow 0} \frac{P(x,y) \Rightarrow (x,y')}{t} = \sigma_{y,y'}$$

$$\sigma_{(x,y),(x',y)} = \lim_{t \rightarrow 0} \frac{P(x,y) \Rightarrow (x',y)}{t} = \sigma_{x,x'}$$

$$\sigma_{(x,y),(x',y')} = \lim_{t \rightarrow 0} \frac{P(x,y) \Rightarrow (x',y')}{t} = 0$$

$$\sigma_{(x,y),(x,y)} = \sum_{x' \neq x} \sigma_{(x,y),(x',y)} + \sum_{y' \neq y} \sigma_{(x,y),(x,y')} = \sum_{x' \neq x} \sigma_{x,x'} + \sum_{y' \neq y} \sigma_{y,y'} = \sigma_{x,x} + \sigma_{y,y}$$

Proof of Proposition 4.3

Assume the generator matrix \mathbf{A} can be written as:

$$\mathbf{A} = \begin{bmatrix} a_{00} & \cdots & a_{0n_1} \\ \vdots & \ddots & \vdots \\ a_{n_1 0} & \cdots & a_{n_1 n_1} \end{bmatrix}$$

Assume the generator matrix \mathbf{B} can be written as:

$$\mathbf{B} = \begin{bmatrix} b_{00} & \cdots & b_{0n_2} \\ \vdots & \ddots & \vdots \\ b_{n_2 0} & \cdots & b_{n_2 n_2} \end{bmatrix}$$

Let's denote the new matrix as \mathbf{C} , $\mathbf{C} = \mathbf{A} \oplus \mathbf{B}$. From the definition of the Tensor sum we know that

$$c_{i \cdot n_2 + j, m \cdot n_2 + n} = \begin{cases} a_{ii} + b_{jj} & i = m, j = n \\ b_{jn} & i = m, j \neq n \\ a_{im} & i \neq m, j = n \\ 0 & i \neq m, j \neq n \end{cases} \quad (\text{P2.1})$$

Let's denote the generator matrix of the joint process as \mathbf{G} . In the joint process, the states are ordered in the lexicography order of \mathbf{A} and \mathbf{B} , i.e., they are ordered as:

$(a_1, b_1), (a_1, b_2), \dots, (a_1, b_{n_2}), (a_2, b_1), (a_2, b_2), \dots, (a_2, b_{n_2}), \dots$. Therefore, $\sigma_{(a_i, b_j), (a_m, b_n)} = G_{i \cdot n_2 + j, m \cdot n_2 + n}$.

From Proposition 4.2 we know that:

$$G_{i \cdot n_1 + j, m \cdot n_1 + n} = \begin{cases} \sigma_{(a_i, b_j), (a_i, b_j)} = \sigma_{a_i, a_i} + \sigma_{b_j, b_j} & i = m, j = n \\ \sigma_{(a_i, b_j), (a_i, b_n)} = \sigma_{b_j, b_n} & i = m, j \neq n \\ \sigma_{(a_i, b_j), (a_m, b_j)} = \sigma_{a_i, a_m} & i \neq m, j = n \\ \sigma_{(a_i, b_j), (a_m, b_n)} = 0 & i \neq m, j \neq n \end{cases} \quad (\text{P2.2})$$

Because $a_{i,m} = \sigma_{a_i, a_m}$ and $b_{j,n} = \sigma_{b_j, b_n}$, (P2.1) is equal to (P2.2). Therefore, $\mathbf{A} \oplus \mathbf{B}$ is the generator matrix of the joint process of \mathbf{A} and \mathbf{B} .

Proof of Proposition 5.1

Let function $\xi_{k, K-\{k\}}(t)$ denote the probability that $t_{ik} < t$ and $t_{ik} < t_{im}, \forall m \in K-\{k\}, K=J-\{j\}$. Then

$$Q_{ij}(t) = \int_0^t f(t_{ij}) \sum_{k \in K} \xi_{k, K-\{k\}}(t_{ij}) dt_{ij}$$

Now we are going to prove that

$$\sum_{k \in K} \xi_{k, K-\{k\}}(t) = e^{-\sum_{k \in K} \sigma_{ik} t} \quad (\text{p3.1})$$

If K_2 contains only 2 elements, assume that $K_2 = \{a, b\}$, we know that: $\sum_{k \in \{a, b\}} \xi_{k, K-\{k\}}(t) = e^{-(\sigma_{ia} + \sigma_{ib})t}$. Assume that when K_n contains n elements, (p3.1) is true. When $K_{n+1} = K_n + \{z\}$, we have

$$\begin{aligned} \sum_{k \in K_{n+1}} \xi_{k, M}(t) &= \sum_{k \in K_{n+1}} \int_t^\infty f(t_{ik}) \cdot \sum_{m \in K_{n+1}-\{k\}} \xi_{m, K_{n+1}-\{k, m\}}(t_{ik}) dt_{ik} \\ &= \sum_{k \in K_{n+1}} \int_t^\infty \sigma_{ik} e^{-\sigma_{ik} t_{ik}} \cdot e^{-\sum_{m \in K_{n+1}-\{k\}} \sigma_{im} t_{ik}} dt_{ik} \\ &= e^{-\sum_{k \in K_{n+1}} \sigma_{ik} t} \end{aligned}$$

Therefore, (P3.1) is true for all K .

$$Q_{ij}(t) = \int_0^t \sigma_{ij} e^{-\sigma_{ij} t} \cdot e^{-\sum_{k \in K} \sigma_{ik} t} dt_{ij} = \sigma_{ij} (1 - e^{-\sum_{l \in J} \sigma_{il} t}) / \sum_{l \in J} \sigma_{il}$$

Proof of Lemma 5.1

$$\sum_{i \in S} \sum_{a_i \in A_i} x_i^{a_i} \gamma_i^{a_i} = \sum_{i \in S} \sum_{a_i \in A_i} (x_i^{a_i} \tau_i^{a_i} c_{ii} + \sum_{j \in J} x_i^{a_i} c_{ij} p_{ij}^{a_i})$$

$$\begin{aligned} \sum_{i \in S} p_i^\pi c_i^{A_i} &= \sum_{i \in S} (p_i^\pi c_{ii} + \sum_{j \in J} p_i^\pi \sum_{a_i \in A_i} c_{ij} \sigma_{ij}^{a_i} p_i^{a_i}) \\ &= \sum_{i \in S} (p_i^\pi c_{ii} + \sum_{j \in J} \sum_{a_i \in A_i} c_{ij} \sigma_{ij}^{a_i} p_i^{a_i} \sum_{a_i \in A_i} x_i^{a_i} \tau_i^{a_i}) \\ &= \sum_{i \in S} (p_i^\pi c_{ii} + \sum_{a_i \in A_i} \sum_{j \in J} x_i^{a_i} \tau_i^{a_i} c_{ij} \sigma_{ij}^{a_i}) \end{aligned}$$

Since $p_i^\pi = \sum_{a_i \in A_i} x_i^{a_i} \tau_i^{a_i}$ and $p_{ij}^{a_i} = \tau_i^{a_i} \cdot \sigma_{ij}^{a_i}$, $\sum_{i \in S} \sum_{a_i \in A_i} x_i^{a_i} \gamma_i^{a_i} = \sum_{i \in S} p_i^\pi c_i^{A_i}$.

Proof of Lemma 5.2

If the Markov process is irreducible, then $\lim_{t \rightarrow \infty} p_i^\pi(t) = p_i^\pi$. From the definition of $C_{i,avg}^\pi$ (2.9) we can derive these equations.

Proof of Theorem 5.0

Consider a request x , the time that this request stays in the queue can be divided into N segments, where N is the maximum capacity of the queue. In the i th segment, there are i requests in the queue. For request x , the length of the i th segment is denoted as $T_i(x)$, note that $T_i(x)$ may equal to zero, which means that during the time that request x is in queue, the number of waiting requests in queue has never been i . During time T , if the number of incoming requests is X , then the average waiting time for each request is:

$$E_T(d) = \sum_{x \in X} \sum_{i=1}^N T_i(x) / X \quad (\text{P.4.1})$$

We also know that, during time $T_i(x)$, there are i requests. Therefore, (P.4.1) can also be written as:

$$E_T(d) = \sum_{i=1}^N i \cdot T_i / X$$

where T_i denotes the sum of time during T that there are i requests in the queue. Let p_i denote the probability that there are i requests in queue, we know that $p_i = T_i / X$, and

$$E_T(q) = \sum_{i=1}^N i \cdot p_i$$

Therefore, $E_T(d) = E_T(q) \cdot T / X$.

Furthermore, if the loss rate of request is small enough, the average request incoming rate $\lambda = T / X$. Therefore,

$$D = Q \cdot \lambda.$$