An Optimal Voltage Synthesis Technique for a Power-Efficient Satellite Application

Dong-In Kang, Jinwoo Suh, and Stephen P. Crago
University of Southern California/Information Sciences Institute
3811 N. Fairfax Dr. Suite 200, Arlington, VA 22203
703-248-6164, 703-248-6160, 703-812-3729
{dkang, jsuh, crago}@isi.edu

ABSTRACT
This paper presents an optimal voltage synthesis technique for a satellite application to maximize system performance subject to energy budget. A period of a satellite's orbit is partitioned into several independent regions with different characteristics such as type of computation, importance, performance requirements, and energy consumption. Given a periodic energy recharge model, optimal voltages for the regions are synthesized such that the overall performance is maximized within the energy budget in the period.

Categories and Subject Descriptors
C.4 [**PERFORMANCE OF SYSTEMS**] Design studies, Modeling techniques, Performance attributes.

General Terms
Algorithms, Management, Performance, Design.

Keywords
Power-aware design, power-efficient design, satellite application, queueing.

1. INTRODUCTION
A satellite such as an LEO (Low-Earth Orbit) satellite circles Earth and captures information of the events on Earth and is used for remote sensing and weather monitoring. Its behavior with respect to energy charge/consumption is highly periodic. Energy is recharged every period using solar panels. The computation is also periodic. A satellite flies over precalculated regions every period and does pre-determined computations depending on the region that it is flying over.

Due to the increasing demand of computations and limited energy source of the satellite, energy-efficient usage of the resources is becoming more important. However, low-power design is not always the best solution because any unused energy within a period is simply wasted due to the periodic recharge. Power-aware design that maximizes system performance within the energy budget is superior to low-power design. The sampling rate of a remote sensing satellite can be increased when there is extra energy left unless the sensor limits the sampling rate. Likewise, the accuracy of the computation can be enhanced with extra energy.

A major source of energy savings is voltage scaling, which scales operational voltages of resources and corresponding maximum clock speeds [1][3][11]. Voltage scaling affects throughput and latency in a nonlinear fashion [2], which makes it difficult to find an optimal trade-off between energy consumption and throughput. Maximum clock frequency depends on the supply voltage. As supply voltage increases, both the maximum clock frequency and energy consumption increase.

In this paper, we solve the optimal voltage synthesis problem for power-aware satellite system design. The voltage synthesis problem is to find the (optimal) voltages of resources in the system that satisfy system-level energy constraints and maximize overall performance. The period of a satellite's orbit is partitioned into several independent regions. Each region has its own event rate, computation demand per input, importance, and energy consumption per computation. We assume the event rate and computation demand are not deterministic but stochastic. With those inputs, our technique synthesizes an optimal voltage for each region that maximizes aggregated throughput in the period while total energy consumption is within the energy budget.

Motivational example:
As an example of an LEO satellite, the weather satellite NOAA10 tracked the following regions on January 26th, 2001 at around 11:00am. Figure 1 is downloaded from the J-Track by NASA [7]. The satellite covered seven different regions during the period. Region R_1 corresponds to the Pacific Ocean, R_2 to North America, R_3 to the Arctic and Greenland, R_4 to the Atlantic ocean, R_5 to Europe and Africa, R_6 to the Atlantic ocean, and R_7 to Antarctica. Based on the characteristics of the similar satellite NOAA12, its period is 101.3 minutes and it's in polar orbit.

The synthetic characteristics in each region are shown in Table 1. The duration d_i at region R_i is the length of the satellite's stay over the region. The arrival rate λ_i at region R_i is the average arrival rate of inputs per millisecond. The service rate $\mu_{i,ref}$ at region R_i is the average number of inputs that can be processed per millisecond at the reference voltage V_{ref} shown in Table 2. The weight w_i of an input for region R_i denotes how valuable an output is when compared to the outputs in other regions. Energy consumption $e_{i,ref}$ for region R_i is the average energy at the reference voltage V_{ref} consumed to process an input in milli-Joules. Energy consumption is assumed to depend on the service time.

DAC 2002, June 10-14, 2002, New Orleans, Louisiana, USA.

We assume that computation in a region may differ from the computation in the other regions. Likewise, the size of input data can be different for each region. The buffer size B denotes how many inputs can be stored in the buffer. The maximum input loss probability Λ_i for region R_i is the maximum allowed input loss probability for the region.

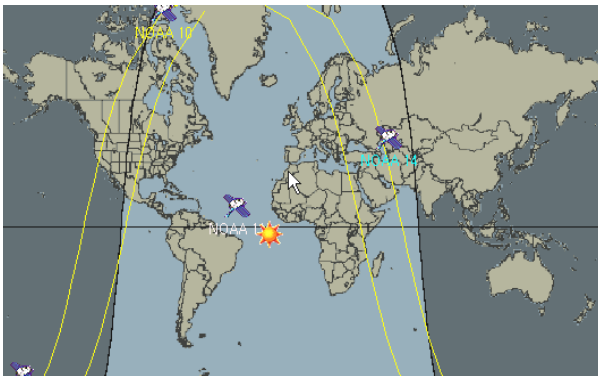

Figure 1. A screenshot of Jtrack

Table 1. Characteristics of the system per region

Region	R_1	R_2	R_3	R_4	R_5	R_6	R_7
d_i (min.)	26.1	11.6	11.6	8.7	14.5	11.6	20.3
λ_i	0.4	1	0.2	0.4	1.2	0.4	0.2
$\mu_{i,ref}$	1	1.5	1	1	1.8	1	1
w_i	1	1.5	0.9	0.7	1.2	0.7	0.7
$e_{i,ref}$ (mJ)	2.1	2.5	1.8	2.1	2.5	2.1	1.8
B	100	50	80	100	70	100	80
Λ_i	0.1	0.02	0.1	0.05	0.01	0.05	0.1

Table 2. System constraints

V_{ref}	V_t	E_{Total}
3.3V	0.6V	3300 J

As an example, the second column labeled R_1 in Table 1 is interpreted as follows. The satellite flies over the Pacific Ocean for 26.1 minutes. During that time 0.4 inputs arrive per millisecond on the average. The satellite can process one input per millisecond on the average in the region at the reference voltage. The average energy consumption for processing one input is 2.1 milli-Joules in the region. The system can hold up to 100 inputs in its buffer. Any more inputs are discarded. Due to the limited buffer space, input loss is expected. However, there is a maximum allowable input loss probability for each region. Up to 10 percent of inputs gathered during the flight over Pacific Ocean can be discarded.

System-wide constraints are described in Table 2. According to Table 2, the satellite can use up to its maximum energy budget E_{Total}, which is 3300 Joules, for computation within a period. The system is assumed to be voltage schedulable between its reference voltage V_{ref}, 3.3 Volts, and its threshold voltage V_t, 0.6 Volt, shown in Table 2.

Our interest lies in how to derive the optimal voltage for each region that produces maximum performance with a given energy constraint E_{Total}. Performance is defined as the sum of inputs processed multiplied by their weights. The solution to the example is derived using our optimal algorithm in Section 3.3.

The rest of this paper is organized as follows. Section 2 presents our model and an overview of the problem. In Section 3, the power optimization techniques and voltage synthesis algorithms are discussed. Simulation results are presented in Section 4. Related work is presented in Section 5. We conclude in Section 6.

2. MODEL AND PROBLEM OVERVIEW

A set of regions $\{R_i \mid i = 1 .. N\}$ exists within a period T. Each region R_i has its own characteristics: data arrival rate (λ_i), service rate at the reference voltage ($\mu_{i,ref}$), energy consumption per input processing at the reference voltage ($e_{i,ref}$), duration (T_i), buffer size (B_i), and weight of an output (w_i) in that region. The weight value associated with a region reflects the importance of an output produced in the region. The performance (Ω_i) of a region R_i is the product of the weight and the number of outputs in the region, which is $\Omega_i = T_i w_i \lambda_i (1 - p_{i,loss})$, where $p_{i,loss}$ denotes the probability of input loss due to the limited buffer space. Performance in a period is described as the sum of the performance of the regions in the period, which is $\sum_i \Omega_i = \sum_i T_i w_i \lambda_i (1 - p_{i,loss})$. Our goal is to find an optimal voltage assignment of each region so that overall performance is maximized within the energy budget.

Periodic Energy Recharge: In the case of satellites with solar panels, the battery is recharged periodically. Both battery charge and discharge occur throughout the entire orbit, although the charging amount at any time may vary. We assume that the amount of the energy usable during a period is predictable and available at the start of a period. This simple model may not capture dynamic behavior of energy supply and consumption.

Heterogeneous computation patterns within a period: As a satellite orbits around the earth, its covering areas change: from land to sea, from desert to mountains, etc. Different areas may have different computation needs and input characteristics. We assume that within a period, there may be multiple regions that have their own input arrival rates, computation requirements, data sizes and energy consumptions per data processing. Those patterns of regions in a period may change as the earth rotates on its axis. We assume that those patterns and their characteristics are known at the start of a period, which is reasonable for most satellites.

Input Arrival and Service Patterns: In this paper we assume inputs arrive in accordance with a Poisson process of rate λ_i within region R_i. The probability distribution of the service time is assumed to be exponential with parameter μ_i. Maximum system capacity is limited to K_i. We assume no correlation among the inputs, and one input produces one output. The system in region R_i is modeled with an $M/M/1/K_i$ queueing system. To simplify the static analysis, we ignore the dependencies between different regions, and treat each region independently. The interference between regions is studied using simulation in Section 4.

Dynamic power consumption of a CMOS digital circuit: In CMOS digital circuits, latency and energy consumption by a task are given by the following equations (Eq. 1) and (Eq. 2) [2]. In the equations, v denotes supply voltage, V_t and V_{ref} denote threshold voltage and reference voltage respectively, which are inherent to the implementation technology, C is a technology dependent constant, and C_{eff} is the effective capacitance.

$$Latency = C\frac{v}{(v-V_t)^2},\quad V_t < v \le V_{ref} \qquad \text{(Eq. 1)}$$

$$Energy \approx (\#\text{ of switches}) \times C_{eff} \times v^2 \qquad \text{(Eq. 2)}$$

When latency $l_{i,ref}$ and average energy consumption $e_{i,ref}$ per input in region R_i at reference voltage $V_{i,ref}$ of the resource are given, latency $l_i(v_i)$ and average energy consumption $e_i(v_i)$ at the supply voltage v_i can be driven by the following equations.

$$e_i(v_i) = e_{i,ref} \times (\frac{v_i}{V_{i,ref}})^2 \qquad \text{(Eq. 3)}$$

$$l_i(v) = l_{i,ref} \times \frac{v \times (V_{i,ref} - V_{i,t})^2}{V_{i,ref} \times (v - V_{i,t})^2} \qquad \text{(Eq. 4)}$$

As the supply voltage decreases, latency increases and the corresponding operating frequency decreases. Likewise, the service rate μ_i at v_i decreases as the voltage v_i decreases shown in (Eq. 5), where $\mu_{i,ref}$ denotes the service rate at the reference voltage $V_{i,ref}$.

$$\mu_i = \mu_{i,ref}\frac{V_{i,ref}(v_i - V_{i,t})^2}{v_i(V_{i,ref} - V_{i,t})^2} \qquad \text{(Eq. 5)}$$

Design Constraints and Problem Overview: Design constraints consist of performance constraints and energy constraints. Performance constraints are given for each region R_i and they are described as (1) maximum input loss probability Λ_i, shown in (Eq. 6) and (2) the sum of the products of weight and the number of outputs per region shown in (Eq. 7). $p_{i,loss}$ denotes the probability of input loss in the region R_i due to the limited buffer space. The weight of an output w_i in region R_i denotes the importance of an output compared with the outputs of other regions. T_i denotes the duration of region R_i. The energy constraint E_{Total} is the total amount of energy that can be used for the processing in a period shown in (Eq. 8), where e_i denotes energy consumption in region R_i. The design goal is to find feasible voltages for the regions that maximize performance while the energy consumption is within the energy budget E_{Total} by tuning voltages and corresponding clock frequencies in the regions.

$$\forall R_i,\quad p_{i,loss} \le \Lambda_i \qquad \text{(Eq. 6)}$$

$$Performance = \sum_i \Omega_i = \sum_i T_i \times w_i \times \lambda_i \times (1 - p_{i,loss}) \qquad \text{(Eq. 7)}$$

$$\sum_i E_i = \sum_i T_i \times \lambda_i \times (1 - p_{i,loss}) \times e_i \le E_{Total} \qquad \text{(Eq. 8)}$$

3. STATIC VOLTAGE SYNTHESIS FOR MAXIMUM PERFORMANCE

In this section, we present an analytical framework for our optimal voltage synthesis technique. We first present the technique for a single region model, and then we extend it for a multiple region model.

3.1 A Period with a Single Region

Throughout a period, we assume the system is homogeneous and has fixed parameters: arrival rate λ, service rate μ_{ref} at the reference voltage V_{ref}, system capacity K, maximum input loss probability Λ and average energy consumption per input e_{ref} at the reference voltage V_{ref}.

The system can be modeled with a single M/M/1/K queue. The steady state probability p_n that there are n data in the M/M/1/K queue is described in (Eq. 9) where ρ is defined in (Eq. 10) [10]. The input loss probability p_{loss} is equal to the steady state probability.

$$p_n = \begin{cases} \frac{(1-\rho)\rho^n}{1-\rho^{K+1}}, & \text{if } \rho \ne 1, n = 0,..,K \\ \frac{1}{K+1}, & \text{if } \rho = 1, n = 0,..,K \end{cases} \qquad \text{(Eq. 9)}$$

$$\rho = \frac{\lambda}{\mu} = \frac{\lambda}{\mu_{ref}\frac{V_{ref}(v-V_t)^2}{v(V_{ref}-V_t)^2}} = \frac{\lambda v(V_{ref}-V_t)^2}{\mu_{ref}V_{ref}(v-V_t)^2} \qquad \text{(Eq. 10)}$$

(Eq. 9) and (Eq. 10) show that the input loss probability depends on the supply voltage. With a higher supply voltage, the input loss probability becomes smaller. Maximum performance is achieved at the smallest p_{loss} value within the energy budget. If the smallest p_{loss} value within the energy budget is larger than the maximum input loss probability Λ, there is no feasible solution. The system produces the largest throughput at the largest v value which satisfies both (Eq. 11) and (Eq. 12). It is trivial to find an optimal solution. We first find v using (Eq. 10), (Eq. 11), and $p_{loss} = \Lambda$. If the resulting v value is not within the feasible range, no feasible solution exists. If there is a feasible v value, the optimal supply voltage is the largest v value satisfying (Eq. 11) and (Eq. 12)

$$E = T \times \lambda \times (1 - p_{loss}) \times e_{ref} \times (\frac{v}{V_{ref}})^2 \le E_{Total} \qquad \text{(Eq. 11)}$$

$$p_{loss} \le \Lambda \qquad \text{(Eq. 12)}$$

3.2 A Period with Multiple Regions

Period T is partitioned into several sub-intervals, and we call them regions denoted by R_1, ..., R_N. Each region R_i has its own parameters: λ_i, μ_i, K_i, T_i, w_i, e_i and Λ_i. We assume no correlation among the inputs, and one input produces one output. The system in region R_i is modeled with a M/M/1/K_i queueing system. We ignore the dependencies between different regions. The assumption of the independence between regions is investigated using simulation in Section 4.

Each region R_i is modeled with an M/M/1/K_i queue. The steady state probability $p_{i,n}$ that there are n data in the M/M/1/K_i queue is described in (Eq. 13) and (Eq. 14). The input loss probability $p_{i,loss}$ is equal to $p_{i,Ki}$ and it must be equal to or smaller than maximum loss probability Λ_i for all regions while the sum of the energy consumption of all regions must be within energy budget, which is shown in (Eq. 7) and (Eq. 8).

$$p_{i,n} = \begin{cases} \frac{(1-\rho_i)\rho_i^n}{1-\rho_i^{K_i+1}}, & \text{if } \rho_i \ne 1, n = 0,..,K_i \\ \frac{1}{K_i+1}, & \text{if } \rho_i = 1, n = 0,..,K_i \end{cases} \qquad \text{(Eq. 13)}$$

$$\rho_i = \frac{\lambda_i}{\mu_i} = \frac{\lambda_i}{\mu_{i,ref}\dfrac{V_{i,ref}(v_i - V_{i,t})^2}{v_i(V_{i,ref} - V_{i,t})^2}} = \frac{\lambda_i v_i (V_{i,ref} - V_{i,t})^2}{\mu_{i,ref} V_{i,ref}(v_i - V_{i,t})^2} \quad \text{(Eq. 14)}$$

It is a trivial process to find the voltages that satisfy the maximum allowable loss probabilities for all regions. However, it is not trivial to find voltages that maximize system performance. We propose an optimal algorithm that finds voltages of the regions producing maximum system throughput within energy constraints. We first define a function $\phi_i(v_i)$ of each region R_i, which is shown in (Eq. 15).

$$\phi_i(v_i) = \frac{dE_i}{d\Omega_i} = \frac{dE_i}{dv_i}\frac{dv_i}{d\Omega_i} = \frac{dE_i}{dv_i}\frac{1}{\dfrac{d\Omega_i}{dv_i}} \quad \text{(Eq. 15)}$$

Function $\phi_i(v_i)$ of each region R_i shows how quickly energy consumption changes as performance changes. Function $\phi_i(v_i)$ increases monotonically as v_i increases in the interval $V_{i,t} < v_i \le V_{i,ref}$, which is shown by Lemma 2. The proofs of the following lemmas can be found in [9][8].

Lemma 1. When the value of v_i is not bounded, $\phi_i(v_i)$ values should be equal for all regions in the optimal solution.

Lemma 2. Function $\phi_i(v_i)$ increases monotonically as v_i increases where $\lambda_i > 0$, $\mu_i > 0$, $V_{i,t} < v_i \le V_{i,ref}$.

Lemma 3. Let n regions R_1, R_2, R_3, .., R_n exist in period T, and their $\phi_i(v_i)$ values be equal. Let energy consumption in period T be E, and performance be Ω. When more energy ΔE is given for more performance and reference voltages are assumed to be infinite, the performance increase $\Delta\Omega$ is maximized where the $\phi_i(v_i)$ values of all regions are the same.

From Lemma 2 and Lemma 3, the maximum throughput gain compared to the setting where their original $\phi(v)$ values are the same is obtained at the same $\phi(v)$ value for all resources. The assumption of infinite reference voltage guarantees the existence of a voltage at the resulting $\phi(v)$ value.

Based on this observation, we propose an optimal voltage assignment algorithm that produces maximum performance subject to energy constraints.

Algorithm 1. Optimal Voltage Assignment for regions in a period

(1) **S** ← $\{i \mid R_i\}$; $E \leftarrow E_{Total}$
(2) For $i \in$ **S [**
(3) Calculate v_i, $\phi_i(v_i)$ and $E_i(v_i)$ where $p_{i,loss} = \Lambda_i$
(4) $V_{i,bot} \leftarrow v_i$
(5) If ($v_t > V_{i,ref}$) **[** No feasible solution exists!! Exit. **]]**
(6) If ($E < \sum_{i\in S} E_i(v_i)$)**[** No feasible solution exists!! Exit. **]**
(7) If ($E \ge \sum_{i\in S} E_i(V_{i,ref})$)**[** Solution found!! Exit. **]**
(8) **Q** ← $\{i \mid v_i = V_{i,ref}\}$; **S** ← **S** – **Q**; $E \leftarrow E - \sum_{i\in Q} E_i$
(9) **T** ← $\{(i, \phi_i(v)) \mid i \in$ **S** and ($v = V_{i,ref}$ or $v = V_{i,bot}$)$\}$
(10) Sort **T** in descending order of $\phi_i(v)$
(11) While (**S** is not empty) **[**
(12) m ← number of unique $\phi_i(v)$ values in **T**
(13) if ($m = 0$) **[**
(14) Maximize $\sum_{i\in S}\Omega_i$ subject to $E = \sum_{i\in S} E_i(v_i)$ using the method of Lagrange multipliers. **]**
(15) Φ ← ($\lfloor m/2 \rfloor$+1)-th largest $\phi_i(v)$ value in **T**
(16) For $i \in$ **S [**
(17) $v'_i \leftarrow v$, where $\phi_i(v'_i) = \Phi$
(18) if ($v'_i > v_{i,ref}$) then **[** $v'_i \leftarrow v_{i,ref}$ **]**
(19) If ($\Psi \ge \sum_{i\in S} E_i(v'_i)$)**[** /* Upper Half */
(20) For ($i_i \in$ **S**) **[** $v_i \leftarrow v'_I$ **]**
(21) If ($E = \sum_{i\in S} E_i(v'_i)$) **[** Solution found!! Exit **]**
(22) $Low \leftarrow \Phi$
(23) **Q** ← $\{i \mid v_i = V_{i,ref}\}$; **P** ← $\{(i, \phi_i(v)) \mid \phi_i(v) \le \Phi\}$ **]**
(24) Else ($E < \sum_{i\in S} E_i(v'_i)$)**[** /* Lower Half */
(25) $High \leftarrow \Phi$
(26) **Q** ← $\{i \mid v'_i = V_{i,bot}\}$; **P** ← $\{(i, \phi_i(v)) \mid \phi_i(v) \ge \Phi\}$ **]**
(27) **S** ← **S** – **Q**
(28) $E \leftarrow E - \sum_{i\in Q} E_i(v_i)$
(29) **T** ← **T** – $\{(i, \phi_i(v)) \mid i \in$ **Q** and ($v = V_{i,bot}$ or $V_{i,ref}$)$\}$
(30) **T** ← **T** – **P]**

The algorithm first finds an initial solution that satisfies input loss constraints of all regions. Each region is evaluated independently and has a single solution. If the resulting solution does not satisfy energy constraints, there is no feasible solution, which is shown in lines (1) through (6). The while loop in lines (11) through (32) is a binary search algorithm that reduces the search space by half, as shown in line (15). The size of the search space is at most 2N where N is the number of regions in the period. At each iteration of the while loop, the median $\phi_i(v)$ value Φ in set **T** is chosen and the corresponding voltage is estimated in lines (15) through (18). If the $\phi_i(v)$ value of a region at the reference voltage is smaller than Φ, its voltage is set at the reference voltage. When the energy consumption at Φ is smaller than the energy constraint E_{Total}, Φ is set to the lower bound of the search space and the regions with current voltages equal to the reference voltage are removed from further consideration, which is shown in lines (22) – (23) and (27) – (30). When the energy consumption is equal to the energy constraint, the optimal solution is found in line (21).

Otherwise, Φ is set to the upper bound, and the regions with current voltages equal to the initial voltages are removed from further consideration as shown in lines (25) – (26) and (27) – (30). When set **T** becomes empty as shown in line (13), the upper bound 'High' and the lower bound 'Low' of $\phi(v)$ value are given. Lagrange's method is used to find an optimal solution as shown in line (14).

We present a simple example to show the behavior of the algorithm. Assume that there are two regions in a period R_1, R_2, $\phi_1(V_{1,bot}) = a_1$, $\phi_1(V_{1,ref}) = b_1$, $\phi_2(V_{1,bot}) = a_2$, $\phi_2(V_{2,ref}) = b_2$, $a_1 < a_2 < b_1 < b_2$. In its first iteration of the while loop, there are four unique ϕ values in **T** and the third largest ϕ value in **T** is chosen as Φ that is a_2. Energy consumption at Φ is smaller than the energy constraint. According to line (23), **Q** = {} and **P** = {(1, a_1),(2, a_2)}, and **T** becomes {(1, b_1), (2, b_2)}. In the second iteration, the algorithm probes b_1. Now, the energy consumption at b_1 is larger than the energy constraint. According to line (26), **Q** = {} and **P** = {(1, b_1), (2, b_2)}. Since set **T** is empty, the algorithm finds optimal voltages $v_{1,opt}$ and $v_{2,opt}$ using Lagrange's method, which is shown in line (14). More explanation of the example and the proof of Theorem 1 are found in [9].

Theorem 1. When there is a feasible design that meets energy constraints and input loss constraints for all regions, Algorithm 1 finds an optimal solution with maximum performance.

When we assume the existence of the inverse function of $\phi_i(v_i)$, the complexity of Algorithm 1 is O(N log N), where N denotes the number of regions in a period and solving for the *Lagrange* multiplier is assumed to be done in O(N) time [4].

3.3 Solution of the motivational example

We used Algorithm 1 to solve the motivational problem shown in Table 1 and Table 2. We used binary search algorithm to find Lagrange's multiplier in line (14) of Algorithm 1 with 0.0001 percent error bound. It takes less than 8 milliseconds to find the solution on a Pentium III-733MHz PC running Linux.

	R_1	R_2	R_3	R_4	R_5	R_6	R_7
V_d	1.795	2.528	1.374	1.844	2.544	1.844	1.340
Λ_i	0.1	0.02	0.1	0.05	0.01	0.05	0.1
$p_{i,loss}$	0.1	0.02	0.02	0.05	0.01	0.05	0.075

4. SIMULATION

The analytical framework assumes there is no interference between two neighboring regions. However, this may not be true. Region R_i can interfere with R_{i+1} when the queue is not empty or the processor is not idle at the time of the region change. In this situation, the processor is working on data collected in region R_i even though the satellite is flying over region R_{i+1}. This situation gets worse when region R_i is highly loaded and the duration of region R_{i+1} is short.

To observe interference between regions, we simulated the behavior of the motivational example and its variants. Three queue management policies at the boundaries of the regions are simulated which include (1) no management, (2) drop all inputs in the queue except the one that is being processed, and (3) drop all inputs including the one that is currently processed.

The simulated result of the example is shown in Table 3. In Table 3(a), $p_{i,loss}$ denotes the loss probability given by the analysis and $p_{i,loss(i)}$ denotes the average packet loss probability observed by the simulation using the *i-th* queue management policy at the boundaries of the regions. Likewise, $\sigma_{i,loss(1)}$ denotes standard deviation of the packet loss probability using the *i-th* queue management policy.

Our analytic result does not depend on how large the arrival rate and service rate are but does depend on their ratio. To observe the influence of a large and small arrival rate and service rate, we scaled the service rate and arrival rate by 10, 0.1 and 0.01 while their ratio remained the same. The results are summarized in Table 4, Table 5, and Table 6. From those tables, we observe that our analysis technique predicts well for high arrival and service rates. At low arrival and service rates, our analytical result deviates from simulation results. As the arrival and service rates decrease as shown in Table 6, interdependency among the regions at the boundaries between regions gets bigger and the technique predicts better for queue management policies (2) and (3) than no management policy (1).

Table 3. Simulation result of the example.

	R_1	R_2	R_3	R_4	R_5	R_6	R_7
$p_{i,loss}$	0.1	0.02	0.02	0.05	0.01	0.05	0.075
$p_{i,loss\ (1)}$	0.10	0.02	0.02	0.05	0.01	0.05	0.075
$p_{i,loss\ (2)}$	0.10	0.02	0.02	0.05	0.01	0.05	0.076
$p_{i,loss\ (3)}$	0.10	0.02	0.02	0.05	0.01	0.05	0.076
$\sigma_{i,loss\ (1)}$	2e-3	1e-3	3e-3	3e-3	7e-3	3e-3	3e-3
$\sigma_{i,loss\ (2)}$	2e-3	1e-3	3e-3	3e-3	7e-4	3e-3	3e-4
$\sigma_{i,loss\ (3)}$	2e-3	1e-3	3e-3	3e-3	7e-4	3e-3	3e-4

(a)

	Energy	$\sigma_{energy,}$	Performance	$\sigma_{performance}$
(1)	3.3e6	1748	3.4e6	1750
(2)	3.3e6	1754	3.4e6	1764
(3)	3.3e6	1752	3.4e6	1763

(b)

Table 4. Simulation result (10 times service rate and arrival rate)

	R_1	R_2	R_3	R_4	R_5	R_6	R_7
$p_{i,loss}$	0.1	0.02	0.02	0.05	0.01	0.05	0.075
$p_{i,loss\ (1)}$	0.1	0.02	0.02	0.05	0.01	0.05	0.075
$p_{i,loss\ (2)}$	0.1	0.02	0.02	0.05	0.01	0.05	0.075
$p_{i,loss\ (3)}$	0.1	0.02	0.02	0.05	0.01	0.05	0.075
$\sigma_{i,loss\ (1)}$	6e-4	3e-4	8e-4	9e-4	2e-4	9e-4	7e-4
$\sigma_{i,loss\ (2)}$	6e-4	3e-4	8e-4	9e-4	2e-4	9e-4	7e-4
$\sigma_{i,loss\ (3)}$	6e-4	3e-4	8e-4	9e-4	2e-4	9e-4	7e-4

(a)

	Energy	$\sigma_{energy,}$	Performance	$\sigma_{performance}$
(1)	33.0e6	5627	34.3e6	5404
(2)	33.0e6	5785	34.3e6	5666
(3)	33.0e6	5785	34.3e6	5666

(b)

Table 5. Simulation result (0.1 times service rate and arrival rate)

	R_1	R_2	R_3	R_4	R_5	R_6	R_7
$p_{i,loss}$	0.1	0.02	0.02	0.05	0.01	0.05	0.075
$p_{i,loss\ (1)}$	0.1	0.02	0.02	0.05	0.01	0.05	0.074
$p_{i,loss\ (2)}$	0.1	0.02	0.02	0.05	0.01	0.05	0.076
$p_{i,loss\ (3)}$	0.1	0.02	0.02	0.05	0.01	0.05	0.076
$\sigma_{i,loss\ (1)}$	5e-3	3e-3	8e-3	9e-3	2e-3	8e-3	8e-3
$\sigma_{i,loss\ (2)}$	5e-3	3e-3	9e-3	9e-3	2e-3	8e-3	8e-3
$\sigma_{i,loss\ (3)}$	5e-3	3e-3	9e-3	9e-3	2e-3	8e-3	8e-3

(a)

	Energy	$\sigma_{energy,}$	Performance	$\sigma_{performance}$
(1)	0.33e6	552	0.34e6	562
(2)	0.33e6	564	0.34e6	564
(3)	0.33e6	565	0.34e6	563

(b)

Table 6. Simulation result (0.01 times service rate and arrival rate)

	R_1	R_2	R_3	R_4	R_5	R_6	R_7
$p_{i,loss}$	0.1	0.02	0.02	0.05	0.01	0.05	0.075
$p_{i,loss\ (1)}$	0.11	0.05	8e-3	0.05	0.03	0.03	0.064
$p_{i,loss\ (2)}$	0.10	0.02	0.04	0.06	0.01	0.06	0.080
$p_{i,loss\ (3)}$	0.10	0.02	0.04	0.06	0.01	0.06	0.081
$\sigma_{I,loss\ (1)}$	2e-3	1e-2	2e-2	3e-2	7e-3	2e-2	0.026
$\sigma_{I,loss\ (2)}$	2e-3	1e-2	1e-2	3e-2	3e-2	2e-2	0.026
$\sigma_{I,loss\ (3)}$	2e-3	1e-2	2e-2	3e-2	7e-3	2e-2	0.026

(a)

	Energy	$\sigma_{energy,}$	Performance	$\sigma_{performance}$
(1)	32.4e3	198	33.7e3	201
(2)	32.9e3	175	34.2e3	177
(3)	32.9e3	175	32.9e3	177

(b)

5. RELATED WORK

There have been many run-time techniques for system-level power reduction. Dynamic voltage scheduling techniques change the voltages of resources and corresponding clock frequencies [5][6][14]. Dynamic power management techniques shut down resources at idle time for additional power savings [11]. Dynamic power management techniques have been extended with a dynamic voltage scheduling technique [13]. In real-time systems, exploiting resource slack time has been studied for reducing power consumption in a single CPU system [1][11] and a distributed system [8]. Low-power techniques have been an important topic for CAD research [1][12]. However, most of them are chip-level VLSI design techniques, such as multiple voltage scheduling among functional units in a chip [3][11]. These techniques address the problem of assigning a supply voltage from a finite number of pre-known supply voltages to each operation in a data flow graph so that the resulting schedule minimizes power consumption and satisfies timing constraints.

Our technique is similar to a dynamic voltage scheduling technique in that it synthesizes the supply voltage of a resource in a region based on stochastic event prediction. However, our technique is unique in solving the power-aware optimization problem subject to a performance constraint rather than the low-power scheduling problem. While low-power scheduling techniques consider only power consumption, our technique considers both system performance and energy consumption and synthesizes optimal voltages that maximize system performance.

6. CONCLUSION

We presented an analytical optimal voltage synthesis technique for a satellite application. The technique optimizes voltages for each region of a satellite's period so that the system performance is maximized under an energy budget. We presented simulation results showing the quality of the packet loss probabilities at the regions estimated by the analysis.

7. ACKNOWLEDGMENTS

This effort is sponsored by Defense Advanced Research Projects Agency (DARPA) through the Air Force Research Laboratory, USAF, under agreement number F30602-00-2-0548.

8. REFERENCES

[1] L. Benini and G. D. Micheli, "System-Level Power Optimization: Techniques and Tools," International Symposium on Low Power Electronics and Design, pp. 288-293, 1999.

[2] A. P. Chandrakasan, S. Sheng, and R. W. Brodersen, "Low-Power CMOS Digital Design," IEEE Journals of Solid-State Circuits, pp. 473-484, vol. 27, no. 4, April 1992.

[3] J. Chang and M. Pedram, "Energy Minimization Using Multiple Supply Voltages," IEEE Transactions on Very Large Scale Integrated (VLSI) Systems, vol. 5, no.4, Dec. 1997.

[4] K. Govil, E. Chan, and H. Wasserman, "Comparing Algorithms for Dynamic Speed-Setting of a Low-Power CPU," Mobile Computing and Networking, November 1995.

[5] J. Dey, J. Kurose, and D. Towsley, "On-Line Scheduling Policies for a Class of IRIS (Increasing Reward with Increasing Service) Real-Time Tasks," IEEE Transactions on Computers, Vol. 45, No. 7, Jul. 1996.

[6] I. Hong, G. Qu, M. Potkonjak, and M. B. Srivastava, "Synthesis Techniques for Low-Power Hard Real-Time Systems on Variable Voltage Processors," IEEE Real-Time Systems Symposium, pp. 178-187, Dec. 1998.

[7] JTRACK, NASA http://liftoff.msfc.nasa.gov/realtime/JTrack/, 2001

[8] D. Kang, S. Crago, and J. Suh, "Power-Aware Design Synthesis Techniques for Distributed Real-Time Systems," ACM Workshop on Languages, Compilers, and Tools for Embedded Systems (LCTES) '01, Jun. 2001.

[9] D. Kang, J. Suh, and S. Crago, "An Optimal Voltage Synthesis Technique for a Power-Efficient Satellite Application," unpublished USC/ISI technical report, also found at http://www.east.isi.edu/publications/satellite.pdf.

[10] Leonard Kleinrock, "Queueing Systems," volume I, John Wiley & Sons, 1975.

[11] Q. Quiu and M. Pedram, "Dynamic Power Management Based on Continuous-Time Markov Decision Processes," Design Automation Conference, June 1999.

[12] D. Singh, J. Rabaey, S. Rajgopal, N. Sehgal, and T. Mozdzen, "Power Conscious CAD Tools and Methodologies: A Perspective," Proceedings of IEEE, pp. 570-594, vol. 83, no. 4, April 1995.

[13] T. Sinumic, L. Benini, A. Acquaviva, P. Glynn, and G. D. Micheli, "Dynamic Voltage Scaling and Power Management for Portable Systems," Design Automation Conference, June 2001.

[14] M. Weiser, B. Welch, A. Demers, and S. Shenker, "Scheduling for Reduced CPU Energy," Operating Systems Design and Implementation, Nov. 1994.