Brief Announcement: Power Management of Devices- When Should I Switch Off?

Ajay Gulati*
VMware Inc., Palo Alto, CA 94304
agulati@vmware.com

ABSTRACT

Increasing power costs has renewed interest in software based power optimizations. One of the key question that any software based mechanism needs to decide is: *How long should I wait before changing the power state of a device?* In this paper, we study the solution space of this problem based on the workload history and device characteristics. We model the workload as a distribution of idle times (collected online) and the device characteristics using break-even time(Δ). Our analysis shows that one should not do any power savings if the mean of the idle time distribution is lower than Δ. We also show that being *greedy* (i.e. wait time = 0) gives the highest power savings for most of the well known distributions. Our model is also able to capture special cases such as bimodal distributions.

Categories and Subject Descriptors: D.4.8[Operating Systems]:Performance - *Modeling and prediction*

General Terms: Algorithms, Design, Management.

Keywords: Dynamic power management, Power states

1. INTRODUCTION

Power is becoming one of the major costs for today's data centers and high performance systems. Virtualization is becoming a key to enabling more flexible workload placement and power savings. Hardware vendors are treating power as the first class citizen in their designs there by supporting multiple power states for each device. Existing works have proposed policy based power management both in hardware and software. A comparison of various power management techniques can be found in [1,2]. Some of the common techniques used are fixed timeout, adaptive exponential average and stochastic modeling. Based on power consumption values for different states and transitions [1, 2], one can define the time threshold Δ (break-even time), which is the idle time interval that leads to zero power savings.

In this paper, we try to solve the following problem: *How long should I wait before powering down the device for a given Δ, so as to get maximum power savings?* We model this as an optimization problem, characterize each workload by a dynamically computing histogram of idle times and computationally solve the optimization for various known distributions. The probability density function for idle times durations for the workload is represented by $f(\textit{idle time})$, where $f(t)$ gives the probability of having an idle time equal to t. We consider a continuous function for modeling but in practice we can have a discrete histogram, which can be computed online based on past history. For a specific idle time t and wait time T, the gain G(t) would be:

$$G(t) = P\ (t\ -\ T\ -\ \Delta) \qquad (1)$$

Here P denotes energy saving per unit time. Let $B(T)$ denote the overall gain obtained using a wait time T, for a given Δ and function $f(t)$. Using $f(t)$, the overall benefit for a particular wait time T, can be written as:

$$B(T) = \int_{T}^{\infty} P\ (t - T - \Delta) f(t) dt \qquad (2)$$

Now for a given $f(t)$, we need to choose the value of T, that gives maximum expected value of benefit $B(T)$. Note that the analysis assumes that there power savings or loss is proportional to $(t - T + \Delta)$.

Our analysis for various distributions (exponential, gaussian, uniform and bimodal) shows that based on the relationship between Δ and mean value of idle time distribution (M), one should either use or discard power savings. In particular, if $\Delta >= M - \epsilon$, one shouldn't do any power savings and if $\Delta < M - \epsilon$, one should do it. Here ϵ is a small constant (1 or 2) for most cases. Furthermore we observe that just using the greedy approach of powering off without any wait gives best savings in most cases except bi-modal distributions. This shows that our model is able to adapt based on the actual distribution of the workload behavior and device characteristics.

2. REFERENCES

[1] L. Benini, A. Bogliolo, and G. Micheli. A survey of design techniques for system-level dynamic power management. In *IEEE Transactions on VLSI Systems*, 2000.

[2] Y.-H. Lu, E.-Y. Chung, T. Šimunić, L. Benini, and G. D. Micheli. Quantitative comparison of power management algorithms. In *DATE*, pages 20–26, 2000.

*Dr. Gulati is a member of resource management team at VMware. An extended version of this paper can be found at: `www.cs.rice.edu/~gulati`

PODC'08, August 18–21, 2008, Toronto, Ontario, Canada.
ACM 978-1-59593-989-0/08/08.