

PICSEL: Measuring User-Perceived Performance to Control Dynamic Frequency Scaling

Arindam Mallik

Jack Cosgrove

Robert P. Dick

Gokhan Memik

Peter Dinda

Department of Electrical Engineering and Computer Science, Northwestern University

Evanston, Illinois, USA

{arindam, i-cosgrove, dickrp, memik, pdinda}@northwestern.edu

Abstract

The ultimate goal of a computer system is to satisfy its users. The success of architectural or system-level optimizations depends largely on having accurate metrics for user satisfaction. We propose to derive such metrics from information that is “close to flesh” and apparent to the user rather than from information that is “close to metal” and hidden from the user. We describe and evaluate PICSEL, a dynamic voltage and frequency scaling (DVFS) technique that uses measurements of variations in the rate of change of a computer’s video output to estimate user-perceived performance. Our adaptive algorithms, one conservative and one aggressive, use these estimates to dramatically reduce operating frequencies and voltages for graphically-intensive applications while maintaining performance at a satisfactory level for the user. We evaluate PICSEL through user studies conducted on a Pentium M laptop running Windows XP. Experiments performed with 20 users executing three applications indicate that the measured laptop power can be reduced by up to 12.1%, averaged across all of our users and applications, compared to the default Windows XP DVFS policy. User studies revealed that the difference in overall user satisfaction between the more aggressive version of PICSEL and Windows DVFS were statistically insignificant, whereas the conservative version of PICSEL actually improved user satisfaction when compared to Windows DVFS.

Categories and Subject Descriptors C.3 [Performance of Systems]: Measurement Techniques, Performance Attributes; B.4.2 [Input/Output and Data Communications]: Input/Output Devices - Image display

General Terms Algorithms, Management, Measurement, Performance, Human Factors

Keywords User-perceived Performance, Dynamic Voltage and Frequency Scaling, Power Management, Thermal Emergency

1. Introduction

Existing architectures and systems software typically optimize for user satisfaction by employing metrics based largely on instruction throughput (e.g., instructions-per-second). These metrics are used because they are easy to access, easy to compare across platforms, and are believed to reflect user demands for performance at a very low level. However, in this paper, we will show that low-level

information is not as good a proxy for user satisfaction with performance as is high-level information actually observed or perceived by the user. We focus on interactive applications and show that it is possible to infer information about user-perceived performance by measuring changes in video output. This provides better information about the performance level necessary to maintain user satisfaction. We demonstrate the utility of this information in on-line power management.

Processor frequency has a strong effect on power consumption and temperature, directly and also indirectly through the need for higher voltages at higher frequencies. Dynamic Voltage and Frequency Scaling (DVFS) is one of the most commonly used power reduction techniques in modern processors. DVFS varies the frequency and voltage of a microprocessor at runtime to trade off power consumption and processor performance. Specifically, existing DVFS techniques in high-performance processors select an operating point (CPU frequency and voltage) based on the utilization of the processor and other information available to the Operating System (OS) kernel. This approach is often pessimistic regarding user satisfaction, setting the processor frequency higher than necessary to ensure user satisfaction with performance. A high level of CPU utilization or a burst of certain OS events leads directly to a high frequency (and high voltage), regardless of the user’s satisfaction with performance. This can produce unnecessary increases in frequency, voltage, power consumption, and temperature.

In response to this observation, we have developed a new power management technique that relies upon a more accurate proxy for user performance needs than CPU- or OS-level events, but that is still inexpensive to measure. We estimate user satisfaction with processor performance using information that is “close to flesh” and apparent to the user rather than information that is “close to metal” and hidden from the user. Interface devices are the logical locations for these measurements since they sit between computation and user perception. The display is particularly useful because it is the user’s primary source of information regarding the performance of the computer.

We must note that a user satisfaction-aware optimization policy does not need an absolute metric for user-perceived performance to make decisions. The policy will only make decisions for the architecture on which it is implemented. Using this idea in the context of DVFS, we can compare the displayed performance of applications that change the display at lower frequencies to the displayed performance at the processor’s highest available frequency. If the two frequencies result in identical sequences and timing of frames on the screen, then we can safely conclude that these two processor states have the same displayed performance. This maximum frequency satisfies user demands for displayed

performance as well as the architecture can, marking a basis for satisfying the user that is fixed for the architecture. Hence, initializing a DVFS policy at the maximum frequency “seeds” the policy with a meaningful level of displayed performance.

To bring this idea to life and evaluate it, we have developed a new power management framework called **PICSEL** (**P**erception-**I**nformed **C**PU performance **S**caling to **E**xtend battery **L**ife) that monitors the rate of change of pixel intensities in the display. An algorithm controlling the processor’s operating frequency then makes decisions based upon these rates of change. The algorithm is tested with two configurations: conservative PICSEL (cPICSEL) and aggressive PICSEL (aPICSEL) (Section 3.2). We focus on the DVFS technique implemented by a commercial OS and show that runtime estimation of user-perceived performance using pixel intensities can enhance the effectiveness of the power management scheme. We also show that this approach can result in optimizations that are not possible otherwise. Our work makes the following contributions:

1. We show that traditional performance metrics do not necessarily represent user-perceived performance,
2. We introduce new metrics that can successfully measure user-perceived performance, and
3. We propose, implement, and evaluate PICSEL, a power management scheme based on estimates of user-perceived performance.

2. User-Perceived Performance

The motivation for including user-perceived performance in any objective function is clear: an optimization ultimately aims to satisfy the user. However, the difficulty in optimizing directly for user-perceived performance is in finding a corresponding metric that can be efficiently measured at runtime. For interactive applications, the events occurring on the input/output devices are good candidates for measuring user satisfaction with performance. However, input events are rare compared to output events. Therefore, considering output to the user is preferable for estimating the performance experienced by the user. Of all the types of output supplied to the user, graphics are used in the highest proportion of applications. Therefore, exploiting properties of the display to estimate user-perceived performance is a good alternative.

Given an application that only changes the display, it is plausible that the frame sequence and frame rate are indications of the user-perceived performance of that application. For example, if there are two architectural alternatives that result in identical frame sequences and frame rates, we can reasonably say that these architectures provide the same user-perceived performance. Ghinea and Thomas [8] have done a perceptual study showing that varying both the color depth and the frame rate has a significant effect on user satisfaction with performance. However, extracting the exact frame rate and color depth information would require changes in the application or OS. Hence, we decided to employ a metric that is independent of the application and easily measurable: the rate of pixel intensity change over time. This captures the combination of these two metrics.

We have performed a set of experiments to understand the relationships between instruction throughput, rate of pixel change, and user-perceived performance. In these experiments, we measured the number of instructions-per-second (IPS) on a 2.13 GHz Intel Pentium M-based laptop (please see Section 4.1 for further details

on the experimental study environment) for three applications: a 3D Shockwave animation, a DVD quality video, and a 3D video game. We also measured the changes in intensity in the red, green, and blue channels of some of the pixels being used to display these applications using the method described in Section 3.1, and averaged these changes together for each time instance to obtain the **Average Pixel Change (APC)**. The procedure to calculate APC is presented in Table 1. We repeated these measurements at all six available processor operating frequencies.




Figure 1. IPS and APC curve

Table 1. User-Perceived Performance Metrics

Metrics	Measurement Procedure
Average Pixel Change (APC)	- Capture the Pixel intensities of the RGB channels of all the pixels in a memory buffer - Calculate the relative changes for all the sampled pixels - The mean of relative changes is the APC
Rate of Average Pixel Change (APR)	$(APC_{T_i} - APC_{T_{i-1}}) / (T_i - T_{i-1})$

Figure 1 illustrates the results of this experiment, with the solid lines representing the APC curve and dotted lines representing the IPS curve. As depicted in the figure, the IPS of the system is closely related to the operating frequency of the CPU and is fairly uniform across the three applications. APC is also dependent on the operating frequency, but this dependence is influenced by the application more than IPS is influenced. For the Shockwave application, the effect on APC due to frequency throttling is below 10% for the highest three frequencies. The Video application shows similar properties. For this task, we could simply set the frequency statically to a lower value without causing noticeable change in the APC. For the game application, the highest two frequency states can sustain the APC value within 10% of its maximum value. However, the lower frequency states cause the APC value to drop suddenly. Most importantly, we see a significant difference between the reduction in IPS and APC. In other words, these results show that the instruction throughput and user-perceived performance are not linearly related. We observe that the APC value of a system can quantify user perceived performance and can be used as a metric for a power management scheme that implements DVFS based on user-perceived performance.

The primary metric we use for user-perceived performance is APC normalized to the total number of pixels in the display. As shown in Figure 1, we observe considerable variation in the APC values across different applications as well as different frequency states.

On the other hand, it is also possible that the reduction in the frequency may result in discontinuities in the display. Previous researchers [10] have found that jitter and latency are the main sources of user discontent in networked multimedia applications. For example, consider an application that starts skipping frames when the computational power is reduced. In such a case, the APC may not be affected significantly: in a sequence of frames, even if some of the intermediate frames are skipped, the pixel difference between the first and the last does not change.

To capture the occurrences of such discontinuities, we record the **Rate of Average Pixel Change (APR)** normalized over the number of pixels. In other words, we calculate the difference between the APC values measured at each time instant. This roughly corresponds to the derivative of the APC. Figure 2 illustrates the APR trends observed in three applications used in this paper. When there are glitches during display, the APR value tends to increase rapidly. This is true for applications where video glitches are observed at lower frequencies, namely the Video and 3D Shockwave animation. For other applications (such as the game), we simply observe an overall slowdown and APR values drop in parallel to APC levels. Such applications reduce game jitter at the price of reducing the frame rate. As a result, for this particular application we actually observe a reduction in APR value at lower frequencies as the game’s average frame rate is reduced.




Figure 2. APR curves for the three applications

APR reveals even more pronounced differential behavior across applications than APC. This behavior can permit a DVFS algorithm to differentiate between two applications with similar computational loads and to assign them to different operating frequencies, one potentially lower than would have otherwise been assigned by existing pessimistic DVFS schemes.

3. PICSEL Framework

User-perceived performance-based frequency scaling has two components. First, we have to measure the rate of change in the pixels displayed on the screen. This measurement tool is described in the next section. Then, we have to make a throttling decision based on these measurements. The algorithm making this decision is described in Section 3.2. In Section 3.3, we describe how PICSEL interacts with the system.

3.1 PICSEL Display Access

PICSEL gathers screen information using the Windows GDI screenshot method, which is simple to implement and can blit any

region of the screen to main memory. However, screen content may be missing from sections of the blitted region if those sections were drawn elsewhere in video memory by a rendering or decoding operation. We set our applications to perform rendering and decoding in software in order to capture these operations with a GDI screenshot. This also places the computational load for those operations on the CPU, thus making them subject to CPU frequency scaling. Ideally we would like to consider all the pixels present in the display while calculating the APC. Furthermore, the rate of APC calculation should be same as the rate of frame change in the system. However, both of these constraints introduce heavy computational overhead on the system. Therefore, it is necessary to reduce the size of the captured screen area so that the capturing process does not occupy too much of the computer’s resources. We decided to limit the overhead to less than 2% CPU utilization. The final captured area is 64 by 51 pixels, or a scaling down of each dimension of a 1280 by 1024 screen by a factor of 20. This area contained 3276 pixels and was fixed at the center of the screen. We chose to capture a contiguous rectangle of pixels rather than a disjoint grid of pixels because capturing the disjoint grid proved to be much more computationally intensive than capturing the contiguous rectangle, holding the number of pixels constant. Because blitting transfers contiguous blocks of data by design, fewer transfers are necessary to capture an area covered by a fraction of the blocks rather than a screen-size grid covered by all of the blocks.

The sampling frequency for calculating APC was chosen to be 10 Hz, the highest frequency with which our framework did not exceed the 2% CPU utilization threshold for the captured pixel area. There is a computational tradeoff between sampling frequency and capture area which we tuned through initial testing on applications including, but not limited to, the three tested applications. As we will show in Section 4, these sampling parameters do not prevent PICSEL from capturing the user-perceived performance for our target applications. Nevertheless, it is possible that applications will not use our focus area; hence it may be desirable to overcome these limitations for other application domains. There are two design alternatives to solve this problem. First, PICSEL can be implemented in hardware (either on the CPU or on the graphics card). The simplicity of the algorithm ensures relatively easy hardware implementation. Second, PICSEL can be executed as software on the graphics hardware. Although such implementations would be desirable, our goal in this work is to provide a proof-of-concept, which is achieved with the current implementation of PICSEL.

After a section of the screen has been captured, it is stored to a memory buffer. This buffer is compared to another buffer containing the previous screen capture, and the intensity differences for the red, green, and blue channels are calculated. Only two buffers are necessary, with each buffer toggling between old and new screen captures. All of the magnitude differences are averaged to obtain a single value for the first time derivative of pixel intensity over the sampling period (APC).

It is important to understand that this method does not capture each frame and that there are unaccounted-for frames between the two frames used to calculate the intensity difference. This introduces noise into the APC metric. However, since we also measure the APR, we can detect trends in pixel intensity that would otherwise be

obscured by the noise. This permits a sampling frequency below the frame rate of the screen.

3.2 PICSEL Algorithm

PICSEL decides on the frequency level by using three state variables: f , the current CPU frequency; μ_{APC} , APC in the last time interval; and μ_{APR} , APR in the last time interval. Pixel data are measured at fixed sampling frequency and stored to a file by a background process. Adaptation is controlled by three constant parameters: ρ , the APC change threshold; γ , the APR change threshold; and α , the threshold difficulty level corresponding to each frequency state.

PICSEL can either be in the initialization or the control state. The idea in the initialization stage is to capture information about the APC and APR values observed at the highest frequency. These values will be used as a base case for comparison during the control stage to make throttling decisions. Therefore, during initialization, the CPU frequency is set at the highest value f_{max} for a time interval T_{init} . The APC and APR values of the system over the time interval T_{init} are obtained from the background process and initialized as APC_{init} and APR_{init} . PICSEL then enters the control state where at the end of each time interval T_i , the APC and APR of the system over the last interval are obtained from the background process. PICSEL then makes a decision as follows:

```

IF ( $APC_{init} - \mu_{APC}$ ) <  $\rho \times (1 - \alpha) \times APC_{init}$ 
  OR  $|\mu_{APR_{init}} - \mu_{APR}| < \gamma \times (1 - \alpha) \times APR_{init}$ 
    Reduce  $f$  by one level
    Reset  $\alpha$  of the last level to 0
ELSE
  Increase  $f$  by one level
  Increment  $\alpha$ 

```

The main idea in this pseudocode is to compare the last observed APC and APR against the APC and APR captured when the processor is executing at the highest frequency. Then, based on the threshold factors defined by ρ , γ , and α , we may conclude that the user-perceived performance is unchanged and try to reduce the frequency and power consumption. Otherwise, out-of-bound values of μ_{APC} and μ_{APR} suggest that user-perceived performance has suffered in the last interval due to low CPU frequency and it is increased accordingly to improve the user-perceived performance.

Factor α introduces hysteresis to eliminate the possible ping-pong effect between two frequency states. If the processor has been at a state several times after which PICSEL had to increase the frequency, α makes it harder to return to that frequency level. Following every third ($n=3$) update to α , PICSEL reenters the initialization state. This feature of the algorithm ensures that PICSEL will gradually adjust to a new set of operating conditions. The constant parameters ($T_i = 7$ seconds, $T_{init} = 10$ seconds) were set based on the experience of the authors using the system. α is initialized to zero for each of the frequency levels and is incremented by 0.1 for each frequency boost. We used two variations of the PICSEL algorithm by fixing the $\rho = 0.05$, $\gamma = 0.15$ and $\rho = 0.10$, $\gamma = 0.30$, which correspond to **conservative PICSEL (cPICSEL)** and **aggressive PICSEL (aPICSEL)**, respectively.

Ideally, we would like to empirically evaluate the sensitivity of PICSEL performance to these parameters. However, it is important

to note that any such study would require having real users in the loop, and thus would be quite slow. Testing three values of five parameters on 20 users would require 243 days (based on 20 users per day and 25 minutes per user). For this reason, we decided to choose the parameters based on qualitative evaluation by the authors and then close the loop by evaluating the whole system with the choices.

3.3 Current Implementation and Integration

For our user studies, we disable the default DVFS policy and give control of the processor frequency to PICSEL. Once PICSEL is active, it executes client software that runs as a Windows toolbar task as well as an API that controls CPU frequency based on user perceived performance. In the client, we log the APC and APR at the background. The API uses these values to control CPU frequency. It is this implementation that we evaluate in the next section.

In its current implementation, PICSEL has some limitations, which will be addressed once it is integrated with the OS. PICSEL should be activated only if the system is executing an application that modifies the display. We detect applications with display output through constant (but infrequent) monitoring of the device. For example, a running process can monitor the APC/APR values every 10 seconds. Then, if this rate is above a threshold, we conclude that the current foreground application modifies the display and we activate PICSEL frequency control. If the APC/APR value drops below a threshold, PICSEL will conclude that the application is completed and give the control back to Windows DVFS. In scenarios when the display is static, PICSEL will detect that the rate of change in the display is below the threshold and give the control back to Windows DVFS. On the other hand, if the machine runs any application that changes a part of the screen, PICSEL will control the frequency. In such a case, the frequency will be reduced if there is no background job but will be kept high if a background job takes CPU resources away from the graphical application. In this way, the application that changes the screen is acting as a ‘‘canary in a coal mine’’ whose performance degradation is readily apparent to PICSEL.

We must note that running background jobs does not cause any problem for PICSEL. In fact, one of our applications targeted in the next section includes a non-interactive background job to prove that our concept is applicable in such cases. If there is a CPU-intensive background job, a reduction in the frequency causes a significant reduction in the APC (even if the interactive application itself is not computationally intensive). Therefore, PICSEL will keep the frequency high. If, on the other hand, the background job is not CPU-intensive, the frequency can be safely reduced, which is exactly the action taken by PICSEL.

4. Evaluation

We now evaluate the cPICSEL and aPICSEL schemes. We compare against the native Windows XP DVFS scheme, displaying reductions in power consumption and temperature. In Section 4.4, we also present user satisfaction results.

Our evaluations are based on user studies, as described in Section 4.1. We trace the user’s activity on the system during the use of the applications and monitor the responses of Windows DVFS, cPICSEL, and aPICSEL. For studies involving PICSEL, the cPICSEL and aPICSEL algorithms are used online to control the

clock frequency in response to APC and APR values. In the rest of this section, we first describe a user study of PICSEL that provides both independent results and traces for later use. Next, we present dynamic CPU power consumption estimates, system power measurements, and temperature measurements.

PICSEL estimates user-perceived performance via APC and APR values and customizes processor frequency to the individual user. This typically leads to significant power savings compared to existing dynamic frequency schemes that rely only on CPU utilization as feedback. The frame buffer readings and the corresponding calculations for measuring user-perceived performance are infrequent, and impose less than 2% computational overhead. We must note that PICSEL performs APC and APR readings during user studies, hence all the results presented for PICSEL (including power and user satisfaction) include this overhead and its potential impact on user satisfaction. A more efficient, GPU-based implementation could be used to further reduce this overhead.

4.1 Experimental Setup

Our experiments were done using an IBM Thinkpad T43p with a 2.13 GHz Pentium M-770 CPU and 1 GB memory running Microsoft Windows XP Professional SP2. The Pentium M uses the second generation of Intel's SpeedStep technology, in which six CPU frequency-voltage operating points are available.

Our base case for comparison, the Windows XP Adaptive scheme, is Microsoft's adaptive DVFS scheme for portables/laptops. Adaptive DVFS uses all of the frequency states in the Intel Speedstep technology. Performance needs are measured from heuristics "such as processor utilization, current battery level, use of processor idle states, and inrush current events" [17]. In our experimental setup, we ran the computer off AC power, the processor was always active, and we ran trials close enough together to prevent hard disk timeout in order to minimize inrush current events. This leaves processor utilization as the main input to the adaptive DVFS, which makes decisions according to the following algorithm. This evaluates either when in the idle loop or after 300 ms have passed since the last evaluation, whichever comes first.

```

IF 150 ms have passed since the last
    frequency state adjustment
AND Performance has increased by 20%
    since the last evaluation
    Increase f by one level within the next 10
    ms
IF 500 ms have passed since the last
    frequency state adjustment
AND Performance has decreased by 30%
    since the last evaluation
AND A decrease of frequency state by
    one operating point will remain
    above 50% of the maximum
    frequency state
    Decrease f by one level within the next 10
    ms

```

In all our studies, we make use of three application tasks, some of which are CPU intensive and some of which frequently block while waiting for user input:

1. Watching a 3D Shockwave animation using the Microsoft Internet Explorer web browser. The animation was stored

locally. Shockwave options were configured so that rendering was done entirely in software on the CPU.

2. Playing the FIFA 2005 Soccer game. FIFA 2005 is a popular sports game. The game was stored locally. There were no constraints on user gameplay.
3. Watching an HD quality movie trailer in Windows Media Player (WMP) while decoding another MPEG movie clip in the background. Both clips were stored locally and decoding was done in software on the CPU.

We conducted a study with twenty users to evaluate PICSEL. We developed a user pool by advertising our studies within Northwestern University. Some participating users were computer science, computer engineering, or electrical engineering students and others were less experienced with computer use. The studies were double-blind and randomized (i.e., the order of schemes during the tests were randomized to eliminate any possible effect of "first-time" execution impact). The studies included intervention by proctors between trials. Each user evaluation lasted about thirty minutes, and consisted of the user doing the following:

- Filling out a questionnaire that asked the user to rate his or her level of experience in the use of PCs, Windows XP, DVD video, 3D animation, and FIFA 2005 from among the following set: "Power User", "Typical User", or "Beginner".
- Listening to an explanation of how to play FIFA 2005 and how to rate his or her satisfaction with each application instance.
- Watching the 3D Shockwave animation three times using cPICSEL, aPICSEL, and Windows DVFS (2 minutes each).
- Playing FIFA 2005 three times using cPICSEL, aPICSEL, and Windows DVFS (3.5 minutes each).
- Watching the movie trailer three times using cPICSEL, aPICSEL, and Windows DVFS (2 minutes each).
- After each application, the users were instructed to assign one of five levels of satisfaction to their experiences with the system performance for each instance of an application. The users were not asked to rank the instances against each other.

4.2 Frequency Results

Figure 3 illustrates the performance of the two algorithms for three applications in our study. Each graph shows the CPU frequency for a randomly selected user as a function of time. Notice that in all the applications both versions of PICSEL reduced processor frequency more than the Windows DVFS policy. The amount of frequency reduction varies across applications. PICSEL is most effective for the 3D animation application. As illustrated in Figure 2, this has the least variation in APC and APR values at lower frequencies. As a result, PICSEL was able to greatly reduce the CPU frequency without affecting user-perceived performance. Similar results were observed for the video application. For the game, we observe less processor throttling. This is also expected as the APC values in Figure 1 degrade very quickly for the game and PICSEL can throttle down the frequency to lower frequency states in few cases. Overall, these results show that PICSEL reduces frequency compared to Windows DVFS while maintaining user satisfaction. In Section 4.4, we also analyze user satisfaction with the default Windows DVFS and PICSEL algorithms and show that the user satisfaction is not adversely affected for any of our target applications.

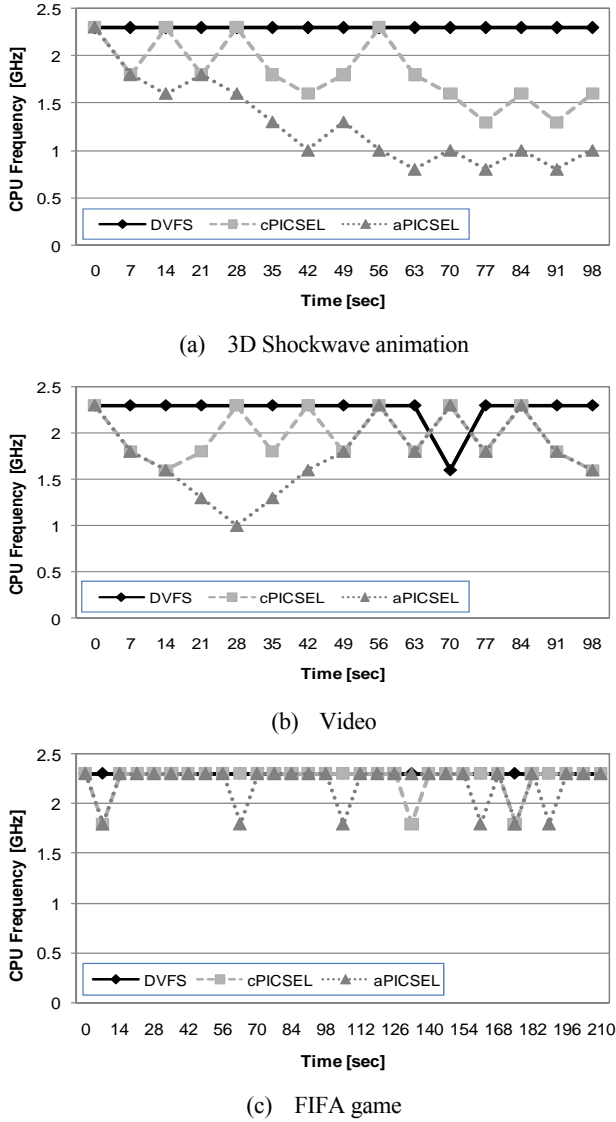


Figure 3. Frequency vs. time for three user trials

4.3 Power Measurements

To analyze the effect of cPICSEL and aPICSEL on the power consumption of the system, we logged the frequency over time during the user studies described in the previous section. We then combine this frequency information with the offline profile to derive power savings for cPICSEL, aPICSEL, and the default Windows XP DVFS policy. In Section 4.3.1 we present the CPU dynamic power savings and in Section 4.3.2 we present the total system power savings. Section 4.3.3 presents the changes in the operating temperatures.

4.3.1 CPU Dynamic Power Reduction

The dynamic power consumption of a processor is directly related to its frequency and supply voltage and can be expressed using the formula $P = V^2CF$, which states that power is equal to the product of voltage squared, capacitance, and frequency. By using the frequency traces and the nominal voltage levels on our target processor [9], we calculated the relative dynamic power consumption. Figure 4 presents the CPU dynamic power reduction

achieved by the PICSEL algorithms (cPICSEL and aPICSEL) for individual users. The rightmost bars correspond to the savings averaged across users.

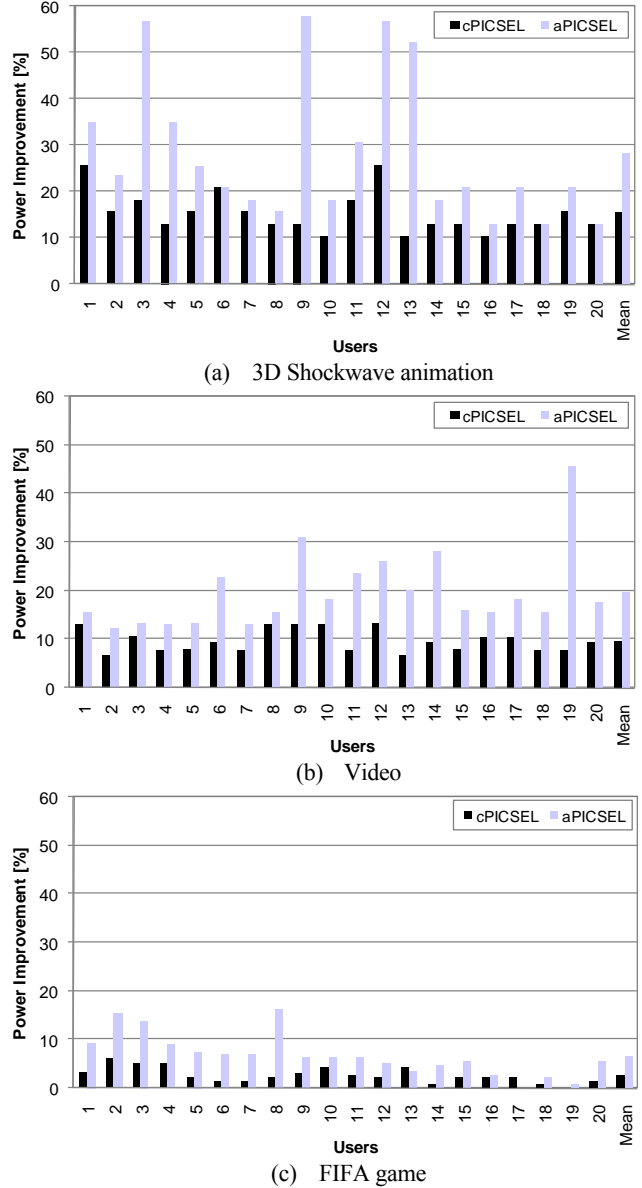



Figure 4. CPU dynamic power reduction with cPICSEL and aPICSEL over Windows DVFS


For the 3D Shockwave animation, we see mixed responses from the users, although on average PICSEL reduces power by 21.8%. On average, cPICSEL and aPICSEL independently reduce the power consumption by 15.3% and 28.2%, respectively. aPICSEL performs better as it allows a larger threshold for APC values over each interval. The results show a considerable variation among different users. This can be explained by the fact that the control agent for APC calculation considers a sampling window of roughly 64x51 pixels at the center of the display window. The relative position of the shockwave player while the user watches the 3D animation plays a role in the calculation of APC and APR. It subsequently affects the decision taken by the PICSEL algorithm. Nevertheless,

as we will show in Section 4.4, such variations do not have an impact on user satisfaction.


For the Video application, cPICSEL and aPICSEL reduce power consumption by averages of 9.6% and 19.7%, respectively. This suggests that the Video application is less conducive to frequency throttling than the Shockwave application. User 19 is the only



(a) 3D Shockwave animation



(b) Video



(c) FIFA game

Figure 5. System power reduction with cPICSEL and aPICSEL over Windows DVFS

exception where aPICSEL results in a power savings of 45.8%, greater than those for the Shockwave application. For the FIFA game, the average power improvements of 2.6% for cPICSEL and 6.7% for aPICSEL were lower than Video and Shockwave applications, suggesting that the FIFA game was the least conducive to frequency throttling. Note that PICSEL does not reduce the frequency for all the users while they play the FIFA game. For example, cPICSEL does not reduce the frequency for user 19. Similarly, aPICSEL does not reduce the frequency for user 17. This is understandable since the game application has the most steeply sloped APC curve (Figure 1), meaning a change in frequency will have a larger effect on the game’s displayed output than on the other applications’ displayed output.

For all three applications, we see that in all cases cPICSEL and aPICSEL lead to power savings compared to Windows DVFS. On average, aPICSEL reduces the dynamic power consumption by 18.2% for all three applications. cPICSEL results in a 9.1% power reduction aggregated over three applications and 20 users.


4.3.2 System power measurement

To further measure the impact of our techniques, we replayed the traces from the user studies described in Section 4.3.1 on the laptop. The laptop was connected to a National Instruments 6034E data acquisition board attached to the PCI bus of a host workstation, which permitted us to measure the power consumption of the entire laptop (including other power consuming components such as memory, screen, hard disk, etc.). The sampling rate was set to 10 Hz. Each of the user studies was replayed five times to average out any variation across trials. Figure 5 presents the system-level power savings of cPICSEL and aPICSEL relative to Windows DVFS. In general, the reduction in system-level power consumption is similar to the estimated processor dynamic power savings. cPICSEL and aPICSEL reduce power consumption by 16.8% and 25.7% on average for the 3D Shockwave animation, by 8.0% and 14.5% on average for the Video application, and by 2.6% and 6.2% on average for the FIFA game, respectively. On average, aPICSEL reduces system-level power consumption by 12.1%, aggregated over 20 users and three applications. cPICSEL reduces the system-level power consumption by 7.1%.


We must note that the dynamic CPU power savings presented in the previous section and the system-level power savings presented in this section cannot be directly compared because the previous section reports the dynamic power consumption of the CPU. This section, on the other hand, reports the measured power consumption of the laptop (which includes leakage power of the CPU as well as all the power consumption of other components in the laptop including memory, screen, hard disk, etc.). However, some conclusions can be drawn from the data in both sections. Applications that result in high CPU dynamic power consumption tend to also observe high system power savings. Clearly, part of the system power reduction comes from the decrease in the CPU dynamic power consumption. Leakage is also reduced due to the decrease in voltage and the decrease in temperature resulting from reduced dynamic power consumption.

4.3.3 Changes in Peak Temperature

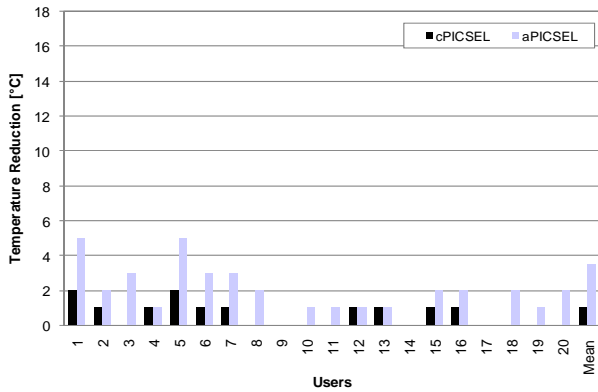
We used CPUCool [18], a Windows-based tool that logs temperatures at processor cores, to measure CPU temperature in the system. Figure 6 shows the reductions in peak temperatures of the system when using the cPICSEL and aPICSEL schemes. In all cases, the cPICSEL and aPICSEL schemes lower the temperature compared to the Windows native DVFS scheme due to the power reductions we have reported in the previous sections. The maximum temperature reduction of 16°C is seen in the case of the aPICSEL scheme used for the Shockwave application. On average, for all three applications, cPICSEL and aPICSEL reduce the peak temperature of the system by 1.7°C and 4.3°C, respectively, aggregated over all 20 users.



(a) 3D Shockwave animation



(b) Video




(c) FIFA game

Figure 6. Peak temperature reduction


4.4 User Satisfaction

We now discuss the satisfaction levels with the Windows DVFS and PICSEL algorithms for three applications as reported by individual users. During the user study, each participant was asked to give a satisfaction level from 1 to 5 (5 being the most satisfactory performance) for each application. Figure 7 illustrates the ranks awarded by each user. Compared to Windows DVFS, cPICSEL results in slightly better satisfaction levels for all three applications aggregated over 20 users. The student t-test analysis of the results reveals that the difference is not due to chance with 90% confidence. aPICSEL and Windows DVFS provide the same satisfaction (a student t-test analysis identifies the two means to be identical with over 99% confidence). On average, aPICSEL is rated highest for the game application (3.8) where it results in the least power reduction. For the Shockwave application, maximum


power reduction for the aPICSEL scheme caused it to have the lowest average user satisfaction score (3.5).



(a) 3D Shockwave animation



(b) Video



(c) FIFA game

Figure 7. User ranking distribution

We noticed cPICSEL was ranked higher than Windows DVFS although Windows DVFS runs the system at higher frequencies. The only time Windows DVFS will throttle the frequency below what CPU utilization would prescribe is in the case of the temperature crossing a thermal trip point [17], and this led to the hypothesis that user dissatisfaction caused by thermal emergencies was the main reason for the decreased user satisfaction with Windows DVFS. We ran an experiment in which FIFA 2005 was played under Windows DVFS until the user observed several distinct processor frequency reductions triggered by thermal emergencies. The results of this experiment are shown in Figure 8.

This figure shows processor temperature and frequency when FIFA 2005 is played until it triggers a thermal emergency (about 16 minutes after starting the game). At that point, the frequency is reduced to the lowest value. This causes a perceivable slowdown in game play and lower instruction throughput. Windows DVFS continues to operate even though it has been over-ridden by the processor, and lowers its frequency to match the lower instruction throughput. Soon after the processor returns frequency control to Windows DVFS, the frequency is again set to the highest available frequency on the processor. This causes the temperature to rise again quickly, leading to consecutive emergencies.

Both cPICSEL and aPICSEL reduced the occurrence of thermal emergencies. As a result, for processor-intensive applications, PICSEL may deliver better user-perceived performance by reducing the probability of thermal emergencies. The satisfaction results also support this claim: aPICSEL provides the highest satisfaction for the game on average, because for this highly compute-intensive application, aPICSEL allows the highest reduction in temperature, and resulting thermal emergencies.

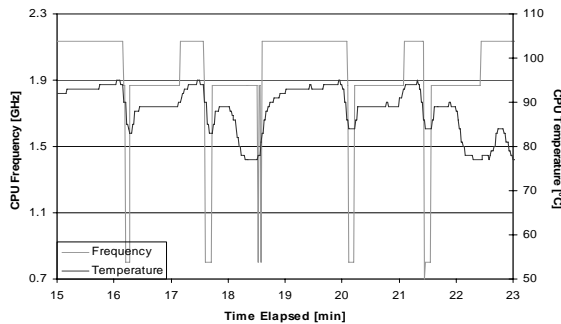


Figure 8. Thermal emergency under Windows DVFS

4.5 Related Work

Dynamic voltage and frequency scaling (DVFS) is an effective technique for microprocessor energy and power control for most modern processors [9], [1]. Energy efficiency has been a major concern for mobile computers. Gurun and Krintz [13] have proposed a new model for estimating energy consumption using hardware and software counters. Fei, Zhong, and Jha [6] proposed an energy aware dynamic software management framework that improves battery utilization for mobile computers. However, this technique is only applicable to highly-adaptive mobile applications. Researchers have proposed algorithms based on workload decomposition [2], but these tend to provide power improvements only for memory-bound applications. Wu et al. [22] presented a design framework of a run-time DVFS optimizer in a general dynamic compilation system. The Razor [5] architecture dynamically finds the minimal reliable voltage level. Dhar, Maksimovic, and Kranzen [4] proposed adaptive voltage scaling that uses a closed-loop controller targeted towards standard-cell ASICs. Intel Foxton technology [20] provides a mechanism for certain Intel Itanium 2 processors to adjust core frequency during operation to boost application performance. However, unlike PICSEL it does not perform any dynamic voltage setting. To the best of our knowledge, no previous DVFS techniques consider user-perceived performance.

Other DVFS algorithms use task information, such as response times in interactive applications [15] and [24] as a proxy for the user. Vertigo [7] monitors application messages and could be used to perform the optimizations implemented in our study. However, compared to Vertigo, our approach uses a much easier metric/framework. Xu, Moss, and Melhem proposed novel schemes [23] to minimize energy consumption in certain real-time embedded systems. However, they try to adapt to the variability of the workload rather than to the users. Gupta, Lin, and Dinda [11] studied user satisfaction with resource borrowing and noted a high variation in user tolerance for any given level of system resources in desktop computing applications. Lin and Dinda [14] developed a CPU scheduling system that used direct user feedback to exploit this variation. Mallik, Lin, Memik, Dinda, and Dick [16] showed that this variation also exists for power management, and presented a successful power management approach based on direct user feedback.

Ranganathan, Geelhoed, Manahan, and Nicholas [19] explored using OS-level knowledge about screen content to reduce the power consumption of the screen itself, however no work has been done using knowledge of screen content to control the voltage and frequency of a processor. Gurun and Krintz [12] looked at OS-level knowledge of user-generated events to control a DVFS scheme but did not use knowledge of screen content. Our work combines these two approaches and uses detailed screen information to control the CPU's voltage and frequency levels.

A study of user perception of audio/video quality found that the loss of video frames would decrease user satisfaction [21]. Frame rate also has a significant effect on user satisfaction, with satisfaction increasing logarithmically with the number of frames displayed per second [3]. Finally, Gulliver and Ghinea found that both video delay and jitter cause a significant reduction in users' perception of the quality of a video [10]. However, none of these results were used to control processor resources.

5. Conclusion

Any architectural optimization ultimately aims to satisfy the user. Its success or failure rests on the accuracy of its performance metrics as proxies for user satisfaction. In this work, we argue that rather than using metrics that are "close to metal", architectures should optimize for metrics that are "close to flesh". To evaluate such an approach, we have developed a new power management technique: **PICSEL (Perception-Informed CPU performance Scaling to Extend battery Life)**. This technique reduces CPU power consumption in comparison with existing DVFS techniques. Extensive user studies show that we can reduce system-level power consumption of our target laptop on average by 7.1% for a conservative approach (cPICSEL) and 12.1% for the aggressive version (aPICSEL) compared to the Windows XP DVFS scheme. Furthermore, CPU temperatures can be markedly decreased through the use of our techniques. User studies also revealed that the difference in overall user satisfaction between the more aggressive version of PICSEL and Windows DVFS were statistically insignificant, whereas the conservative version of PICSEL improved the users' overall satisfaction when compared to Windows DVFS.

Acknowledgements

This work is in part supported by DOE Awards DE-FG02-05ER25691 and DE-AC05-00OR22725 (via ORNL), NSF Awards CNS-0720691, CNS-0721978, CNS-0715612, IIS-0613568, CNS-0551639, CNS-0347941, CCF-0541337, IIS-0536994, CCF-0444405, ANI-0093221, ANI-0301108, and EIA-0224449, by SRC award 2007-HJ-1593, by Wissner-Slivka Chair funds, and by gifts from Symantec, Dell, and VMware. The authors would like to thank the anonymous reviewers for their helpful comments and suggestions.

References

- [1] Brock, B. and Rajamani, K. 2003. *Dynamic Power Management for Embedded Systems*. In Proc. of the IEEE SOC Conf. (SOC'03).
- [2] Choi, K., Soma, R., and Pedram, M. 2004. *Dynamic Voltage and Frequency Scaling based on Workload Decomposition*. In Proc. of the 2004 Int. Symp. on Low Power Electronics and Design (ISPLED'04), 174-179.
- [3] Claypool, M., Claypool, K., and Damaa, F. 2006. *The Effects of Frame Rate and Resolution on Users Playing First-person Shooter Games*. In Proc. of ACM/SPIE Multimedia Computing and Networking (MMCN'06).
- [4] Dhar, S., Maksimovic, D., and Kranzen, B. 2002. *Closed-Loop Adaptive Voltage Scaling Controller for Standard Cell ASICs*. In Proc. of the 2005 Int. Symp. on Low Power Electronics and Design (ISPLED'05.), 103-107.
- [5] Ernst, D., Kim, N. S., Das, S., Pant, S., Rao, R., Pham, T., Ziesler, C., Blaauw, D., Austin, T., Flautner, K., and Mudge, T. 2003. *Razor: A Low-Power Pipeline Based on Circuit-Level Timing Speculation*. In Proc. of the 36th ACM/IEEE Int. Symp. on Microarchitecture (MICRO-36), 7-18.
- [6] Fei, Y., Zhong, L., and Jha, N. K. 2004. *An Energy-aware Framework for Coordinated Dynamic Software Management in Mobile Computers*. In Proc. of the IEEE Computer Society's Int. Symp. on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems (MASCOTS'04), 306-317.
- [7] Flautner, K. and Mudge, T. 2002. *Vertigo: Automatic Performance-Setting for Linux*. ACM SIGOPS Operating Systems Review 36, SI (Winter 2002), 105-116.
- [8] Ghinea, G. and Thomas, J. P. 2005. *Quality of Perception: User Quality of Service in Multimedia Presentations*. IEEE T. Multimedia 7, 4 (Aug. 2005), 786-789.
- [9] Gochman, S., Ronen, R., Anati, I., Berkovits, A., Kurts, T., Naveh, A., Saeed, A., Sperber, Z., and Valentine, R. C. 2003. *The Intel Pentium M Processor: Microarchitecture and Performance*. Intel Technology J. 7, 2 (May 2003), 21-36.
- [10] Gulliver, S.R. and Ghinea, G. 2007. *The Perceptual and Attentive Impact of Delay and Jitter in Multimedia Delivery*. IEEE T. Broadcast 53, 2 (June 2007), 449-458.
- [11] Gupta, A., Lin, B., and Dinda, P. A. 2004. *Measuring and Understanding User Comfort with Resource Borrowing*. In Proc. of the 13th IEEE Int. Symp. on High Performance Distributed Computing (HPDC'04), 214-224.
- [12] Gurun, S. and Krintz, C. 2005. *AutoDVS: an Automatic, General-purpose, Dynamic Clock Scheduling System for Hand-held Devices*. In Proc. of the 5th ACM Int. Conf. on Embedded Software (EMSOFT'05), 218-226.
- [13] Gurun, S. and Krintz, C. 2006. *A Run-Time, Feedback-Based Energy Estimation Model for Embedded Devices*. In Proc. of the Int. Conf. on Hardware/Software Codesign and System Synthesis. (CODES+ISSS'06).
- [14] Lin, B. and Dinda, P. A. 2006. *Towards Scheduling Virtual Machines Based on Direct User Input*. In Proc. of the 1st Int. Workshop on Virtualization Technology in Distributed Computing (Tampa, FL, USA, November 17, 2006). VTDC'06. See also technical report NWU-EECS-06-07, Northwestern University, EECS.
- [15] Lorch, J. and Smith, A. 2003. *Using User Interface Event Information in Dynamic Voltage Scaling Algorithms*. In Proc. of the IEEE Computer Society's Int. Symp. on Modeling, Analysis, and Simulation of Computer and Telecommunications Systems (MASCOTS'03), 46-55.
- [16] Mallik, A., Lin, B., Memik, G., Dinda, P. A., and Dick, R. P. 2006. *User-Driven Frequency Scaling*. IEEE Computer Architecture Letters 5, 2 (July 2006), 16. A summary of this work also appeared in ACM SIGMETRICS 2007.
- [17] Microsoft Corporation. 2003. *Windows Native Processor Performance Control*. Windows Platform Design Notes (May 2003). Retrieved from <http://www.microsoft.com/whdc/system/pnppwr/powermgmt/rocPerfCtrl.msp>.
- [18] Podien, W. CPUCool. Retrieved from <http://www.cpu-cool.de/index.html>.
- [19] Ranganathan, P., Geelhoed, E., Manahan, M., and Nicholas, K. 2006. *Energy-Aware User Interfaces and Energy-Adaptive Displays*. Computer 39, 3 (March 2006), 31-38.
- [20] Wei, J. *Foxton Technology Pushes Processor Frequency, Application Performance*. Technology@Intel Mag. (July 2007). Retrieved from <http://www.intel.com/technology/magazine/computing/foxton-technology-0905.htm>.
- [21] Wijesekera, D., Srivastava, J., Nerode, A., Forrsti, M. 1999. *Experimental Evaluation of Loss Perception in Continuous Media*. Multimedia Systems 7, 6 (Nov. 1999), 486-499.
- [22] Wu, Q., Martonosi, M., Clark, D. W., Reddi, V. J., Connors, D., Wu, Y., Lee, J., and Brooks, D. 2005. *Dynamic Compilation Framework for Controlling Microprocessor Energy and Performance*. In Proc. of the 38th IEEE/ACM Int. Symp. on Microarchitecture (MICRO-38), 271-282.
- [23] Xu, R., Moss, D., and Melhem, R. 2005. *Minimizing Expected Energy in Real-time Embedded Systems*. In Proc. of the 5th ACM Int. Conf. on Embedded Software (EMSOFT'05), 251-254.
- [24] Yan, L., Zhong, L., and Jha, N. K. 2005. *User-perceived Latency-based Dynamic Voltage Scaling for Interactive Applications*. In Proc. of ACM/IEEE Design Automation Conf. (DAC'05), 624-627.