

Power and Performance Optimization: A Case Study with the Pentium M Processor

Heather Hanson* Stephen W. Keckler[‡]

Computer Architecture and Technology Laboratory
cart@cs.utexas.edu - www.cs.utexas.edu/users/cart

*Department of Electrical and Computer Engineering

[‡]Department of Computer Sciences
The University of Texas at Austin

IBM Technical Contact : Ron Kalla, IBM Systems Group

Abstract

*Power and thermal constraints limit computer system performance, and although current and next-generation commercial computing systems have access to several power management mechanisms, the management techniques designed to ensure safe operation often degrade performance. We conducted a study of power management techniques during a summer internship at IBM's Austin Research Laboratory as part of a larger project to investigate application-specific power/performance trade-offs. This report describes one component of the study, a characterization of two power management techniques on a Pentium M system, dynamic voltage and frequency scaling (DVFS) and clock throttling. We analyzed microbenchmark power and performance for three dataset sizes corresponding to core-bound, intermediate, and memory-bound programs. We found that DVFS is more effective than clock throttling at reducing power while preserving performance, and that the two techniques applied together offer a range of power management options to suit application behavior and system requirements.*¹

1 Introduction

The increasing capacity and density of components in computer systems have reached a point where power and thermal constraints are barriers to further growth of computing performance. Leakage current causes a baseline level of static power, independent of program execution. However,

¹This report summarizes a study conducted during an internship at the IBM Austin Research Laboratory during the summer of 2005. Karthick Rajamani was the intern mentor; the project team includes Juan Rubio, Soraya Ghiasi, Freeman Rawson, and manager Tom Keller.

the switching activity that causes dynamic power consumption depends on the workload characteristics. Power management techniques such as clock throttling and dynamic voltage and frequency scaling (DVFS) alter power and performance levels. Application characteristics influence the effectiveness of the management techniques. For example, reducing the voltage and frequency degrades performance, but the extent of degradation is worse for compute-bound applications than for memory-bound applications. This study evaluates power management techniques individually and in combination for use in commercial server systems. This report describes one component of the study characterizing the Pentium M DVFS and clock throttling power-management options, which was performed during a summer internship in 2005. A forthcoming IBM technical report will describe the full study in detail. The project team has continued the study and is in the process of submitting a conference paper on recent work.

The long-term goal of the project is to use information from this study in a real-time power management control scenario. The controller will evaluate the estimated effects of changing power management settings and make an informed decision for when and how to change settings.

Section 2 describes the power management techniques in this study and Section 3 describes the measurement infrastructure and application suite. Section 4 presents the results of the power-management characterization on the Pentium M system. Section 5 describes an example of power management with a fixed power cap and Section 6 concludes the report.

2 Power Management Techniques

We chose to study the Pentium M processor because it is readily available and offers two power management op-

Supply Voltage (Volts)	Frequency (MHz)
0.988	600
1.052	800
1.100	1000
1.148	1200
1.196	1400
1.244	1600
1.292	1800
1.34	2000

Table 1. Frequency and Voltage Pairs

tions, clock throttling and frequency-voltage scaling, that are controlled independently and may be used together in any combination of settings. We expect that the information gathered from a Pentium M system would provide insight useful for designing dynamic power management for other commercial computing systems, as well.

2.1 Frequency and Voltage Scaling

Table 1 lists the 8 frequency-voltage pairs used in this study, which span the full frequency range of 600 MHz to 2.0 GHz and the corresponding supply voltage range of 0.988 V to 1.34 V, as recommended in the Pentium M documentation [1]. During frequency transitions, the processor stalls for up to 10 μ s.

2.2 Clock Throttling

Clock throttling gates the main clock with a throttling signal to form *run* and *hold* regions in the clock signal. When the clock is enabled, the clock signal runs freely with a standard clock period; when it is throttled, the clock is held at a zero-voltage level. Figure 1 [3] illustrates a simplified view of the run and hold window. Eight clock throttling levels indicate the fraction of running time in increments of 1/8ths. For example, throttling level 8 is unthrottled (running 8/8 of the time) and throttling level 3 runs 3/8 of the time. Note that the clock is not throttled 3 of each 8 cycles, but rather runs continuously for 3/8ths of the *run + hold* window and then is held idle for 5/8ths of the window. The *run + hold* window size is approximately 3 μ s.

3 Infrastructure

3.1 Processor and system board

The experiments in this report were conducted on a Pentium M-based system running Windows XP. The 90 nm

Pentium M processor “Dothan” has a 32 KB primary instruction cache, 32 KB primary data cache, and a 2MB, 8-way unified secondary cache [1]. The processor chip is paired with an Intel 855GME chipset and 512 MB of DDR SDRAM memory on a Radisys uniprocessor motherboard [2]. In our experimental setup, the motherboard resides in a modified tower enclosure, lying flat with the top panel removed to allow access for probe cables. The standard heat sink and fan, combined with the open-air configuration in a cool room allows the processor to operate throughout its full range of frequency, voltage, and throttling settings without tripping the temperature sensor that would intervene and change settings automatically.

3.2 Power Measurement

We added a sense resistor between each voltage regulator module and the processor and placed data acquisition probes to monitor processor supply voltage and current levels. We collect and analyze power data on a separate computer to avoid interference with workloads executing on the system under test. A National Instruments data acquisition system samples current and voltage values and interfaces with a Pentium III system that executes a custom program in LabView software to capture the data in a trace file.

3.3 Workloads

We characterized application behavior with a suite of microbenchmarks, *MS-Loops*. The suite contains several tests, each one a kernel of code that repeats multiple times. The application behavior for each test is monophasic, which allows a single power measurement and a single performance measurement to describe each benchmark for analysis. Table 2 describes each test in the MS-Loop suite.

4 Characterization

This section presents the results of application characterization on the Pentium M. First, we illustrate the power and performance impact of each technique independently, then characterize the system with both techniques applied simultaneously.

4.1 Voltage and Frequency Scaling

Figure 2 shows the effect of scaling voltage and frequency on power consumption and performance for the 4 KB (L1-resident), 128 KB (L2-resident), and 4 MB (main memory resident) data footprints for unthrottled operation. Power consumption ranges from about 2 watts to nearly 18 watts throughout the range of footprints and

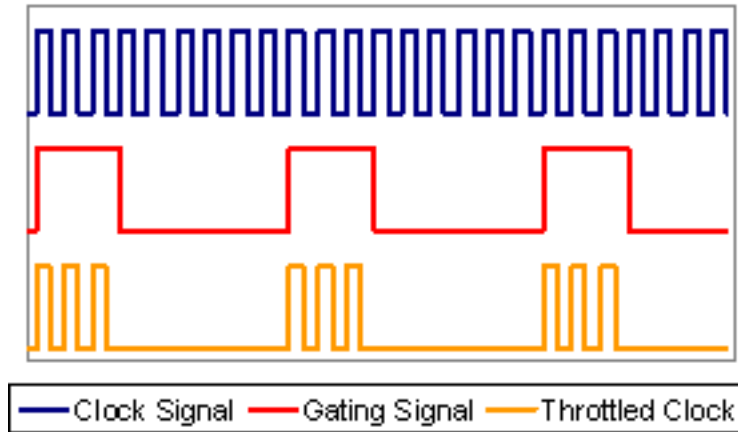


Figure 1. Pentium M clock throttling (not to scale)

MS LOOP Test	Description
DAXPY	Double-precision calculation of $aX + Y$
FMA	Floating point multiply and accumulate
MCOPYY	Copy arrays from one memory location to another
MLOAD-RAND	Random memory accesses

Table 2. Microbenchmarks: MS-Loops

voltage-frequency settings. Benchmarks with L1- and L2-resident footprints display a quadratic relationship between the frequency-voltage pair and power, with more variation among applications at high frequencies. The memory-bound 4 MB footprint also reflects the quadratic relationship, with a wider range of application behavior throughout the frequency spectrum than the core-bound and intermediate footprints, with increasing spread in power consumption at higher frequencies. In all footprints, the MLOAD-RAND test exhibits lower power consumption than the other tests, in part because its random behavior does not benefit from pre-fetching and consequently spends more time waiting for data. In a system such as Pentium M with aggressive clock-gating for idle components, the longer stall time translates to lower power consumption.

Charts in Figure 2 display performance in terms of inverse program execution time, normalized to the maximum performance (minimum execution time) for each benchmark/footprint. The data are normalized for a fair comparison of the microbenchmarks, which vary in program length.

The core-bound and intermediate footprints show very little application-specific behavior variation throughout the DVFS spectrum and normalized performance is linear with frequency. The memory-bound footprint, however, is greatly influenced by frequency settings. As the core frequency is reduced, the relative speed of the core with re-

spect to memory slows and the core waits fewer cycles for data from memory, effectively reducing the performance penalty of lower frequencies. The extent of degradation is application-specific, with the FMA test incurring the largest degradation of the test suite.

4.2 Clock Throttling

Figure 3 shows the effect of clock throttling on power consumption and performance for microbenchmarks executing in each memory footprint at 2 GHz. The experimental data indicate that power consumption for throttle level 4 is approximately the same as throttle level 5. An examination of the power and performance data suggested that the implementation of throttle level 4 on this system does not conform to the expected throttling behavior for a 50% duty cycle. Throttling level 5 is also somewhat more throttled than expected for a 5/8 duty cycle.

Aside from the mid-range throttling abnormalities, power consumption reflects an approximately linear relationship with clock throttling level. Clock-throttling power trends are similar to voltage-frequency scaling power trends. MLOAD-RAND consumes less power than the other tests at all clock throttling settings and the power data exhibit more variation under less-throttled conditions and with larger footprints. However, extensive clock throttling

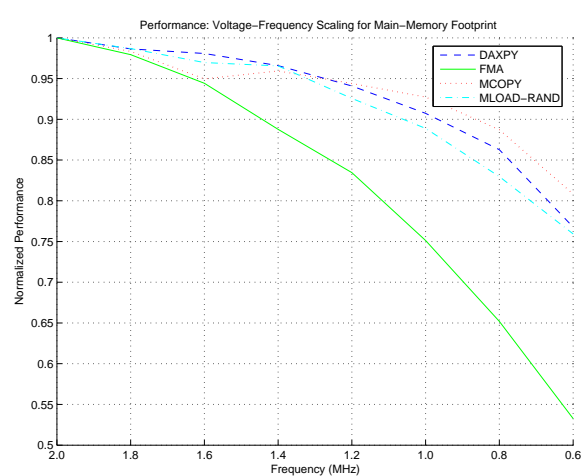
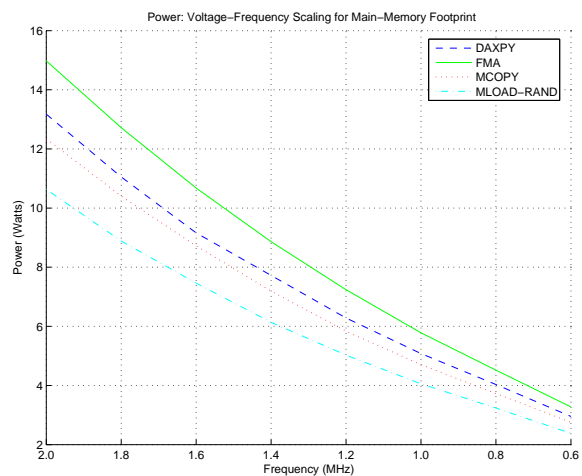
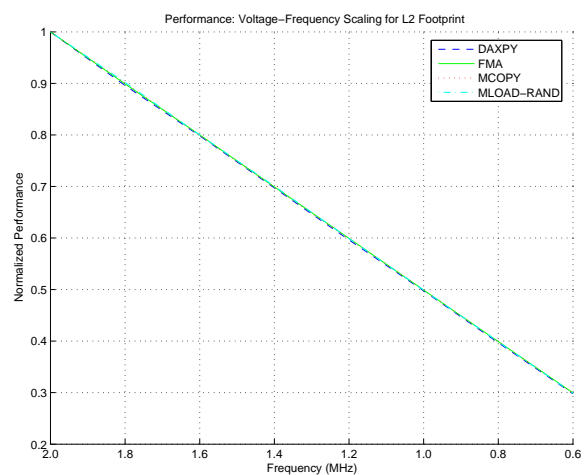
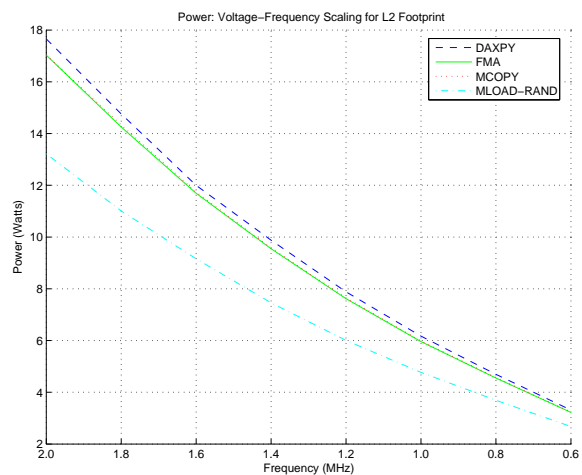
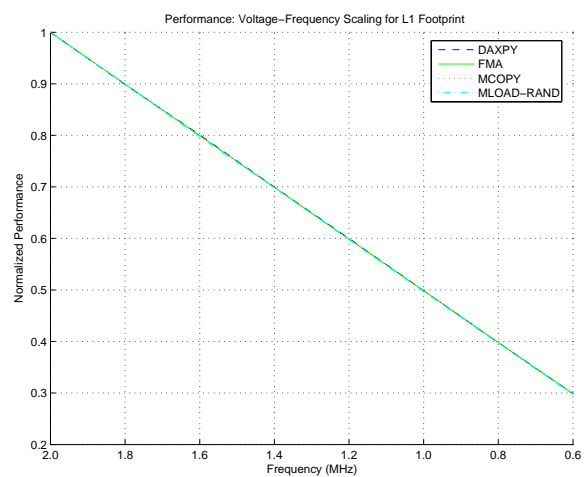



Figure 2. Effects of Voltage-Frequency Scaling on Power and Performance

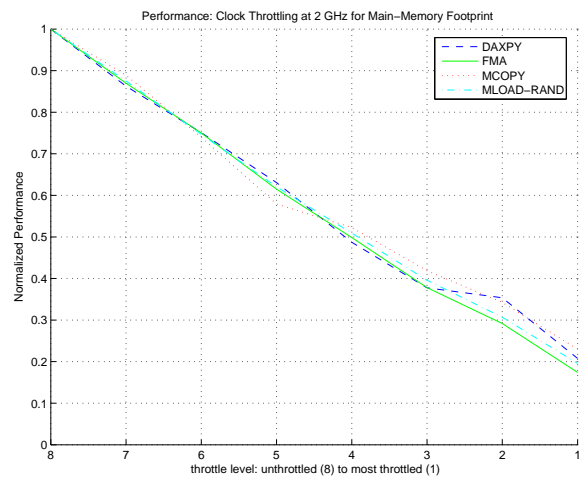
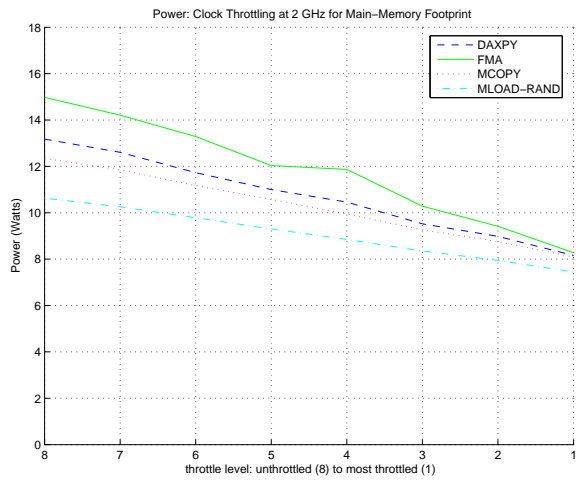
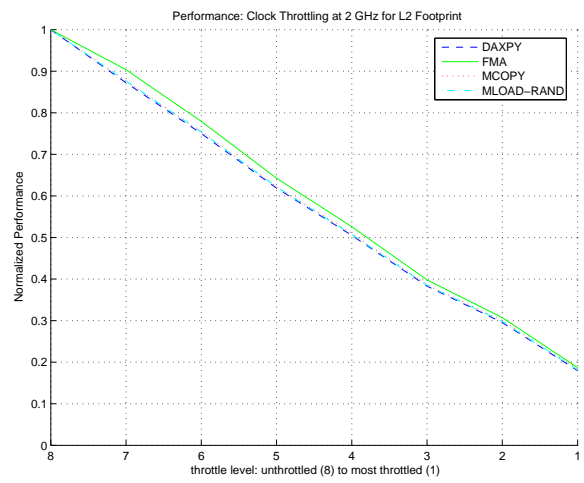
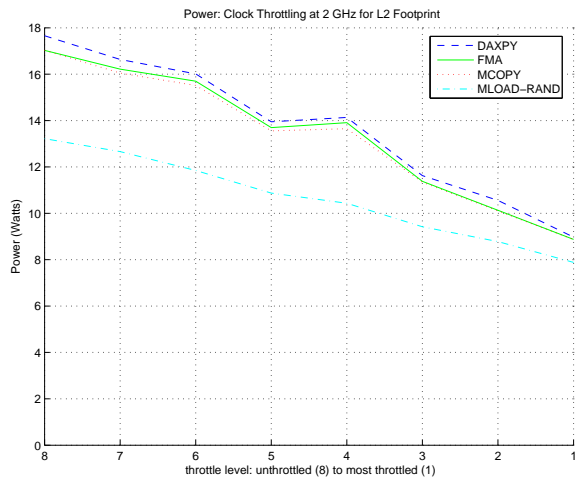
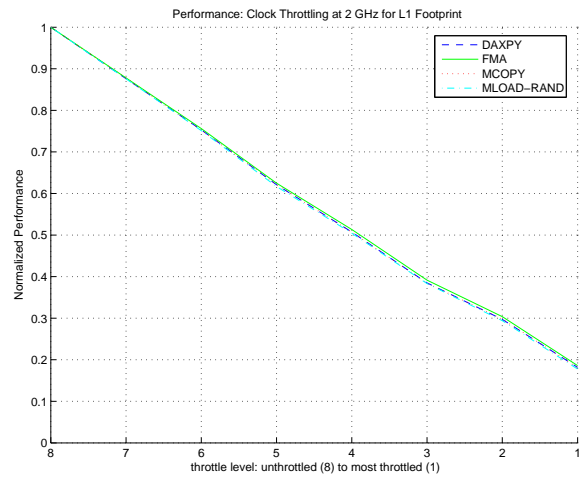
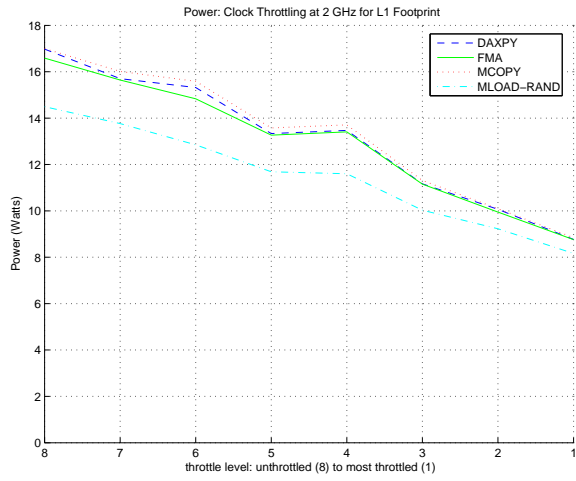


Figure 3. Effects of Clock Throttling on Power and Performance

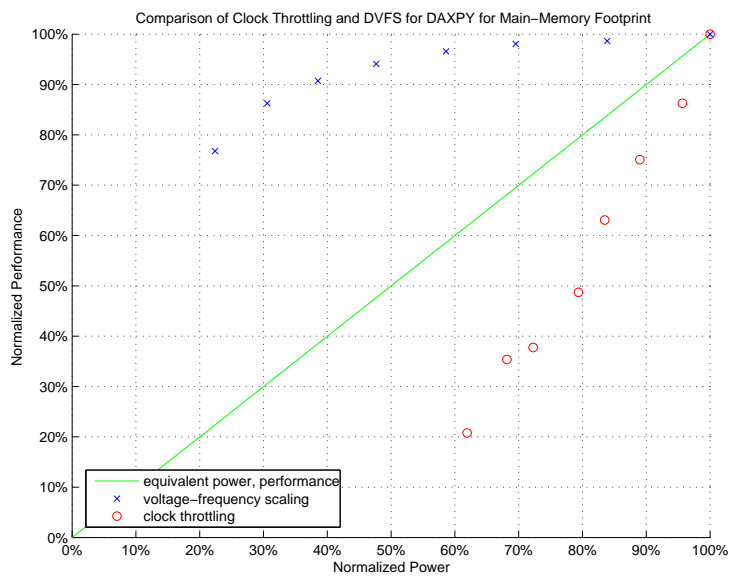

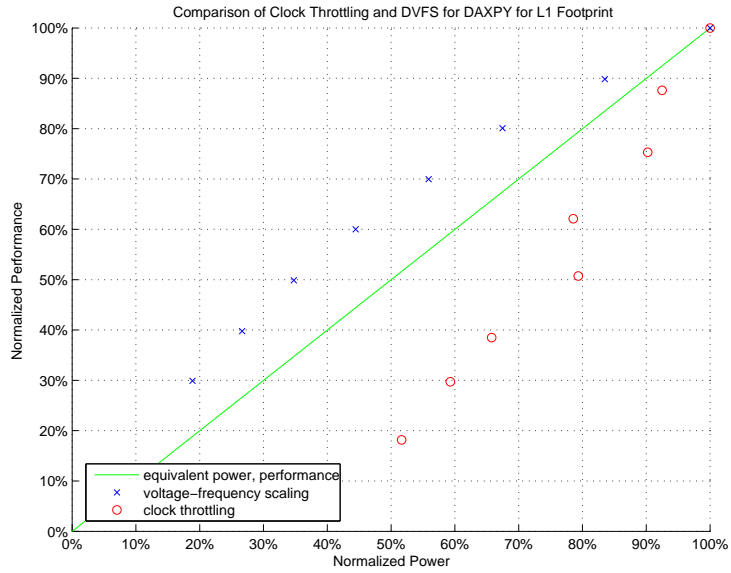


Figure 4. Comparison of Clock Throttling and Voltage-Frequency Scaling

does not reduce power to the same degree as low frequency-voltage settings. Power consumption ranges from about 7.5 watts at the maximum throttling to about 18 watts unthrottled.

Clock throttling with large *run+hold* windows does not alter the relative speed of memory with respect to the core frequency; therefore, performance trends for clock throttling are similar for tests of all footprint sizes.

4.3 Technique Comparison

Figure 4 shows normalized data for the DAXPY microbenchmark executing with 3 footprint sizes: L1, L2, and main memory with each power-management technique applied separately. The diagonal line indicates the break-even point where performance and variable power are affected equally by a technique. Circle markers indicate measurements at each of the 8 clock throttling levels and the x markers show data for each of the 8 frequency-voltage pairs.

Clock throttling results in greater performance degradation per power reduction, with data points below the break-even point in the graphs. The best power reduction for DAXPY with a main-memory resident data footprint is approximately 60% of the maximum power at a cost of degrading performance to only 20% of maximum performance. The clock throttling data points exhibit similar power though different performance for throttling levels 4 and 5, as previously discussed.

Voltage-frequency scaling is able to reduce power with less performance degradation than clock throttling, especially for the memory-bound workloads, which benefit from lower frequency in both reduced power and also fewer core clock cycles stalled for memory accesses.

4.4 Combined Techniques

Although Figure 4 shows that DVFS provides more effective power-performance management than clock throttling, it can incur a relatively long stall while adjusting to new frequency and voltage settings. The Pentium M stall time is approximately $10\mu s$ for any change in frequency-voltage pairs. Switching clock throttling levels does not incur a stall, so it can be used for quick emergency response or to tailor power management settings to rapidly changing application behavior.

Clock throttling in combination with DVFS can provide a range of power-management options. For each test in the microbenchmark suite, we measured performance and power at each of the 64 combinations of DVFS and clock throttling levels. Graphs in Figure 5 display the trends for performance and power for the DAXPY microbenchmark test throughout the setting space. The grid intersection points indicate combinations of DVFS and clock throttling

combinations, such as 1 GHz and throttle level 6 or 600 MHz and throttle level 2. Each isoline is labeled as the percentage of maximum power or performance that it traces.

The isolines demonstrate that a performance or power target can be achieved with multiple settings, and that the settings' effect varies with working set size. For example, the compute-bound DAXPY test reaches approximately 50% of maximum performance for equivalent settings of clock throttle level 4 and a frequency of 2 GHz, or unthrottled 1 GHz. For the memory-bound DAXPY test, the 2 GHz/throttle level 4 combination is 50% of maximum performance but the unthrottled 1 GHz setting achieves 90% of maximum performance.

While improving performance causes greater power consumption for compute-bound applications, memory-bound applications can benefit from substantial power reduction at lower frequencies and voltages. For the same 50% performance, the memory-bound application at 2 GHz and throttle level of 4 requires approximately 80% of maximum power, while an equivalent performance setting of 800 MHz and throttle level 5 reduces power to approximately 25% of maximum power.

The combination of DVFS and clock throttling provides a power manager with options to suit the program behavior and system constraints. In the compute-bound case, the 800 MHz/throttle 5 choice is more power-efficient than 2 GHz/throttle 4, but due to the overhead to change frequency-voltage settings, a quick change to throttle level 4 preserving the 2 GHz frequency may be a better choice for a short power-saving excursion from unthrottled, maximum frequency operation.

5 Power Limit

To illustrate the power-performance tradeoffs for DVFS and clock throttling, we applied a fixed power limit for the system. In this example, the limit is 6 watts. We created a MATLAB script that evaluated the power consumption for each of the 64 combinations of clock throttling and DVFS settings, eliminated choices that would exceed 6 watts, and plotted performance for the remaining settings. Figure 6 illustrates the performance of settings that meet the 6-watt limit for the L1-resident DAXPY test. The graph shows clock throttling and frequency scaling settings on the x and y axes and resulting performance levels on the z (vertical) axis. Legal settings within the power budget form a 3-dimensional sheet of performance values. The surface shows the impact on performance for applying clock throttling and DVFS settings. Better performance values are higher in the z axis; lower-performing settings are closer to the z origin. In this example, the best performance within the 6-watt limit would be achieved with unthrottled conditions and a core frequency of 1 GHz, the highest point on the

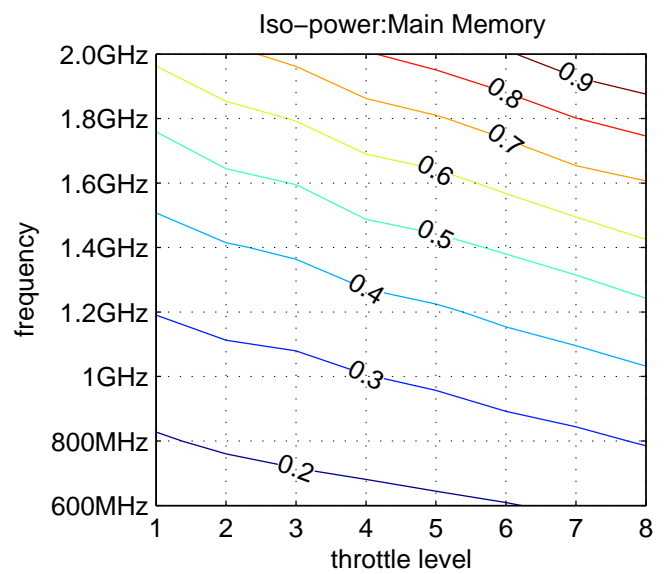
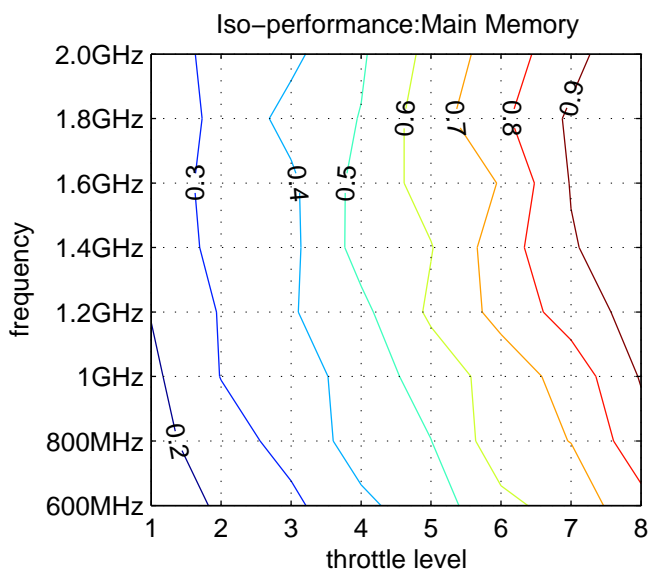

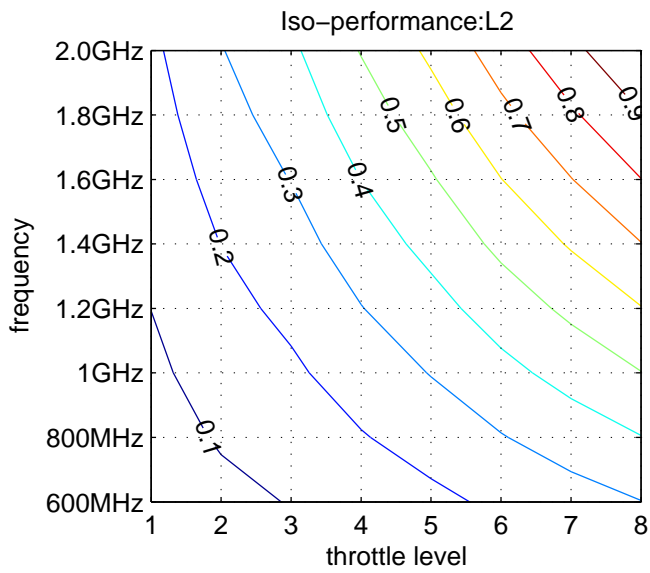
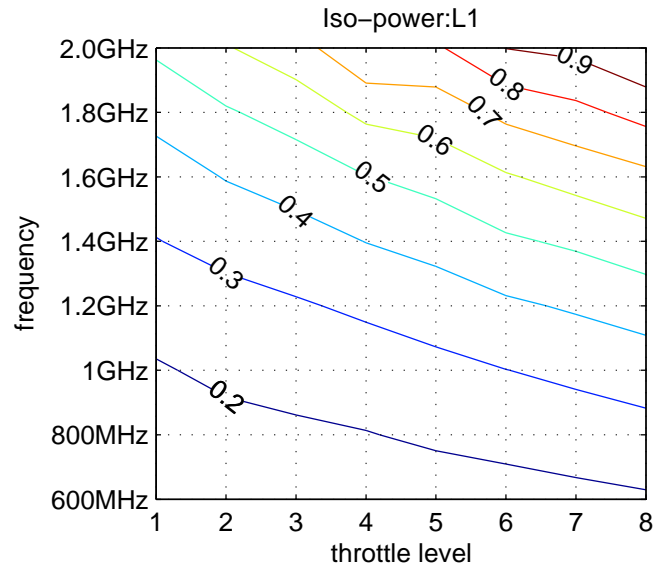



Figure 5. Microbenchmark Performance and Power Isolines




Figure 6. DAXPY L1 Footprint: Performance for Settings That Fit a 6-watt Power Budget

legal performance surface. A simple power manager could choose the maximum point on the surface as the best performance within the power budget. A more complex manager could evaluate the performance surface in conjunction with a similar mapping for power to determine an efficient power-management setting for the application considering both power constraints and performance requirements.

6 Conclusion

Current and next-generation commercial computing systems have access to several power management techniques, and to employ them effectively, we need to understand their effects on both power and performance.

We characterized application response to two power management techniques, dynamic voltage and frequency scaling (DVFS) and clock throttling, applied individually and in combination on a Pentium M system with a set of microbenchmarks designed to capture core and memory subsystem activity. We analyzed microbenchmark power and performance for three dataset sizes corresponding to core-bound, intermediate, and memory-bound programs. We found that for techniques applied individually, DVFS is more effective than clock throttling at reducing power while preserving performance.

We also applied both techniques together in each of the 64 combinations of 8 frequencies and 8 clock throttling lev-

els and demonstrated that a performance or power target can be achieved with multiple combinations of DVFS and clock throttling settings. The combined-technique data also indicate that memory access patterns are an important consideration in choosing appropriate settings. An example of a fixed 6-watt power budget illustrates the spectrum of performance available from the DVFS and clock-throttling level combinations within limited power resources.

Adjusting DVFS settings incurs a stall latency of up to $10\mu s$, while changing the clock throttling level does not incur a stall. A dynamic power manager can tailor its response to program behavior by choosing among multiple options that suit performance requirements and power constraints, and evaluating the overhead of changing power management settings.

This report summarizes one component of a study conducted during a summer internship in 2005. Project members have continued other aspects of the study since the internship and we are currently preparing a conference paper on a method to predict the performance and power response to changing the power-management settings for use in choosing optimal settings in a real-time system. Our future work will continue development with characterization and prediction mechanisms for use in real-time dynamic power management scenarios.

7 Acknowledgments

We would like to thank the IBM Austin Research Laboratory and the ARL power team for the internship opportunity and a special thanks to Karthick Rajamani, Juan Rubio, Soraya Ghiasi, Freeman Rawson, and Tom Keller for ongoing collaboration.

References

- [1] Intel. Pentium M processor on 90 nm process with 2-MB L2 cache datasheet. <http://www.intel.com/design/mobile/datashts/302189.htm>, Jan. 2005.
- [2] Radisys. Endura LS855 product data sheet. http://www.radisys.com/oem_products/ds-page.cfm?productdatasheetsid=1158, Oct. 10 2004.
- [3] J. Rubio, K. Rajamani, and T. Keller. Power management with the Pentium M processor. In *ACEED 2005*, 2005.