

Energy Savings Through Compression in Embedded Java Environments

G. Chen, M. Kandemir, N. Vijaykrishnan,
M. J. Irwin
Microsystems Design Lab
Pennsylvania State University
University Park, PA 16802, USA
{gchen,kandemir,vijay,mji}@cse.psu.edu

W. Wolf
Department of Electrical Engineering
Princeton University
Princeton, NJ 08544, USA
wolf@princeton.edu

ABSTRACT

Limited energy and memory resources are important constraints in the design of an embedded system. Compression is an useful and widely employed mechanism to reduce the memory requirements of the system. As the leakage energy of a memory system increases with its size and because of the increasing contribution of leakage to overall system energy, compression also has a significant effect on reducing energy consumption. However, storing compressed data / instructions has a performance and energy overhead associated with decompression at runtime. The underlying compression algorithm, the corresponding implementation of the decompression and the ability to reuse decompressed information critically impact this overhead.

In this paper, we explore the influence of compression on overall memory energy using a commercial embedded Java virtual machine (JVM) and a customized compression algorithm. Our results show that compression is effective in reducing energy even when considering the runtime decompression overheads for most applications.

Keywords

Compression, Leakage Energy, Embedded Java

1. INTRODUCTION AND MOTIVATION

Java has become a popular vehicle for portable network programming, spanning not just resource-rich server and desktop environments, but resource constrained environments as well. It is estimated that the market for Java-enabled devices for resource-constrained environments such as cell-phones, PDAs and pagers will grow from 176 million in 2001 to 721 million in 2005 [12]. Various embedded Java virtual machines (JVMs) and Java accelerators have been proposed to target this potential market over the past year.

A Java system for an embedded/portable environment needs to meet an entirely different set of constraints as compared to executing on a high-performance or desktop environment. Three important aspects to which current embedded JVMs such as Sun's KVM [11] and HP's ChaiVM [3] conform are soft real-time, restricted memory size, and long-duration sessions requirements. Energy consumption is also an important design consideration for such battery-driven sys-

tems. However, currently, there is little support for analyzing and optimizing energy behavior of such embedded JVMs. In particular, the energy consumption in the memory system is a significant portion of overall energy expended in execution of a Java application [15]. Thus, it is important to consider techniques to optimize memory energy consumption. There are two important components of memory energy: dynamic energy and leakage energy. Dynamic energy is consumed whenever a memory array is referenced. Leakage energy is consumed as long as the device is powered and is consumed even when the device is not being accessed. While dynamic energy has been the traditional focus of most optimizations, leakage is becoming an equally important portion as supply voltages and thus threshold voltages and gate oxide thicknesses continue to scale [4]. Recent energy estimates for 0.13 micron process indicate that leakage energy accounts for 30% of L1 cache energy and as much as 80% of L2 cache energy [9]. Leakage energy is of particular concern in the dense memory structures as it increases with the size of the memory. In contrast, the effect of larger SRAM sizes on dynamic energy can be controlled by partitioning large SRAMs into smaller structures. Also, it is customary to use multiple levels of memory hierarchy to confine most accesses in the smallest memory.

In this work, we use a system-on-a-chip (SoC) with two-level memory hierarchy where a software-managed memory known as scratch pad memory (SPM) is used between the memory and the processor core. The SPM, due to its smaller size, has a lesser per access dynamic energy cost associated with it. Hence, confining most accesses to the smaller SPM (instead of large main memory) reduces the overall dynamic energy. However, the increased memory space due to the two-level memory hierarchy can increase the overall leakage energy of the system. Various compression schemes have been widely used to reduce the memory requirements. In this work, we use compression to reduce the size of the required memory. Specifically, we store the code of the embedded JVM system and the associated library classes in a compressed form in the memory. Thus, the effective number of active transistors used for storage and the associated leakage are reduced. We employ a mechanism that turns off power supply to the unused portions of the memory to control leakage. Whenever the compressed code or classes are required by the processor core, a mapping structure stored in a reserved part of the SPM serves to locate the required block of data in the compressed store. Then, the block of data after decompression is brought into the SPM. Thus, the use of scratch pad memory in conjunction with a compressed memory store targets the reduction of both dynamic and leakage energy of a system. The focus of this paper is on investigating the influence of different parameters on the design of such a system. The issues addressed in this work are listed below:

- Storing compressed code or data has an associated decompression cost from both the energy and performance aspects. To obtain any energy savings, the energy overhead of decompression must be smaller than the leakage energy savings obtained through storage of compressed code. The underlying compression algorithm and the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CODES'02, May 6-8, 2002, Estes Park, Colorado, USA.

Copyright 2002 ACM 1-58113-542-4/02/0005...\$5.00.

corresponding implementation of the decompression critically impact the energy and performance overhead. We explore this idea using a specific hardware compression scheme and also experiment with different decompression overheads to account for a range of possible implementations from customized hardware to software.

- The size of the compressed block influences both the compression ratio and the overhead involved in indexing the compressed data. A larger granularity of compression, typically, provides a better compression ratio. In turn, this provides an ability to turn off power supply to more unused memory blocks, thereby providing larger leakage energy savings. Also, it reduces the mapping overhead for indexing into the compressed store. However, a larger block also occupies a larger space in the SPM and increases the storage pressure. This can lead to more frequent conflicts in the scratch pad memory resulting in more frequent decompressions.

The rest of this paper is organized as follows. Section 2 introduces the virtual machine used in this study, our SoC, and our embedded applications. Section 3 discuss our strategy for saving leakage energy through compression. Section 4 presents our simulation environment and Section 5 gives our experimental results. Finally, Section 6 concludes the paper with a summary of our major contributions.

2. KVM, SOC ARCHITECTURE, AND APPLICATIONS

2.1 Virtual Machine

In this study, we used K Virtual Machine (KVM) [11], Sun's virtual machine designed for resource-constrained (e.g., battery-operated) environments. It targets embedded computing devices with as little as a few kilobytes total memory, including actual virtual machine and Java class libraries themselves. These devices include smart wireless phones, pagers, mainstream personal digital assistants, and small retail payment terminals.

2.2 Architecture

A system-on-a-chip (SoC) is an integrated circuit that contains an entire electronic system in a single sliver of silicon. A typical SoC contains a library of components designed in-house as well as some cores from chipless design houses also known as intellectual property. In this work, we focus on an SoC-based system that executes KVM applications. Figure 1(a) depicts the high level (logical) view of the relevant parts of our SoC architecture. This architecture has a CPU core, a scratch-pad memory (SPM), and two main memory modules. The processor in our SoC is a microSPARC-IIep embedded core. This core is a 100MHz, 32-bit five-stage pipelined RISC that implements the SPARC architecture V8 specification. It is primarily targeted for low-cost uniprocessor applications. Both main memory and SPM are SRAMs which are organized as blocks. Each main memory block can be mapped into one SPM block. Each SPM block has a tag register indicating which main memory block is currently mapped into this SPM block. The tag registers are set by SPM manager, which is implemented in hardware.

When the CPU generates an address, the high-order bits of this address are compared with each tag in parallel. If one of the tags generates a match, the corresponding SPM block is selected and low-order bits of the address are used to access the contents of the block. If no tag match occurs, then the "Hit" signal line (shown in Figure 1(b)) is disabled and an interrupt is generated. The corresponding interrupt service routine activates the SPM manager which brings the faulted block from main memory to the SPM. In case no free SPM block is available, a timer-based block replacement policy is used. Specifically, for each SPM block, there is a timer and an access bit. Whenever the block is accessed, its access bit is set and its timer is reset. When a block is not accessed for a certain period of time, the timer goes off and the access bit is reset. When a block replacement is to be performed, the SPM Manager always tries to select a block whose access bit is reset. If no such block exists, the manager selects a

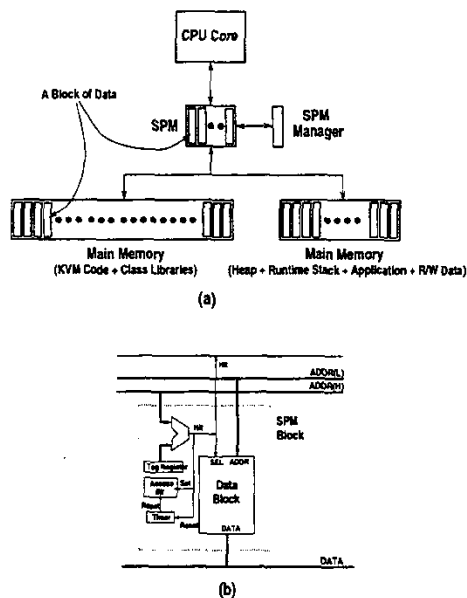


Figure 1: (a) High-level view of the SoC memory architecture. (b) Details for an SPM block.

block in a round-robin fashion. The main memory is composed of two parts: one part which contains the KVM code and class libraries and the other part which contains all writable data including heap and C stack as well as application code.

It should be mentioned that a number of parameters in this architecture are tunable. For example, the capacities of SPM and main memory can be modified. Also, by playing with the width of the timers associated with each block, we can modify the behavior of the block replacement policy. Finally, the SPM block size can be changed. Note that changing the block size affects the block replacement rate as well as the overhead (per block) when a replacement occurs.

Instead of an SPM, a cache could also have been employed. In our experiments, we found that using 2-way associate 32KB instruction and data caches both with line sizes of 32 bytes consumed 11% more energy than an equivalent 64KB SPM configuration. Hence, for lack of space, we show only SPM results.

2.3 Applications

In this subsection, we describe the applications used in this study. To test the effectiveness of our energy saving strategy, we collected twelve applications shown in Figure 2. These applications represent a group of codes that are executed in energy-sensitive devices such as hand-held computers and electronic game boxes, and range from utilities such as calculator and scheduler, embedded web browser to game programs. These applications represent a good mix of codes that one would expect to run under KVM-based environments.

3. COMPRESSING KVM CODE AND CLASS LIBRARIES

As noted earlier, leakage energy consumption of SRAM blocks is proportional to their size as well as the duration of time that they are powered on. In this work, we try to reduce the number of active (powered on) memory blocks by storing read-only data, including KVM binary codes and Java class libraries, in compressed form. To

Application	Brief Description	Source
Calculator	Arithmetic calculator	[Omitted for anonymity]
Crypto	Cryptography	www.bouncycastle.org
Dragon	Game program	comes with Sun's KVM
Elite	3D rendering	home.rochester.rr.com/ohommes/Elite/
Kshape	Electronic map	www.jshape.com
Kvideo	KPG decoder	www.jshape.com
Kwml	WML browser	www.jshape.com
ManyBalls	Game program	comes with Sun's KVM
MathFP	Math lib	home.rochester.rr.com/ohommes/MathFP/
Missiles	Game program	comes with Sun's KVM
Scheduler	Weekly/daily scheduler	[Omitted for anonymity]
StarCruiser	Game program	comes with Sun's KVM

Figure 2: Brief description of the benchmarks used in our experiments.

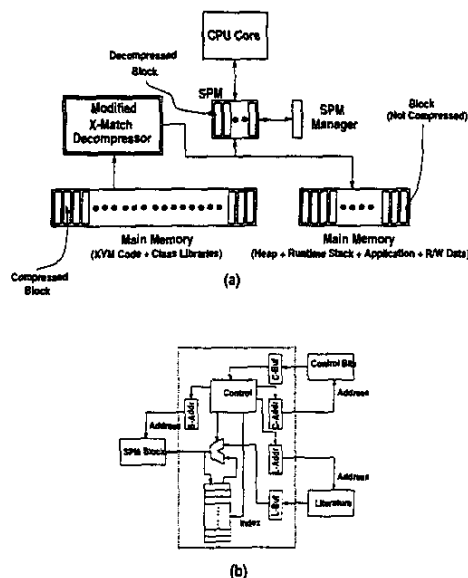


Figure 3: (a) High-level view of the SoC memory architecture with decompressor. (b) Details for the modified X-Match decompressor.

avoid incurring the cost of runtime compression, writable data are stored in original form. The high-level view of our architecture with the decompression support is shown in Figure 3(a). When a data item belonging to a compressed memory block is requested by the processor, the whole block is decompressed by the decompressor and is then written into the SPM. The advantage of this strategy is that since data in read-only memory is in the compressed form, it occupies fewer memory blocks. This, in turn, reduces the leakage energy consumption in read-only part of the main memory system. Note that the amount of this saving is determined by the compression rate of the algorithm used. The drawback is that decompressing a block (at runtime) incurs both time and energy penalty. The magnitude of this penalty depends on how frequently decompression is required and how much time/energy it takes to decompress a block of data. The number of decompressions is directly related to the number of misses in the SPM (which is a characteristic of application behavior and SPM configuration). The time/energy expended in decompressing the block depends on the decompression method used and can be reduced by an efficient implementation.

The compression/decompression algorithm to be used in such a framework should have the following important characteristics: (1) good compression ratio for small blocks (typically less than 4KB); (2) fast decompression; and (3) low energy consumption in decompression. Since compression is performed offline, its speed and energy overhead are not constrained. The first characteristic is desirable because the potential leakage energy savings in memory are directly related to the number of memory blocks that need to be powered on. Note that a compressed memory block needs to be decompressed before it can be stored in the SPM. Consequently, a decompression overhead is incurred in every load to SPM and, in order for this scheme to be effective, we should spend very little time and energy during decompression.

Kjelsø et al. [8] presented a dictionary-based compression/decompression algorithm called X-Match. This algorithm maintains a dictionary of data previously seen, and attempts to match the current data element (to be compressed) with an entry in the dictionary. If such a match occurs, the said data element is replaced with a short code word indicating the location of data in the dictionary. Data elements that do not generate a match are transmitted in full (literally), prefixed by a single bit. Each data element is exactly 4 bytes in width and is referred to as a tuple. A full match occurs when all characters in the incoming tuple fully match a dictionary entry. A partial match occurs when at least two of the characters in the incoming tuple match exactly a dictionary entry; the characters that do not match are transmitted literally. The coding function for a match encodes three separate fields: (1) match location, (2) match type indicating which characters from the incoming tuple matched the dictionary entry, and (3) any characters from the incoming tuple which did not match the dictionary entry at the match location (i.e., those transmitted without encoding).

In the original X-Match algorithm, the dictionary is maintained using a move-to-front strategy, whereby the current tuple is placed at the front of the dictionary and other tuples move down by one location. If the dictionary becomes full, the tuple occupying the last location is simply discarded. The move-to-front operation is implemented with content addressable memory, which is expensive from the energy consumption perspective. In our implementation, we replaced the move-to-front strategy with a simple round-robin strategy, i.e., the new tuple is always appended to the end of current dictionary entries. When the dictionary is full, the replacement pointer is moved to the first dictionary entry and the entry becomes the one that will be replaced next time. The elimination of the move-to-front strategy may cause a slight degradation in the compression ratio, but the implementation is simpler and energy-efficient. We also separate the literal bytes from the control bits (i.e., prefixes, match types, and dictionary locations), which allows the control bits and literal bytes to be fed into the decompressor as separate streams. We refer to this modified algorithm as the modified X-Match algorithm in the rest of the paper. The hardware block diagram of the modified X-Match decompressor is shown in Figure 3(b).

While this modified X-Match implementation is used in our evaluation, the idea of trading additional decompression energy with reduced memory leakage energy is applicable using other compression schemes such as Lempel-Ziv and Huffman.

4. SIMULATION METHODOLOGY

4.1 Energy Model

The energy numbers reported in this paper are obtained by a simulator implemented on the SPARC simulation tool-set, Shade [6], augmented with energy models. The simulator takes as input the KVM system executing a Java application and computes performance as well as energy data. The current implementation runs on SPARC systems and simulates the SPARC V8 instruction set of our target processor. Our simulator tracks energy consumption in the processor core, SPM, and main memory blocks. The energy consumed in the processor core is estimated by counting (dynamically) the num-

ber of instructions of each type and multiplying the count by the base energy consumption of the corresponding instruction. The base energy consumptions of the different instruction types are obtained using a customized and validated version of our in-house cycle accurate energy simulator [14]. The simulator is configured to model a five-stage pipeline similar to that of the target microSPARC-IIep architecture.

The energy consumption in SPM and main memory is divided into two components: dynamic energy and leakage energy. In computing per access dynamic energy consumptions for SPM and main memory, we used the CACTI tool Version 2.0 [10] assuming a 0.10 micron technology. In computing the leakage energy, we assumed that the leakage energy per cycle of the entire main memory is equal to the dynamic energy consumed per access. This assumption tries to capture the anticipated importance of leakage energy in the future. It should be stressed that leakage becomes the dominant part of energy consumption for 0.10 micron (and finer) technologies for the typical internal junction temperatures in a chip [4]. Note that, as opposed to dynamic energy which is expended only when an access occurs, leakage energy is spent as long as memory is powered on.

In computing the overall energy consumption in main memory and SPM, we assumed that a memory block (or an SPM block) can be in one of three states (modes) at any given time: *R/W*, *active*, or *inactive*. In the *R/W* (read/write) mode, memory is being read or written and consumes full dynamic energy as well as full leakage energy. In the *active* state, on the other hand, the memory is powered on but not being accessed. In this state, it consumes no dynamic energy but full leakage energy. Finally, the memory modules that are not needed by the system are not powered on, i.e., in the *inactive* state, consequently, no energy consumption at all. Obviously, one would want to place as many memory blocks as possible to the *inactive* state so that the energy consumption can be minimized. One way of achieving this is to reduce the amount of data stored in memory, which can be achieved using compression.

4.2 Base Configuration and Energy Distribution

Figure 4 gives the simulation parameters used in our base configuration. Figure 5 shows (in columns two through five) the energy consumptions (in micro-joules) for our applications executing on base configuration without decompression. The energy consumption is divided into four components: dynamic energy in SPM, leakage energy in SPM, dynamic energy in main memory, and leakage energy in main memory. The contribution of the processor energy to the overall (main memory + SPM + processor) energy is around 10% and is not much affected by decompression. Consequently, we focus only on main memory and SPM energies. A memory block that contains no valid information throughout the application execution is turned off so that it does not consume any leakage energy. We see from these results that the memory leakage energy consumption (shown in the third column) constitutes a large percentage of the memory system (main memory + SPM) energy budget (61.74% on the average) and is a suitable target for optimization. The sixth column in Figure 5 gives percentage of energy consumption due to read-only part of the memory. We see that, on the average, the read-only part of the memory is responsible for 62.42% of the overall memory energy consumption. Finally, the last two columns give the number of SPM misses and the number of execution cycles (in millions) for each application.

5. RESULTS

In this section, we present data showing the effectiveness of our strategy in saving energy and also measure the sensitivity of our strategy to different parameters such as SPM capacity, block size, and cost of decompression. All energy numbers reported here are values normalized to the energy consumption in the base case without any decompression (Figure 5). Also, when a simulation parameter is modified, the remaining parameters maintain their original values given in Figure 4.

Parameter	Value
SPM capacity	64KB 4KB SPM management 40KB for read-only data 20KB for writable data
SPM block size	1KB for read-only data 512 bytes for writable data
Main memory capacity	512KB
SPM access time	1 cycle
Main memory access time	3 cycles
SPM dynamic energy/read	0.5216 nJ
SPM dynamic energy/write	0.6259 nJ
Main memory dynamic energy/read	1.334 nJ
Main memory dynamic energy/write	1.601 nJ
Main memory leakage energy/byte/cycle	2.54×10^{-6} nJ
SPM leakage energy/byte/cycle	2.54×10^{-6} nJ
SPM access bit reset time	6000 cycles for read-only 4000 cycles for <i>r/w</i> clean 8000 cycles for <i>r/w</i> dirty

Figure 4: Simulation parameters and their values for our base configuration.

The top part of Figure 6 gives the normalized energy consumptions in read-only portion of the main memory and the SPM. It can be observed from this figure that the energy saving is 20.9% on the average. The bottom part of Figure 6, on the other hand, shows the overall (normalized) energy consumption in main memory and SPM, including the energy expended during decompression. We see that most of the applications achieve an overall energy saving of 10% (an average of 7% across all applications). In two applications (*Calculator* and *Scheduler*), the decompression overhead (energy) plays a larger role and the overall energy consumption becomes worse than the original case. We also experimented with a 50% reduction in leakage energy per main memory cell to account for design variations that permit the slower main memory cells to operate using a higher threshold voltage. In this case, the overall memory system energy saving across all applications is 5.4% on the average.

In general, there are two application-related factors that determine the effectiveness of our energy saving strategy: (1) the overall running time of the application, (2) the number of SPM misses. Since the major energy gain in our strategy comes from the memory leakage energy, the longer the application runs, we can expect more energy benefits. Recall that each SPM misses invokes a decompression. Therefore, the number of SPM misses is an important factor in determining the energy spent in decompression during the course of execution. The reasons that *Calculator* and *Scheduler* do not get benefit from our strategy are different. In *Calculator*, the execution time is rather short (only 5.60 million cycles) and the energy spent on decompression does not pay off. On the other hand, although the execution time of *Scheduler* is not short (105.5 million cycles), it suffers from a high number of SPM misses (a total of 96033).

Since our strategy focuses on energy savings in the read-only part of the main memory, in the rest of this section, we mainly present results pertaining only this part and the SPM (unless otherwise stated). However, to evaluate the impact of decompression, we also show the energy consumed during the decompression process.

5.1 Sensitivity to the Decompression Cost

To see how a more efficient or a less efficient implementation of the X-Match decompressor would impact our results, we conducted another set of experiments. Assuming a decompression rate of 4 bytes/cycle and that the energy consumed in each stage is equal to one SPM access, we determined that the energy consumption for decompressing one word (4 bytes) is equal to three SPM accesses. We normalize this energy cost to 1 and experiment with its multiple as

Application	Memory Energy		SPM Energy		Read-Only Contribution	Number of SPM Misses	Number of Cycles
	Dynamic	Leakage	Dynamic	Leakage			
Calculator	1.81 (13.70%)	7.46 (56.58%)	3.08 (23.31%)	0.87 (6.61%)	66.00%	5258	5.59
Crypto	7.30 (3.37%)	137.94 (63.75%)	54.97 (25.40%)	16.17 (7.47%)	60.46%	19848	103.58
Dragon	1.10 (1.05%)	68.70 (65.41%)	27.18 (25.88%)	8.05 (7.67%)	60.86%	3064	51.58
Elite	1.10 (1.12%)	64.42 (65.36%)	25.49 (25.86%)	7.55 (7.66%)	60.88%	3206	48.37
MathFP	2.42 (1.50%)	104.96 (65.04%)	41.70 (25.84%)	12.30 (7.62%)	60.40%	6735	78.81
ManyBalls	5.90 (5.05%)	73.08 (62.59%)	29.22 (25.03%)	8.56 (7.33%)	59.98%	11455	54.87
Missiles	2.51 (2.67%)	60.34 (64.27%)	23.96 (25.53%)	7.07 (7.53%)	61.61%	7288	45.30
KShape	16.18 (5.01%)	201.90 (62.57%)	80.93 (25.08%)	23.66 (7.33%)	61.22%	45321	151.61
KVideo	2.21 (10.24%)	12.69 (58.87%)	5.17 (23.99%)	1.49 (6.90%)	64.87%	6396	9.53
KWML	121.56 (11.97%)	584.96 (56.58%)	240.78 (23.70%)	68.55 (6.75%)	64.35%	339395	439.26
Scheduler	32.96 (13.30%)	140.51 (56.71%)	57.83 (23.34%)	16.47 (6.65%)	66.50%	96035	105.50
StarCruiser	3.22 (4.16%)	48.96 (63.18%)	19.58 (25.27%)	5.74 (7.40%)	61.91%	9356	36.76

Figure 5: Energy consumptions and execution cycles for our applications under the base configuration.

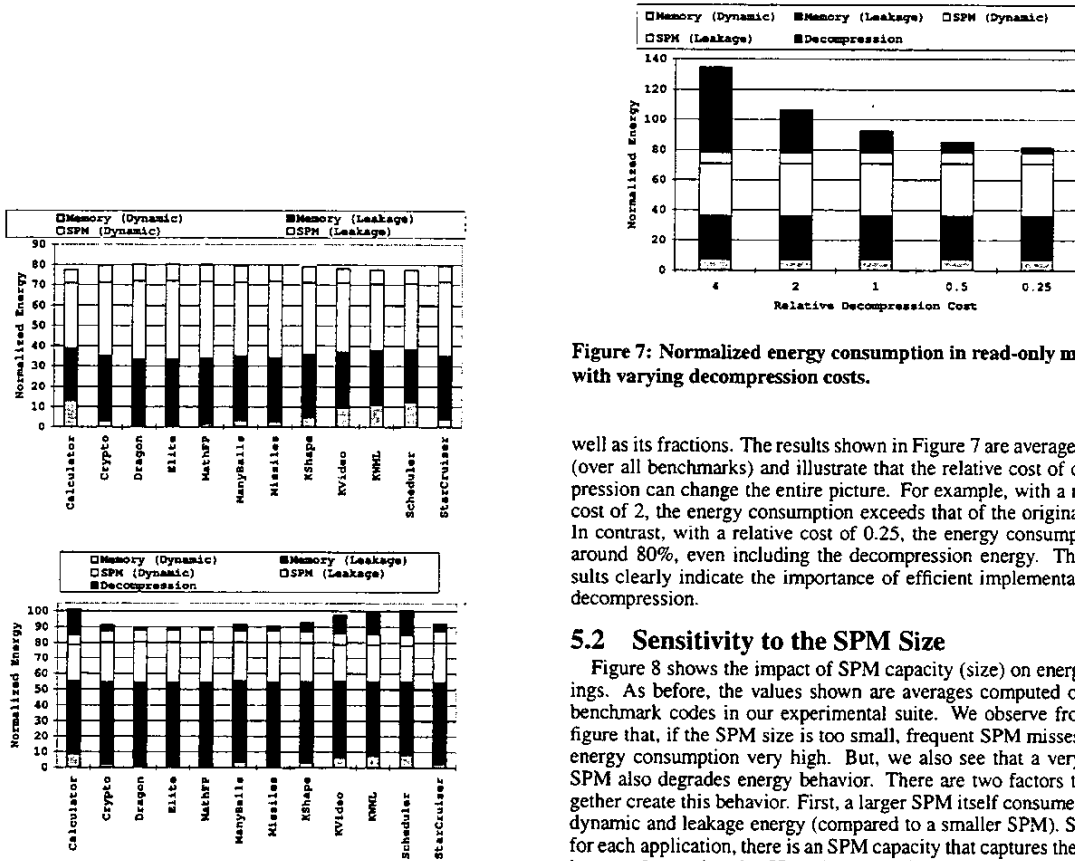


Figure 6: Normalized energy consumption in read-only memory and SPM (top) and in overall memory and SPM (bottom).

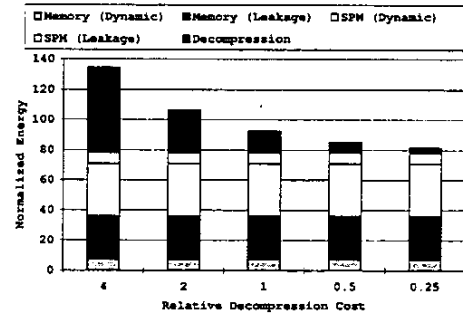


Figure 7: Normalized energy consumption in read-only memory with varying decomposition costs.

well as its fractions. The results shown in Figure 7 are average values (over all benchmarks) and illustrate that the relative cost of decompression can change the entire picture. For example, with a relative cost of 2, the energy consumption exceeds that of the original case. In contrast, with a relative cost of 0.25, the energy consumption is around 80%, even including the decompression energy. These results clearly indicate the importance of efficient implementation of decompression.

5.2 Sensitivity to the SPM Size

Figure 8 shows the impact of SPM capacity (size) on energy savings. As before, the values shown are averages computed over all benchmark codes in our experimental suite. We observe from this figure that, if the SPM size is too small, frequent SPM misses make energy consumption very high. But, we also see that a very large SPM also degrades energy behavior. There are two factors that together create this behavior. First, a larger SPM itself consumes more dynamic and leakage energy (compared to a smaller SPM). Second, for each application, there is an SPM capacity that captures the working set. Increasing the SPM size beyond this capacity does not reduce the number of misses further. So, this stability in the number of misses, combined with the first factor, leads to an increase in energy consumption. In an embedded system design, the maximum possible SPM size is determined by the chip budget. Our experimentation indicates that the best SPM size depends on the application at hand. So, embedded system designers should select a suitable SPM size considering the applications in question as well as the impact of SPM size on energy consumption.

5.3 Sensitivity to the Block Size

In this set of experiments, we tried to measure the sensitivity of our energy savings to the block size used. Recall that our default

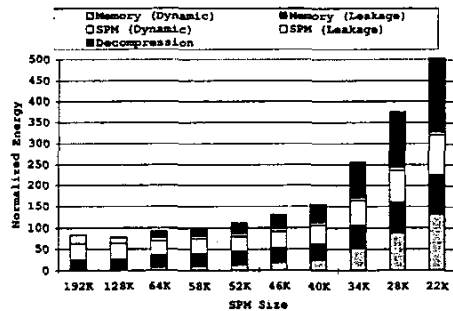


Figure 8: Normalized energy consumption in read-only memory with varying SPM sizes.

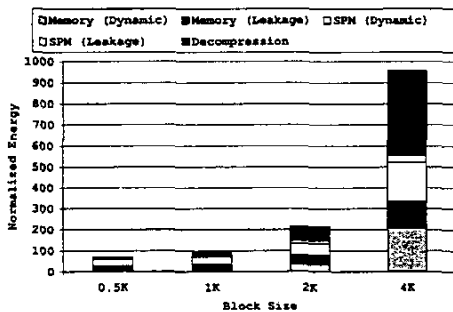


Figure 9: Normalized energy consumption in read-only memory with varying block sizes.

block size was 1KB. The results shown in Figure 9 indicate that, given the SPM capacity, the size of each block has a great impact on the energy consumption. For most compression algorithms in the literature, a larger block size has, in general, a better compression ratio. However, it should be noted that, a very large block size might increase both SPM miss rate and miss penalty (decompression cost). That is exactly the behavior we observed during our experiments. As shown in Figure 9, a block size 0.5KB generated better results than our default 1KB blocks. In contrast, increasing the block size to 2KB increased the original energy consumption by more than a factor of two.

6. CONCLUSIONS

Storing compressed code or data has an associated decompression cost from both the energy and performance aspects. However, compression itself helps to reduce the portion of the memory to be powered on and the consequent leakage energy of the memory system. Our experiments with a set of embedded applications using a commercial embedded JVM and a specific compression scheme show that the proposed technique is effective in reducing system energy. We expect our findings to be applicable to other compression algorithms and implementations as well.

Acknowledgment

This work was supported in part by NSF CAREER Awards 0093082 & 0093085; NSF Awards 0073419, 0082064, 0103583 and an award from GSRC.

7. REFERENCES

- [1] L. Benini and G. De Micheli. System-level power optimization: techniques and tools. *ACM Transactions on Design Automation of Electronic Systems*, 5(2), pp.115-192, April 2000.
- [2] F. Catthoor, S. Wuytack, E. D. Greef, F. Balasa, L. Nachtergaele, and A. Vandecappelle. *Custom Memory Management Methodology - Exploration of Memory Organization for Embedded Multimedia System Design*. Kluwer Academic Publishers, June 1998.
- [3] ChaiVM for Jornado. <http://www.hp.com/products/1/embedded/jornado/index.html>
- [4] A. Chandrakasan, W. J. Bowhill, and F. Fox. *Design of High-Performance Microprocessor Circuits*. IEEE Press, 2001.
- [5] G. Chen, R. Shetty, M. Kandemir, N. Vijaykrishnan, M. J. Irwin, and M. Wolczko. Tuning garbage collection in an embedded Java environment. In *Proc. the 8th International Symposium on High-Performance Computer Architecture*, Cambridge, MA, February 2-6, 2002.
- [6] B. Cmelik and D. Keppel. Shade: A Fast Instruction-Set Simulator for Execution Profiling. In *Proc. ACM SIGMETRICS Conference on the Measurement and Modeling of Computer Systems*, pp. 128-137, May 1994.
- [7] J. Flinn, G. Back, J. Anderson, K. Farkas, and D. Grunwald. Quantifying the energy consumption of a pocket computer and a Java virtual machine. In *Proc. International Conference on Measurement and Modeling of Computer Systems*, June 2000.
- [8] M. Kjelso, M. Gooch, and S. Jones. Performance evaluation of computer architectures with main memory data compression. *Elsevier Science, Journal of Systems Architecture*, 45 (1999), pp. 571-590.
- [9] M. D. Powell, S.-H. Yang, B. Falsafi, K. Roy, and T. N. Vijaykumar. Gated-Vdd: a circuit technique to reduce leakage in deep-submicron cache memories. In *Proc. the ACM/IEEE International Symposium on Low Power Electronics and Design*, August 2000.
- [10] G. Reinman and N. Jouppi. An integrated cache timing and power model. COMPAQ Wester Research Lab, Palo Alto, CA, 1999. <http://www.research.compaq.com/wrl/people/jouppi/CACTI.html>
- [11] R. Riggs, A. Taivalsaari and M. VandenBrink. *Programming Wireless Devices with the Java 2 Platform*. Addison Wesley, 2001.
- [12] D. Takahashi. Java chips make a comeback. *Red Herring*, July 12, 2001.
- [13] The future of SoC design. <http://www.eetasia.com/ART.8800141212.HTM>.
- [14] N. Vijaykrishnan, M. Kandemir, M. J. Irwin, H. Y. Kim, and W. Ye. Energy-driven integrated hardware-software optimizations using SimplePower. In *Proc. the International Symposium on Computer Architecture*, Vancouver, British Columbia, June 2000.
- [15] N. Vijaykrishnan, M. Kandemir, S. Tomar, S. Kim, A. Sivasubramaniam and M. J. Irwin. Energy Characterization of Java Applications from a Memory Perspective. In *Proc. USENIX Java Virtual Machine Research and Technology Symposium*, April 2001.