

Synergistic Integration of Dynamic Cache Reconfiguration and Code Compression in Embedded Systems*

Hadi Hajimiri, Kamran Rahmani, Prabhat Mishra
Department of Computer & Information Science & Engineering
University of Florida, Gainesville, Florida, USA
{hadi, kamran, prabhat}@cise.ufl.edu

Abstract—Optimization techniques are widely used in embedded systems design to improve overall area, performance and energy requirements. Dynamic cache reconfiguration is very effective to reduce energy consumption of cache subsystems which accounts for about half of the total energy consumption in embedded systems. Various studies have shown that code compression can significantly reduce memory requirements, and may improve performance in many scenarios. In this paper, we study the challenges and associated opportunities in integrating dynamic cache reconfiguration with code compression to retain the advantages of both approaches. Experimental results demonstrate that synergistic combination of cache reconfiguration and code compression can significantly reduce both energy consumption (65% on average) and memory requirements while drastically improve the overall performance (up to 75%) compared to dynamic cache reconfiguration alone.

1. INTRODUCTION

Energy conservation has been a primary optimization objective in designing embedded systems as these systems are generally limited by battery lifetime. Several studies have shown that memory hierarchy accounts for as much as 50% of the total energy consumption in many embedded systems [1]. Dynamic cache reconfiguration (DCR) and code compression are two of the extensively studied approaches in order to achieve energy savings as well as area and performance gains.

Different applications require highly diverse cache configurations for optimal energy consumption in the memory hierarchy. Unlike desktop-based systems, embedded systems are designed to run a specific set of well-defined applications. Thus it is possible to have a cache architecture that is tuned for those applications to have both increased performance as well as lower energy consumption. Since too many cache configurations are possible, the challenge is to determine the best cache configuration (in terms of total size, associativity, and line size) for a particular application. Studies have shown that cache tuning can achieve 53% memory-access-related energy savings and 30% performance improvement [2].

The use of high-level programming languages coupled with RISC instruction sets leads to a larger memory footprint and increased area/cost and power requirements, all of which are important design constraints in most embedded applications. Code compression is clearly beneficial for memory size reduction because it reduces the static memory size of executable code. Several code compression techniques have been proposed for reducing instruction memory size in low cost

embedded applications [3]. The basic idea is to store instructions in compressed form and decompress them on-the-fly at execution time. More importantly, code compression could also be beneficial for energy by reducing memory size and the communication between memory and the processor core [4].

Design of efficient compression techniques needs to consider two important aspects. First, the compressed code has to support the possibility of starting the decompression during execution at several points inside the program (i.e., branch targets). Second, since decompression is performed on-line, during program execution, decompression algorithms should be fast and power efficient to achieve savings in memory size and power, without compromising performance. We explore various compression techniques (including dictionary-based compression, bitmask-based compression and Huffman coding) that represent a trade-off between compression performance and decompression overhead.

It is expected that by compressing instructions the cache behavior of programs is no longer the same. Thus in order to have the optimal cache configuration, more analysis should be done including hit/miss behavior of the compressed programs. In other words, cache reconfiguration needs to be aware of code compression to obtain best possible area, power and performance results. In this paper, we present an elaborate analysis of combining two optimization techniques: dynamic cache reconfiguration and code compression. Our experimental results demonstrate that the combination is synergistic and achieves more energy savings as well as overall performance compared to DCR and code compression alone.

The rest of the paper is organized as follows. Section 2 provides an overview of related research activities. In Section 3 we describe our compression-aware cache reconfiguration methodology. Section 4 presents our experimental results. Finally, Section 5 concludes the paper.

2. BACKGROUND AND RELATED WORK

2.1 Dynamic Cache Reconfiguration (DCR)

In power constrained embedded systems, nearly half of the overall power consumption is attributed to the cache subsystem [1]. Applications require vastly different cache requirements in terms of cache size, line size, and associativity. Research shows that specializing the cache to application's needs can significantly reduce energy consumption [2]. Fig. 1 illustrates how energy consumption can be reduced by using inter-task (application-based) cache reconfiguration in a simple system supporting three tasks. In application-based cache tuning, DCR happens when a task starts its execution or it resumes from an interrupt (either by preemption or when exe-

* This work was partially supported by NSF grant CNS-0915376.

cution of another task completes) and the same cache for the application gets chosen no matter if it is starting from the beginning or resuming anywhere in between. Fig. 1 (a) depicts a traditional system and Fig. 1 (b) depicts a system with a reconfigurable cache. For the ease of illustration let's assume cache size is the only reconfigurable parameter of cache (associativity and line size are ignored). In this example, Task1 starts its execution at time P1. Task2 and Task3 start at P2 and P3 respectively. In a traditional approach, the system always executes using a 4096-byte cache. We call this cache as **base cache** throughout the paper. *Base cache is the best possible cache configuration optimized for all the tasks.* With the option of reconfigurable cache, Task1, Task2, and Task3 execute using 1024-byte cache starting at P1, 8192-byte cache starting at P2, and 4096-byte cache starting at P3 respectively. Through proper selection of cache size for each task the system can achieve significant amount of energy savings as well as performance gains compared to using only the *base cache*.

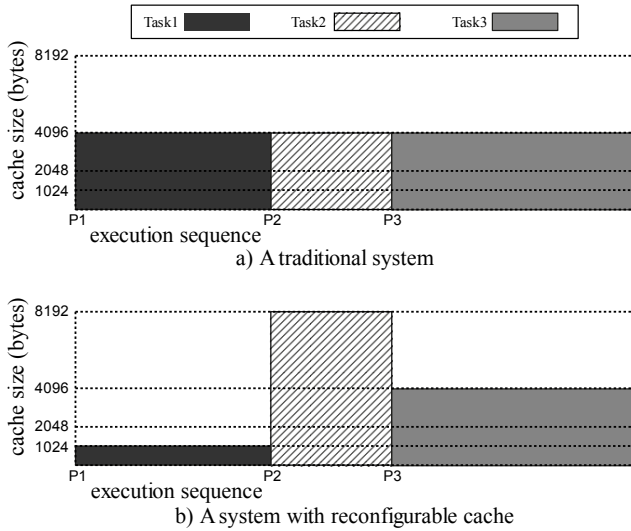


Fig. 1: DCR for a system with three tasks

The inter-task DCR problem is defined as follows. Consider a set of n applications (tasks) $\mathbf{A} = \{a_1, a_2, a_3, \dots, a_n\}$ intended to run on a configurable cache architecture capable of supporting m possible cache configurations $\mathbf{C} = \{c_1, c_2, c_3, \dots, c_m\}$. We define $e(c_j, a_i)$ as the total energy consumed by running application a_i on the architecture with cache configuration c_j . We also define $c_o \in \mathbf{C}$ as the optimal cache configuration for application a_i , such that $e(c_o, a_i) \leq e(c_j, a_i), \forall c_j \in \mathbf{C}$. Through exhaustive exploration of all possible configurations of $\mathbf{C} = \{c_1, c_2, c_3, \dots, c_m\}$, best energy optimal cache configuration for each application can be found.

Dynamic cache reconfiguration has been extensively studied in several works [5] [6] [7] [8]. The reconfigurable cache architecture proposed by Zhang et al. [6] determines the best cache parameters by using Pareto-optimal points trading off energy consumption and performance. Their method imposes no overhead to the critical path, thus cache access time does not increase. Chen and Zou [9] introduced a novel reconfiguration management algorithm to efficiently search the large design space of possible cache configurations for the optimal one. None of these approaches consider the effects of compressed code on cache reconfiguration.

DCR can be viewed as a technique that tries to squeeze cache size with other cache parameters to reduce energy consumption without (or with minor) performance degradation. Smaller caches contribute less static power but may increase cache misses which can lead to increased dynamic power and performance degradation (longer execution time thus higher energy consumption). Therefore, the smallest possible cache may not be a feasible solution in many cases. DCR techniques find the best cache that fits the application by exploring cache configurations using various schemes. In this paper, we show that code compression which significantly reduces the code size can also help the cache reconfiguration technique to choose relatively smaller cache sizes, smaller associativity, or smaller line size without performance degradation, therefore, reduces cache energy consumption significantly.

The configurable caches used in our work are based on the architecture described in [10]. The underlying cache architecture contains four separate banks that can operate as four separate ways. Special configuration registers are used to inform the cache tuner – a custom hardware or a lightweight process – to concatenate ways such that the associativity can be altered. The special registers may also be configured to shut down ways to vary the cache size. Similarly, by configuring the fetch unit to fetch cache lines in various lengths, we can adjust the line sizes. Cache reconfiguration time and energy overhead of the reconfigurable hardware is negligible [10].

2.2 Code Compression in Embedded Systems

Various code compression algorithms are suitable for embedded systems, i.e., provide good compression efficiency with minor (acceptable) or no decompression overhead. Wolfe and Chanin [11] were among the first to propose an embedded processor design that incorporates code compression. Xie et al. [12] introduced a compression technique capable of compressing flexible instruction formats in VLIW architectures. Seong et al. [13] modified dictionary-based compression (BMC) technique using bitmasks which improved compression efficiency without introducing any additional decompression overhead. Lin and Xie [14] proposed LZW-based algorithms to compress branch blocks. Recently, Rawlins et al. [15] used compressed programs in their approach of combined loop caching with DCR. Their approach has several limitations. They primarily focus on loop caching which may not be applicable in many embedded systems due to intrusive addition of another level of cache. Furthermore, due to emphasis on loop caching, interactions between compression and DCR was not explored in detail. In this paper we provide comprehensive analysis of how compression and DCR synergistically interact with each other as well as energy-performance trade-offs available for system designer.

Traditional code compression and decompression flow is illustrated in Fig. 2 where the compression is done offline (prior to execution) and the compressed program is loaded into the memory. The decompression is done during the program execution (online) and as shown in Fig. 7 it can be placed before or after cache. It is possible to place the decompression unit between two levels of cache as well, if the system has multi-level cache hierarchy.

In this paper we explore three compression techniques: dictionary-based compression (DC), bitmask-based compression (BMC) [13], and Huffman coding. DC and Huffman coding represent two extremes. DC is a simple compression technique and therefore produces moderate compression but decompression is very fast. On the other hand, Huffman coding is considered to be one of the most efficient compression techniques but has higher decompression overhead/latency. DC and Huffman are widely used but BMC is a recent enhancement of DC that enables more matching patterns. Fig. 3 shows the generic encoding formats of bitmask-based compression technique for various numbers of bitmasks. Compressed data stores information regarding the bitmask type, bitmask location, and the mask pattern itself. The bitmask can be applied in different places in a vector and the number of bits required for indicating the position varies depending on the bitmask type. Bitmasks may be sliding or fixed. A fixed bitmask can be applied to fixed locations, such as byte boundaries. However, sliding bitmasks can be applied anywhere in the code vector.

The main advantage of bitmask-based compression over traditional dictionary-based compression is the increased matching patterns. In dictionary-based compression, each vector is compressed only if it completely matches with a dictionary entry. Fig. 4 illustrates an example of bitmask-based compression in which it can compress up to six data entries using bitmask-based compression, whereas using only dictionary-based compression would compress only four entries. The example in Fig. 4 uses only one bitmask. In this case, vectors that match exactly a dictionary entry are compressed with 3 bits. The first bit represents whether it is compressed (using 0) or not (using 1). The second bit indicates whether it is compressed using bitmask (using 0) or not (using 1). The last bit indicates the dictionary index. Data that are compressed using bitmask requires 8 bits. The first two bits, as before, represent if the data is compressed, and whether the data is compressed using bitmasks. The next three bits indicate the bitmask position and followed by two bits that indicate the bitmask pattern.

In this example, the compression ratio is 80%. *Compression ratio* (CR), widely accepted as a primary metric for measuring the efficiency of code compression, is defined as:

$$CR = \frac{\text{Compressed program size}}{\text{Original program size}}$$

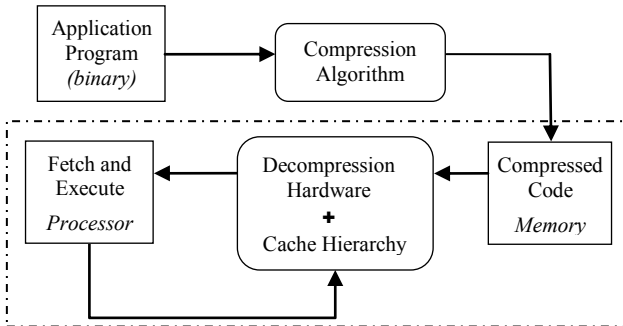


Fig. 2: Traditional code compression methodology

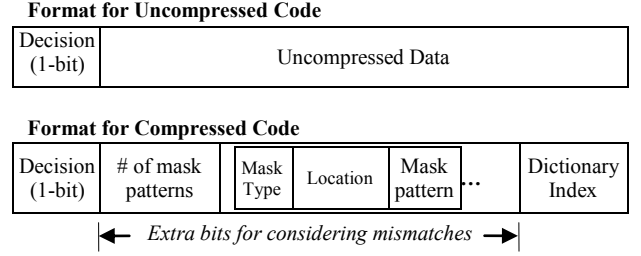


Fig. 3: Encoding format for incorporating mismatches

Bitmask selection and dictionary selection are two major challenges in bitmask-based code compression. Seong et al. [13] have shown that the profitable bitmasks to be selected for code compression are 1s, 2s, 2f, 4s, and 4f (s and f stand for sliding and fixed bitmasks respectively). Since the decompression engine must be able to start execution from any of jump targets, branch targets should be aligned in the compressed code. In addition, the mapping of old addresses (in the original uncompressed code) to new addresses (in the compressed code) is kept in a jump table.

3. COMPRESSION-AWARE DCR

It is a major challenge to optimize both performance and energy consumption simultaneously. In case of DCR, tradeoffs between performance and energy consumption should be considered in order to choose the most profitable cache configuration for each application. Fig. 5 shows an example of performance-energy consumption tradeoff using *Anagram* benchmark. Each dot represents a cache configuration showing its corresponding energy consumption and total execution time of the task. By plotting all cache configurations in performance-energy consumption graph (based on time and energy consumption from simulation results) we can determine Pareto optimal points representing feasible alternatives. For instance, increasing cache line or associativity can improve performance and may increase energy consumption as well. High performance alternatives will sacrifice some amounts of energy while selecting energy saving options would have lower performance. The remainder of this section describes how to combine the advantages of both compression and dynamic reconfiguration.

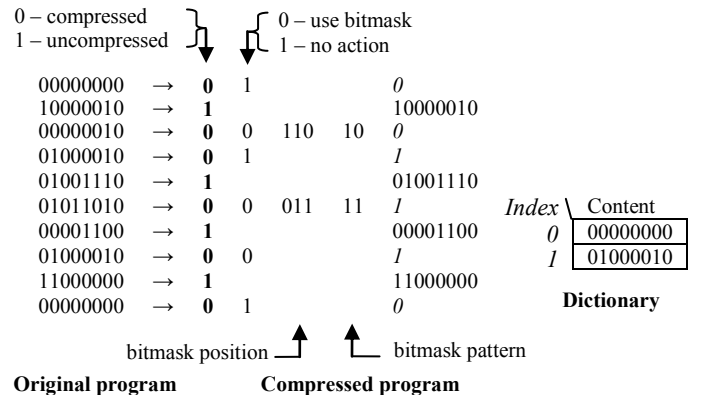


Fig. 4: An example of bitmask-based code compression

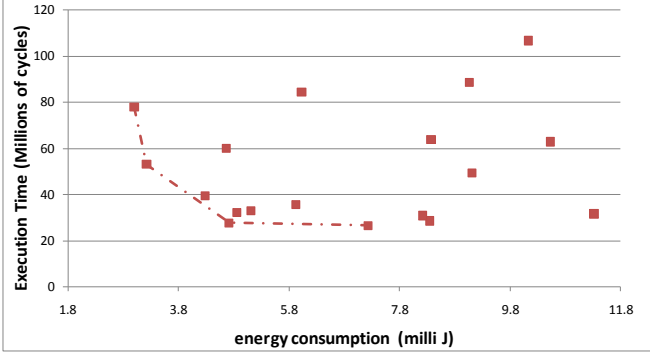


Fig. 5: An example of performance-energy consumption tradeoff using *Anagram* benchmark (Pareto optimal alternatives are connected using dashed lines)

3.1 Motivation

A reconfigurable cache can be viewed as an elastic cache with flexible parameters such as cache size, line size, and associativity. The dynamic reconfiguration technique exploits the elasticity of such caches by selecting a profitable cache configuration which is capable of maintaining the critical portion of the application to reduce energy consumption. Choosing smaller caches that fail to store the critical portion of the program may lead to increased cache misses thus longer execution time and eventually escalation in energy consumption. However, it is possible that the cache reconfiguration method may find a cache configuration that increases the execution time of the application in spite of reduced energy consumption. This may not be an issue for systems without real-time constraints but timing constraints in real-time applications limit use of such cache reconfiguration techniques. Integrating code compression with cache reconfiguration resolves this problem by effectively shrinking the program size in order to fit the critical portion of the application into a smaller cache.

Fig. 6 illustrates different caches for a real-time embedded system with a set of applications. Associativity is ignored for the ease of illustration. The horizontal and vertical axis show different possibilities of cache size and line size, respec-

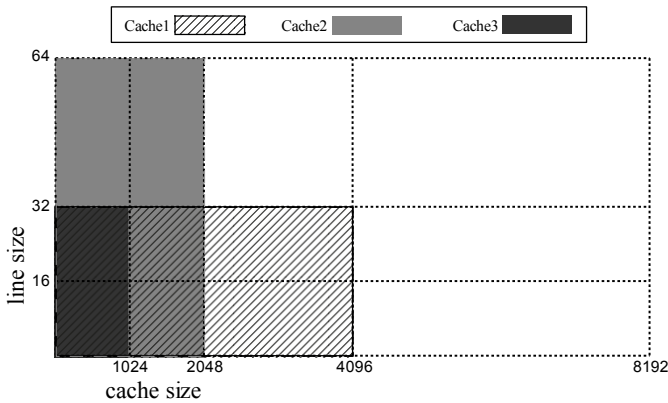


Fig. 6: Different caches used in different scenarios
Cache1: conventional system without reconfiguration,
Cache2: only dynamic reconfiguration (no compression),
Cache3: both dynamic reconfiguration and compression

tively. *Cache1* is the optimized *base cache* chosen for the system which is used for all applications and will have the minimal energy consumption while ensuring that no deadlines will be missed. *Cache2* is the cache selected by dynamic reconfiguration technique (with no compression) to reduce the energy consumption of this application. But to ensure real-time (deadline) constraints, low energy cache alternatives may get rejected because of longer execution times (critical portion of applications may not fit, for example). Incorporating compression into DCR would lead to selection of *Cache3*. Applying compression will help dynamic reconfiguration to perfectly fit the critical portion of the application into smaller cache thus gaining even more energy savings without increasing the execution time.

3.2 Compression-Aware DCR

Algorithm 1 outlines the major steps in our cache configuration selection in the presence of compressed applications. The algorithm collects simulation results for all possible cache configurations (cache sizes of 1KB, 2KB, 4KB, and 8KB; associativity of 1, 2, 4-way; cache line sizes of 16, 32, 64). It finds the best energy optimal cache configuration for each application through exhaustive exploration of all possible cache configurations of $C = \{c1, c2, c3, \dots, c_m\}$. Number of simulation cycles for each run is collected based on the simulation results. The energy model of [6] is used to calculate the energy consumption using the cache hit and miss statistics. The algorithm finally constructs the Pareto optimal alternatives and returns it in a list. The most energy efficient cache configuration among all Pareto optimal alternatives which satisfies timing requirements of the application is chosen next. Suppose there are two cache configurations, C1 with execution time of 2 million cycles and energy consumption of 5 mJ and C2 with execution time of 1.8 million cycles and energy consumption of 6 mJ, available in the Pareto optimal list of alternatives. If the task has to be done in 1.9 million cycles, the faster alternative (C2) gets chosen. If the timing requirement of the task is not constrained by 2 million cycles, the more energy efficient cache alternative (C1) gets selected.

The algorithm is similar to traditional DCR but uses compressed code. Therefore the simulation/profiling infrastructure needs to have decompression unit to provide the

Algorithm 1: Finding Pareto optimal cache configurations

Input: Compressed code

Output: List of Pareto optimal cache alternatives

Begin

 li = an empty list to store cache alternatives

for s = cache sizes of 1KB, 2KB, 4KB, and 8KB **do**

for a = associativity of 1, 2, 4-way **do**

for l = cache lines of 16, 32, 64 **do**

 do cycle accurate simulation for cache $C_{s,a,l}$;

$t_{s,a,l}$ = simulation cycles;

$e_{s,a,l}$ = energy consumption of the cache subsystem;

 add the triple $(C_{s,a,l}, t_{s,a,l}, e_{s,a,l})$ to li;

end for

end for

end for

return Pareto optimal points in li;

end

ability of decoding compressed instructions. For example, in our case, we implemented and placed the required decompression routines/functions for respective compression algorithms in *SimpleScalar* simulator [16].

Since we consider systems with only one level of reconfigurable cache architecture, number of cache reconfigurations is small. So we can exhaustively explore all possible configurations in a reasonable time. Since the reconfiguration of associativity is achieved by way concatenation as described in Section 3.1, 1KB L1 cache can only be direct-mapped as other three banks are shut down. For the same reason, 2KB cache can only be configured to direct-mapped or 2-way associativity. Therefore, there are 18 ($=3+6+9$) configuration candidates for L1.

3.3 Placement of Decompression Hardware

Fig. 7 shows two different placement of the decompression unit. In pre-cache placement the memory contains compressed code and instructions are stored in cache in original form. Whereas, in the post-cache placement the decompression unit is placed between cache and processor thus both memory and cache contain compressed instructions.

Our studies show that having the pre-cache placement has very little effect on energy and performance of cache. In this case uncompressed instructions are stored in the cache and when cache miss occurs, the cache controller asks the decompression unit to provide a block of instructions. In majority of the cases the decompression hardware requires one clock cycle in pipelined mode (as shown in Fig. 7), so one clock cycle will be added to the latency of entire block fetch. In rare cases, e.g., when the first instruction of the block is not compressed, it will introduce two cycle penalty since it will take two cycles to fetch and decompress the instruction [17]. As demonstrated in Fig. 8, the energy consumption of cache in the pre-cache placement is almost the same as the case when there is no compression involved. So the best choice is to use post-cache placement to achieve maximum performance as well as minimum energy consumption.

Incorporating compression, cache miss penalty caused by memory fetch latency is reduced because of improved bandwidth (since compressed code is smaller). In addition, off-chip access energy (the buses to main memory and memory access) is also reduced since the decompression engine reads compressed code from memory resulting in lower traffic to main memory. However, post-cache placement can introduce significant performance overhead to the system. Seong et al. [13] presented a bitmask-based compression technique that adds no penalty to the system performance using pipelined one-cycle decompression engine with negligible power requirement. Using this decompression engine makes it practical to place

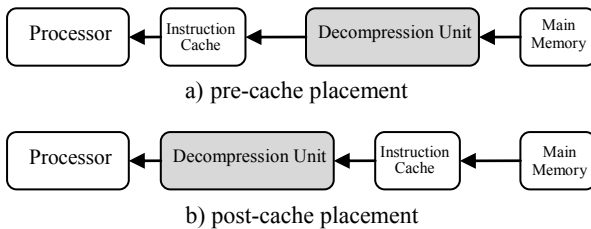


Fig. 7: Different placement of decompression unit

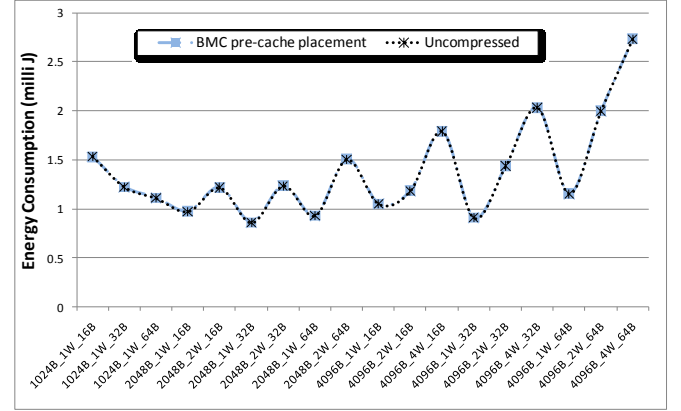


Fig. 8: The impact of pre-cache placement of decompression engine on cache energy – *djpeg* benchmark

the decompression unit after cache (post-cache placement) and benefit from the compressed code stored in the cache.

In the context of embedded systems one of the main goals is maximizing energy savings while ensuring the system will meet applications requirements. Usually, choosing a cache configuration for energy savings may result in performance degradation. However, the synergistic combination of cache reconfiguration and code compression enables energy savings without loss of performance. Our proposed methodology provides an efficient and optimal strategy for cache tuning based on static profiling using compressed programs.

4. EXPERIMENTS

4.1 Experimental Setup

In order to quantify compression-aware cache configuration tradeoffs, we examined *cjpeg*, *djpeg*, *epic*, *adpcm* (*rawaudio* and *rawaudio*), *g721* (*encode*, *decode*) benchmarks from the MediaBench [18] and *dijkstra*, *patricia*, *crc32*, *bitcnts* from MiBench [19] compiled for the Alpha target architecture. All applications were executed with the default input sets provided with the benchmarks suites.

Three different code compression techniques including bitmask-based, dictionary-based and Huffman code compression were used. To achieve the best attainable compression ratios, in bitmask-based compression, for each application we examined dictionaries of 1 KB, 2KB, 4KB, and 8 KB. Similar to Seong et al. [13] we tried three mask sets including one 2-bit sliding, 1-bit sliding and 2-bit fixed, and 1-bit sliding and 2-bit fixed masks. Similarly for dictionary-based and Huffman compression we used 0.5 KB, 1KB, 2KB, 4KB, and 8 KB dictionary sizes with 8 bits, 16 bits and 32 bits word sizes. We found out that dictionary size of 2 KB and word size of 16 bits are the best choices for this set of benchmarks. The reason is that using 8 bits words increases the number of compression decision bits and using 32 bits word size decreases the words frequencies significantly. Hence, as simulation results showed, 16 bits word size is the best choice.

Code compression is performed offline. In order to extract the code (instruction) part from executable binaries, we used ECOFF (Extended Common Object File Format) header files provided in SimpleScalar toolset [16]. We placed the

compressed code back into binary files so that they can be loaded into the simulator.

We utilized the configurable cache architecture developed by Zhang et al [6] with a four-bank cache of base size 4 KB, which offers sizes of 1 KB, 2 KB, and 4 KB, line sizes ranging from 16 bytes to 64 bytes, and associativity of 1-way, 2-way, and 4-way. For comparison purposes, we used the *base cache* configuration set to be a 4 KB, 4-way set associative cache with a 32-byte line size, a reasonably common configuration that meets the average needs of the studied benchmarks.

To obtain cache hit and miss statistics, we modified the *SimpleScalar* toolset [16] to decode and simulate compressed applications. We implemented and placed the required decompression routines/functions for respective compression algorithms in *SimpleScalar* simulator. We considered the latency of decompression unit carefully. Decompression unit can decompress the next instruction in one cycle (in pipelined mode) if it finds the entire needed bits in its buffer. Otherwise, it takes one cycle (or more cycles, if cache miss occurs) to fetch the needed bits into its buffer and on more cycle to decompress the next instruction. Correctness of the compression and decompression algorithms was verified by comparing the outputs of compressed applications with uncompressed versions. The performance overhead of decompression includes decompression unit buffer flush overhead due to jumps, and variable latency of memory reads in each block fetch (because of variable length compressed code). These overhead are negligible according to the experimental results.

We applied the same energy model used in [6], which calculates both dynamic and static energy consumption, memory latency, CPU stall energy, and main memory fetch energy. The energy model was modified to include decompression energy. We updated the dynamic energy consumption for each cache configuration using CACTI 4.2 [20]. Utilizing Perl scripts, the design space of 18 cache configurations is exhaustively explored during static analysis to determine the performance, and energy-optimal cache configurations for each benchmark.

4.2 Energy Savings

Energy consumption for several benchmarks from the MediaBench and MiBench in different approaches are analyzed: a fixed *base cache* configuration, bitmask-based compression without utilizing DCR (BMC only), DCR without compression (DCR only), dictionary-based compression with DCR (DC+DCR), Huffman coding with DCR (Huffman+DCR), and bitmask-based compression with DCR (BMC+DCR). The most energy efficient cache configuration

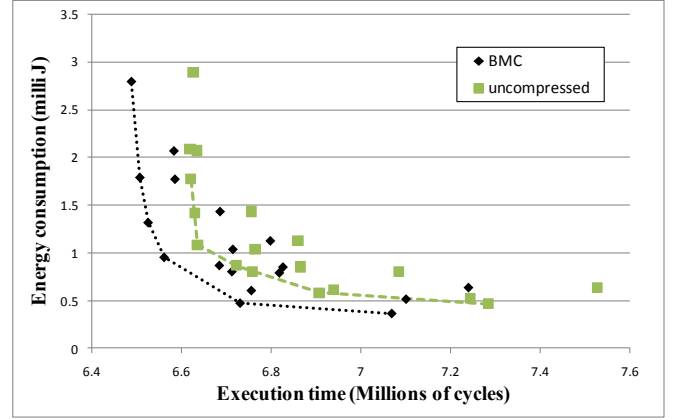


Fig. 10: Performance-Energy consumption tradeoff for compressed and uncompressed codes using rawaudio (adpcm-enc) benchmark

found by exploration in each technique is considered for comparison. Fig. 9 presents energy savings for the instruction cache subsystem. Energy consumption is normalized to the fixed *base cache* configuration such that value of 100% represents our baseline. Energy savings in the instruction cache subsystem ranges from 10% to 76% with an average of 52% for utilizing only DCR. As we expected, due to higher decompression overhead, Huffman (when combined with DCR) achieves lower energy savings compared to BMC virtually for all benchmarks. Energy savings in DC+DCR approach are even lower than Huffman+DCR as a result of moderate compression ratio by DC. Incorporating BMC in DCR increases energy savings up to 48% – on top of 10% to 76% energy savings obtained by DCR only – without any performance degradation. Our methodology achieves on average 65% energy savings of the cache subsystem.

Fig. 10 illustrates an example of performance-energy consumption tradeoffs for both uncompressed and compressed (using BMC) cases for *rawaudio* (*adpcm-enc*) benchmark. It can be observed that for every possible configuration for the uncompressed program there is an alternative which has a better performance and lower energy requirement if the program is compressed. This observation shows that compression-aware DCR leads to better design choices.

Another observation we have made is that without DCR, applying compression on an application (which executes using *base cache* configuration that already fits the critical portion of the application) will not gain noticeable energy savings. However, compression-aware DCR effectively uses the ad-

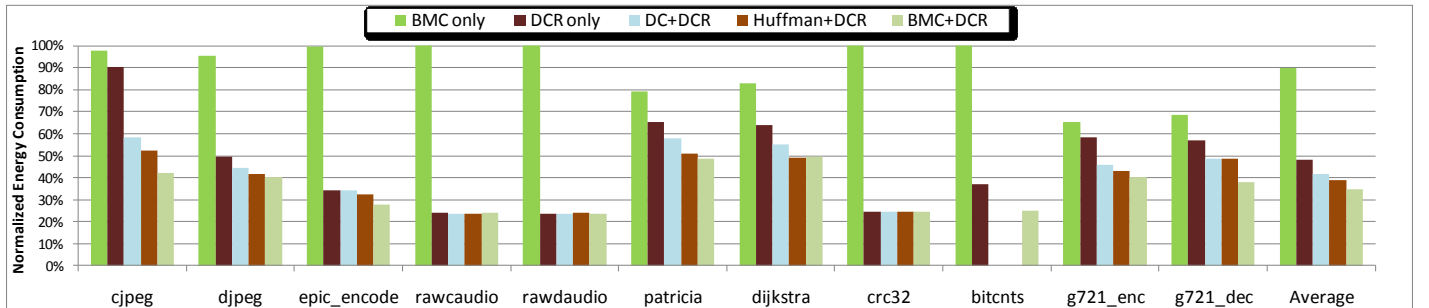


Fig. 9: Energy consumption of the selected "minimal-energy cache" normalized to the base cache

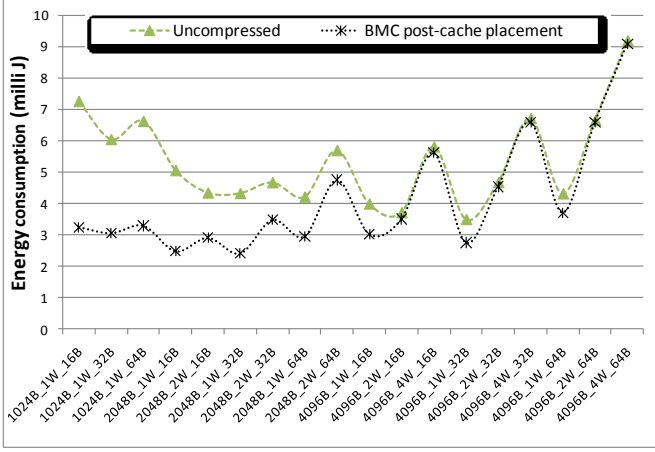


Fig. 11: The impact of cache/line size on energy profile of cache using cjpeg benchmark

vantage of reduced program size achieved by compression to choose smaller cache size, associativity, or line size and yet fit critical portion of programs. Therefore, compression aware-DCR can achieve more energy savings compared to DCR alone. Fig. 11 illustrates comparison of energy profile for different caches for compressed (using BMC) and uncompressed cjpeg benchmark. Using a 4KB cache with associativity of 4 and 64-bit line size, energy consumption of cjpeg benchmark is nearly the same for compressed and uncompressed programs.

In the post-cache placement, compression has a significant effect when combined with small cache sizes. In this case compressed instructions are stored in the cache. Since the compressed code size is 30 to 45 percent less than uncompressed code it can fit in smaller cache sizes. However, when size of the selected cache increases, the critical portion of program (regardless of whether compressed or not) will fit into cache entirely. Therefore by utilizing large cache sizes energy consumption of the compressed code is very close to uncompressed one. It should be noted that the main objective of exploration is to find the most energy efficient cache configurations so we are not interested in large cache sizes since they require more energy.

4.3 Performance Improvements

Fig. 12 shows performance of applications for different schemes normalized to the *base cache*. Applying DCR alone for the purpose of energy saving, results in 11% performance loss on average. We observe that code compression can im-

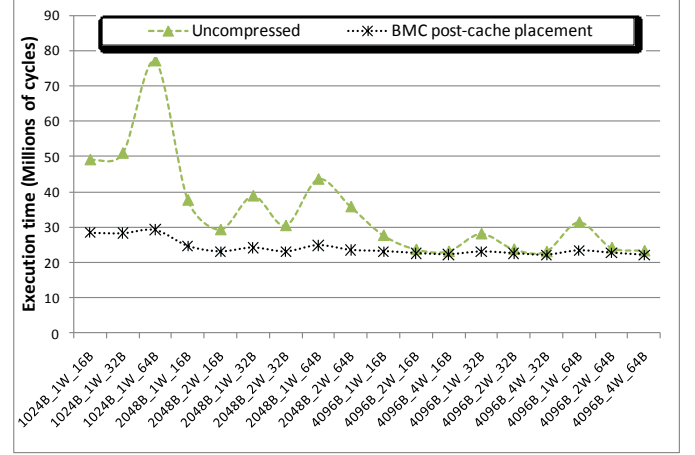


Fig. 13: performance trend of different cache configurations using cjpeg benchmark

prove performance in many scenarios while achieving significant reduction in energy consumption. For instance, in the case of the application *patricia*, applying only DCR would result in 12% performance degradation with 34% energy savings. However, incorporating BMC boosts performance by 33% while gaining extra 17% energy savings on top of DCR achieving 51% energy savings compared to the *base cache*. Results show that synergistic integration of BMC with DCR achieves as much as 75% performance improvement for *g721_enc* (21% improvement on average) compared to DCR alone. Thus it is possible to have a cache architecture that is tuned for applications to have both increased performance as well as lower energy consumption.

Fig. 13 shows performance trend of all cache configurations for both uncompressed and compressed codes for *cjpeg* benchmark. It is interesting to note that compression also improves performance. The compressed program can fit in smaller cache because of 30 to 45% reduction in code size. This decreases cache misses significantly for small caches. Reduced number of misses can lead to reduced stalls and improved performance. As it can be observed in Fig. 13, without compression, reducing the cache size may lead to major performance degradation so DCR is forced to discard many cache alternatives due to timing constraints. For instance, having timing constraint of 25 million cycles for *cjpeg* benchmark will force to discard all cache configurations of size 2048KB or lower. However, compression improves the performance significantly when small cache sizes are used. Thus combination of cache reconfiguration and code compression enables energy savings while improving overall performance.

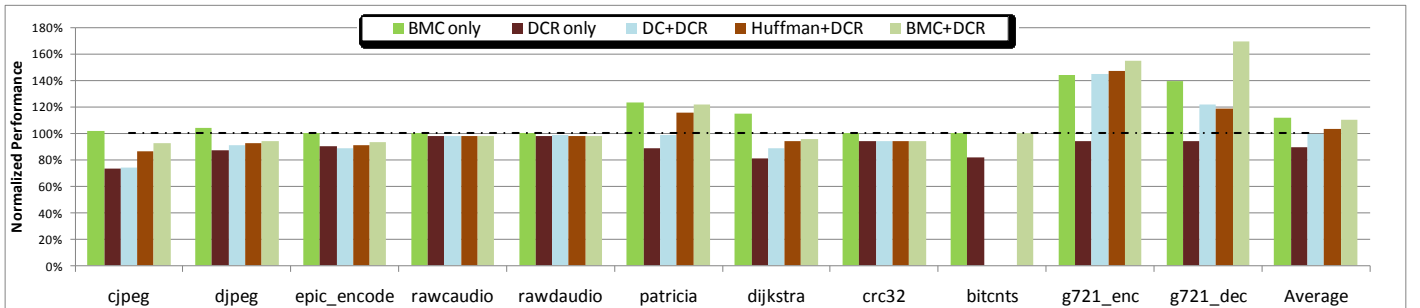


Fig. 12: Performance of the selected "minimal-energy cache" normalized to the base cache

5. CONCLUSION

Optimization techniques are widely used in embedded systems to improve overall area, energy and performance requirements. Dynamic cache reconfiguration (DCR) is very effective to reduce energy consumption of cache subsystem. Code compression can significantly reduce memory requirements, and may improve performance in many scenarios. In this paper, we presented a synergistic integration of DCR and code compression for embedded systems. Our methodology employs an ideal combination of code compression and dynamic tuning of cache parameters with minor or no impact on timing constraints. Our experimental results demonstrated 65% reduction on average in overall energy consumption of the cache subsystem as well as up to 75% performance improvement (compared to DCR only) in embedded systems.

REFERENCES

- 1 A. Malik, B. Moyer, and D. Cermak. A low power unified cache architecture providing power and performance flexibility. *ISLPED* (2000).
- 2 A. Gordon-Ross, F. Vahid, N. Dutt. Automatic tuning of two-level caches to embedded applications. *DATE* (2004).
- 3 C. Lefurgy. Efficient execution of compressed programs. *Ph.D. Thesis, University of Michigan* (2000).
- 4 L. Benini, F. Menichelli, M. Olivieri. A Class of Code Compression Schemes for Reducing Power Consumption in Embedded Microprocessor Systems. *IEEE Transactions on Computers* (April 2004), p.467-482.
- 5 A. Gordon-Ross, F. Vahid, N. Dutt. Fast configurable-cache tuning with a unified second level cache. *International Symposium on Low Power Electronics and Design* (2005).
- 6 C. Zhang, F. Vahid, W. Najjar. A highly-configurable cache architecture for embedded systems. *30th Annual International Symposium on Computer Architecture* (June 2003).
- 7 P. Vita. Configurable Cache Subsetting for Fast Cache Tuning. *Design Automation Conference, DAC* (2006).
- 8 D. H. Albonesi. Selective Cache Ways: On-Demand Cache Resource Allocation (2000).
- 9 L. Chen, X. Zou, J. Lei, Z. Liu. Dynamically Reconfigurable Cache for Low-Power Embedded System. *Third International Conference on Natural Computation* (2007).
- 10 C. Zhang, F. Vahid, and R. Lysecky. A self-tuning cache architecture for embedded systems. *DATE* (2004).
- 11 A. Wolfe, A. Chanin. Executing compressed programs on an embedded RISC architecture. *Proc. of the Intl. Symposium on Microarchitecture* (1992), 81–91.
- 12 Y. Xie, W. Wolf, and H. Lekatsas. A Code Decompression Architecture for VLIW Processors. *MICRO-34. Proceedings. 34th ACM/IEEE International Symposium on Microarchitecture* (2001).
- 13 S. Seong, P. Mishra. Bitmask-Based Code Compression for Embedded Systems. *IEEE Trans. CAD* (2008), 673–685.
- 14 C. Lin, Y. Xie, and W. Wolf. LZW-Based Code Compression for VLIW Embedded Systems. *Proc. DATE* (2004), 76–81.
- 15 M. Rawlins, A. Gordon-Ross. On the Interplay of Loop Caching, Code Compression, and Cache Configuration. *ASP-DAC* (2011).
- 16 D. Burger, T. Austin, S. Bennet. Evaluating future microprocessors: the simplescalar toolset. *University of Wisconsin-Madison. Computer Science Department Technical Report CS-TR-1308* (July 2000).
- 17 C. Murthy and P. Mishra. Lossless Compression Using Efficient Encoding of Bitmasks. *isvlsi, pp.163-168, 2009 IEEE Computer Society Annual Symposium on VLSI* (2009).
- 18 C. Lee, M. Potkonjak, and W. H. Mangione-smith. MediaBench: A Tool for Evaluating and Synthesizing Multimedia and Communications Systems. *International Symposium on Microarchitecture* (1997).
- 19 M.R. Guthaus, J.S. Ringenberg, D. Ernst, T.M. Austin, T. Mudge, and R.B. Brown. MiBench: A free, commercially representative embedded benchmark suite. *International Workshop on Workload Characterization (WWC)* (2001).
- 20 CACTI. HP Labs, CACTI 4.2, <http://www.hpl.hp.com/>.
- 21 T. Sherwood, E. Perelman, G. Hamerly, S. Sair, and. Discovering and exploiting program phases. *Micro* (2003).
- 22 W. Wang and P. Mishra. Dynamic Reconfiguration of Two-Level Caches in Soft Real-Time Embedded Systems. *IEEE International Symposium on VLSI* (2009).
- 23 S. Ramprasad, N.R. Shanbhag, and I.N. Hajj. A Dictionary-Based En/Decoding Scheme for Low-Power Data Buses. *Very Large Scale Integration (VLSI) Systems* (1999).
- 24 H. Lekatsas and W. Wolf. SAMC: A Code Compression Algorithm for Embedded Processors. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* (1999).