Adaptive Cache Infrastructure: supporting dynamic program changes following dynamic program behavior

Fabian Nowak, Rainer Buchty, and Wolfgang Karl
{nowak|buchty|karl}@ira.uka.de

Institut für Technische Informatik (ITEC)
Universität Karlsruhe (TH)

Abstract. Recent examinations of program behavior at run-time revealed distinct phases. Thus, it is evident that a framework for supporting hardware adaptation to phase behavior is needed. With the memory access behavior being most important and cache accesses being a very big subset of them, we herein propose an infrastructure for fitting cache accesses to a program's requirements for a distinct phase.

1 Introduction

When regarding a program's long-time behavior, it is evident that each program consists of at least three different phases: the first one can be called the *initialization phase*, the second one the *main* or *computational phase*, and the last one the *final* or *termination phase* [1]. It can even be shown that after millions of instructions in the so-called *main phase*, new phases of program execution commence [2]. Metrics changing from phase to phase include, but are not limited to, branch prediction, cache performance, value prediction, and address prediction [3, 4].

By providing an architecture tailored at only one phase, as is done usually, this very phase is executed with best results concerning the aspired enhancements, i.e. performance or energy efficiency. By means of reconfiguration, however, we are able to support a program during its whole run-time when adapting the hardware in every distinct phase.

In this paper, we address the basic concept, present our latest implementation results, and show in detail how much cache reconfiguration can possibly speed up program execution by giving benchmark [5] results. Furthermore, we show a simple means to handle phase changes more dynamically.

This is especially interesting for scientific super-computing where slight enhancements of the whole system can result in some fewer days of computation or less energy consumption, and therefore, cooling and money savings. Another interesting aspect is to further speed up execution of parallel computing nodes sharing some memory levels, like L2 cache and main memory, by dynamically partitioning a cache's area to the different processors' needs.

The rest of this paper is organized as follows: as a start, an overview of the finished and ongoing work is given. Upon that base, in Section 3 we present a novel architecture for supporting cache reconfiguration at reunite. The current state of our implementation is then summed up in Section 4 giving a first impression of how much speed-up can be achieved. This is explained in detail in Section 5. In order to round the whole work off, an application supporting the tool-chain is illustrated. The paper finishes with the conclusion.

2 Related work

Much work was done already in the vast field of cache partitioning and cache adaptation. As far as cache partitioning is concerned, dividing into instruction and data caches is a very well-known method for increasing cache hit-rates. Other methods are partitioning into scalar and array data caches [6] or separating by temporal and spatial locality [7]. The last approach was extended by Johnson and Hwu by means of memory address tables yielding a speed-up of up to 15% [8]. They request a framework for intelligent run-time management of the cache hierarchy. Partitioning can also be used for offering instruction reuse buffers based on cache technology as is done by Ranganathan et al. in [9]. Unfortunately, some software changes have to be made for the system to work. Suh et al. examined dynamic partitioning of shared cache memory and come to the conclusion that re-partitioning only needs to take place when a context switch occurs.

Cache adaptation, which may benefit of cache partitioning, requires the micro-system to monitor its memory system performance, detect changing demands, and initiate reconfiguration of at least a subpart of the memory system. Benitez et al. evaluate performance and energy consumption of a cache system, which is very limited concerning possible parameter changes. With their micro-architecture of the Field-Programmable Cache Array (FPCA), they also introduce basic phase detection where branches are counted as a simple means to recognize possible changes related to the memory system [10]. When reconfiguring, any cache content is lost, thus reconfiguration has to be controlled with the processor initiating a cache flush in advance.

More sophisticated phase detection is achieved by Sherwood et al. based on basic block vectors and clustering [11]. Similarly, Balasubramonian and Albonesi managed phase detection by measuring branch frequencies and cache misses. Upon this information, cache parameter adjustment is carried out, such as sizes, associativities and access latencies [4]. In their latest work, Huffmire and Sherwood applied wavelet-based phase detection onto L1 accesses and are able to accurately predict L2 miss rates, and thus, phases [12].

With their work towards reconfigurable cache memory, Ogasawara et al. proved that varying certain cache parameters indeed makes sense [13]. Despite delivering first results, their system is only tailored at simulation and very limited with respect to variety of dynamic parameters.

In the embedded field, Vahid and Gordon-Ross among others are developing configurable cache controllers with a strong focus on embedded applications and energy consumption [14–17].

A concept of the required tool-chain for a reconfigurable system was worked out by Mei et al. [18]. It consists of a two-level approach that is based on profiling and mapping and should be adaptable to the hardware infrastructure under development at our chair.

3 Reconfigurable Cache Architecture

The first goal of our architecture was the creation of a cache capable of acting both as Level-1 and Level-2 cache. Secondly, it has to be reconfigurable. As a last aspect, we want the cache controller to be synthesizable for hardware usage.

While most reconfigurable caches depend on the reconfiguration capabilities of the underlying FPGA chip, our implementation handles the reconfiguration logically, only by hardware logic in the controller itself. This decision offers a great degree of freedom in choosing different sizes for both the cache memory and its control and replacement information, in running with different associativities and in changing replacement strategies.

The cache only caches accesses to the ordinary main memory, which consists of DDR-SDRAM on the Xilinx ML310 and ML403 evaluation boards available at our institute.

3.1 Cache Structure

As already mentioned, we decided to split the whole cache into three distinct areas: *cache memory*, *control memory* and *replacement memory*. The *control memory* indicates whether the according cache line is valid, modified, partly valid and additionally stores the tag. The *replacement memory* is only needed, when an associativity of more than one and a more sophisticated replacement strategy like LRU is used. In Figure 1, the layout of a 2-way set-associative cache with eight lines and a replacement strategy like FIFO, Pseudo-LRU, or LRU is illustrated.

3.2 Cache Controller

The *cache controller* (cf. Figure 2) handles several things: not only is it responsible for serving read/write requests to the cache memories explained above, but also for initiating the reconfiguration process, which in return is executed by the *reconfiguration controller*.

3.3 Interfaces

The new cache controller is surrounded by four different external interfaces, which will be explained in more detail below.

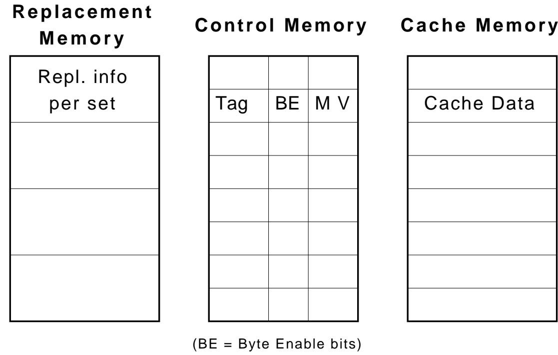

Fig. 1. Exemplary cache layout

- **DCR Bus:** The DCR Bus is used for reconfiguration purpose. A master device, e.g. the processor, sends parametrization values after a successful handshake.
- **PLB:** This is where the data/instruction requests are received from and served.
- **DDR-SDRAM:** Data unavailable in the caches is found in the DDR-SDRAM.
- **Monitoring:** For implementing phase detection and in order to create traces, an interface indicating hits/misses, the affected cache line, and the originating memory address is available.

3.4 Reconfiguration

Reconfiguration of the cache is achieved by requesting the controller to enter the reconfiguration state, then writing new parametrization values, explicitly indicating the end of the reconfiguration request, and waiting for the reconfiguration process to finish. Meanwhile, the processor could do different tasks not involving the cached main memory. But for ease of implementation and in order to avoid additional code for busy waiting, we decided to simply halt the processor. The complete process is illustrated in Figure 3.

In [19], we already proved that several traditionally fixed cache parameters are reconfigurable and presented some implementation approaches using heuristics where necessary.

3.5 Monitoring

By providing a 64 bit wide output register filled with information about cache hits and misses, the conflict-causing line number, and additionally the complete memory address, it is possible to both collect information for on-line/offline analysis and achieve phase prediction such that the processor can initiate cache reconfiguration itself or have the system automatically adopt to the detected phase changes. The monitor output register is depicted in Figure 4.

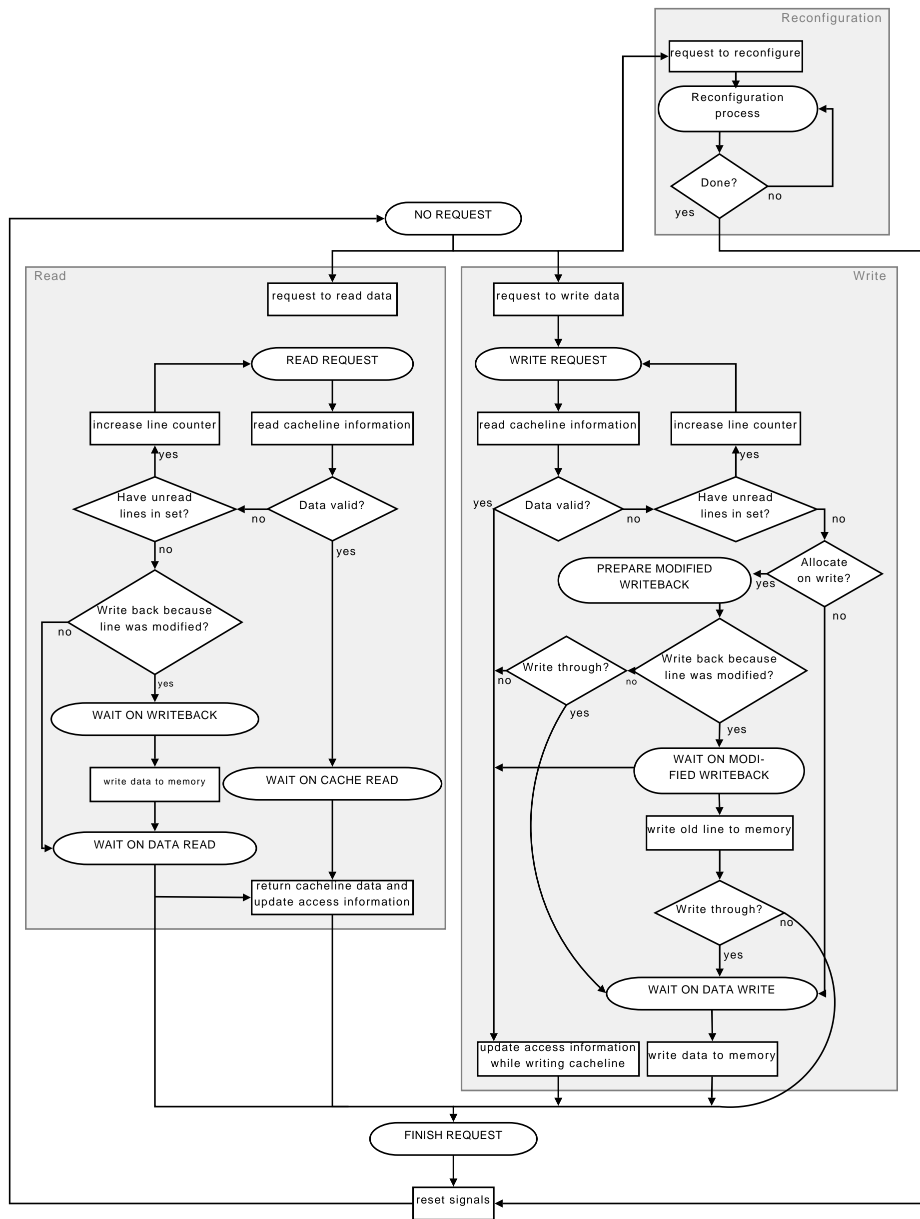

Fig. 2. Cache Controller

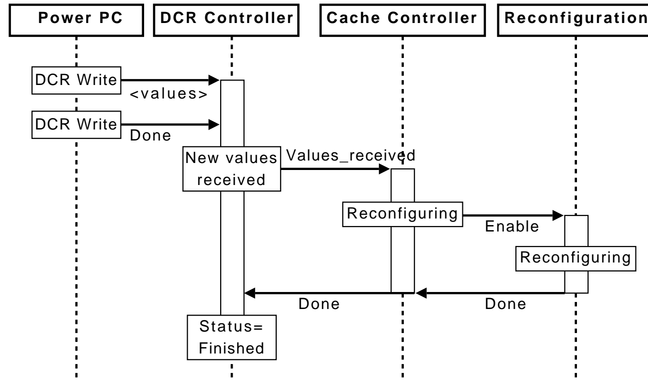

Fig. 3. Reconfiguration flow

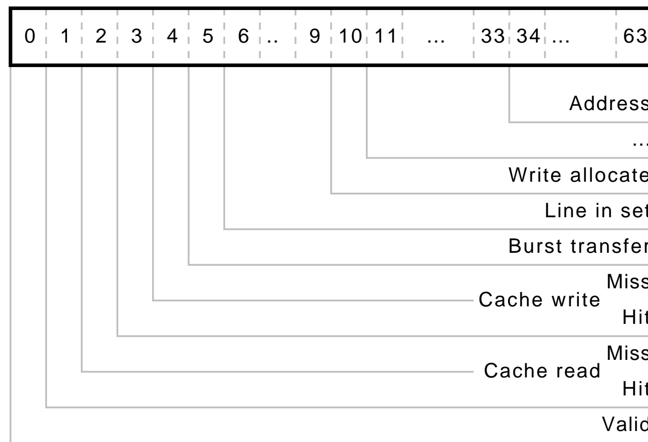

Fig. 4. Monitor output register: the unused bits are intentionally reserved for further extensions to the current monitoring interface.

3.6 Putting It All Together

In Figure 5, the complete extended cache/memory controller is depicted with both, the four interfaces, and the sub-modules controlling these interfaces or reading from them.

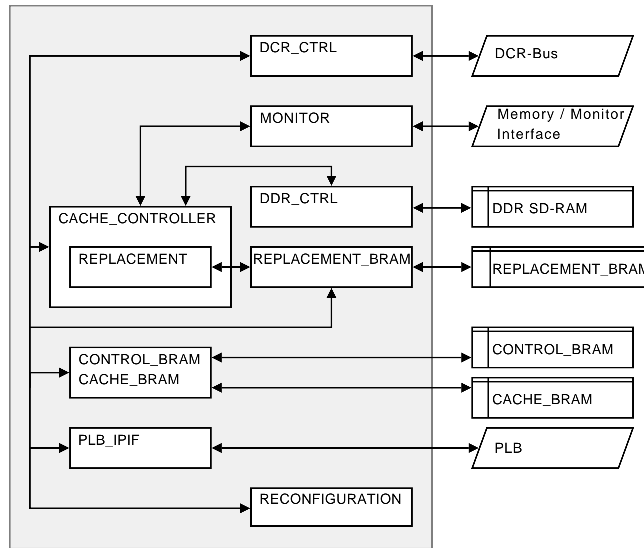

Fig. 5. The big picture: the cache controller is the central part of the reconfigurable cache, communicating by means of supplemental modules with its external interfaces.

4 Implementation Results

We are currently in process of synthesizing the whole design for the Virtex-II Pro and the Virtex-4 FX12 in order to benchmark the system in real hardware. First results show that some modifications still have to be made to the design; especially, the maximum clock rate is only at 9.779 MHz. This is due to the fact that the processor clock is also used for lots of comparisons for memory accesses dependent on the chosen associativity and number of sets.

In Table 1, we present the hardware resource usage of parts of the current implementation tailored at the Virtex-II Pro, that is to say of the cache controller, the reconfiguration module, the DCR controller, the monitoring component and the basic PLB DDR Components. Due to long synthesis times, further results will still require some weeks of synthesis. Despite all, the results look quite promising, although we expect that we have to further limit the design for real application. In fact, synthesizing for bigger chips, which unfortunately are yet unavailable, would already lead to working hardware.

Logic Utilization	**Used**	**Available**	**Utilization**
Total Number Slice Registers	2089	27392	7.6%
Number used as Flip Flops	1826		
Number used as Latches	263		
Number of 4 input LUTs	37980	27392	138.7%
Logic Distribution			
Number of occupied Slices	20212	13696	147.7%
Total Number of 4 input LUTs	38660	27392	141.1%
Number used as logic	37980		
Number used as a route-thru	680		
Number of MULT18X18s	2	136	1.471%
Number of GCLKs	10	16	62.5%
Total equivalent gate count for design	268036		
Additional JTAG gate count for IOBs	79488		

Table 1. Hardware resource usage of synthesized components.

5 Application Speed-Up

Given the benchmark results from [5], we will show how much speed-up can be obtained.

First, we write down the rather simple formula for calculating whether reconfiguration is appropriate and for comparing the overall memory access time with reconfiguration to the overall memory access time without adaptation to program phases.

$$\begin{aligned} & n_{wm} * t_{wm} + n_{wh} * t_{wh} + n_{rm} * t_{rm} + n_{rh} * t_{rh} > \\ & n'_{wm} * t_{wm} + n'_{wh} * t_{wh} + n'_{rm} * t_{rm} + n'_{rh} * t_{rh} + t_{reconfiguration} \end{aligned} \quad (1)$$

where n_{wm} indicates the number of write misses, t_{hr} the time required for serving a read request when a hit in the cache occurs, and n' the numbers achieved after reconfiguration. As can be seen easily, the overall memory access time is the sum of the cycles needed in all program phases plus their respective reconfiguration times.

Then we want to outline the time needed for both cache accesses and reconfiguration time measured in clock cycles. Table 2 gives a comparison of the implemented memory access times with and without our cache. Obviously, most speed-up can be achieved by increasing the read hit rate, while the write accesses do not contribute too much.

Already upon that basis, we are able to extend Equation 1 to the following for a one- and two-way set-associative cache:

$$\begin{aligned} & c * (p_{wm} * 10 + p_{wh} * 8 + p_{rm} * 15 + p_{rh} * 8) > \\ & c' * (p'_{wm} * 10 + p'_{wh} * 9 + p'_{rm} * 15 + p'_{rh} * 9) + t_{reconfiguration} \end{aligned} \quad (2)$$

where $n_\alpha = p_\alpha * c$, c the number of memory access cycles “per phase”, and $\sum_i p_i = 1$. Of course, we have to pay attention to use the best case access times

Access type	Duration w/o cache	Duration w/ cache
Read Hit	15	$8 + \lfloor n/2 \rfloor$
Read Miss	15	$15 + \lfloor (a-1)/2 \rfloor$
Write Hit	9	$8 + \lfloor n/2 \rfloor$
Write Miss	9	$10 + \lfloor (a-1)/2 \rfloor$

n denotes the line in the set where the hit occurs; a denotes the level of associativity.

Table 2. Access times to main memory without and with cache (using write-through)

for the first part, while in contrast the worst case has to be regarded for the second part.

Regarding reconfiguration time, it must be noted that for area and memory concerns, the reconfiguration of associativity only saves half of the cache's content. This decision makes reconfiguration up to 37% faster. The process is thus split into two phases with the first one being responsible for assuring cache consistency of the "rear half" by executing write-back where necessary and the second one moving cache lines to their new locations. The first step takes either 2 or 12 cycles per cache line, depending on whether write-back is needed. If write-through is used for write accesses, this step is only responsible for invalidating the "rear" cache lines. The second phase then takes 3 cycles per cache line for rearranging cache line data (this is indeed where reconfiguration would need double the time if the whole cache content was kept).

Hence, for doubling associativity, a one-way set-associative cache with 1024 lines and write-through strategy requires $512*2+512*3 = 2560$ cycles (ignoring the few setup cycles of each reconfiguration step).

Now, the memory access count stays the same when executing a disinct phase with a different configuration, thus, $c = c'$. In addition, we assume an enhancement of cache hit rate of 20% as stated in [14] when going from one-way to two-way associativity. We further assume $p_{wm} = p_{wh} = p'_{wm} = p'_{wh} = 0.25$, $p_{rh} = p_{rm} = 0.25$. With this enhancement, we get $p'_{rh} = 0.3$ and accordingly $p'_{rm} = 0.2$. Equation 2 therefore becomes

$$\begin{aligned} c*(0.25*10+0.25*8+0.25*15+0.25*8) &> \\ c*(0.25*10+0.25*9+0.2*15+0.3*9) &+2560 \qquad (3) \\ &\Leftrightarrow \\ c*(0.25*7+0.5*15-0.3*9) &> 2560 \\ &\Leftrightarrow \\ c > 2560/(1.75+7.5-2.7) &= 390.84 \qquad (4) \end{aligned}$$

Hence, after only 391 memory accesses, the reconfiguration effort proved sensible.

6 Toolchain Integration

In order to simplify the correct and consistent parameterization of the rather complex-to-configure cache-controller, we developed a program with a graphical user interface (cf. Figure 6), which automatically adjusts all parameters with respect to the user's demand.

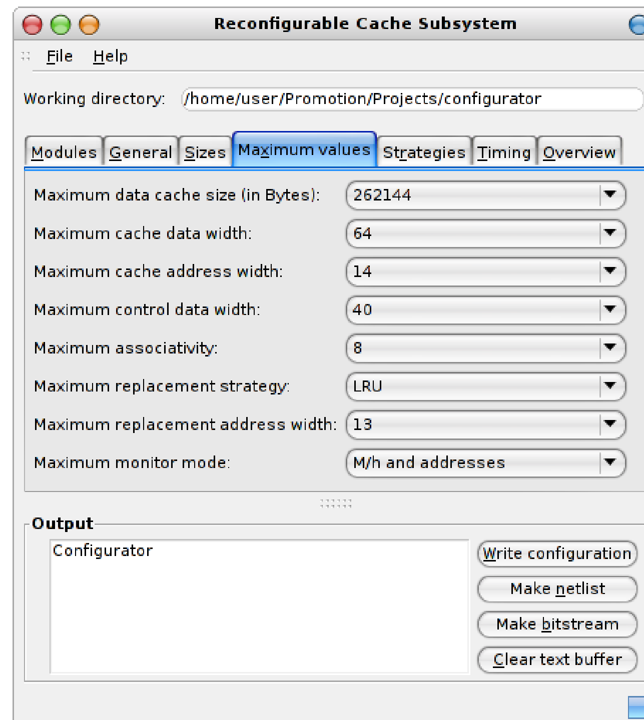

Fig. 6. Graphical User Interface for a-priori configuration of the cache/memory system

The *configurator* program has to be invoked manually after having designed the system-on-chip in *Xilinx Platform Studio (XPS)* [20]. It then enables the user to e.g. specify a maximum associativity of 4, while in the beginning of the execution, an associativity of 2 is chosen. Furthermore, it ensures that the data widths of the replacement and control information are wide enough to store all information required. Having written the changes, the user can return to *XPS* and undertake remaining changes, synthesize the description, and download the bitstream to the device. The configuration task flow is depicted in detail in Figure 7.

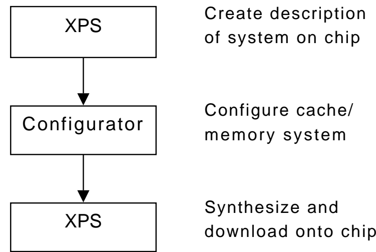

Fig. 7. Task flow for initial configuration of the cache/memory system

7 Conclusions and Future Work

We have presented an architecture concept for supporting exploitation of program phase behavior by adapting a subset of the memory system – the cache system – to a program's needs. With such an architecture, it is possible to not only statically reconfigure the system every n instructions, but to have the operating system decide itself whether it seems advantageous to use another configuration. This is where the particular strength of our approach becomes obvious: the time needed for reconfiguration can be clearly estimated between lower and upper bounds, which do not differ too much. In contrast, when using dynamic FPGA reconfiguration, parts of the system have to be halted and due to side effects, the reconfiguration time may vary unpredictably.

Future work includes the development of a monitoring component, which is capable of indicating new phases, and operating system functions for phase and system evaluation by use of the gathered monitoring information and application-embedded hints. These hints can easily be determined manually by visualization tools like in [21] and inserted into binary code.

We also intend to work on faster main memory access by a more direct connection of the memory to the processor and on using wider cache-lines. Additionally, offering hardware acceleration units in the unused cache area seems interesting [22].

References

1. Lim, H.B., Yew, P.C.: Parallel program behavioral study on a shared-memory multiprocessor. In: ICS '91: Proceedings of the 5th international conference on Supercomputing, New York, NY, USA, ACM Press (1991) 386–395
2. Sherwood, T., Perelman, E., Hamerly, G., Sair, S., Calder, B.: Discovering and exploiting program phases. IEEE Micro **23**(6) (2003) 84–93
3. Sherwood, T., Calder, B.: The time varying behavior of programs (August 1999) Technical Report UCSD-CS99-630, University of California, San Diego.
4. Balasubramonian, R., Albonesi, D.H., Buyuktosunoglu, A., Dwarkadas, S.: A dynamically tunable memory hierarchy. IEEE Trans. Computers **52**(10) (2003) 1243–1258

5. Guthaus, M.R., Ringenberg, J.S., Ernst, D., Austin, T.M., Mudge, T., Brown, R.B.: Mibench: A free, commercially representative embedded benchmark suite. In: WWC '01: Proceedings of the Workload Characterization, 2001. WWC-4. 2001 IEEE International Workshop on, Washington, DC, USA, IEEE Computer Society (2001) 3–14
6. Naz, A., Kavi, K., Oh, J., Foglia, P.: Reconfigurable split data caches: a novel scheme for embedded systems. In: SAC '07: Proceedings of the 2007 ACM symposium on Applied computing, New York, NY, USA, ACM Press (2007) 707–712
7. Gonzàlez, A., Aliagas, C., Valero, M.: A data cache with multiple caching strategies tuned to different types of locality. In: ICS '95: Proceedings of the 9th international conference on Supercomputing, New York, NY, USA, ACM Press (1995) 338–347
8. Johnson, T.L., mei W. Hwu, W.: Run-time adaptive cache hierarchy management via reference analysis. In: ISCA '97: Proceedings of the 24th annual international symposium on Computer architecture, New York, NY, USA, ACM Press (1997) 315–326
9. Ranganathan, P., Adve, S., Jouppi, N.P.: Reconfigurable caches and their application to media processing. In: ISCA '00: Proceedings of the 27th annual international symposium on Computer architecture, New York, NY, USA, ACM Press (2000) 214–224
10. Benitez, D., Moure, J.C., Rexachs, D.I., Luque, E.: Evaluation of the field-programmable cache: performance and energy consumption. In: CF '06: Proceedings of the 3rd conference on Computing frontiers, New York, NY, USA, ACM Press (2006) 361–372
11. Sherwood, T., Perelman, E., Hamerly, G., Calder, B.: Automatically characterizing large scale program behavior. In: ASPLOS-X: Proceedings of the 10th international conference on Architectural support for programming languages and operating systems, New York, NY, USA, ACM Press (2002) 45–57
12. Huffmire, T., Sherwood, T.: Wavelet-based phase classification. In: PACT '06: Proceedings of the 15th international conference on Parallel architectures and compilation techniques, New York, NY, USA, ACM Press (2006) 95–104
13. Ogasawara, Y., Tate, I., Watanabe, S., Sato, M., Sasada, K., Uchikura, K., Asano, K., Namiki, M., Nakajo, H.: Towards reconfigurable cache memory for a multi-threaded processor. In Arabnia, H.R., ed.: PDPTA, CSREA Press (2006) 916–924
14. Zhang, C., Vahid, F., Najjar, W.: A highly configurable cache architecture for embedded systems. In: ISCA '03: Proceedings of the 30th annual international symposium on Computer architecture, New York, NY, USA, ACM Press (2003) 136–146
15. Gordon-Ross, A., Vahid, F., Dutt, N.: Automatic tuning of two-level caches to embedded applications. In: DATE '04: Proceedings of the conference on Design, automation and test in Europe, Washington, DC, USA, IEEE Computer Society (2004) 10208
16. Gordon-Ross, A., Vahid, F.: Dynamic optimization of highly configurable caches for reduced energy consumption. Riverside ECE Faculty Candidate Colloquium (March 2007) Invited Talk.
17. Gordon-Ross, A., Zhang, C., Vahid, F., Dutt, N.: Tuning caches to applications for low-energy embedded systems. In Macii, E., ed.: Ultra Low-Power Electronics and Design, Kluwer Academic Publishing (June 2004)
18. Mei, B., Vernalde, S., Man, H.D., Lauwereins, R.: Design and optimization of dynamically reconfigurable embedded systems (2001) citeseer.ist.psu.edu/mei01design.html.

19. Nowak, F., Buchty, R., Karl, W.: A run-time reconfigurable cache architecture. (2007) accepted for Proceedings of ParCo 2007.
20. Xilinx, Inc.: Platform Studio and EDK Details (2007) Web site: http://www.xilinx.com/ise/embedded/edk_pstudio.htm.
21. Tao, J., Karl, W.: Optimization-oriented visualization of cache access behavior. In: Proceedings of the 2005 International Conference on Computational Behavior (Lecture Notes in Computer Science 3515), Springer (May 2005) 174–181
22. Kim, H., Somani, A.K., Tyagi, A.: A reconfigurable multifunction computing cache architecture. IEEE Trans. Very Large Scale Integr. Syst. **9**(4) (2001) 509–523