An Educational Environment
for Designing and Performance Tuning of Embedded Systems

Roberto Giorgi and Cosimo Antonio Prete

Dipartimento di Ingegneria della Informazione
Università di Pisa, Via Diotisalvi 2, I-56126, Italy

Abstract - *Teaching how to design and tune an embedded system is indeed a difficult task, since the student has to learn the many trade-offs that lead to the final system configuration.*

Existing tools are often too complex, or do not stress the basic steps in the design path. These steps are very useful during the first training sessions.

The environment Csim2, which is used at our university, permits the student to become familiar with concepts of program locality, cache structure and performance tuning, while analyzing actual data produced by the actual software that has to be tied with the embedded system.

The student can analyze program behavior by means of locality graphs, or run extensive parametric simulations in order to find the best configuration that minimize either system cost, power consumption, or execution time. Further optimizations allow the designer to explore more sophisticated features like selective cacheing, cache locking, scratch memory, and code mapping for better cache exploitation.

In this paper we show the basic capabilities of the environment, and some example of training sessions. By means of graphs about program locality and performance metrics, the student is readily conducted to learn how to select an adequate embedded system configuration.

1 Introduction

The aim of Csim2 educational environment is to combine two complementary exigencies in embedded system architecture design courses: on the one side, the availability of a tool, which can provide practical example sessions in order to show the main concepts about embedded system architecture; on the other side, fostering the students to the actual design activity of embedded application oriented systems [Gajski95].

The idea of building this environment came up from the observation that actual computer systems and/or commercial design tools are generally not suitable to be used as didactic tools and to present the basic concepts of architecture design in both basic and advanced Computer Engineering courses. Their structure is often too complex and usually prevents the detection of all the events occurring in the activity of the machine. Moreover, the high number and frequency of these events may require a too expensive acquisition system or, also, several events may not be directly observable, since they occur within the chip.

The design of embedded systems, through new methodologies like co-design, and *system-in-a-chip* approach, implies the demand for specific architectural knowledge by computer engineers. A typical design path starts from the definition of the hardware/software requirements needed to implement the specified function through the embedded system [Gajski95]. After that, the designer usually has a prototype program and a prototype hardware configuration that has to be tuned in order to meet the low power consumption and/or low cost requirements or to minimize execution time. At this point, it is very important that the designer understands how to modify the system in order accomplish that task. This understanding usually rely on a good knowledge about the memory hierarchy behavior and the program locality effects on that hierarchy.

Also, due to the conflicting requirements, a trade-off has to be chosen. For example, low power consumption and low cost requirements suggest the adoption of slow memory devices, so that designers often turn to on-chip cache memories to both provide high processing power and allow using large slow main memory. In this case, typical question to be answered are: i) given a system and an application, is it necessary to add a cache memory to obtain the desired performance? ii) if so, which is the optimal cache configuration? iii) given a specific on-chip cache, which is the cheapest main memory satisfying the performance requirements of the application? Without an accurate tool for system configuration and simulation, reliable answers are hard to come by.

All these reasons motivated us to develop a new tool which could combine the different needs of student and the actual designer. A relative easiness is guaranteed concerning the practical use of the environment, so that the student, when switching from one phase to another, does not need to get familiar with a different – possibly much more complex – tool.

As a mere didactic environment, Csim2 offers a wide range of opportunities for investigating the system performance and system structure up to a relatively complex level of depth. The concept of program locality is particularly emphasized, since it is one of the critical issues in this branch of computer architecture. To this purpose, Csim2 provides an advanced program-locality analysis and a close evaluation of all the quantities which affect the execution time. The student is actively involved in making authentic choices that affect the target system, such as changing the parameters of a simulation and analyzing the immediate response of that system to the user actions; otherwise, elaborate graphics and simulations may result not effective.

As a design tool, the package allows the user: i) to project his own embedded application within a proper software development environment; ii) to explore different strategies for the memory hierarchy, including different level of caching, split

caches, write-buffering, by means of a graphical design tool; iii) to carry out the performance evaluation, in order to choose the system configuration which can guarantee the best performance for the target application.

For a further parameter space exploration the designer can perform a parametric simulation to evaluate system performance while varying the timing and architecture features of each component of the system. The final results, shown by means of easy-to-read graphs, help him to find rather quickly an acceptable trade-off solution. Thus, focusing on these simple crucial steps of the system design, we foster the student attitude to select appropriate memory hierarchy and tune the influencing parameters of the architecture.

Second level optimizations include the investigation of more sophisticated techniques like *selective cacheing* to cache only particular memory areas, *cache locking* to leave some data in cache thus avoiding replacement, *scratch memory* a small on-chip memory for allocating frequently used data, *code mapping* for reducing conflict cache-misses. All these techniques are supported by adequate tools that help the designer to select the best strategy for the application code and data [Lorenzini98]. Etherogeneous multicore architecture is also permitted, thus allowing the designer to try combinations of processors and DSPs.

The environment has been made up by our University [Prete97b] and the full version is integrated in a toolkit (JumpStart) distributed by VLSI Technology, Inc., for the design of ARM-based applications. (ARM [Jaggar97], [Furber93] is a 32-bit microprocessor designed by ARM Ltd. and largely used in embedded products. It uses RISC technology and a fully-static design approach to obtain both high performance and very low power consumption.)

2 The Csim2 environment

To facilitate the students, the package has a very friendly point-and-click graphical interface, by which the teacher can easily show and discuss, using practical examples, the basic concepts of cache architecture and behavior. The environment consists of four phases shown in Figure 1.
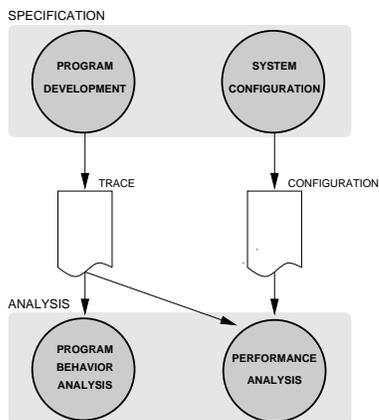

Figure 1. Structure of a Csim session.

In the *Program Development* phase, the user builds an application, debugs it and produces a trace file. Applications can be executed and debugged on a dedicated ARM instruction set simulator or loaded in an ARM CPU-based board for a native evaluation. Once that the application has been developed, the user can generate a trace file by simply pointing-and-clicking while the program is running in emulator mode.

In the *System Configuration* phase, the user defines the system architecture and the features of each component. The user first draws the schematic of the system architecture at functional level. The system may include the following components: an ARM core, cache memory that can be combined in different levels or in split instruction/data architecture, a system bus, memory banks, and a number of I/O devices. For each component, the designer has to specify the timing and the other custom parameters (Figure 2).
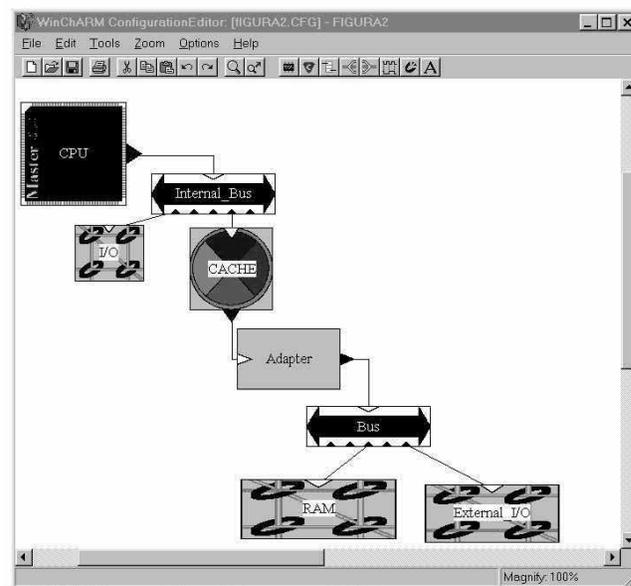

Figure 2. Embedded system configuration example. The system includes an ARM core, an internal I/O device, a cache memory, a bus adapter, and an external RAM and an external I/O device. The configuration is produced by means of the configuration editor of the environment.

The *Program Behavior Analysis* phase allows the student to perform two types of trace analysis. The first one uses traditional program statistics such as the percentages of data/code, read/write accesses. The second one regards program locality. An accurate knowledge of locality features plays a crucial role in understanding cache concepts. The locality statistics include: the number of unique blocks, the locality surface [Grimsrud96], and the spatial locality.

The *Performance Analysis* phase allows the user to plan, perform a single simulation or a performance evaluation *experiment*, and finally analyze the results. An experiment is defined by: i) the trace file; ii) the system configuration; and iii) the varying parameters (one or two). Csim2 may initially simulate an adequate number of memory references without an outcome. This allows the cache to exit its cold state [Easton78] and to reach a steady condition. The results con-

sist of: *global system performance* (execution time, lost time in waiting, and word transfer ratio); *cache behavior* (miss, code miss, data miss, read miss, data read miss and data write miss ratios and cumulative cold misses); and *bus traffic* (occupation rate, number of read-block operations, number of write operations for write-through cache models and number of update-block operations for copy-back cache models).

3 Exploring the program behavior

In a didactic approach to computer architecture, one of the key concepts that the student has to deal with is *program locality*. Hence, the first phase of a typical Csim2 didactic session, concerns the analysis of locality features of a program written directly by the student or chosen within a set of predefined, very simple programs; in the example shown in detail in this Section, the program is the implementation of the median filter algorithm applied to a 34×34 pixel image with a 3×3 pixel window [Gallagher81].

As shown in the previous Section, during the *Program Development* phase (Figure 1) a trace can be produced to allow a detailed program locality analysis (number of unique blocks, locality surface, spatial locality).

If we define $T[i]$ as the i-th reference of a trace T, for each couple $\{T[i], T[j] \; such \; that \; j > i\}$ we can also define the *distance* (d) as the number $j - i$ of intervening references, and the *stride* (s) as the offset $T[j] - T[i]$ between the two references.

The concept of *spatial locality* refers to the fact that address locations close to the "currently" referenced location $T[i]$ are more likely to occur in the next few references than locations far away. Similarly, the concept of *temporal locality* reflects the fact that the address of the "current" reference $T[i]$ is very likely to occur again in the next few references.

A quantitative approach to the locality analysis was first proposed by Archibald *et al.* [Grimsrud96] by means of the introduction of the *locality surface*. They proposed a 3D-graph where stride and distance are the base axes. The magnitude of locality surface for a specific couple (s, d) is defined as the probability that $T[i] + s = T[i + d]$, where $T[i] + s \notin \{T[i + 1], ..., T[i + d - 1]\}$ and i assumes all the values between 1 and the length of the trace T minus one. From the locality surface, the designer may derive information about locality features like *sequentiality*, *striding*, *temporality* and *loops*. Sequentiality is typically due to the fetching of consecutive instructions. It is visible as a ridge along the diagonal region with $s = d$, in which the length reflects the distribution of sequential run lengths in the reference stream, while the amplitude reflects sequential run frequency. Striding is produced by a series of references with a fixed step and is typical of numerical algorithms, such as matrix operations where the elements are accessed in row order instead of in column order. It is characterized by a ridge in the region with $s > d$. The temporality region, i.e., the region with $s = 0$, shows the distribution of distances between repeated accesses at the same addresses. Finally, loops are characterized by ridges which are parallel to the stride axis, in the region with $-d < s < 0$. Figure 3 gives an example of a locality surface concerning the median filter program. In par-

ticular, the temporal locality window shows that, in this case, the latest referenced address has a very high probability (more than 80%) to be referenced again within the next 128 memory accesses.
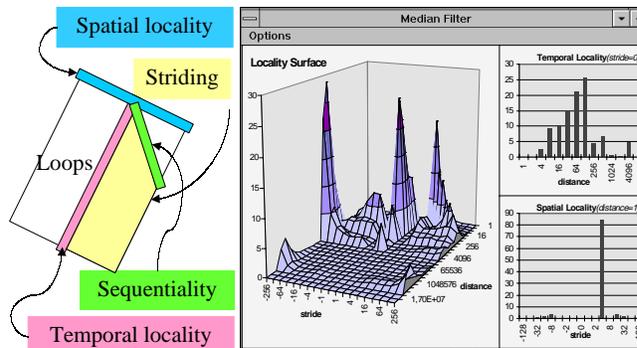

Figure 3. The locality surface for median filter program.

According to the definition, spatial locality is the distribution of the offset between two consecutive addresses in the trace. It can be obtained from the locality surface in the case of $d = 1$. Csim2 can show spatial locality in a specific graph as distribution of: i) all accesses and ii) data and code accesses separately.

The user can optionally analyze the *unique blocks* graph. The number of unique blocks for a given number i of references is the number of distinct blocks used by a program before the reference i-th. These blocks only cause misses in an infinite cache, the number of unique blocks delineates a lower bound for the miss ratio. Csim2 shows a family of curves, where the number of unique blocks is given as a function of the number of references and the parameter is the block size (Figure 4).
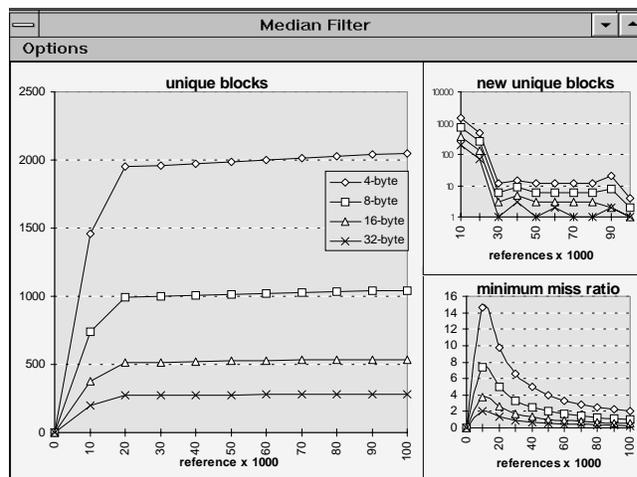

Figure 4. Number of unique blocks (total and incremental) and lower bound of miss ratio for median filter program.

The next step in a didactic path is to show how the presence of a cache memory can exploit program locality in order to im-

prove the system performance. For this purpose, the student has to select the system parameters concerning cache organization. A cache scheme is defined by the following parameters: i) the mapping policy; ii) the replacement algorithm; iii) the update policy; iv) the cache size, v) the block size; vi) the number of blocks for each set (in the case of a set associative cache); and vii) the presence and the length of a write buffer.
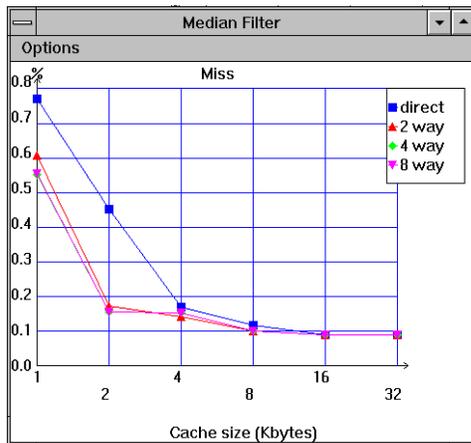

Figure 5. Miss percentage for median filter program.

Figure 5 shows the miss percentage of the median filter program as a function of the cache size and degree of associativity. For this program, a 2-way set associative cache is recommended since it supplies the best balance between cost and performance. Also, we notice that a 4-way cache produces quite the same result. At a more detailed level of analysis, the student is called to search, for a given cache structure, the cache and block sizes which can provide the best results in terms of global performance. Again a new graph can be produced (like the one in Figures 7 and 8) to find the best value for the cache block size.

4   Fostering students to actual design

In the design of embedded systems, a key point is the optimization of each component, which needs to meet, as better as possible, the specific application for which the whole system is designed. Also, it should be noticed that often an embedded system runs only one program for all its life. We are going to present an example of design training path, and we will show how Csim2 can help a student to find out the optimal cache and system configuration for a specific application. We consider the cjpeg program, a jpeg image compression/decompression tool [Wallace91] which is frequently used in commercial embedded systems.

First, the student should wonder about the following question: for a 20 MHz ARM running an application using the cjpeg program, is it necessary to add a cache memory in order to achieve the required performance? We suppose that the product requires that the image compression be completed in less than 1s. First, the student traces the execution of the cjpeg program while compressing an image stored in, e.g., "ppm" format, using the JumpStart trace facility. The ppm image is a 101-KByte image consisting of 227x149 pixels. (The sample picture is a red rose and the image produced by cjpeg occupies 5 KBytes.)

Then, the student defines the system configuration including a 20 MHz ARM core, a system bus, a 1-MByte memory DRAM bank, a 128-KByte memory PROM bank, and a memory-mapped graphical I/O device. The student needs to specify the timing and architecture of each component. In order to obtain a low cost and a low power-consumption solution, a slow system bus and slow memory devices are selected.

In the case of the ARM core, the student provides the timing for both read and write operations. For these operations the simulator requires i) the minimum time necessary for the ARM core to perform the bus operation and ii) the maximum available to a slave to complete an operation without requiring waiting time for the CPU.

A specific window shows the ARM timing plot, derived from the ARM data book, in order to drive the student to find the proper values. In our example, considering that the ARM processor employs a pipelined bus, the student obtains and sets these values: 50 ns, as minimum time for read/write operations, and 64 ns as the maximum time usable by a slave to complete an operation.

The simulator models a generic bus, which is capable to accommodate the typical memory operations. The student has to specify the data bus width and the time for each type of bus operation. In our example, the data bus width is 32 bits and the time is 200 ns for both read (T_{read}) and write (T_{write}) operations.

Finally, the student specifies the features of memory and I/O devices. For each module, the configuration parameters include the module type, the starting address and the size. In our example, the system includes three modules: a DRAM bank, a ROM bank and a memory-mapped graphical I/O device. The simulator requires to know the delays (additional time with respect to the bus time) introduced by a component to complete each bus operation. Among the three modules considered, only the PROM module needs additional time (200 ns) to complete read operations.

With the system configuration just examined, a simulation shows that, without cache memory, the application takes up 2.898s to execute the program. The addition of a cache memory proves to be necessary, therefore, to meet the time requirement.

As shown in the previous Section, the designer has to define the cache structure in terms of cache size, block size, number of blocks per set, and replacement policy; furthermore, for simulation to be possible, cache timings have to be specified. Figure 6 shows the scheme used to set the cache timings: T_{rd} and T_{wt} are the times for reading data from or writing data to a cache block; T_{tag} is the tag access time; T_{cmp} is the compare time; and T_{orq} is the time needed to initiate an operation involving an attached module, after a miss has been detected.

Finally, the student has to specify the timing of bus block-operations used by the cache to interact with the main memory. In particular, the cache uses the *read-block* operation to fetch a memory block when a miss condition occurs. An *update-block* operation allows a copy-back cache to update a memory block when its cached copy is dirty and has to
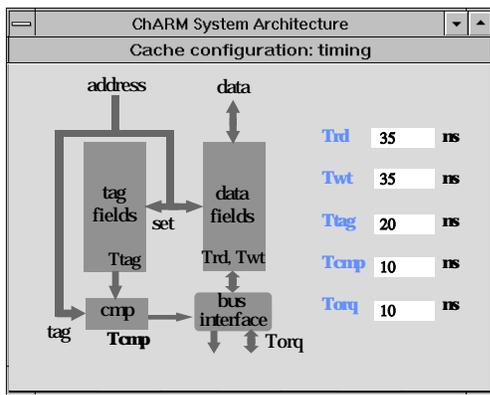
Figure 6. Setting the cache timings.

Figure 7. Miss ratio and time execution for `cjpeg` *program.*

be replaced. The time required by a bus block-operation is calculated by considering the bus width, the block size and four timing values: T_{read}, T_{write}, T_{sread} and T_{swrite}. A block operation is described as a single operation followed by fast transfer operations. T_{read} (T_{write}) is the time needed to perform a single read (write) operation, and T_{sread} (T_{swrite}) is the time needed to perform a subsequent sequential read (write) transfer. In this example, the timings are: $T_{read} = T_{write} = 200$ns, $T_{sread} = T_{swrite} = 160$ns.

The student can now execute a parametric simulation in order to search the optimal cache configuration. Figure 7 shows the miss ratio and the execution time versus block size (from 8 to 64 Bytes) and cache size (from 2 to 32 KBytes) for two cache configurations. The first cache is a simple write-through, direct access cache without a write buffer; the second one is a more complex copy-back, two-way set associative cache with a two-deep write buffer. The cache uses the LRU technique as replacement policy. In both cases, the timings are: $T_{rd} = 35$ns, $T_{wt} = 35$ns, $T_{tag} = 20$ns, $T_{cmp} = 10$ns, $T_{orq} = 10$ns.

The designer can observe that, in both configurations, execution time and miss ratio exhibit different values and behaviors. For cache sizes greater than 16 KBytes, the execution time is constant and independent of the block size. In this way the student can select a configuration that best meets cost-effectiveness and performance requirements (execution time $\leq$ 1s). For example, an optimal choice is a 16-KByte, write-through, direct access cache with 16-Byte block size without a write buffer.

Now, the student can also answer the question: for the selected cache configuration, which is the cheapest main memory meeting the time requirement? The designer can find the solution by executing simulations having the RAM bank access time as parameter. The simulation shows that the memory bank delay can be increased by no more than 30ns with respect to the values specified in the configuration.

Now, if the student uses the same design path for a different embedded application, he/she finds a different cache and system configuration. For example, we trace the execution of the `rawcaudio` program [CCITT84] while it converts a 6-KByte ADPCM sound sample to a 24-KByte raw 16-bit PCM format. The audio sample is the voice of a man saying "hello, world." Figure 8 shows the miss ratio and the execution time
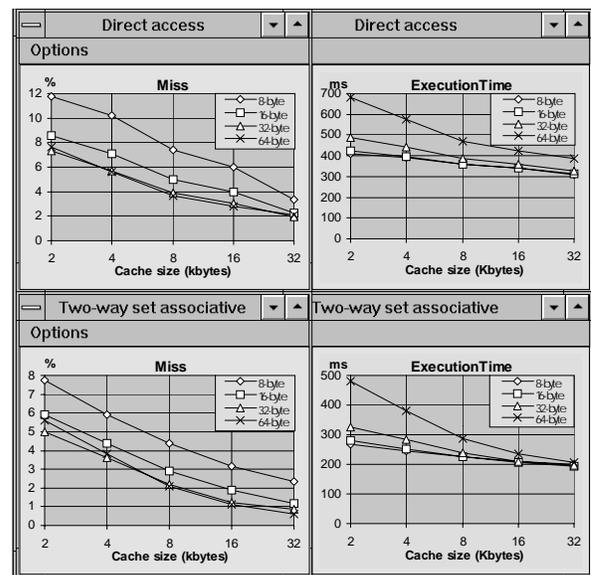
for the same configurations considered in the first example. The differences between the graphs of Figures 7 and 8 are due to the different locality characteristics of the programs.

We assume that the application requires that the conversion should be completed in less than 90ms, in order to show a configuration tuning session. The system takes up 359ms to execute the program without cache memory, therefore cache is necessary. Table 1 lists some examples of cache configurations that allow the system to satisfy the time requirement. If the designer selects cache configuration 2, the memory bank delays can be increased by 260ns with respect to the values specified above.

	mapp.	update policy	cache size (KB)	block size (B)	write buffer len.	exec. time (ms)	max delay (ns)
1	direct	copy-back	8	8	0	88.97	40
2	direct	copy-back	16	16	0	86.40	260
3	direct	copy-back	8	8	2	88.75	60
4	direct	copy-back	16	16	2	86.34	260
5	2-way	copy-back	4	16	0	88.96	40
6	2-way	copy-back	8	64	0	88.62	50
7	2-way	copy-back	4	16	2	88.86	50
8	2-way	copy-back	8	64	2	88.61	60

Table 1. Some cache configurations that allow the system to satisfy the time requirements for `rawcaudio` *program.*

These two examples show that meeting time requirements of different applications yields different cache configurations. The student can observe that, in the second example, the write-through cache configurations never guarantee the fulfillment of time requirements. The two examples also show that write and miss ratios affect the performance of systems with cache memories in a different way. The `cjpeg` and
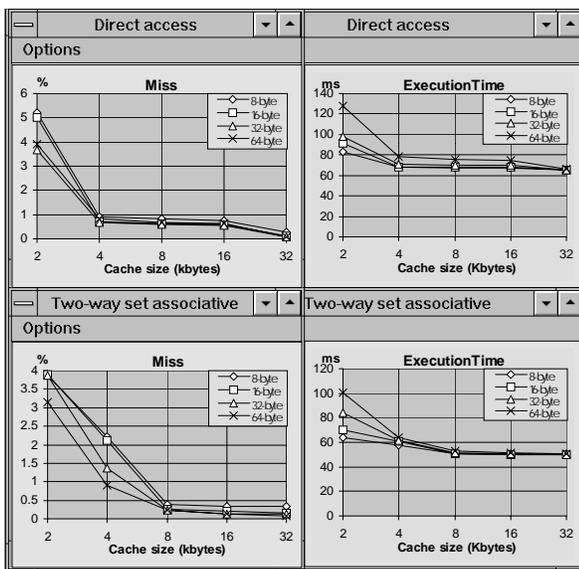
Figure 8. Miss ratio and execution time for `rawcaudio` *program.*

`rawcaudio` programs have similar write ratios, but exhibit different miss ratios due to different locality features. A lower miss ratio enhances the influence of write operations on global performance. In this case, the choice of an optimal update policy becomes critical.

5 Conclusions

The growing demand for embedded products requires highly sophisticated computing functions. Designers must select the most efficient cache/system configuration in order to resolve complex – even conflicting – requirements for low-power/high-speed and component cost. This makes accurate and reliable system/cache memory simulation and performance analysis crucial. We have presented an educational environment based on a trace-driven system simulator that can help students in the design activity of cache memory to be employed in ARM-based embedded systems. By means of practical examples, we have shown how the student can successfully use the tool in two typical schemes of a didactic path.

6 Acknowledgments

References

[CCITT84]   *32 Kbit/s Adaptive Differential Pulse Code Modulation (ADPCM)*, CCITT (Int'l Telegraph and Telephone Consultative Committee), 1984.

[Easton78]   M. Easton, "Computation of cold-start miss ratio," *IEEE Trans. Computers*, vol. C-27, no. 5, pp. 404–8, May 1978.

[Furber93]   S. B. Furber, P. Day, J. D. Garsidex, N. C. Paver, and J. V. Woods, "A Micropipelined ARM," *Proc. IFIP TC 10/WG 10.5 Int'l Conf. on Very Large Scale Integration (VLSI '93)*, Sept. 1993.

[Gajski95]   D. D. Gajski and F. Vahid, "Specification and Design of Embedded Software-Hardware Systems," *IEEE Design & Test of Computers*, vol. 12, no. 1, Spring 1995.

[Gallagher81] N. Gallagher and G. Wise, "A theoretical analysis of the properties of median filters," *IEEE Trans. Acoustics Speech and Signal Proc.*, vol. 29, pp. 1136–1141, 1981.

[Grimsrud96]   K. Grimsrud, J.Archibald, R. Frost, and B. Nelson, "Locality as a Visualization Tool," *IEEE Trans. Computers*, vol. 45, no. 11, pp. 1319–1326, Nov. 1996.

[Jaggar97]   D. Jaggar, "ARM Architecture and Systems," *IEEE Micro*, pp. 9–11, July 1997.

[Lorenzini98] S. Lorenzini, G. Luculli, and C. A. Prete, "A Fast Procedure Placement Algorithm for Optimal Cache Use," *Proc. MELECON'98, Tel Aviv, Israel*, May 1998.

[Prete97b]   C. A. Prete, M. Graziano, and F. Lazzarini, "The ChARM Tool for Tuning Embedded Systems: Selecting and tuning system configurations to meet cost, performance, and power consumption requirements," *IEEE Micro*, vol. 17, no. 4, pp. 67–76, July/Aug. 1997.

[Wallace91]   G. K. Wallace, "The JPEG Still Picture Compression Standard," *Communications of the ACM*, vol. 34, no. 4, pp. 30–44, Apr. 1991.