

**The U.S.C. Multiprocessor  
TestBed: Project Overview**

L. Barroso, S. Iman, J. Jeong  
K. Oner, K. Ramamurthy and Michel Dubois

CENG Technical Report 94-15

Department of Electrical Engineering - Systems  
University of Southern California  
Los Angeles, California 90089-2562  
(213)740-4475

August 1994

## **THE U.S.C. MULTIPROCESSOR TESTBED: PROJECT OVERVIEW<sup>1</sup>**

**Luiz Barroso, Sasan Iman, Jaeheon Jeong,  
Koray Öner, Krishnan Ramamurthy and Michel Dubois**

USC CENG Technical Report  
CENG-94-15  
August 1994

### **Abstract**

In multiprocessor systems, processing nodes contain a processor, some cache and a share of the system memory, and are connected through a scalable interconnect. The system memory partitions may be shared (shared-memory systems) or disjoint (message-passing systems). Within each class of systems many architectural variations are possible. Comparisons among systems are difficult because of the lack of a common hardware platform to implement the different architectures.

The U.S.C. Multiprocessor Testbed is a hardware emulator for the rapid prototyping of vastly different multiprocessor systems. In the testbed the hardware of the target machine is emulated by reprogrammable controllers implemented with Field-Programmable Gate Arrays (FPGAs). The processors, memories and interconnect are off-the-shelf and their relative speed can be modified to emulate various component technologies. Every emulation is an actual incarnation of the target machine and therefore all software written for the target machine can run directly on it without modification.

In this paper, we will describe the architecture of the testbed, its performance and the prototyping methodology, and we will compare the hardware emulation approach to simulation and prototyping.

**Keywords:** Field-Programmable Gate Arrays (FPGAs), message-passing multicomputers, shared-memory multiprocessors, design verification, performance evaluation, simulation.

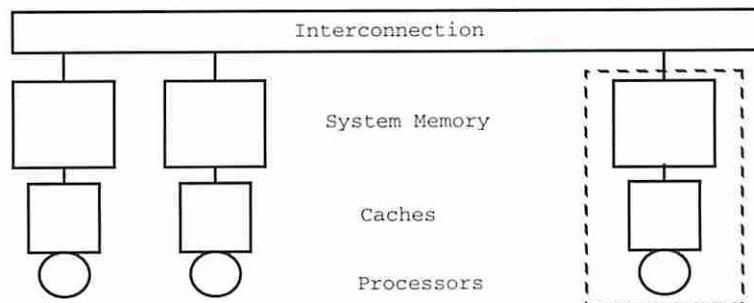
---

1. This work is supported by the National Science Foundation under Grant MIP-9223812

## 1. INTRODUCTION

Multiprocessor systems are becoming common place in the computing industry. The consensus among machine designers is to favor asynchronous MIMD (Multiple Instructions Multiple Data streams) systems, in which processors execute their own instructions and run on different clocks. In these systems, processing elements contain a processor, some cache and a share of the system memory, and are connected through a scalable interconnect which facilitates machine packaging, such as a bus or a mesh (see Fig. 1). Whereas this physical model dominates, disagreement exists as to the interprocessor communication mechanism. There are two dominant models. One is based on disjoint memories and message-passing and the other is based on shared-memory. In a message-passing system, processors communicate by exchanging explicit messages through *send* and *receive* primitives. A received message is put into a local buffer, which is accessed by the processor when it executes a receive. In the shared-memory model, processors communicate through load and store instructions and some form of explicit synchronization among processors is required at times to avoid data access races [5].

FIGURE 1. Physical Organization of Most MIMD Systems



The shared-memory model facilitates fine grain (word level) communication but requires a large number of instructions to transmit large chunks of data, whereas the message-passing model can transmit large amounts of data in a single message. In terms of ease of programming, the shared-memory model has been so far the favored transition path from uniprocessors to multiprocessors. On the other hand, message-passing systems are generally perceived as more scalable than shared-memory systems. One concern in both kinds of systems is the growing disparity

between the processor speed and the speed of communication. In message-passing systems, messages have traditionally suffered from a large software overhead, which explains the comparatively large latencies of sends and receives as compared to the latencies of loads and stores in shared-memory machines. Ways to combat these large latencies are to implement light-weight, hardware-based message primitives and to overlap message passing with computation. In shared-memory systems, the large latencies of loads and stores on shared data is also a problem, which is usually solved by complex shared-memory access mechanisms.

Some researchers advocate private caches [5, 11] with hardware- or software-based consistency maintenance, such as in the Stanford DASH [8] prototype. The consistency protocol, the constraints on the ordering of memory accesses [5], the cache parameters, as well as the interconnect latency and bandwidth are all factors affecting the performance and programming ease of multiprocessors. Machines such as the DASH have been called CC-NUMA (for Cache-Coherent Non-Uniform Memory Access) architectures to differentiate them from COMAs (Cache-Only Memory Architectures) exemplified by the Data Diffusion Machine (or DDM) [7]. A COMA has the same architecture as in Fig. 1 and communication is through shared variables; however, no main memory is present and, instead, the memory in each processor node implements a huge cache called *attraction memory*.

There is a trend today towards integrating the message-passing and shared-memory paradigms in order to draw from the strengths of both. The implementation of message-passing on top of shared-memory is straightforward but special hardware is however needed for efficient implementations. Recently [9] it was shown that shared-memory can be implemented on top of a message-passing system. It is widely expected that future multiprocessors will be hybrid systems, supporting both the shared-memory and message-passing paradigms. To understand which form such an integrated system should take comparisons among systems are required. Presently, comparisons are difficult to make and hard to validate because of the lack of a common hardware platform to implement the different models.

Because multiprocessors are complex and powerful the correctness of a design and its expected performance are very difficult to evaluate before the machine is built. Traditionally two approaches have been taken to verify a design: breadboard prototyping and software simulation. A prototype is costly, takes years to build and explores a single or a few design points. Discrete-event simulation is very flexible but very slow if the design is simulated in details and it is subject to validity problems because the target system must be considerably abstracted in order to keep simulation times reasonable. In some industrial projects where a detailed and faithful simulation of a target system has been done in software, the simulation runs at the speed of a few cycles of the target system per second of simulation. The parallelization of discrete-event simulation is an ad-hoc procedure and usually exhibits low speedup [6]. Most simulators [4, 10] rely on the direct execution of each target instruction on the host and, because the code (either source or binary) must be instrumented, it is difficult to simulate efficiently the execution of interesting workloads, other than scientific programs.

The major objective of the U.S.C. Multiprocessor Testbed is to develop a common, configurable hardware platform to emulate faithfully the different models of MIMD systems with up to eight execution processors. Emulation is orders-of-magnitude faster than simulation and therefore an emulator can run problems with the real data set sizes (the ones for which the target machine is designed). Emulation is more faithful to the target implementation than simulation and therefore more reliable performance evaluation and design verification can be done with emulation. Finally, an emulator is a real computer with its own I/O and the code running on the emulator is not instrumented. As a result, the emulator “looks” exactly like the target machine to the programmer and runs any binary including binaries from production compilers, operating systems, and software utilities.

In the following, we first introduce the hardware emulation approach for multiprocessors; then we describe the architecture of our testbed. In Sections 4 to 6 we explain the methodology to keep track of simulated time, to measure performance, and to program the testbed. Finally, in Sec-

tions 7 and 8 we show the expected performance of the testbed and compare the emulation approach to simulation and prototyping.

## **2. THE HARDWARE EMULATION APPROACH**

At the University of Southern California we have been experimenting for a year with a new approach to the rapid prototyping of multiprocessor systems. The approach is based on *hardware emulation*. Emulators have been used in the past to experiment with instruction sets. At the time when most processors were microcoded, it was convenient to have a machine which could execute various instruction sets in order to run software developed for different machines. Moreover, emulators were used in the development of new instruction sets in order to quickly obtain an implementation, to verify the correctness of the instruction set and to start writing the software so that both the machine and the software could be ready at the same time. No attempt has ever been made to apply the emulation approach to the design and verification of multiprocessor systems.

Several technologies are currently converging, making it possible to build flexible, low cost multiprocessor emulators in a short time. These technologies are FPGAs (Field-Programmable Gate Arrays), open software and hardware standards, efficient Computer-Aided Design (CAD) tools, and the existence of services for the rapid design and fabrication of printed-circuit boards.

### **2.1. FPGAs (Field Programmable Gate Arrays)**

Field-Programmable Gate Arrays (FPGAs) [12] are high-density, user configurable ASIC devices. FPGAs are in-circuit programmable by software. They have evolved over the past 15 years from simple chips that could replace a small number of gates (300 gates in a 20-pin package) to complex arrays that can be used to implement large circuits of up to 20,000 gates in 232-pin packages (Xilinx Inc. has announced the XC4025 with the equivalent of 25,000 gates for the Fall of 1994.) Meanwhile maximum operating speeds have soared from 20 to over 100 MHz. The current trend for FPGA technology is expected to continue unabated for some time to come. To

take advantage of this rapid technological improvement and to increase the lifetime of their design, designers must engineer the rest of their board for the highest possible clock rate and take advantage of the pin compatibility of new FPGA products to upgrade their design.

These recent advances in FPGA technology have turned hardware emulation into a more practical design verification/analysis technique. Solutions for hardware emulation have been made available by commercial companies in form of turnkey solutions. Turnkey solution vendors such as *Quickturn* use generalized prototype boards where all programmable devices and interconnects are pre-placed in a fixed manner. Tightly integrated software takes care of logic mapping and partitioning of multiple FPGAs. Prototypes of new microprocessors such as the Intel Pentium have been built and debugged with their actual software, using *Quickturn* Enterprise Emulation Systems.

## **2.2. Open Standards**

Open standards for instruction sets, software systems and interconnects make it possible to obtain easily and inexpensively the key components making up a multiprocessor system. For example, the processors in our testbed use the SPARC instruction set [2] for which we have the detailed functional simulator from Sun Microsystems. The chip set from Newbridge implements the Futurebus+ standard for the interconnection and relieves us from developing our own bus. In the future, when we will port system code, we will use a standard for which we can acquire the source code.

## **2.3. Efficient CAD Tools**

The new generation of design automation tools allows designers to move to a higher level of abstraction. Designers no longer need to directly deal with the gate-level design of digital circuits. They can specify their design at the Register Transfer Level (RTL) using High-level Description Languages (HDL) such as VHDL and Verilog. The availability of such tools allows designers to complete more complex designs in a much shorter time frame. In particular, the programming of

FPGAs is greatly simplified. Testing and simulation tools also help designers to verify the functionality of their design before it is implemented. Integrated tool sets that encapsulate all such design automation technologies provide a fast and reliable methodology to implement and verify a proposed design.

#### **2.4. Services for the Rapid Design and Fabrication of Custom PC-boards:**

We have been using the services of EZFAB at the Information Science Institute (ISI), which is part of the ARPA-sponsored Systems Assembly Project. Our task was limited to providing a correct netlist. EZFAB designed the board layout and placement and produced all the specifications needed by the board manufacturer in the form of a gerber format output. As a result, our involvement in board design and fabrication was minimal.

#### **2.5. Our Emulation Approach Using FPGAs**

The U.S.C. multiprocessor testbed is built mostly from off-the-shelf components (including processors, SRAMs, DRAMs, FIFOs, bus interface and drivers, and backplane), but the cache, memory, coherence and communication controllers are built with FPGAs. The emulation of a particular machine model is done through the FPGAs and a part of the memory to which they are attached.

The clock rate is 10MHz, which is about 10 times slower than the rate permitted by current board and PLD technologies (estimated at 100MHz). This compromise on the emulation efficiency results from two trade-offs. First, the design and fabrication of the PC-boards of the emulator are greatly simplified, because the boards are at the mature end of the technology curve. Second, the lower clock rate facilitates the configuration of the FPGAs. FPGAs are slower than other programmable logic devices or custom circuits --especially when they are programmed with VHDL synthesizers, which are often less than optimum but which greatly reduce the design time-- and clocking them at low speed promotes the mapping of more complex designs. In order to have better flexibility, to emulate complex mechanisms and to further simplify the design, each



processor clock (pclock) is emulated in several testbed clocks. (Currently one pclock is eight clocks, but this number can be changed.) So, overall, the emulator currently runs 80 times slower than the target system which could be built with the best current technology. This low processing rate allows us to use a standard interconnection fabric and still be able to eliminate conflicts, which can cause distortions in the collection of performance data (see Section 4).

### **3. TESTBED ARCHITECTURE**

The testbed architecture has been geared towards the evaluation of multiprocessors with the general architecture shown in Fig. 1. The interconnections of the possible targets are limited to FIFO (First-In-First-Out) interconnections with uniform access latencies such as crossbars or busses; however, other interconnections, such as rings, can be modeled approximately. A FIFO interconnection is an interconnection such that messages sent between two nodes arrive at the destination in the same order in which they were sent.

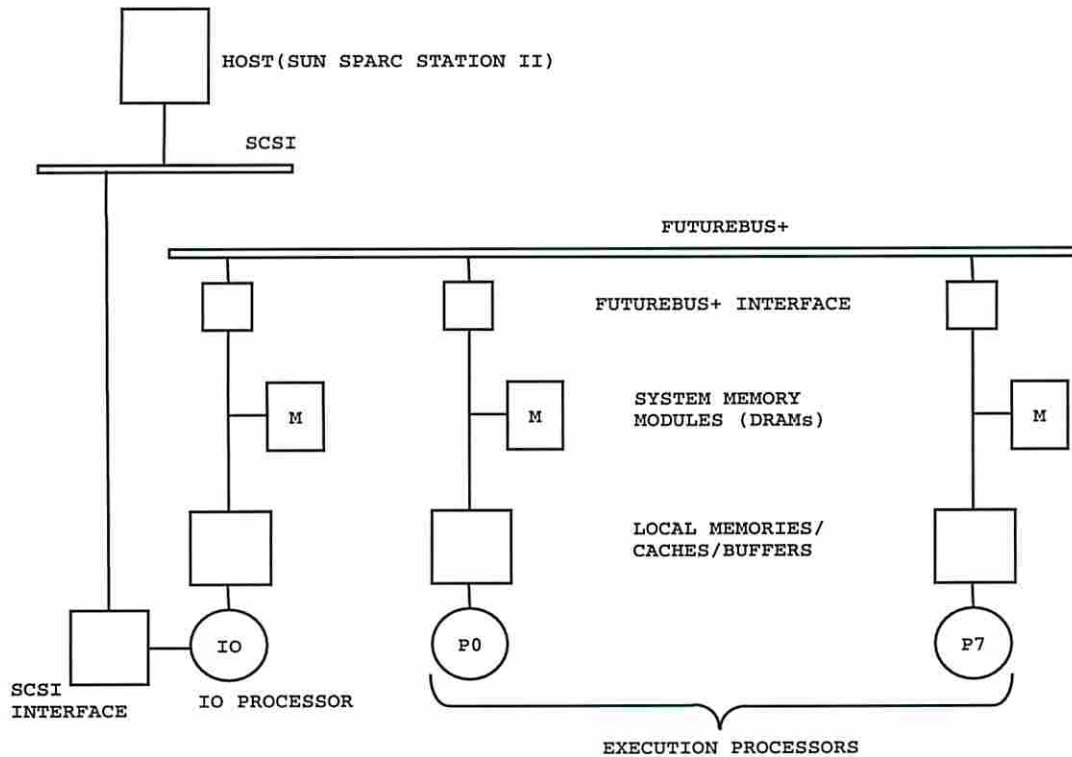
#### **3.1. Hardware Organization**

The testbed organization is shown in Fig. 2. It is made of nine SPARC processors (eight execution processors and one I/O processor) connected to a Futurebus+ backplane. The FutureBus+ is 64-bit wide (data) and packet-switched. It supports data broadcasting and interprocessor interrupts. The arbitration is distributed and takes between 200 and 600 nsec. The peak transfer rate is 20 Mwords/second or 80 Mbytes/second.

The number of clocks in each processor clock (pclock) is variable (it depends on the complexity of the mechanisms to emulate) but is currently set at eight clocks, which gives the testbed a peak emulation rate of 10 MIPS (i.e., eight processors at 1.25 MIPS each) of the target system or 1.25 million cycles of the target per second. An I/O processor whose configuration is identical to the configuration of the execution processors is connected to a SUN SPARC Station II through a SCSI interface [2]. This workstation serves as the console for the testbed and, additionally, executes its I/O requests. The peak I/O bandwidth is 1.25 Mbytes per second, which is more than suf-

ficient for a 10 MIPS machine.

FIGURE 2. Overall Configuration of the Testbed



The architecture of each processor board of the testbed is shown in Fig. 3. The nine processors are LSI Logic L64831 SPARC IU/FPU [2]. These single-chip processors can be clocked at up to 25MHz and execute both integer and floating-point instructions. They have no on-chip cache and therefore all instruction fetches and data accesses are visible on the pins of the chip.

There are three memories on each board, each controlled by a set of Xilinx FPGAs:

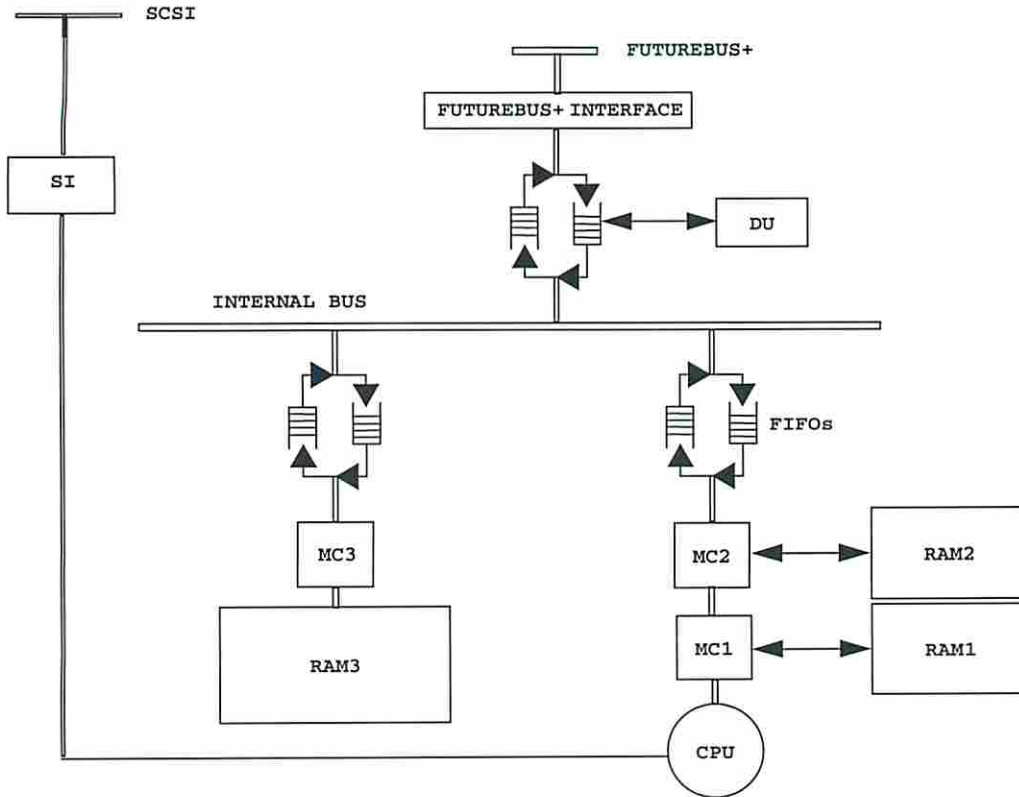
- MC1/RAM1: Each processor is attached to 2 Mbytes of static RAM (RAM1) controlled by two Xilinx XC4013<sup>2</sup> (called MC1). The major function of MC1 is to control the cycle-by-cycle execution of the processor. At the beginning of each processor access, MC1 blocks the processor, executes the sequence of steps needed to satisfy the access and unblocks the processor when the

---

2. Each XC4013 contains the equivalent of 13K gates. We plan to upgrade to the Xilinx XC4025 as soon as it becomes available (this Fall). This will double the number of equivalent gates in each controller. The XC4025 is pin-to-pin compatible with the XC4013.

access is completed. MC1 also manages RAM1 as a cache and interacts with the second-level memory.

FIGURE 3. Block Diagram of the Testbed Processor Node



- MC2/RAM2: An additional 8 Mbytes of static RAM (RAM2) controlled by three Xilinx XC4013 called MC2 makes up the second-level memory. MC2 interfaces the processor to the rest of the system and usually acts as a second-level cache controller.
- MC3/RAM3: System memory is emulated by 96 Mbytes of dynamic RAM (RAM3) controlled by two Xilinx XC4013 and one Cypress CYM7232 DRAM controller (called MC3).

The internal bus is a synchronous bus with a protocol similar to Sun Microsystems MBUS protocol [2]. It is a 32-bit wide packet-switched bus which transfers packets of sizes 16 to 128 bytes. Controllers MC2 and MC3 connect to the internal bus through very large, two-way FIFO buffers, which are there to prevent deadlocks and relieve the controllers from managing the data transfers. All on-board datapaths are 32-bit wide.

The Delay Unit (DU) is a programmable unit which emulates variable interconnection delays. It is built with a FIFO controlled by one AMD MACH 210 chip. The FIFO (8 kbytes) contains blocks and messages which are sent to the bus interface after a programmable delay depending on the target machine's interconnect latencies and packet size. This delay is computed by the formula:

$$\text{Latency} = T_{\text{start}} + (\# \text{ of words in packet}) \times T_{\text{word}}$$

The FutureBus+ interface is made of off-the-shelf chip sets from Newbridge and National Semiconductors. It includes bus transceivers, plus the Newbridge LIFE chip and distributed arbiter chip. The processing speed of the testbed is very low relative to the bandwidth available on the FutureBus+ and the effect of conflicts is negligible. We have run simulations of the testbed configured as a CC-NUMA architecture using the SPLASH benchmarks<sup>3</sup> with extremely small data set sizes (and therefore a high communication-to-computation ratio) and the bus utilization is lower than 10% even under the worst-case hypotheses. Therefore the testbed will yield accurate performance readings for situations where the network has infinite bandwidth. Interconnection traffic between any pair of processors can be measured in order to observe whether a particular switch in the target system's interconnect could be a hot spot. Note that the contention for the caches, the internal bus and the memory inside each processor node is modeled in full detail.

### 3.2. Emulation of CC-NUMAs with Central Directory Protocol

The first emulator that we have developed is a system with hardware-enforced cache-coherence under strong or weak ordering of memory accesses [5]. The protocol is directory-based and each memory block has a home node where a directory records the presence and state of copies in every cache [8, 11]. The memory and the cache directories have pending states so that transactions for different blocks can be executed concurrently. In this emulator, MC1/RAM1 is a first-level write-through cache (containing both data and instructions), MC2/RAM2 is a second-level

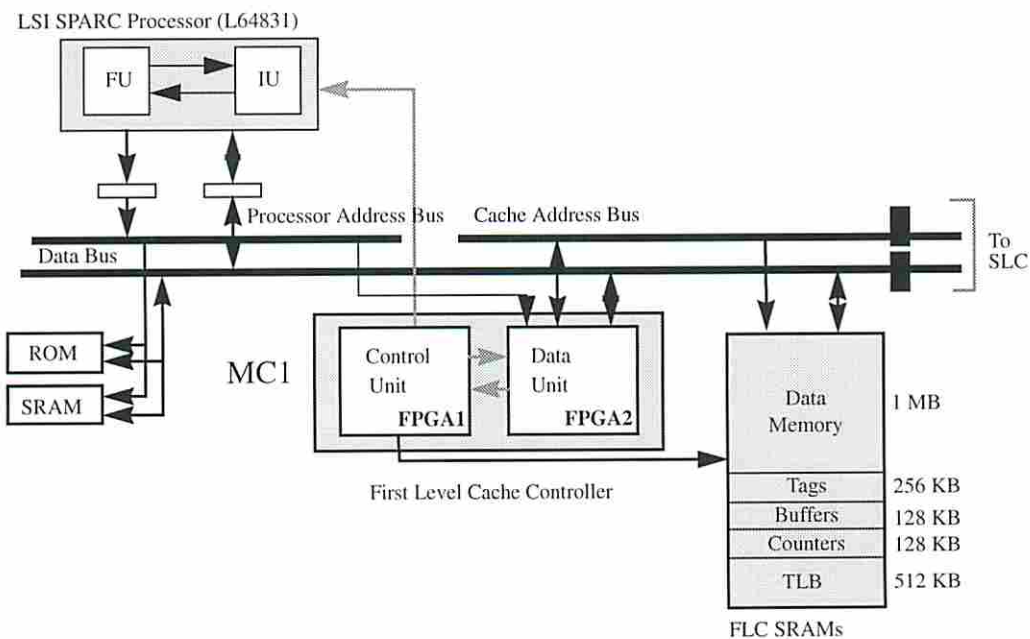
---

3. The SPLASH benchmarks are scientific benchmarks commonly used to compare the effectiveness of architectural features of multiprocessors [3]

write-back cache and MC3/RAM3 is the main memory.

Non-blocking loads are supported in the second-level cache. These non-blocking loads are issued by the compiler to direct the second-level cache to prefetch cache blocks before the data is needed. A virtually unlimited number of prefetches can be pending at any time in the second-level cache. When a store is issued by the processor the store is always propagated to the second-level cache. If the store misses in the first-level cache no block is allocated (no allocation on store misses). Under strong ordering of memory accesses (enforcing sequential consistency), the first-level cache and the processor block on a write access that misses or that requires coherence activity in the second-level cache. Under weak ordering of memory accesses there can be a write buffer between the first- and second-level caches (first-level write buffer) and between the second-level cache and the internal bus (second-level write buffer). The second-level write buffer can be assisted by a write cache (WC), which is a small cache keeping track of partially modified blocks [3]. The first protocol is a pure write invalidate protocol but we intend to implement write-update and competitive-update protocols as well as hardware-based prefetching, described in [3].

FIGURE 4. The Processor and its First-Level Cache

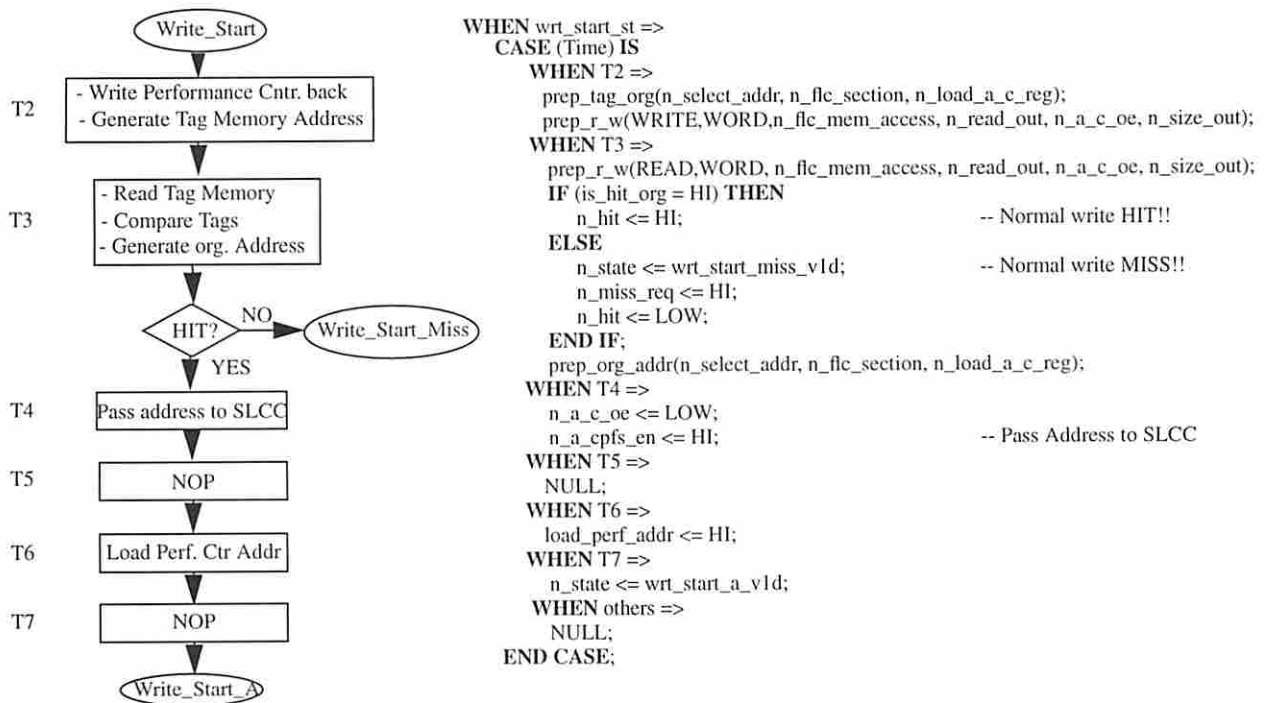


The emulation of each pclock is implemented by a combination of control in the FPGAs

and buffer space in the RAMs attached to them. This buffer space is used to emulate special memories, such as cache directories, TLB (Translation Lookaside Buffer), which is part of the hardware for virtual-memory support [2], or special buffers such as write buffers between the first- and the second-level caches and between the second-level cache and the internal bus. Additional memory space (called *count memory*) is reserved for event counters (see Section 5).

The SRAM implementing the **first-level cache** (FLC) is divided into five parts (see Fig. 4): the data memory (up to 1 Mbyte), the cache directory, the Translation Lookaside Buffer (TLB), the space for the emulation of prefetch and write buffers, and the space dedicated to the collection of performance statistics. The controller is partitioned across two FPGAs: one for the control unit and the other for the data unit. Currently the first-level cache is write-through and direct-mapped with a block size of 16 bytes.

FIGURE 5. Partial Flowchart and its VHDL Code for a Write Cycle in the First-level Cache

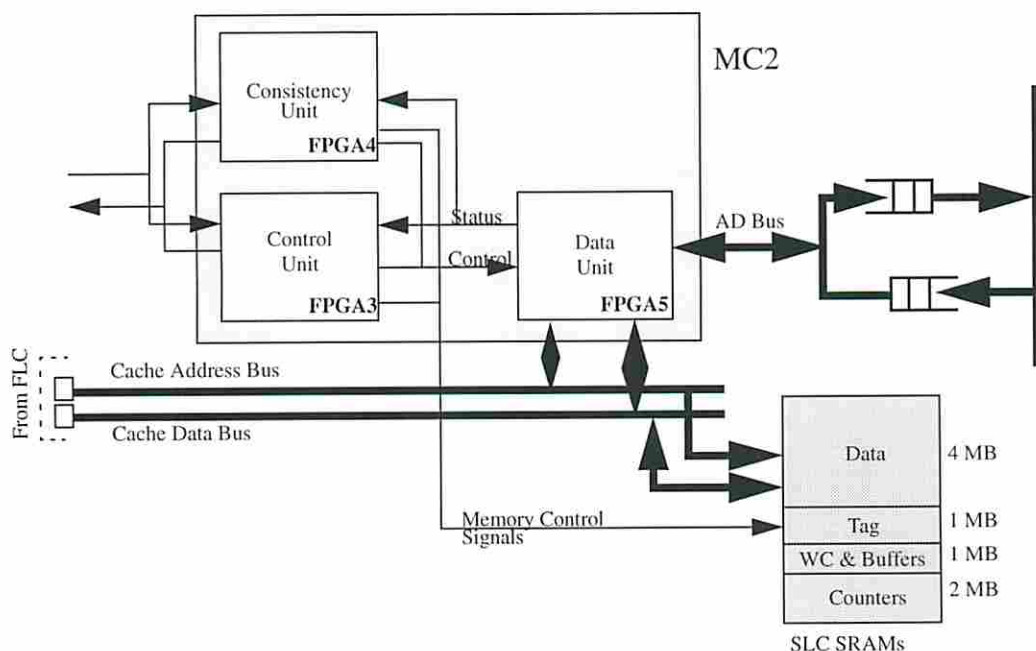


A typical data read cycle in the first-level cache consists of receiving an address from the processor, translating the address in the TLB space, accessing the cache directory space, fetching

the data from the cache data space, fetching the counter in count memory for that event, updating the counter and returning the data to the processor (all this is done in eight cycles). Fig. 5 shows a partial flowchart for a write access in the first-level cache and its VHDL implementation. (This sequence has no TLB access.) This sequence emulates one pclock in the processor write cycle. As can be seen, the updates of event counters in count memory are implemented within the control sequence and are pipelined. T0 and T1 are missing in the flowchart because they are common to all accesses.

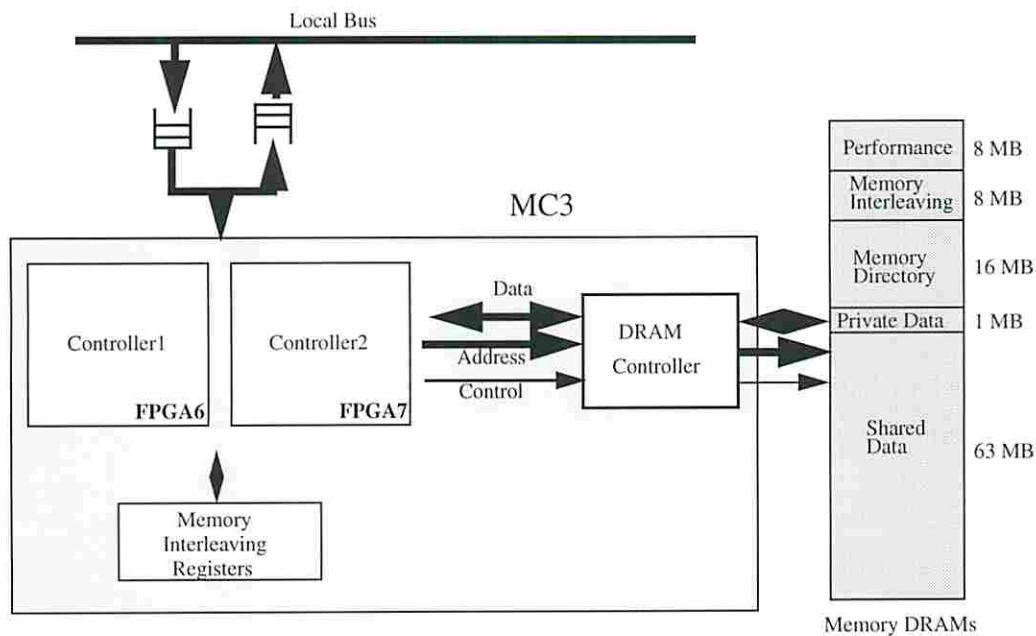
The **second-level cache** (SLC) is implemented by MC2/RAM2. The current configuration is a two-way set-associative write-back cache and a block size of 16 bytes (see Fig. 6). Half of the memory (up to 4 Mbytes) is for the second-level cache data memory. The other half is dedicated to the cache directory, various buffers (second-level write buffer and write cache) and count memory. The second-level cache controller is by far the most complex controller of the machine. It is implemented by three FPGAs: the data unit is the datapath for the controller, the control unit implements the basic cache control mechanisms and the consistency unit contains the control for the mechanisms enforcing access ordering and for the write buffers and prefetching hardware [3].

FIGURE 6. The Second-Level Cache Architecture.



The **main memory** is implemented in MC3/RAM3. 64 Mbytes are devoted to the data memory (1 Mbyte for private data and 63 Mbytes for shared data and code), 16 Mbytes to the directory (one 32-bit word per 16-byte block), 8 Mbytes for performance counters and 8 Mbytes to support the emulation of various hardware mechanisms, including the *virtual interleaving* of on-board memory (Fig. 7).

FIGURE 7. The Main Memory and its Controllers



If current trends continue and the speed gap between processor and DRAM speeds keeps widening, the conflicts at the memory modules will be so high that some form of on-board memory interleaving will be required (by interleaving we mean that several independent memory banks with their own controller can process memory requests concurrently). In the testbed, this interleaving effect can be obtained by multiplexing in time the memory controller. *Virtual interleaving* of memory relies on the fact that a large number of cycles are available in the emulator to emulate memory transactions and is supported by a set of eight interleaving registers (for up to eight interleaved banks) and some buffer space in memory. Each interleaving register contains a counter which is decremented at every pclock. When the register content is null the memory bank is free; otherwise the register content indicates the number of plocks until the bank is released. A



request to a busy bank is simply rejected by the controller.

A typical memory transaction has three phases: *prelude*, *suspension* and *completion*. During the prelude the packet is received and decoded and the directory is accessed to find the state of the block. To suspend the request, a transaction completion record containing all the information needed to resume and complete the request is stored in a memory location corresponding to the memory bank and the interleaving register is filled with a value in plocks corresponding to the suspension time. While the request is suspended, a different (free) virtual memory bank can be accessed. When the interleaving register reaches zero, the memory controller is interrupted; it then fetches the transaction completion record in memory and completes the transaction by (possibly) sending some messages (completion phase). Virtual memory interleaving is further illustrated in Section 4.

Many parameters can be changed easily in the systems above, within limits. All cache and TLB parameters can be changed. Latencies and bandwidths of various components can be changed. We can explore latency tolerance techniques such as non-blocking second level cache with write buffers and write caches as well as data prefetching hardware. We can experiment with various directory structures [11]. We can change the protocols (write-invalidate, write-update, competitive-update), or we can implement multiple protocols and apply different protocols on a per-page basis.

### 3.3. FPGA Statistics

Table 1 gives the current statistics on the FPGAs (These statistics are for strong ordering of memory accesses with no latency tolerance hardware except for prefetch support in the second-level cache.) CLBs or Configuration Logic Blocks [12] are the logic building blocks in the Xilinx XC4013 FPGAs. Each block contains hardware to implement random logic. A CLB is packed when all of its logic is used and it is simply occupied when any of its logic is used. The number of block (CLB) levels is indicative of the delay through the chip. The delay through one CLB is variable. Finally, the total number of pins in the XC4013 is 232. As can be seen, the CLB utilization

and the delays of a few FPGAs are dangerously close to their maximum.

Table 1: FPGA Statistics

Controller	FPGA Name	Occupied CLBs (%)	Packed CLBs (%)	Block Levels	# of IO pins used
MC1	Control Unit	82	65	16	107
	Data Unit	51	44	6	173
MC2	Control Unit	77	58	15	118
	Consistency Unit	not used	not used	not used	not used
	Data Unit	64	56	11	171
MC3	Controller1	88	78	9	124
	Controller2	64	54	8	81

There are two ways to improve on this situation. The first one is to replace the Xilinx XC4013 with Xilinx XC4025 as soon as they become available. This will double the capacity of each FPGA. Higher capacity implies that more complex or faster designs can be mapped. The second way is to upgrade the CAD tools to obtain more optimized mappings. Our current tools are far from optimum for compiling and mapping VHDL designs.

### 3.4. Emulation of Other Architectures

Given its generic board architecture, the testbed is capable of emulating in full detail very complex target multiprocessors with vastly different architectures, provided they fit the generic block diagram of Fig.1. Examples are:

- **CC-NUMAs with Linked-List Directories:** Instead of a centralized directory located at the home memory, the directory is distributed by linking caches containing a copy of the block through a set of hardware pointers in each cache entry. This organization is adopted in the Scalable Coherent Interface (SCI) standard [11].
- **COMAs (Cache-Only Memory Architectures):** This is the architecture of machines such as DDM [7]. In this case there is no system memory and MC3/RAM3 acts as a huge cache (also called *attraction memory*) and the DRAM directory is replaced by the *attraction memory state*. In this

configuration, MC1/RAM1 and/or MC2/RAM2 may be configured as caches.

- **MPS (Message-Passing Systems):** In this configuration MC3/RAM3 acts as the local (private) memory of the processor, MC2/RAM2 acts as a Message Passing Controller (MPC), and MC1/RAM1 acts as the processor cache. The functions of the MPC are to buffer messages sent or received by the processor, to format out-going packets according to the protocol in the target, to decode the messages received, and to interrupt the processor when messages are received. This message-passing architecture can also include hardware primitives to support virtual shared memory efficiently [9].

- **Mixed Shared-memory and Message Passing Systems:** All the shared-memory organizations can be augmented with a message-passing facility for bulk transfers of data among processors.

#### 4. KEEPING TRACK OF TIME: TIME SCALING

Since the speeds of the emulator and the testbed are different, timings measured on the emulator must be related to the timings in the target machine. Rather than keeping track of simulated time through event-driven mechanisms and timestamping (as is done in software simulators [6]), time is *scaled*. *Time scaling* preserves the *relative* timing of components in the emulator and in the target, and absolute times in the target are derived from executions on the host by simple scaling arguments. For example, it is intuitively obvious that the processor utilization in a system with processors running at 100 MHz and with average memory latencies of 100 nanoseconds is equal to the processor utilization in a system with the same architecture but with processors running at 1 MHz and with average memory latencies of 10 microseconds. All performance metrics can be scaled this way, and therefore we do not have to build the system with the most up-to-date and fastest technology provided we scale memory, interconnection and processor speeds appropriately.

Any component is characterized by two fundamental performance measures: latency and bandwidth. These two measures can be independent. For example, two networks can have the same latency but one may have more bandwidth because it has more links; similarly, the band-

width of a memory can be increased (while its latency remains the same) by interleaving it. Another important factor is the width of the data paths. In the testbed, all data paths on board are 32-bit wide and the system bus (FutureBus+) is 64-bit wide. We can emulate one cycle of a data path with 64, 128 or even 256 bits by 2, 4, or 8 cycles of a 32-bit testbed data path.

A convenient unit for all timings is the *pclock*-- the clock period of the processor. If the latencies of all components are expressed in terms of pclocks and if all component bandwidths are expressed in terms of bytes per pclock, then systems with components of equal latencies and bandwidths are equivalent.

In the testbed, a pclock is currently eight cycles. This gives the emulator eight cycles to simulate all the activities occurring in one pclock in the target system. To simulate variable latencies, we delay requests. To simulate variable bandwidth of a given resource, we must vary the number of cycles that each request keeps the resource busy. For example, the latency of local memory can be increased by delaying the requests; its bandwidth can be decreased by inserting dummy cycles in the control sequence of the memory controller or it can be increased through memory interleaving.

The Delay Unit (DU) simulates variable latencies in the interconnection (emulated by the FutureBus+). We do not attempt to simulate variable interconnection bandwidths; rather, the pclock and the interconnect speed of the testbed are such that conflicts in the interconnect are negligible. Therefore, for all practical purposes the bandwidth can be considered infinite. The accuracy of measurements requires that the FutureBus+ traffic is monitored so that excessive traffic leading to significant conflicts and delays at the bus are detected.

To illustrate time scaling and virtual memory interleaving, consider the simple case where a miss occurs in the second-level cache, the block address maps to the local on-board memory (local miss), and the block is uncached elsewhere. In the target system,  $T_m$  is the latency of the miss and the memory is  $n$ -way interleaved so that the memory can deliver  $n$  blocks every  $T_m$  pclock. The testbed has a monolithic memory controller, which must emulate the operations of

the  $n$  parallel memory controllers of the target system. Therefore the memory transaction must keep the controller busy for  $T_m/n$  pclocks by idling the controller before suspending the miss request (see Section 3.3). On the other hand, the total latency of the miss in the testbed must be  $T_m$  pclocks, the same as in the target system. Let  $T_p$ ,  $T_i$ ,  $T_s$  and  $T_c$  be respectively the prelude time, the idle time, the suspension time, and the completion time of the miss request in the testbed. (All times are in pclocks.)  $T_p$  and  $T_c$  are known from the testbed implementation. To find  $T_i$  and  $T_s$  we have the following constraints:

$$T_m = T_p + T_i + T_s + T_c \quad \text{and} \quad (1)$$

$$T_m/n = T_p + T_i + T_c \quad (2)$$

(1) enforces the same latency in the testbed and in the target whereas (2) enforces the same utilization of the controllers in both systems.

Unknown  $T_i$  and  $T_s$  are then given by:

$$T_i = T_m/n - T_p - T_c \quad \text{and} \quad T_s = T_m(n - 1)/n \quad (3)$$

The target system has the timing characteristics shown in Figure 8 (left column). This example uses reasonable parameter values. Designs using such parameters have been simulated with the SPLASH benchmarks and work well. In this example, the memory is assumed to work in page mode: the access time to a 64-bit word takes 100nsec and the access time to a 128-bit block takes 140nsec in the target. Because the target pclock is 5 nsec and the testbed pclock is 800 nsec, a 140nsec block access time in the target translates into  $140/5$  pclocks or  $28 \times 8 = 224$  clocks in the testbed. This is a huge number of cycles and therefore the memory controller of the testbed can emulate very complex directory mechanisms during this time. The timing is tighter in the second level cache because it is built with SRAMs. In the example, the second-level cache has 6 pclocks or 48 clocks to emulate a block transfer between second- and first-level caches. As the block size increases, the number of extra cycles available for emulation also increases.

FIGURE 8. Simple Example Illustrating Time Scaling

**Target System:**

- 200 MHz PROCESSORS; 1 PCLOCK = 10 NSEC
- INSTRUCTION FETCH: 1 PCLOCK
- LOAD: 2 PCLOCKS
- 100 NSEC (20 PCLOCKS) PER WORD ACCESS IN MEMORY
- 64-BITS DATA PATHS
- 16-BYTE BLOCKS
- MEMORY AND 2-ND LEVEL CACHE ARE 64-BIT WIDE
- 6 PCLOCKS PER MISS FROM SLC TO FLC
- 28 PCLOCKS PER BLOCK FETCH IN MEMORY
- MEMORY IS TWO-WAY INTERLEAVED
- MAX MEMORY BANDWIDTH: 1 BLOCK PER 14 PCLOCK

**Testbed Configuration:**

- FLC INSTRUCTION FETCH: 8 CLOCKS
- FLC DATA FETCH: 8 CLOCKS
- 28 PCLOCKS OR 224 CLOCKS IN MEMORY PER MISS
- 6 PCLOCKS OR 48 CLOCKS AVAILABLE TO SLC TO TRANSFER BLOCK TO FLC.
- PRELUDE TIME: 30 CLOCKS
- COMPLETION TIME: 30 CLOCKS
- IDLE TIME:  $112 - 30 - 30 = 52$  CLOCKS
- SUSPENSION TIME:  $224/2 = 112$  CLOCKS

The scaling of I/O is very natural in the testbed. The I/O bandwidth is 1.25 Mbytes per second. This bandwidth matches the I/O requirements of the testbed, which has a peak processing rate of 10 MIPS. Note that when we simulate target systems with faster processors, we do not have to adjust I/O bandwidth because the testbed runs at 10 MIPS peak whatever the target system speed is. So, I/O bandwidth scales automatically. The other issue is I/O latency, which must be scaled. The service of I/O requests issued by the testbed must be delayed as faster target processors are emulated. This scaling of I/O latency should be done in software.

## 5. COLLECTING PERFORMANCE DATA

The primary mechanism to collect performance data involves event counters stored in a special area of each memory called *count memory*. In each of the three on-board memories, a set of counters keeping track of the occurrence of mutually exclusive events are updated every time a transaction is completed in the controller. These events are mutually exclusive, so that only one counter is updated at a time in each memory. Thousands of counters are present in each memory, meaning that thousands of different events can be counted. At the end of the emulation run, the counters are uploaded and post-processed (basically they are added together) to obtain the required performance data. This counting mechanism can be started and stopped under software control.

FIGURE 9. Generation of Addresses for Event Counting in the First-Level Cache's Count Memory

Counter Address	Private/ Shared	Read/ Write	Hit/ Miss	Basic Event
0	0	0	0	Shared-Write-Miss
1	0	0	1	Shared-Write-Hit
2	0	1	0	Shared-Read-Miss
3	0	1	1	Shared-Read-Hit
4	1	0	0	Private-Write-Miss
5	1	0	1	Private-Write-Hit
6	1	1	0	Private-Read-Miss
7	1	1	1	Private-Read-Hit

Fig. 9 illustrates how the addresses in count memory are generated for mutually exclusive events in the first-level cache. In this simple example, three signals --programmed in MC1-- correspond each to one property of an access in the first-level cache. One signal indicates whether the access is to private or to shared data. The second line is the read/write signal from the processor and the third signal is high when the access hits and low when the access misses in the first-level cache. The combination of these signals forms a three-bit address, which can be used to address a counter in the area of RAM1 allocated to count memory. In this example there are only eight addresses, but in a practical situation up to 20 signals can be defined in each controller. For instance one signal may also distinguish between instructions and data, and, in the case of instructions, the opcode could also be a field of the address of the performance counters. (In this case, instruction types are histogrammed and, at the end of the program execution, we can obtain a dynamic instruction mix of the program by summing together the counters with addresses in the count memory having the same opcode field.)

## 6. PROGRAMMING THE TESTBED

In the testbed, seven FPGAs implement the controllers for the caches and the main memory. The testbed behavior is dictated by the protocols that these controllers execute. Therefore the testbed can be programmed by mapping the controllers for the target design into the FPGAs.

Currently, the data path and the RTL description of each controller are specified in

VHDL. The netlist for each controller is then generated by synthesizing the VHDL description using ViewLogic's ViewSynthesis tool. The netlists are then mapped into the FPGAs using the Xilinx mapping tools.

Design parameters such as block sizes, cache sizes, latencies and bandwidth are specified as constants in the VHDL descriptions. Therefore the target machine parameters can easily be changed by modifying these constants in the designs and then recompiling the VHDL codes. This approach requires that each design is recompiled for each parametric change. An alternative approach would be to design the controllers such that the machine parameters are stored in registers inside the FPGAs. Changing parameters would only require that the content of these registers be changed. This approach would lead to slower and more complex designs, but would save time in recompiling the design for parametric changes.

In order to change the functionality of the target machine, controllers that implement different protocols have to be mapped into the FPGAs. Using the current approach, modifications to the functionality of each FPGA require the design and specification of the RTL description of each controller in VHDL. In reality, however, a large number of common functions are shared by all the possible designs for a given FPGA. Nevertheless, the task of re-programming the FPGAs is error-prone and increases the turnaround time for emulating different machines on the testbed. In the future, we expect that the turnaround time for this design stage will be reduced as behavioral compilers become more widely available. A behavioral compiler is a high-level synthesis tool which accepts an algorithmic description of a circuit to create the hardware. This algorithmic description could be derived directly from the high-level description of each protocol. In the short term we will develop libraries of parameterized designs for every possible configuration of each FPGA.

## **7. EXPECTED PERFORMANCE OF THE TESTBED**

Table 2 shows the slowdown factor (which is the processing speed ratio between the target system and the emulator) for various uniprocessor technologies in the target systems. These slow-



down factors can be predicted accurately in our emulation approach and they are independent of the number of processors (from 2 to 8) since the number of processors in the testbed and in the target is the same.

Table 2: Slowdown Factors Between Target and Testbed

Target Uniprocessor Speed	50 MIPS	100 MIPS	200 MIPS	500 MIPS	1 GIPS
Slowdown	40	80	160	400	800
Testbed Time per Second of Target Execution	40 seconds	1 minute and 20 seconds	2 minutes and 40 seconds	6 minutes and 40 seconds	13 minutes and 20 seconds
Testbed Time per Minute of Target Execution	40 minutes	1 hour and 20 minutes	3 hours and 40 minutes	6 hours and 40 minutes	13 hours and 30 minutes

Table 2 also shows the testbed time needed for different execution times of the target. Clearly we can reasonably expect to obtain experimental points for serious workloads for systems with processors of up to 1 GIPS. Some of these experiments would take months to run on a current software simulator. Currently it is difficult to simulate more than a fraction of a second of target execution time on a software simulator. These simulations must be very simplified and abstracted even to reach one second of target execution. On the other hand the testbed emulates the target system in all its details. It is an actual implementation of the target and it demonstrates the feasibility and hardware complexity of the target. The numbers in Table 2 also show that the testbed can provide valuable information for systems built until the end of this decade.

## 8. COMPARISON WITH OTHER APPROACHES

The methodologies for evaluating multiprocessor systems have slowly evolved in the past. Initially, analytical models or trace-driven simulations were used. These approaches evaluate hardware systems at a coarse level of detail. Recent breakthroughs in simulation methodology have opened the possibility of efficient, detailed and flexible evaluations.

### 8.1. Software Simulation

Event-driven simulators need some mechanism to schedule events, which include instruction exe-

cution and memory system operations. In an execution-driven simulator such as Tango [4], or the Wisconsin Wind Tunnel (W.W.T.) [10] each instruction is run directly on the host machine. The processors and the components of the memory systems are simulated as processes or threads, and each event requires a context switch. The source or the binary codes are instrumented to avoid scheduling an event at every instruction execution. Instrumentation code is added at basic block boundaries and at each “sensitive” data access (usually shared-data accesses). This code reduces the frequency of context switches. By contrast, program-driven simulators such as Cache-Mire [1] interpret each instruction in software. Cache-Mire also relies on activity scanning (rather than an event list) for scheduling activities from processors and memory. The pay-offs are that 1) the whole simulation runs in the same context, saving expensive context switches, and 2) instrumentation of the code is not required. The performance impact of interpreting target instructions in software depends on the complexity of the memory model. If the memory system is complex, our experience with Cache-Mire shows that instruction interpretation represents less than 10% of the total simulation time.

Software simulation, whether it is execution- or program-driven, is slow. A simulator such as Tango or Cache-Mire can execute in the order of 10,000 instructions of a target multiprocessor with a complex memory model per second on a current 50 MIPS workstation. This low performance is due to multiple factors:

- the overhead due to event scheduling (e.g., context switching and event list management or activity scanning);
- the code expansion due to code instrumentation to keep track of target instruction execution times in execution-driven simulators (which was reported to be between 2 and 3 in [4]) or the overhead of decoding and executing instructions in program-driven simulations;
- the management of timestamps associated with events;
- the collection of performance data;
- the semantic gap between hardware mechanisms and their execution on the simulator (the fact

that each basic activity which takes one cycle in the hardware of the target takes several instructions to simulate on the host); and

- the speedup of the target multiprocessor.

To keep simulation times reasonable, the data set sizes of the workload must be drastically reduced. Observations made on the small data set sizes must then be extrapolated to the workload with the actual data set size, a difficult task which has never really been validated. Nonetheless, because this is the only possible approach with current simulation technology, the research community has embraced the results from such studies.

A common drawback of all simulations is that they abstract the behavior of the target multiprocessor. Many effects are ignored or approximated, on the premise that they are negligible. In some cases some key hardware components and physical events are totally removed from the simulation. For example, it is not uncommon that a simulator avoids simulating the caches and the data transfers among them. The validity of these simplifications is usually not verified and relies on the experience and judgement of the evaluator. Overall, simulators can give a good idea of the qualitative differences among various software or hardware mechanisms and are suitable for exploratory investigation of the design space, but they are less reliable for evaluating a particular design quantitatively and for verifying the design. Besides performance, design verification is a critical issue in these complex systems; simulation and formal techniques can help but they are so abstracted that they cannot detect all design errors. Furthermore, simulators often yield little insight in the complexity of the actual implementation.

## **8.2. Parallel Software Simulation**

A classical way to speed up simulation is to try to parallelize it. However, the parallelization of discrete-event simulation is a notoriously hard problem [6, 10]. The key problem is the reception of a message by a simulation process with a higher timestamp than the timestamp of the message. Approaches to solving this fundamental problem are classified into conservative and optimistic approaches. Conservative approaches synchronize processors to avoid the reception of

messages with lower timestamps. Optimistic approaches allow the reception of such messages by keeping track of the simulation events generated by the process and by rolling back the execution. In the case of the simulation of multiprocessor systems, a natural decomposition consists of allocating one processor of the host to one or multiple processors of the target, as was done in the Wisconsin Wind Tunnel (W.W.T.) [10], which adopts a conservative approach to parallel simulation. Added overheads of W.W.T. over sequential simulation include:

- barrier synchronization among processors every 100 (or so) cycles of the target and
- the processing of messages and their timestamps.

### 8.3. Comparison Between the Testbed and Software Simulation

Trying to compare the efficiency of the testbed with existing simulators is a hazardous task at best. However, we have used two state-of-the-art software simulators, Cache-Mire [1] and TangoLite (an optimized multithreaded version of Tango [4mai]) to try to quantify the relative performance of the two approaches.

We have run Cache-Mire simulations of some SPLASH benchmarks on a Sun Sparc Station 10 Model 30 with 128 Mbytes of main memory and no off-chip cache. This machine is rated at about 40 SPEC MIPS (36 MHz CPU). The benchmarks we have run are: MP3D with 10K molecules, for 10 iterations, Water with 64 molecules, and Cholesky with matrix bcsstk14. In each case we have simulated an 8-processor system with a complex memory system. The network in the simulation has constant latency and infinite bandwidth; the memory accesses are weakly ordered [5] and there are write buffers with multiple outstanding requests in both the first-level and the second-level caches [3]. The *simulation rate* of Cache-Mire is the number of cycles of the target system simulated per second. These numbers are shown in Table 3.

We have also experimented with TangoLite and the results were quite similar to the ones obtained using Cache-Mire. The TangoLite simulations were executed on a SGI Indigo workstation which is about twice as fast as the SparcStation 10 used for the Cache-Mire simulations (75

MHz MIPS 4400, 1 MByte secondary cache). Moreover, the memory model used in these simulations was considerably simpler than the one used in the Cache-Mire simulations. These simplifications included no simulation of instruction fetches, no simulation of the memory/directory modules or the local busses, strong (instead of weak) ordering of memory accesses (which practically means no write buffers), and very little gathering of performance data. The simulation rate figures for TangoLite and MP3D were 7877 cyc/sec for eight processors, 3495 cyc/sec for sixteen processors and 1186 cyc/sec for thirty two processors. That makes it about a factor of two faster than the numbers for Cache-Mire in Tables 3 and 4, on a machine twice as fast.

Table 3: Comparison Between the Testbed and Software Simulation

Benchmark (Number of processors)	Number of References (Inst + Data)	Simulation Rate (Cache-Mire) (cycles/sec)	Simulation Rate (Testbed) (cycles/sec)	Speedup (Testbed/Cache- Mire)
MP3D (8)	18.5M	3,786	1.25M	330
Water (8)	136.5M	3,960	1.25M	315
Cholesky (8)	79.5M	3,426	1.25M	365

Of course all these simulations are much more abstracted and simplified than the emulation (for instance, in the simulation, data movements are not simulated.) The level of implementation details in the testbed is actually closer to a cycle-by-cycle, register-transfer level simulation. In practical cases, such simulations run at the rate of a few cycles per second and the speed up of the testbed over these detailed simulators is around one million.

#### 8.4. Prototypes

Multiprocessor testbeds were developed for research purposes in the 70's. However, at that time, the goal was to experiment with parallel software and the hardware was not configurable. Therefore, the experimental evidence obtained with them was valid only for the particular hardware.

Breadboard prototypes are also extensively used in industry to validate a new architecture and in academic research projects to explore architectures not pursued by industry. One basic problem with breadboard prototypes is their cost and the relatively low amount of research infor-

mation they provide besides proving one point. For example, recently, one such prototype, the Stanford DASH [8] was built to demonstrate the feasibility of cluster-oriented, directory-based CC-NUMAs. DASH was one of the rare successful academic building projects and it succeeded by making many compromises in order to use existing hardware as much as possible. Besides forming students and keeping some sanity checks on the simulation experiments, little was learned by actually building the prototype.

### **8.5. The Hardware Emulation Approach**

The approach advocated in this project and demonstrated by the testbed is intermediate between software simulation and prototyping. It does not replace these approaches but it complements them. Emulation is faster, more reliable and more faithful to the target system than software simulation. An emulator is a possible hardware implementation of the architecture. It is a fully functional machine on which any workload of the target machine can run. An emulator is also more flexible and easier to build than a prototype and, in terms of research, we can expect to learn more from building an emulator.

The limitations of the current emulator are the small number of processors and the relatively low pclock rate. We were somewhat conservative in our design because of our limited experience with FPGAs and with FPGA CAD tools. With current CAD tools and FPGA technologies the clock rate could easily be raised to 20MHz, and it is obvious from past trends that this clock rate will increase rapidly every year. Moreover the number of clocks in each pclock could also be cut in half if we used a processor with an on-chip instruction cache so that we do not have to emulate instruction fetches. Such an emulator would emulate 5 Millions cycles of the target per second. Large speedups can be expected for bigger machines, with 128 or 256 processors. Table 4 shows the expected speed up over simulation for a 16- and a 32-processor emulator. As can be seen from the table, the speedup over simulation is superlinear. This happens because of two main factors: (1) as the number of processors increases, so does the number of caches, memory modules and other resources in the system that have to be checked at every global event; (2) as the size

of the target system grows the effect of contention for global resources is higher and adds to the overhead of both activity scanning and event-list simulation mechanisms.

Table 4: Performance of an Emulator clocked at 20MHz and with 4 clocks per pclock

Benchmark (Number of processors)	Number of References (Inst + Data)	Simulation Rate (Cache-Mire) (cycles/sec)	Simulation Rate (Emulator) (cycles/sec)	Speedup (Emulator/Cache- Mire)
MP3D (16)	18.5M	1,856	5M	2,694
Water (16)	136.6M	1,868	5M	2,677
Cholesky (16)	113.8M	1,635	5M	3,058
MP3D (32)	18.7M	525	5M	9,524
Water (32)	136.6M	411	5M	12,165
Cholesky (32)	197M	402	5M	12,438

With respect to flexibility, emulations with the testbed is limited to machines with the overall organization of Fig. 1. However, in the future, programmable interconnect technology may be used to increase the flexibility of emulators. Like the crossbar switches used for years in communication and test systems, Field Programmable Interconnect Components (FPICs) at the chip level allow multiple input signals to be directed to multiple output pins. In the future it will be possible to build emulators with FPGAs and FPICs; these emulators will have greater flexibility to be configured for different target systems.

## 9. CONCLUSION

The first goal of the U.S.C. multiprocessor testbed is to develop, demonstrate, and exploit a novel approach for the rapid prototyping of multiprocessor systems. The approach is based on hardware emulation. It relies on emerging technologies such as Field Programmable Gate Arrays (FPGAs) and advanced CAD tools (e.g., VHDL and its hardware synthesizers). The second goal is to provide a vehicle for the verification of complex multiprocessor systems. The third goal is to explore various multiprocessor models on the same, configurable hardware platform for processor speeds up to 1 GIPS (i.e., for processors which will be available during this decade).

To this end a processor emulator system with eight processors has been built. The processors are SPARC processors clocked at 1.25 MHz, which means that the testbed can emulate complex systems in full detail at the speed of 1.25 million cycles per second for systems of up to 8 processors. The testbed can be configured into various multiprocessor systems, including CC-NUMA (Cache-Coherent Non-Uniform Memory Access) architectures, COMAs (Cache-Only Memory Architectures), Message-Passing Systems, and Virtual Shared Memory Systems. Among these configurations many different hardware mechanisms can be implemented. The testbed will be the first hardware platform on which we can compare the effectiveness of the different architectures and their variants for realistic workloads such as operating system kernels, database systems, and multitasked scientific workloads with the real data set sizes for which the target machine is built.

## 10. Bibliography

- [1] Brorsson, M., Dahlgren F., Nilsson, H., and Stenstrom, P. "The CacheMire Test Bench: A Flexible and Effective Approach for Simulation of Multiprocessors", In *Proceedings of the 26th Annual Simulation Symposium*, pp. 41-49, March 1993.
- [2] Catanzaro, B., *Multiprocessor System Architectures*. Prentice-Hall, 1994.
- [3] Dahlgren, F., Dubois, M., and Stenstrom, P., "Combined Performance Gains of Simple Cache Protocol Extensions," *Proceedings of the International Symposium on Computer Architecture*, pp. 187-199, April 1994.
- [4] Davis, H., Goldschmidt, S.R., and Hennessy, J., "Multiprocessor Simulation and Tracing Using Tango," *Proceedings of the Parallel Processing Conference*, pp. II99-II107, August 1991.
- [5] Dubois, M., Scheurich, C., and Briggs, F.A., "Synchronization, Coherence, and Event Ordering in Multiprocessors," *IEEE Computer*, Vol. 21, No. 2, pp. 9-21, February 1988.
- [6] Fujimoto, R.M., "Parallel Discrete Event Simulation," *Communications of the ACM*, Vol. 33, No. 10, pp. 30-53, October 1990.
- [7] Hagersten, E., Landin, A., and Haridi, S., "DDM -- A Cache-Only Memory Architecture," *IEEE Computer*, Vol. 25, No. 9, pp. 44-54, September 1992.
- [8] Lenoski, D., Laudon, J., Gharachorloo, K., Weber, W.-D., Gupta, A., Hennessy, J., Horowitz, M., and Lam, M.S., "The Stanford DASH Multiprocessor," *IEEE Computer*, pp. 63-79, Vol. 25, No. 3, March 1992.
- [9] Li, K. and Hudak, P., "Memory Coherence in Shared Virtual Memory Systems," *ACM Transactions on Computer Systems*, Vol. 7, No. 4, pp. 321-359, November 1989.



- [10] Reinhardt, S., Hill, M.D., Larus, J.R., Lebeck, A.R., Lewis, J.C., and Wood, D.A., "The Wisconsin Wind Tunnel: Virtual Prototyping of Parallel Computers," *Proceedings of the ACM Sigmetrics Conf. on Measurements and Modeling of Computer Systems*, pp. 48-60, May 1993.
- [11] Stenstrom, P., "A Survey of Cache Coherence Schemes for Multiprocessors," *IEEE Computer*, Vol. 23, No. 6, pp. 12-24, June 1990.
- [12] Trimberger, S., "A Reprogrammable Gate Array and Applications," *Proceedings of the IEEE*, Vol. 81, No. 7, pp. 1030-1041, July 1993.