

FPGA Framework for CMP

Kenneth Østby

Master of Science in Computer Science

Submission date: June 2007

Supervisor: Morten Hartmann, IDI

Co-supervisor: Marius Grannæs, IDI

Problem Description

The project's main goal is to develop a platform supporting experiments of Chip Multiprocessors on a Field Programmable Gate Array. The NCAR Group performs research on the simulation of architectural variations in Chip Multiprocessors. It's desirable to extend the experiments to run on parametrizable architectures designed for Field Programmable Gate Arrays. The solution must account for several processors and their interconnection, cache structures, and if time permits operating systems and the ability to communicate with existing simulators. Knowledge about FPGA, VHDL and the C programming language is required.

The subtasks are (to the extent time allows them to be addressed):

Background research on related projects

- Multicore solutions in FPGAs, with emphasis on cache solutions
- potential operating systems

Research on cache solutions on an FPGA

- assess the option of customizable solutions and automated synthesis of cache systems

Make several cores work on existing lab equipment in a simple testing environment

Research and identify suitable cores for the system and possibly an operating system

Cache design with emphasis on customizability and suitable cache hierarchy (L1/L2) on available FPGA cards (Nalle)

Investigate possibilities towards integration with simulation environments (M5 and/or SimpleScalar) and experiments at NCAR

Run suitable experiments and analyse results

Assignment given: 20. January 2007

Supervisor: Morten Hartmann, IDI

Abstract

The single core processor stagnated due to four major factors. (1) The lack of instruction level parallelism to exploit, (2) increased power consumption, (3) complexity involved in designing a modern processor, and (4) the performance gap between memory and the processor. As the gate size has decreased, a natural solution has been to introduce several cores on the same die, creating a chip multicore processor.

However, the introduction of chip multicore processors has brought a new set of new challenges such as power consumptions and cache strategies. Although thoroughly researched in context of super computers, the chip multiprocessor has decreased in physical size, and thus some of the old paradigms should be reevaluated, and new paradigms found.

To be able to research, simulate and experiment on new multicore architectures, simulators and methods of prototyping are needed by the community, and has traditionally been done by software simulators. To help decrease the time between results, and increase the productivity a hardware based method of prototyping is needed.

This thesis contributes by presenting a novel multicore architecture with interchangeable and easily customizable units allowing the developers to extend the architecture, rewriting only the subsystem in question. The architecture is implemented in VHDL and has been tested on a Virtex FPGA, utilizing the MicroBlaze microcontroller. Based upon FPGA technologies, the platform is close in nature to that of a chip multiprocessor. The thesis also shows that a hardware based environment will significantly decrease the time to results.

Preface

This Master's Thesis was written as a part of the degree as "Sivilingeniør" in Computer Engineering. The Master's Thesis is founded in an earlier project in the subject TDT4720 – Computer Design and Architecture. The goal of TDT4720 was to give the student an introduction to the current state of Computer Architecture, while introducing him to the tools used in hardware construction.

The work is done for the Norwegian University of Science and Technology (NTNU) at the Faculty of Information Technology, Mathematics and Electrical Engineering, and the Department of Computer and Information Science. The group which hosted the project was the NTNU Computer Architecture and Design Group under the supervision of Associate Professor Morten Hartmann.

The advisor for this thesis was Associate Professor Morten Hartmann. Co-advisor was Research Fellow Marius Grannæs.

Kenneth Østby
June 17, 2007

Acknowledgments

I would like to thank the following persons for their support and input throughout this project. Associate Professor Morten Hartmann for being my advisor, and thus allowing me to partake in this journey through Chip Multiprocessors. Marius Grannæs for being my co-advisor, reading through my thesis several times and always available for technical discussions.

I would also like to thank the people at NTNU Computer Architecture Research Group for including me in their much interesting meetings. Also important are the people at my study room, ITV-458, May Linda Martinsen, Knut Imar Hagen, Christian Larsen, Idar Borlaug, Rolf Anders Syvertsen and Jan Peder David-Andersen.

Finally I would like to send my thanks and thoughts to all of my friends and loved ones not mentioned by name.

Contents

List of Tables	xiii
List of Figures	xv
1 Introduction	1
1.1 Introduction	1
1.2 Problem Description	4
1.3 Project Motivation	5
1.4 Contribution	5
1.5 Scope	6
1.6 Outline	6
1.6.1 Appendices	6
2 Background	7
2.1 Related Work	7
2.2 CMP Architecture	8
2.2.1 Cores	9
2.2.2 Cache	12
2.2.3 Interconnect	14

2.3	Introduction to FPGAs	16
2.4	Cores	17
2.4.1	MicroBlaze	19
2.4.2	PowerPC	20
2.5	Environment	21
2.5.1	Hardware	21
2.5.2	Field Upgradable Systems Environment - FUSE	22
2.5.3	Virtex-E	22
2.6	Tools	23
2.6.1	GCC	23
2.6.2	ISE	23
2.6.3	XPS	23
3	Methodology	25
3.1	Introduction	25
3.2	OMA	26
3.3	Core	27
3.3.1	Wrapping	28
3.3.2	CPU Identification Component	29
3.3.3	Implementation Challenges	30
3.4	Cache	31
3.4.1	Storage Unit	32
3.5	PCI Communication	33
3.6	Software	36

3.7	Benchmark	37
3.7.1	BogoMIPS kernel	37
3.7.2	Load-Store kernel	38
4	Results	39
4.1	Introduction	39
4.2	Benchmark results	40
4.2.1	BogoMIPS Kernel	40
4.2.2	Load/Store kernel	42
5	Discussion	45
5.1	Cores	45
5.2	Cache	46
5.3	Time	47
5.4	3rd Party Challenges	48
5.4.1	GNU Compiler Collection for MicroBlaze	48
5.4.2	Xilinx IP Glue for OPB	49
5.5	Benchmarks	50
6	Conclusion	53
6.1	Conclusion	53
6.2	Future Work	54
	Bibliography	57
	Appendices	61

A	Benchmarks	61
A.1	BogoMIPS kernel	61
A.2	Sample Linker script	62
A.3	Sample Makefile	63
A.4	LoadStore kernel	64
B	Hardware	65
B.1	Core OMA Interface	65
B.2	CPU Identifier	71
B.3	RAM Block	75
B.4	Cache Block	79
B.5	Arbiter	87
B.6	PCI COM	89
B.7	Toplevel - mCache	94
B.8	Toplevel - mCore	99
B.9	Toplevel - sCore	103
C	Software	107
C.1	Controller.c	107
C.2	Controller.h	112
C.3	memory.h	113
C.4	memory.c	114
C.5	cpu.c	118
C.6	cpu.h	122

C.7	util.h	123
C.8	util.c	124

List of Tables

3.1	Signal interface	27
4.1	Processor Configurations	40

List of Figures

1.1	CPU Memory Gap	2
1.2	Project Overview	5
2.1	Memory Hierarchy	9
2.2	Core number 3 is disabled	10
2.3	Fully Associative Cache	13
2.4	Set Associative Cache	13
2.5	Direct Associative Cache	13
2.6	Shared Bus	15
2.7	Crossbar interconnect	15
2.8	FPGA routing	17
2.9	Sample CLB	18
2.10	LUTs per Core	19
2.11	Cores per FPGA	20
2.12	BenERA diagram	21
3.1	Overall Architecture	26
3.2	Core	29

3.3	Single Memory block	33
3.4	2-way set associative block	34
3.5	PCI Unit	34
3.6	Com States	35
3.7	Control Word	35
4.1	BogoMIPS Performance	41
4.2	BogoMIPS memory Access	41
4.3	BogoMIPS Speedup	42
4.4	MemTest Performance	43
4.5	MemTest memory Access	43
5.1	BRAM Location on the FPGA.	47
5.2	Correct assembly code	49
5.3	Output from GCC	49

Chapter 1

Introduction

1.1 Introduction

During the last couple of years, the traditional single core monolithic CPU has failed to further increase its performance proportionally to the decrease in transistor size, as the popular although not correct interpretation of Moore's Law[34] implies. This is, according to Dr. Patterson in his President's Letter[30], due to 3 main factors. (1) Power dissipation, (2) lack of Instruction Level Parallelism to exploit and (3) the long known memory gap which has steadily grown since the beginning of computer science. A fourth factor mentioned by Spracklen et al.[38] is the inherent complexity of designing a processor. A high performance single core CPU requires vast amounts of chip resources to implement the control logic, ensuring that operations don't interfere with each other. This leaves less room for implementing the computational logic, which in turn influences the overall performance of the processor. The Chip multiprocessor (CMP) tries to solve these problems by utilizing several cores inside a single processor. Having several cores on a single chip introduces several new problems, some which have been encountered before in the world of super computing and others new. This includes handling cache in an attempt to reduce the off-chip access, different topologies to allow inter-communication between on-chip processing elements and finally the power dissipation.

The memory gap, as shown in figure 1.1, has long troubled the computer engineers and has been the subject of several research projects. This is also one of the major focal points when researching CMP architectures. The reasons why the memory gap has appeared is a product of several factors. First there is the sheer distance the data signals must travel between the processor and the memory. Data on-chip have a shorter path to travel, and thus have a noticeable decrease in latency compared to off-chip access. Second, the technology and the larger size of memory that exists outside the processor adds to the latency by requiring more time to deliver the requested data onto the bus itself. All of these factors makes it important to limit redundant memory accesses, storing the frequently used

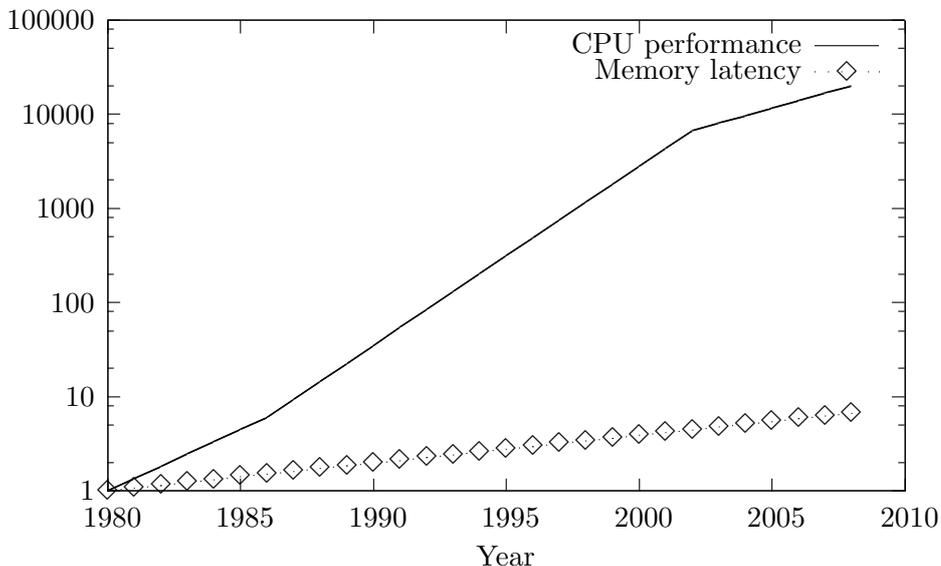


Figure 1.1: CPU/Memory performance[30].
Logarithmic plot for readability.

data in cache banks on the processor, keeping recently used data close in locality. This is done by employing different strategies which forces the processor and the cache banks to cooperate, trying not to evict the frequently used lines of data off the chip. This is an effort to reduce memory access outside the processor itself. Besides reducing the off-chip access much of the cache research studies ways of keeping data requested by the different cores near in locality close to the core which is most likely to request it. This might involve duplication and spreading cache lines around internally on the chip.

Another reason which forced the change of paradigm from single core processors to multi-core processors is the increased difficulty of exploiting Instruction Level Parallelism(ILP). Exploiting ILP is to exploit the fact that certain combinations of instructions can be rearranged without changing the outcome. This is done to avoid stalls in the processor's pipeline, and hence reduce the time spent by the CPU idling. The main challenge exploiting ILP is that it gets exponentially harder per fraction of parallelism level, and thus stagnating at a certain level[15]. To achieve a perfect level of instruction parallelism, unrealisable hardware must be constructed. This to account for perfect branch prediction, unlimited access of registers and all the memory locations must be known beforehand in order to rearrange load and store operations[15]. All needed to achieve perfect ILP. Seeing how the work load on most of the commercial servers are not focused on computationally intensive tasks, but supporting several concurrent requests have forced forth a change where Thread Level Parallelism(TLP) has gained focus. By adding several cores true TLP can be achieved by running the different threads on different cores.

The power consumption of the computer processor has always been the subject of inquiries, although it has not been the major point of focus in desktop computers. Here

performance has been the main focal point. However, as the frequency using traditional CMOS technology reaches its practical limit, ways of reducing the power dissipation have gained popularity. The CMP contributes mainly in two ways to help solve this problem. The first is due to the nature of cubic dependency between the operating frequency of the processor and its power dissipation as presented by Jerraya et al.[18] and Gochman et al.[13]. The practical implication of this means that by halving the frequency of a single core, the power consumption of a single core will be reduced to a mere quarter of its original use. Then by doubling the existing cores on a chip with half the frequency it is possible to retain the same performance with less dissipation of power. This is though a simplified version of the power usage. When calculating the overall consumption on a chip, the mechanisms for communication between different processing elements, cache and etcetera, must be taken into consideration. As shown by Kumar et al.[21], depending on the surrounding infrastructure, it is not only the cores themselves that dissipates power on a modern processor. Another interesting capability that comes with the inherent modularity of a CMP is the ease of resource management. If the load on a CMP is low, cores that are idling can be disabled to reduce the total power consumption. The ability to scale down the number of operational cores, and halving the total power dissipation are both important steps in battling the consumption of power in modern processors.

The final point which is mentioned by Spracklen et al.[38] and Olukotun and Hammond[27] is the inherent complexity of designing a modern processor and its cores. A high performance single core processor requires serious amount of effort into ensuring safe computations and a coherent environment. Since the cores employed by CMP's can afford a reduction in performance per each individual core and still have the same overall throughput, a single CMP core can be simplified compared to a core in a traditional single core processor. Also by allowing a scaled down version of the cores simplifies the development process, and thus reducing the time to market and hence decreasing the production cost. Beside the purely economic reason, a simplified model decreases the chance of bugs in the core itself. Hence when gate size decreases, its performance can be improved by increasing the number of cores on-chip.

To help developers address new architectural challenges, software based simulator such as SimpleScalar[4] and M5[3] have been used. Here the developer specifies the different components on the processor and their behaviour. The problems with software based simulators are that they are not completely accurate in the sense that it has to sacrifice accuracy in some field to gain in another[16]. E.g., a simulator which perspires to be instruction set accurate will not attempt to claim timing accuracy[16]. Another problem is that simulating an entire processor is a slow process, taking quite some time to produce results. The time it takes to produce results could be reduced by decreasing the level of details as proposed by Hines et al.[16], but this would lead to less accurate simulations. Also, by having to manually specify the different levels of detail in the model leaves room for inconsistencies between the different run levels. This might produce different results depending on the detail level when simulating.

The Field Programmable Gate Array(FPGA) is a natural evolution of the Programmable Logic Arrays which allows the developer to program the behavior of the logic hardware.

The extension which made the FPGA popular to prototyping hardware is its Field-Programmable attribute. Unlike many implementations of the Programmable Logic Array, the behaviour of the FPGA is fully reprogrammable. In modern FPGA's this is solved by connecting each programmable unit, i.e., its logic elements and routing resources, to the corresponding bit in memory. Then by changing the bit-stream the behaviour of the FPGA changes. An example would be a mathematical unit, the Arithmetic Logic Unit(ALU). Pending on its corresponding bit, the ALU might work as either a multiplier, divisor or a different mathematical operation. Then by connecting several logic elements it's possible to get the desired behaviour. The major drawback of FPGA's is that due to their flexible nature, they cannot have the same high clock frequency as a dedicated chip, Application Specific Integrated Circuit. The advantage is that the calculations which the software based simulator must calculate serially per simulated cycle, will happen in true parallel on a FPGA. This means that the reprogrammable hardware will have an increase in performance compared to a software based simulator.

To be able to prototype new multicore architectures and architectural parameters trying to improve CMP performance, a platform which allows for rapid changes is needed. It is also important that the time to produce results is reduced to increase productivity. To help decrease the period between a new architecture or parameter is decided upon, and until the result is known, a flexible FPGA platform is used. To achieve the flexibility in hardware, an architecture must be modular to allow for interchangeable elements. This to prevent time demanding and challenging rewrites of the entire system.

1.2 Problem Description

One of the new trends in computer architecture to battle the performance wall by introducing several cores on the same chip. By doing so, the computer engineers were put in front of old challenges in a new setting. Techniques which were meant for large cluster suddenly had to be scaled into a single die and novel ideas were needed. Seeing how the cost in terms of money and time imprinting the first chip is too high to allow prototyping, simulators and logic analyzers have been the mainstream technology of prototyping processors. However software based simulators such as SimpleScalar[4] are notoriously slow, and in nature not able to fully accurately simulate the inner workings of a CPU[16].

The goal of the Master's Thesis is to investigate and develop a hardware based platform for testing Chip Multiprocessors using Field Programmable Gate Arrays, figure 1.2, as an alternative to the software based simulator. The platform should allow for change of architectural parameters and components.

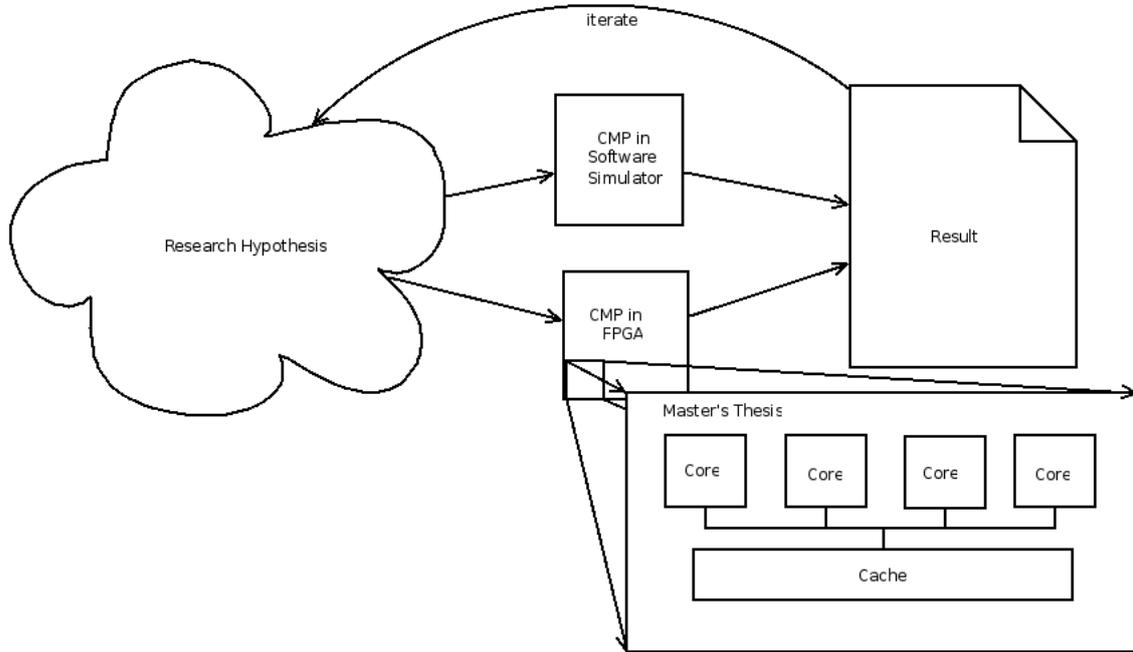


Figure 1.2: Project Overview

1.3 Project Motivation

The NTNU Computer Architecture Research Group has focused their research on Chip Multiprocessor Challenges[25]. More specific, the research group have focused their attention on intercommunication, caches, pre-fetching and scheduling. At the time this thesis was written, their main method of prototyping multicore architectures was based upon software simulators such as SimpleScalar and M5.

Since the software based simulations used to much time in producing results, the NCAR group wanted an supplementary method of prototyping their proposed multicore architectures. They also wanted a platform through which they could rapidly change the architectural parameters in order to test several configurations. This without changing the underlying behavioural code.

1.4 Contribution

Through his work with the Master's Thesis the author has contributed with the following set of items. A thoroughly background search has been done in the field of Chip Multiprocessors. The author has contributed by presenting a hardware based platform for prototyping novel computer architectures. Through the hardware platform the author has

contributed by allowing fellow researchers to produce results more accurate, and decreased the time waiting for results.

1.5 Scope

This thesis will describe a multicore architecture and the methodology used. The thesis will present relevant research and background information needed to understand chip multicore architecture. It will not describe in detail the inner workings of processor core technology beyond what's needed to implement it in a multicore environment.

1.6 Outline

In chapter 2, the thesis will present the relevant background information for reading this paper. It will start by displaying the relevant work, before it will continue by introducing the central concepts in computer architecture relevant to chip multiprocessors. Chapter 2 will end with a presentation of the tools used.

Chapter 3 presents the methodology, where the hardware architecture is described. Chapter 3 will end by a presentation of the controlling software and the application used for testing the hardware design.

In chapter 4 the thesis will show the result from the benchmark suite introduced in chapter 3, before the results will be further discussed in chapter 5.

Finally the thesis will conclude with chapter 6 with the conclusion and further work.

1.6.1 Appendices

Appendix A contains the benchmarks created for the processor. Appendix A also contains the makefile and linker script required in order of building the binary application.

Appendix B contains the hardware developed for this project, with the toplevel configurations in B.9, B.8 and B.7. Besides the code found outside the architecture section in appendix B.1 and B.2, which is glue code generated by the Xilinx tools, everything has been written by the author. Some bugs has also been fixed in the glue code, as discussed in section 5.4.2.

Appendix C contains the software developed by the author to control the FPGA and act as main memory as described in section 3.6.

Chapter 2

Background

This chapter will discuss different architectural challenges designing a CMP, and present some of the proposed solutions. It will introduce the reader to the field of processor architecture and FPGA's by first presenting the already existing work before introducing central concepts, and finally the tools used to develop the FPGA framework.

2.1 Related Work

One of the central problems designing new hardware architectures is to be able to test the architecture before sending it to production. Being able to both test the hardware and software before production, is of grave importance to create a stable and optimal platform. Traditionally the software developers have lagged behind the hardware manufacturers due to traditional simulators producing results too slow. This is one of the challenges that the Research Accelerator for Multiple Processors(RAMP) project is addressing[41]. The RAMP project is a collaboration between different universities and cooperations such as Berkly, Xilinx, IBM and others[2]. Having a FPGA framework allows the software developers to test different implementations in operating systems and compilers before the architecture reaches production. To serve as the developing platform, the RAMPants¹ have opted for the Virtex5 FPGAs to provide the hardware platform. The processor cores used are a mixture of MicroBlaze soft cores, and PowerPCs hard cores. The reference designs available from the RAMP project focuses on either transactional memory, distributed systems or distributed shared memory[41].

While the RAMP project is the first one where which a well defined interface for communication between different parts, allowing for exchangeable parts, there has been different implementations of CMP in FPGAs. Amongst these, Socrates is worth mentioning. Socrates, being one of the early adapter, has focus upon proving that a CMP can be de-

¹Member of the RAMP group

signed on a FPGA, and that it can be done rapidly[7]. The cores in Socrates are based upon an own implementation of ARM cores which in turn are connected to an crossbar mechanism to provide communication between the different elements. This allows the design to increase in size, leaving room for growth in FPGA transistor counts. The downside to the Socrates implementation is that the crossbar interconnect, as shown by Kumar et al.[21], wont scale as gracefully with the number of cores connected to it. This severely limits the amount of cores available to the Socrates platform, both in terms of FPGA resources, and in a real life situation where also factors like power consumption plays a major part.

Although the FPGA gains approval for prototyping hardware, there are still project prototyping using commodity hardware. The Stanford Hydra project is a research project testing out novel architectures for Chip Multiprocessors, basing their processor cores on four MIPS based processors[14]. To be able to simulate the processor design they have built a circuit board using four MIPS R3000 processors cores[28]. Each of the cores are connected to a floating point unit and a Virtex FPGA in order to form a single processor tile. By configuring the FPGAs, the Hydra project can simulate an array of different configurations.

2.2 CMP Architecture

Having several processor cores on a single chip brings forth new and challenging problems related to power consumption, cache strategies with multiple cores, communication between cores to ensure cache coherence and off-chip access. Although the physical scale has reached a new level, currently at a gate level of 45nm[10]², these problems have been addressed before by the world of supercomputers. Seeing how supercomputers traditionally have several processors, distributed/shared memory, and a network between the processing nodes, many of the concepts can be transposed down to fit the scale of a single chip. Furthermore, some of the rejected ideas have resurrected seeing how the scale have changed from several processing nodes, to a single chip, moving the limits on latency and bandwidth. This have led to interesting CMP strategies, where tried and tested supercomputer paradigms have been reused.

More specific, as shown in figure 1.1, during the last years, the core frequency have stagnated, which has left the developers looking for other ways to improve the overall performance, while improving the power consumption. Since the gate size has had a drastic drop, it has been possible to put several cores on a single processor. Then by decreasing the frequency per core, the power consumption have drastically decreased. Other interesting features of the CMP, being modular, is that it allows for easier resource management. This in turn makes it possible to turn of inactive parts, in an attempt to reduce consumption of power. Another important point of focus is that of off-chip access to remote peripherals, e.g. memory. The gap between the CPU and memory has steadily increased, and has

²In 2007

evolved into a complex hierarchy, known as the Memory Hierarchy, figure 2.1. This states that the further memory are from the CPU registers, the higher access cost in terms of latency, but more memory it has allocated. This challenge has introduced the notion of having banks of cache distributed around the system, allowing for temporary storage closer to the processing unit. Since the storage capabilities, in effect, dictates the latency several levels of cache can be found on the CMP. In a multicore environment the challenge lies upon having the cores cooperating on retrieving lines of external data, avoiding retrieving the same data twice. This in attempt to try reduce the amount of unnecessary duplicate off-chip communication.

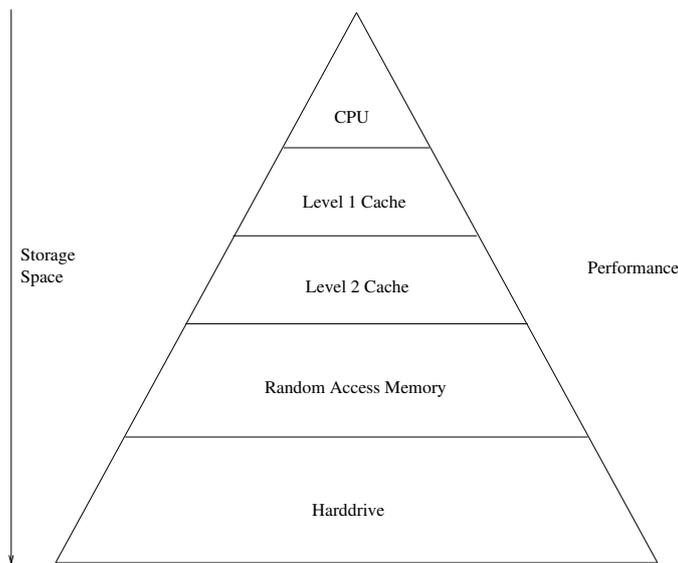


Figure 2.1: Example memory hierarchy

Although if the cores weren't to cooperate on the data retrieval, they would still need a way of invalidating cache lines. Each core might have their own copy of data in cache, and if one of the cores writes to the memory, the system need some way of notifying the other cores holding a copy of the data. If not the task would lie upon the software developer to lock memory accesses, which would be a tedious and error prone job.

2.2.1 Cores

The core is the main processing element on a traditional processor, performing instructions which it loads from memory. Traditionally the performance of a single core CPU has been given by the frequency of the internal clock, and the internal architecture of the core. The frequency has been dependant of the underlying feature size and has been steadily rising until year 2002, see figure 1.1. Combined with the fact that extracting performance per clock cycle gets exponentially harder³, has led for a new way of extracting performance.

³e.g., achieving a perfect level of ILP is impossible due to the demands of perfect hardware[15]

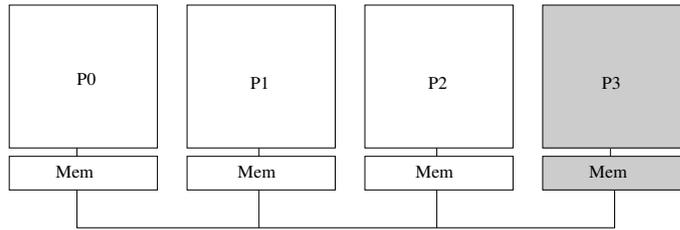


Figure 2.2: Core number 3 is disabled

Besides the raw performance challenges in designing a modern CPU, lowering the power consumption has increasingly gained popularity. The power consumed by a CPU is directly influenced by the frequency on which its internal core is set as shown in formula 2.4. As the generated heat is a function of the power consumed, the manufacturers have had problems cooling the CPU with apparatuses acceptable to the general public. This, and seeing how there is a lower bound to the latency of a signal traveling across the chip has led IBM to predict that their Power6 cores will extract the last amount of clock cycles available at circa 5Ghz[9] using the current available technology.

This has led to a search of other methods for increasing performance, acknowledging that the single core technology will stagnate at a given point using current technologies. Although at the time where this report was written a company named D-Wave in cooperation with NASA presented a co-processor using quantum technologies[1], this technology is still long from perfected, and has a long way before it is common household equipment. This has led to the manufacturers decreasing the complexity of a single core and increasing the total number of cores on each chip. This allows each core to decrease its individual throughput, but the system as a whole will retain its performance[38].

Scaling down the complexity of a single core while decreasing the core frequency influences the power dissipation of the processor as a whole. This has led to a decrease in both the Dynamical and Static power dissipation. Static power dissipation is power dissipating due to transistor leakage[5], while Dynamical power dissipation refers to the rise and fall in current when the transistors changes state. A simplified equation for the total power dissipation is shown below in equation 2.1.

$$P_{total} = P_{static} + P_{dynamic} \quad (2.1)$$

Dynamic Power Dissipation

The dynamical power dissipation is attributed the change of state in the transistor. When a transistor is set low, it has to discharge to ground, which is the main source of dissipation. The following formula (2.2) as described by Jeraya et al.[18] and Gochman et al.[13] shows

the dynamical power dissipation. Here F_0 is the clock frequency. The C_0 is the effective capacitance of the circuit, while V_0 is the voltage and α is the activity factor.

$$P_{dynamic} = \alpha * C_0 * V_0^2 * F_0 \quad (2.2)$$

In the same article by Gochman et al.[13], they show that the frequency can be approximated to be proportional of the core voltage V_0 , which leads forth to formulas 2.3 and 2.4

$$F_0 \approx K_f * V_0 \quad (2.3)$$

$$P_{dynamic} = \alpha * C_0 * V_0^3 * K_f \quad (2.4)$$

As shown in equation 2.4, $P_{dynamic}$ is cubic dependant of the frequency. Thus by halving the frequency per core, the core will dissipate of one quarter of its original power. Then by doubling the number of cores on the chip, the chip will retain its performance while halving its dynamical power dissipation.

Static Power Dissipation

The static power dissipation, “leakage”, is an effect of the current gate technology[12], mainly due to subthreshold and oxide leakage[12]. In their article, Ghiasti and Grunwald[12] presents equation 2.5. This models the static power loss as a product of the Voltage current, leakage current, the number of gates N and a scaling factor k . The scaling factor is dependant of the inherit complexity of the design itself.

$$P_{static} = V_{cc} \cdot I_{leak} \cdot N \cdot k_{design} \quad (2.5)$$

A overall measure of the effectively per power dissipated is shown in equation 2.6. This shows the static power leakage over the million instructions per second, giving a rough estimate of how resourceful the processor is.

$$Eff = \frac{P_{static}}{MIPS} \quad (2.6)$$

To reduce the static power dissipation, Muthana et al.[22] suggests that by reducing the I_{leak} factor, would have a great impact. One of the methods would find an architecture which allowed for disabling caches and cores, as shown in figure 2.2.1.

Spracklen et al.[38] also mentions that in a single core processor, much complexity and logic is used in a controlling context, not in performance issues. This leads to a high N in formula 2.5, and thus the effectively according to formula 2.6 decreases. Spracklen et al. also mentions that a high performance core is a complex design, which leads to a higher k_{design} and further decreases the efficiency.

2.2.2 Cache

As the gap between the CPU performance and the memory latency have grown as shown in figure 1.1 has grown, the need for temporary storage of data has increased. This has led forth to a hierarchy of memory structures where frequently used data is located near the CPU in terms of access time, allowing faster access to more popular data. A sample structure is shown in figure 2.1, where the data with the lowest access time could be stored inside the CPU itself and its registers. Duplicates of data that is frequently used would be placed in the level 1 cache, less frequently used in the level 2 cache and so forth, whereas the data stored on the disk or other external devices would take the longest to access. This is due to external memory have a higher latency before the requested data reaches the bus, the distance the signals have to travel, and the obstacles getting there. E.g., a DDR2 has the memory clock set at 400Mhz, whilst a traditional hard drive has a seek time given in milliseconds. Both considerably higher than the internal registers to the core which operates on clock frequencies measured in gigahertz and where the registers can be accessed in a few clock cycles.

To solve this challenge, cache banks have been introduced into the computer architecture storing a subset of available memory close to the processor, and thus reducing the access time for a set of frequently used data. Which lines of data and the amount of data that the cache can store internally is given by an amount of different parameters. When the cache gets full and a new item is to be stored, the cache bank must choose which one to evict, and different strategies exists to choose the right one, such as Least-Recently-Used[35] and Random-evict[40] depending on how the cache stores its data in memory. Each cache unit can store a certain amount of datum, cache lines, in memory. What differentiate one cache organisation from another is how its chooses store its cache lines, and the different parameters controlling the behavior.

One of the easiest conceptual ways of storing cache lines would be to store the cache line at any given spot in the array of available cache lines. This strategy is known as fully associative, figure 2.3. To be able to keep track of which cache lines that is stored where in the cache, each cache lines' tag, i.e. the part of its associated meta data that describes which data that is stored in a given location, would have to be its full address. This would lead to an massive amount of overhead per cache line stored, seeing how for each line the cache must keep track of its corresponding address. Also during a lookup the cache must traverse through each line, matching its address to the requested address. This leads to a increase in latency when performing cache lookup. Another, faster way of organizing cache is directly associative cache, figure 2.5. Here a given number of the least significant bits of

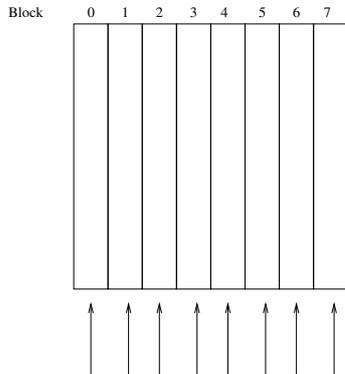


Figure 2.3: Fully Associative Cache

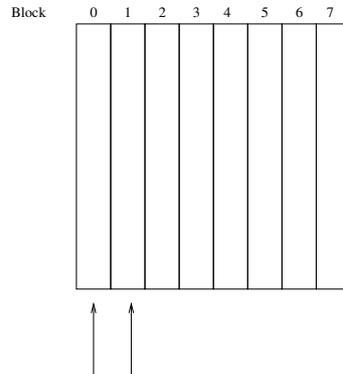


Figure 2.4: Set Associative Cache

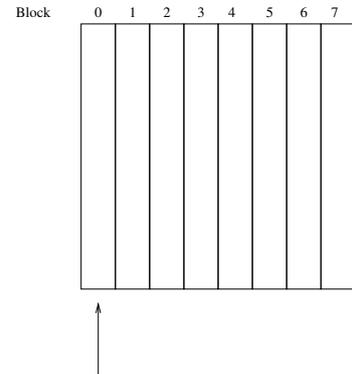


Figure 2.5: Direct Associative Cache

the address can determinate which index that the cache line will be stored in. This would be conceptually equivalent of the mathematical modulo operation. By varying the amount of bits used to determine the index, i.e. varying the number in the modulo operation, the cache can keep a different amount of cache lines. The two obvious benefits by this method contra the fully associative is, one, the cache can reduce the number of bits in its tag. If X number of bits is used to determine the index where the data is stored, address - X bits would be needed to provide the cache line's tag. The second advantage to directly associative cache is that a lookup requires significant less logic when determining a cache hit. A mixture between fully and directly associative cache is set associative, figure 2.4. The basic notion is that a set associative cache can keep several cache lines per index, i.e. the degrees of set associativity. A 1-set associative cache is the same as a directly associated cache, whereas a 2-set associative cache can store 2 cache lines per index.

Multicore Cache Architecture

Having several cores on a single chip introduces a new challenge, namely resource sharing. The basic challenge with the multicore cache architectures is the same as with a single core CPU, efficient use of off chip communication. Although the same problem, the environment has changed. Duplicate the number of cores on the chip, and the memory access will be duplicated using a naive single core cache strategy. Although the CMP is a relative new product in computer science, having several processing units in a computer isn't a new paradigm[32]. Challenges seen in CMP, e.g., communication, cache strategies, etcetera, have been addressed by earlier work. One of the more interesting effects in the CMP world is to see implementations which have been discarded in traditional supercomputing being reused in multicore CPUs. This might be strategies which have been discarded due to problems with latency, low bandwidth and so forth. Since all of the components on a CMP are placed on a single chip, old or discarded research can be re-evaluated seeing how the physical scale has changed.

To hide the gap between primary memory and the CPU, efficient off-chip communication is required. In a single core environment this can be solved using advanced prefetching, cache eviction schemes and etcetera. However, introducing multiple core on the chip have further brought new challenges. Seeing how another on-chip cache bank might hold the data requested from different core, some sort of cache cooperation is needed. Second, having several cache banks brings forth another phenomenon from the supercomputers, the Non-Uniform Memory Access(NUMA) effect, named Non-Uniform Cache Access(NUCA) [20] in CMP terminology. Although, to the core, each line of cache appears to be located in one uniform area of cache each cache bank stores a certain part of the whole. Due to the wire latency, pending on the physical location of the cache bank, accessing different parts of the memory will have different access times.

To help battling the problem several schemes have been proposed. Even though the implementations differs, they all have the same goal. Increase the off-chip communication efficiency, by making the cache banks cooperate. Chang et al. introduces an elaborate scheme based upon ideas from software[6]. Here all the cache banks are aggregated. When a cache bank evicts its data it will first try to “spill” the data over into another bank. However, if the cache is full with own data the bank will reject the “spilled” data. Another strategy is proposed by Dybdahl et al. where the instead of the LRU-scheme[35] a frequency counter is used[11]. Each cache is then allowed to grow shrink cache sets. In doing so, they allow cores with more frequently accessed data to dominate the cache.

2.2.3 Interconnect

As soon a processor has several cores on the same chip, it starts to require an interconnect network between the cores, and other resources on the chip itself such as cache banks. Depending on a various amount of underlying architectural features, such as the number of cores available, the wanted performance in terms of latency, bandwidth and finally the power consumption, the topology of the underlying interconnect varies[21]. Thus depending on the requirements of the chip in production, different topologies will suit different needs. However, there are two main groups of interconnects which will be discussed, the crossbar interconnect and a shared bus as seen in figure 2.7 and 2.6. These two represents two completely different strategies, and thus they have two different sets of characteristics.

The shared bus, figure 2.6 is a network where all of the resources, i.e. the ones on the same network, are connected to the same set of buses. Having several resources connected to the same bus presents the problem with arbitration. If several resources tries to communicates on the same time, the signal would be ruined and the transfer would have to restart. Hence the need for a mechanism which arbitrates either the signal from the resource itself onto the bus, or a device which tells which resource that are allowed to send signals onto the bus at a given time. Kumar et al.[21] discusses a mechanism in which the cores requests access to the address and data bus by communicating with a arbitration device. However, since the medium through which the devices communicates is a shared one, only one signal can be active on the bus at one time. Even so, Kumar et al. presents methods of pipelining

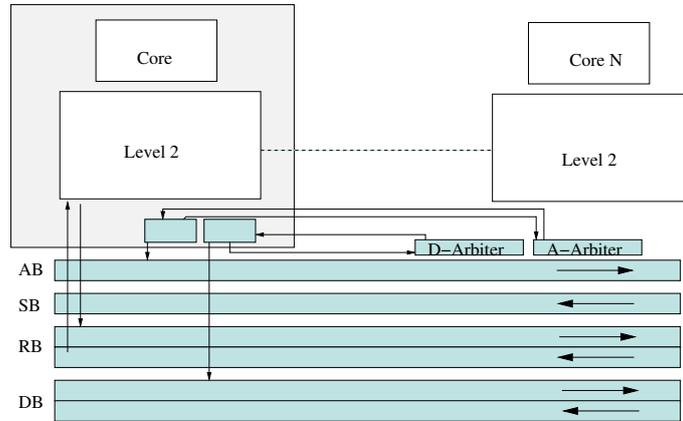


Figure 2.6: Shared Bus

As described in Kumar et al.[21]

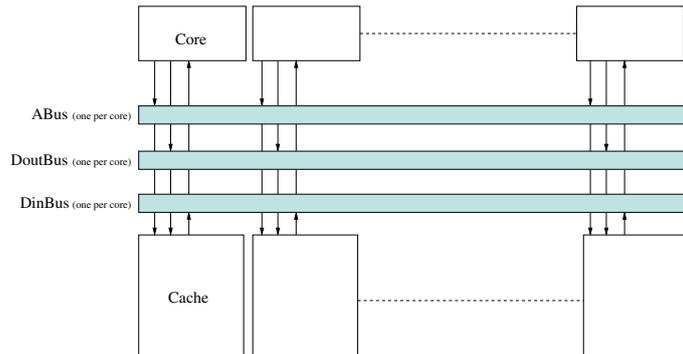


Figure 2.7: Crossbar

As described in Kumar et al.[21]

the process. One method is to include one bus for the control signals, address and data signal. Adding several buses, opens for the possibility to have one resource driving the address bus, while another one drives the data bus, fulfilling the previous request sent out on the address bus.

Unlike the shared bus, the crossbar topology, figure 2.7, relies on a direct connection between the resources on the chip. This helps reduce the time waiting for the data requested to arrive, but has a much higher cost in areal and energy consumption. An example shown in Kumar et al. shows that a crossbar mechanism introduces an area overhead of 11.4%, 22.8% and 46.8% with respectively 2-, 4-, or full sharing on a 8 core processor using a $400mm^2$ die[21]. However, cores will have to stall less waiting for the data to arrived using the crossbar. The same paper shows that with a 8 core, fully connected crossbar and shared cache, the power consumption will match that of 3 cores in just the interconnect

alone.

2.3 Introduction to FPGAs

Due to their flexibility in nature, the FPGA have gained status as “reconfigurable” hardware. Having flexible hardware is a great advantage in situations where the developer don’t want to be locked down by the restrictions put forth by the ASIC. Examples would be in computer architecture research and embedded devices, where reconfigurability is the key. Having a configurable hardware unit is preferred compared to performance when developing hardware.

Although traditionally, the FPGA has been viewed upon as a device for prototyping new hardware, it has recently gained approval for usage in computationally heavy areas such as DNA string matching and several cryptology algorithms[8]. The FPGA being a reconfigurable device has always lagged behind the CPU in terms of clock frequency, and thus it has been ignored when raw clock frequency is preferred. However, the ability to act as a true parallel device outperforms the serial processor in areas where parallelism is the key feature. The reconfigurability comes on cost of frequency, and when this thesis was written, Xilinx produced FPGA’s which operated at about 500 megahertz[46]. This in contrast to the modern processors provided by Intel and AMD which operated in the range of 3000 megahertz.

The way that FPGA’s achieve such flexibility are based upon, two attributes. First the ability to configure each individual Configurable Logic Block(CLB), and second a configurable network connecting CLB’s to each other. To program the FPGA, the developer loads a bitfile into the memory of the FPGA. Portions of the memory is connected to each resource in the FPGA, being either the routing switch, or the CLB. This leaves the developer able to create virtually any circuit by loading the right bitfile into memory, given that the FPGA have enough resources available.

Each CLB, figure 2.9, consist of several logic elements[8]. This might range from D-Flip Flops and lookup tables to coarser units as Arithmetic Logic units. Having several logic blocks in the same block gives the developer the ability to create more advanced units without having to use an excessive amount of the available resources. However, it is important to retain the possibility to have a fine grained output. An arithmetic logic unit will outperform the lookup table at arithmetic operations, but it will not give the developer enough flexibility to design more specialized units.

The routing on a FPGA is controlled by utilizing pass-through structures[8], controlling their behaviour through the bitfile. On modern FPGA’s, the CLB is connected to a nearby connect box as shown in figure 2.8, forming a Island Style network[8]. At each intersection, the control bit decides if the signal should continue in the same vertical or horizontal direction, switch direction or be routed to a neighbouring CLB. That way it is

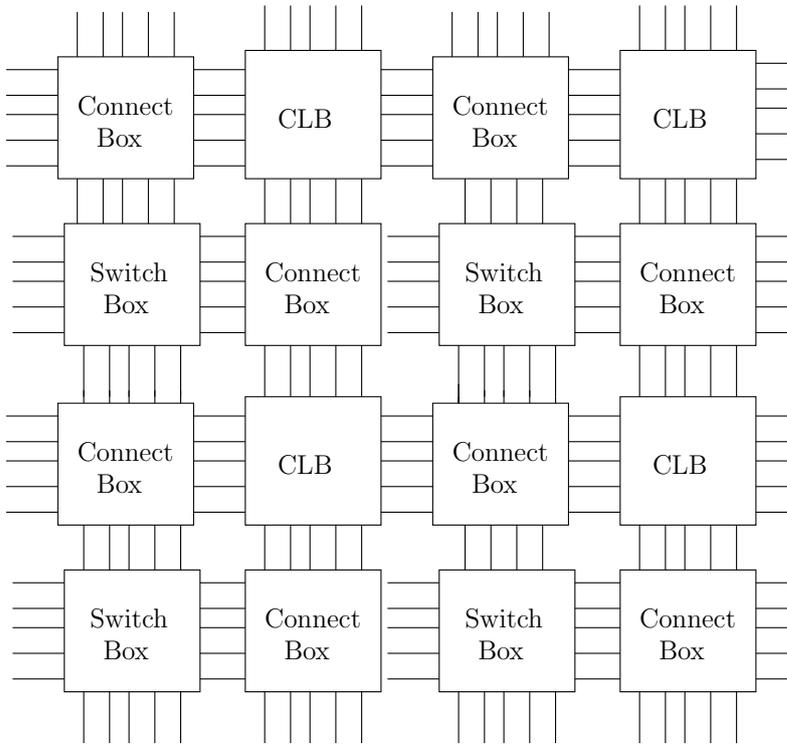


Figure 2.8: FPGA routing

possible to route the signal between different parts of the FPGA, forming complex designs utilizing more than one CLB.

When deciding upon which FPGA to use for this project, several factors are important. A modern CMP is a complex design, and complex designs utilizes a lot of logic cells. Of the FPGA's available to the author, the VirtexE 2000 has 43,200 of logic cells, while the Virtex 1000 has a number of 27,648. Greater amount of logic cells available allows for more complex implementations, using more complex structures or having more soft cores. Figure 2.10 shows a sample of soft cores and their resource in terms of slices on the VirtexE FPGA, each slice being an aggregate of two logic blocks. Other structures that are important are the interconnect, on-chip block RAM and off-chip communication. A integral part of a multicore processor is its cache banks. Thus if the FPGA has on-chip block RAM, additional slices can be saved not having to implement storage blocks using logic units.

2.4 Cores

When designing a multiprocessor on a FPGA, one of its most fundamental features is its cores. Of the cores available to a designer there are two different sub-types, depending

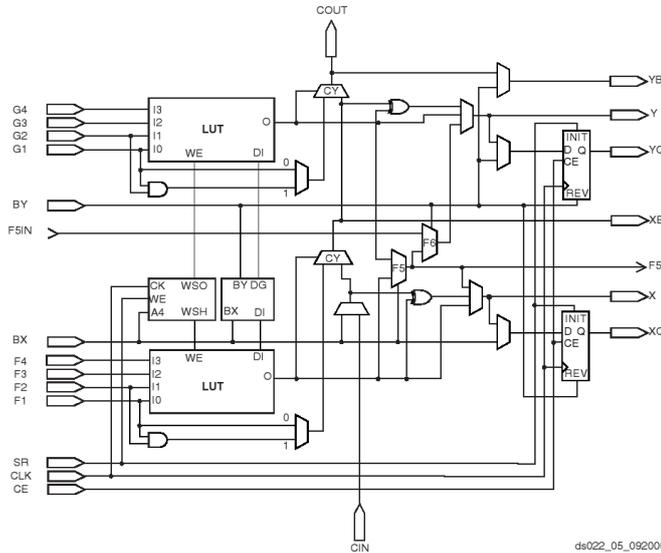


Figure 2.9: Sample CLB[42]

on (1) the wanted qualities and (2) what the hardware supports. What differentiates the hard core from a soft core is if it is a physical implementation, e.g., core molded onto the FPGA, or if it has to be implemented together with the rest of the code.

The soft core, is a separate project delivered in a form that allows the developer to synthesize it together with the rest of the code. The core might be delivered in form of VHDL, or an already implemented netlist mapped to the underlying hardware. Having an independent core gives the developer more freedom to experiment, and decouples the core from the underlying hardware. The hard core is a new phenomenon seen where the FPGA manufactures integrates existing processor cores on the FPGA itself. The hard core allows for greater performance compared to an implementation in VHDL, being a specialized circuit. However by utilizing a hard core the design gets more bound to the underlying hardware and thus scarifies flexibility for performance.

Of the soft cores available there are mainly two groups, commercial available such as the MicroBlaze[45] and freely available microcontrollers as the NanoBlaze and other experimental cores developed by hobbyists found at sites as OpenCores. The PicoBlaze[43] core is a product of the Xilinx Cooperation, and is a small 8-bit controller. For the PicoBlaze controller, Xilinx have opted for a unit which leaves a small footprint in the design, and thus have optimized away several complex instructions which would have made the controller increase in areal. Instructions such as multiply, divide and floating point calculations are non-existing. There exists schematics which shows how to connect several PicoBlaze controllers together to gain support for such instructions. However, in doing so would render the main feature of the PicoBlaze void. Being a small microcontroller, one of stressed points about PicoBlaze is that it provide an alternative to hardware circuits without the areal overhead normally associated with introducing a microcontroller.

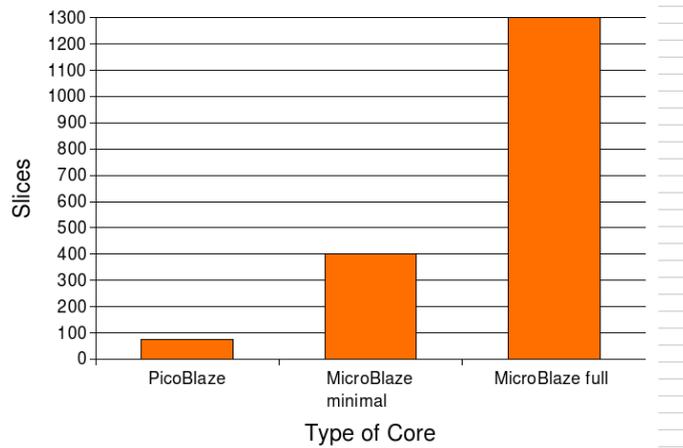


Figure 2.10: Number of Slices per Core.

2.4.1 MicroBlaze

The MicroBlaze microcontroller is the grown up version of the PicoBlaze. Being a 32-bit processor, it matches the modern day processor both in data and address width, and thus it is more suited for modern applications than its predecessor, the PicoBlaze. MicroBlaze implements a wide array of instructions, adding support for floating point, multiplication and several others on the expense of areal and resources used. However, one of the major selling points is that the MicroBlaze controller can be customized to provided the needed functionality. Support for such instructions as floating point might be omitted if the design doesn't require it, freeing space for other components. The controller also, in attempt to match the PPC-core found on some Xilinx FPGAs, conforms to IBM's CoreConnect[17] architecture[45]. This allows the hardware developer to extend the controller with extra peripheral units. Some of the most important buses includes the Local memory bus (LMB) and On-chip peripheral bus (OPB). With these two buses the controller can attach units through which it can communicate using a memory mapped scheme. The major difference between the two buses is that the LMB is a much simpler interface, connecting units which are to guarantee a one cycle response. This bus is where the designer normally would connect on-chip memory such as Block RAM controllers. OPB is a more complex bus, allowing slower peripheral units. Typical examples would be peripherals which are not memory, including units such as media access controllers and off-chip units. The final way of attaching units to the controller is through a link called Fast Simplex Link(FSL). The MicroBlaze can have 8 FSL interfaces which provides a low latency interface through which hardware accelerators can be attached.

2.4.2 PowerPC

Of the hard cores available, Xilinx have embedded IBM's PowerPC 405 on their Virtex-II Pro line of FPGAs. The 405 is an embedded core developed by IBM to suit the embedded market. This includes a more specialized system for memory management and specialized registers for debugging etcetera. The external interface matches that of the MicroBlaze controller, so hardware developed for one would suit the other without, in theory, code rewrite. What differs the PPC405 from the MicroBlaze is that, in being a PowerPC, it must conform to the PowerPC standard. Each PowerPC must correctly implement the User Instruction-Set Architecture(USIA). The USIA guarantees that the controller will behave exactly the same as all other PowerPC cores when in userspace. This leads forth to the PowerPC having a better support for compilers, and operating systems than the MicroBlaze core. Although a version of microcontroller-Linux has been ported to MicroBlaze, several others including NetBSD have been ported to the PowerPC platform[31, 29, 26].

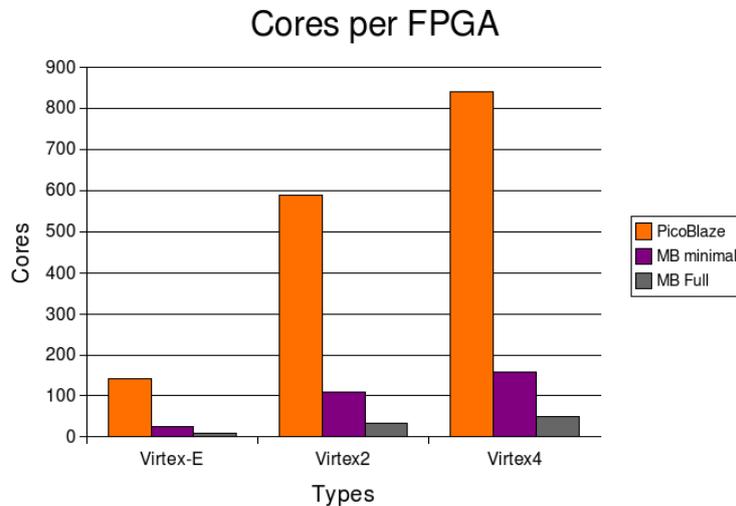


Figure 2.11: Number of Cores on a FPGA

One major implication the type of core imposes, is the number of cores available on a single FPGA. Figure 2.11 shows the amount of cores available on various types of FPGAs. However, it's worth mentioning that it is an approximation. The number of cores are a function of the number of slices available on the FPGA over the number of slices occupied by a single core. In real life other resources should be taken into consideration. Even so, figure 2.11 gives an rough estimate over the number of cores available.

2.5 Environment

2.5.1 Hardware

The server which hosted the project was a CompactPCI IBM Compatible server running the Debian Gnu/Linux operating system. On board the host computer were two BenERA CompactPCI DIME-II motherboards[24]. Each of the two BenERA motherboards had two Virtex-E FPGAs as shown in figure 2.12, respectively marked red and green.

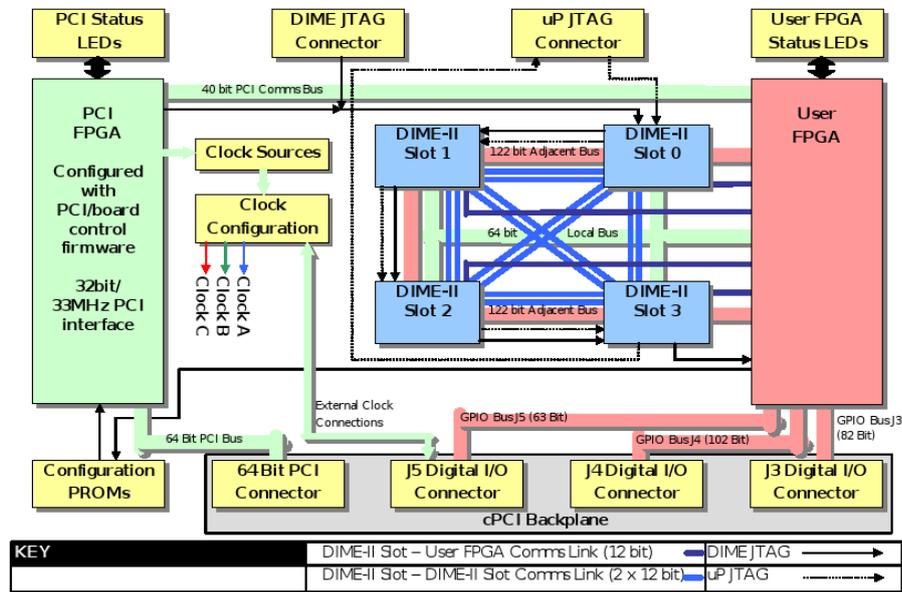


Figure 2.12: BenERA functional diagram [24]

Having two FPGAs on the same motherboard, Nallatech could occupy one to perform solely administrative functions. That FPGA is marked in the functional diagram(2.12) with green and is the PCI FPGA. The PCI FPGA acts as the bridge between the User FPGA and the host computer, simplifying the communication between the software and hardware. To achieve this, Nallatech loads the PCI FPGA with firmware that communicates through the PCI bus on the host computer. To provide communication between the host computer and the User FPGA, the PCI FPGA has a FIFO buffer which can be read and written to from both User FPGA and the software through the FUSE library described in section 2.5.2.

The user programmable FPGA was a single Virtex controlled through the PCI FPGA[24] and JTAG[33]. Further, the same motherboard could be extended with four new modules conforming to the DIME-II industrial standard. Modules that fit the DIME-II standard includes memory modules such as SDRAM and DRAM or FPGAs, including the new Virtex II-Pro.

2.5.2 Field Upgradable Systems Environment - FUSE

To provide communication between the User FPGA and the controlling software, Nallatech ships a FUSE library with their motherboards. This creates an abstraction layer between the software developer and the design running on the User FPGA[23]. The library gives the software developer the opportunity to control the clock frequency, loading bitfiles containing the FPGA design and configuring the DMA communication channel.

2.5.3 Virtex-E

The programmable FPGA hosted on the BenERA motherboard is a Virtex-E. The Virtex-E is a SRAM based FPGA[42], which means that the behaviour is defined by loading the generated bit stream into SRAM memory. The bitfile loaded into memory will control the different logic elements and routing resources, behaving close to the generic FPGA described in section 2.3.

The configurable logic block is implemented as shown in figure 2.9, with four logic cells pared into two slices. Each logic cell has in turn one four bit input lookup table, which acts as the function generator. Each CLB can aggregate their Lookup Table(LUT) s, and that way it is possible to get function generators with a greater width, totalling at 5-6. However, in the modern lines of Virtex FPGAs, the native width of the LUT is 6, such as in the Virtex 5[46]. Besides the role as a function generator, and thus as a pure logical unit, it is possible to configure the logic blocks such that it will act as memory. Done correctly, the developer can use the CLB s to create memory banks. Each CLB having a four bit input is able to store either 1x16 bit RAM or connect two CLB to generate either a 2x16 or 1x32 bit RAM block. A more resourceful way of storing RAM on the Virtex-E FPGA is to use the already existing Block RAM structures. Placed evenly spread between each row of logic elements, Xilinx has placed 96 BlockRAM on the xcv1000E model. Each BlockRAM being a 4096 bit dual port RAM gives the FPGA a total amount of 393,216 bits dedicated memory. Another feature of the BlockRAM is that it might be used as a large LUT taking the memory address as input, giving the data stored as the output.

The routing network on the Virtex FPGA is close to the one described in section 2.3. However what differs is that each connect box is connected to several other lines running a various amount of distances. Each CLB is connected to, what Xilinx calls, a General Routing Matrix(GRM) which acts as both the connect and switch box. Each GRM is connected to their direct neighbours, the neighbours with a Manhattan distance of one. To help decrease the latency, each GRM is also connected to a longer wire, which stretches over 6 GRMs. However this wire is driven from the end GRMs, but is accessible from the boxes in between. The last set of wires is those who runs from one side of the chip to the other, allowing for rapid global communication. Finally all GRMs can access global signals such as clock signals and reset.

2.6 Tools

The tools used can be divided into two main groups, the ones used for hardware development, and the ones used to develop software. The hardware tools for this project is mainly those delivered with the Virtex-brand of FPGA s. This includes Integrated Studio Environment(ISE) which works as a IDE providing everything from syntactical analysis of the VHDL code, down to the creation of the FPGA specific bitfile. The tool to configure the MicroBlaze core is XPS, which easily allows for configuration. The software development tools used is mainly various ports of GCC to create both the running environment on the host machine, and compiling applications to run on the FPGA implemented processor.

2.6.1 GCC

GCC is the Gnu Compiler Collection, and has been ported to several platforms. One of the major advantages to the compiler, which makes it cross compatible is that they have separated the different functional layers[39]. The first layer is the language layer, converting from a programming language to an internal tree structure. Afterwards the compiler will optimize the tree, before it passes it to an machine dependant layer. By having done so, all which is needed to support a new platform is to extend the back end to support a new architecture. This has made the GCC-compiler the preferred choice for many embedded producers. Both the embedded PowerPC core and the MicroBlaze controller have a version of GCC ported by Xilinx.

2.6.2 ISE

Integrated Software Environment is the development environment created by Xilinx[37]. The application is a front end to the entire chain of tools needed from synthesising the written VHDL files, to the generation of the chip specific bitfile to be loaded onto the FPGA. The general flow in creating a bitfile starts with the developer writing VHDL code which specifies the behavior of the hardware. Then XST, will generate a Xilinx specific netlist. Afterwards the applications will perform map, translate and placement and routing(par). This will map the resources needed by the netlist to the resources found on the actual FPGA. Afterwards the bitfile used to configure the chip is generated by the bitgen command. All this is taken hand of by the ISE, and thus decreases the development cycle.

2.6.3 XPS

To allow the user to make customizations to the MicroBlaze microcontroller, Xilinx packs XPS with their Embedded Development Kit[44]. Through the XPS application, the de-

veloper can extend the MicroBlaze or PowerPC core by attaching different modules to one of the many buses. Through the application the developer can tune the core to include certain features such as floating point unit, hardware support for multiply and division.

Chapter 3

Methodology

This chapter will describe the design of the multicore processor. It will start with a short description of the overall design notion, before it will present the most important components. After having presented the multicore processor design it will present some of the software developed to control the CPU. Finally it will present some of the developed benchmarks and test applications used to test the working CPU.

3.1 Introduction

The architecture is designed for extendibility, allowing the developer to test new and novel architectural changes, only rewriting or changing the parameter in question. The partitioning of the system is shown in figure 3.1¹. The processor's main memory is implemented in software, while the rest of the design is implemented in hardware using VHDL. Having the memory in software allows the designer to easily calculate the number of cache misses, each software memory access not being in the on-chip cache banks. Software based memory also has the advantage that it is much cheaper in terms of areal, thus freeing resources from the FPGA which can be used for cache or other logic.

The software communicates to the hardware through the PCI FPGA, using a memory mapped scheme. On the hardware side the software communicates to the PCI FPGA, a specialized FPGA mounted to provide an interface between the programmable FPGA and the software running on the host computer. On the PCI FPGA the software accesses a FIFO buffer which is connected to the User FPGA, containing the data to be communicated to or from the processor. Being a 32 bit wide FIFO, the software will fetch 2 words, before deciding if the last transfer should be a write or read to the buffer pending on the nature of the memory access. Once transferred to the buffer, it is communicated through a bus dubiously called the Peripheral Component Interconnect which is connected to the

¹Red marking the systems developed, blue the 3rd party cores, and gray FPGA specific modules

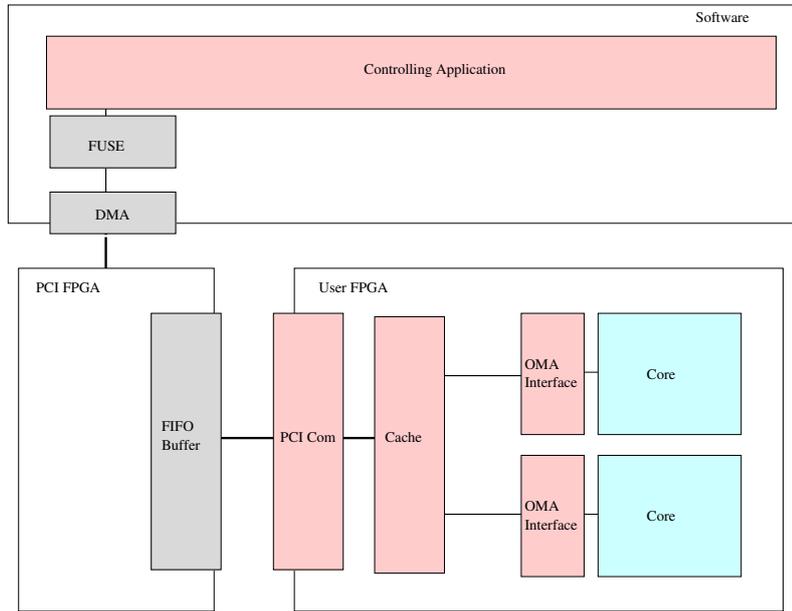


Figure 3.1: Overall Architecture

User FPGA and the PCI Communication module as shown in figure 3.1. This is where the data gets handled by the implemented design, and translated into a format which suits the internal components. Limited by the width of the buffer, the communication module has to multiplex 96 bits of data² into 3 words which can be transferred between the FIFO buffer and the User FPGA.

Once the data reaches FPGA, the data is communicated between the different modules, e.g. arbiter, cores and cache through the unified signal interface described in section 3.2. Having a unified interface allows all the modules to be replaced as long as they confirm to the same interface. Having this allows for the replacements of the cores themselves, freeing the design from the restrictions put forth by the MicroBlaze core and the Xilinx team, which in turn relinquishes the complexity associated to developing with the Xilinx tool chain. This leaves an overhead in the communication latency between the cores themselves and the rest of the system, having to translate from the CoreConnect architecture to an customised interface.

3.2 OMA

To uphold the demand for extendability, allowing for easily exchangeable units, a interface which all components must adhere was created. The interface itself was loosely based upon the CoreConnect architecture, although some of the fancier mechanisms removed

²Address, Data and Control Word

Name	Description
Valid	If the data presented by the data lines are either valid, in the case of a read or have been written to memory.
Data_I	The data to be read when the RW flag is set to high.
Data_O	The data to be written when the RW flag is set high.
Address	The address to which the data either should be read or write
RW	The direction of the transfer, will read on asserted.
Active	When this signal is high a transaction is ongoing.

Table 3.1: Signal interface

to achieve a simpler model for communication. The signal interface was based on a pure memory mapped processor architecture where all communication between the processor and the off-chip peripheral units were done through an elaborate memory scheme. The downside was the inability to communicate meta data in a separate channel, such as internal statistics, cache hit ratios and etcetera. However this could be solved using other mechanisms as JTAG interfaces and ChipScope, providing a much cleaner and more customisable mechanism to probe the hardware.

The signal interface that each component needs to confirm to is shown in table 3.1 and contains the bare minimum to provide communication. However, it is meant as interface between the different layers in the architecture, and not between specialized components, leaving room for advanced features where it is needed.

The interface defines 4 signals which are driven by the source, namely the Data Out bus, Address bus, Active flag and finally the Read not Write flag. The two foremost signals are the Data to be transferred out from the component and the address of the data that are either to be written or read. When a transaction is wanted the source will assert the active flag until the transaction is done. The direction of the transfer is signaled through the RW-flag. If the RW flag is asserted the following action is a read transfer, and the client side should drive the signal on the Data In bus, ignoring the Data Out. On the other hand, if the RW-flag is set low, the operation in question is a write transfer, and thus the source should drive the signal on the data out bus. Independent of the direction of the transfer, when the destination side is finished handling the transfer, it should assert the valid flag for one cycle and thus signal that the transfer is completed.

3.3 Core

The core, being the processing element of the CPU, is the main component of the processor. However to this project the core alone is just another component needed to be able to perform benchmarks on different architectures. Important criteria is the number of cores that can be made available in the design, its documentation and to some degree its feature

set, i.e. a core can't have a too scarce instruction set or else it won't be able to run the most basic benchmarks and tests.

Of the different cores investigated, MicroBlaze was the one that satisfied all the 3 different criteria. By opting for a soft core such as MicroBlaze the design has the ability to scale beyond the hard cores available on the FPGA, which in the case of a Virtex-II Pro is limited to two PowerPCs. Also by utilizing a soft core the implemented multiprocessor is not bound to the specific implementation of the FPGA itself and thus if the resources available on the FPGA can't support the design, the VHDL code can be synthesized to fit on a new type of FPGA. As to the feature set, the core has a RISC based ISA which is most able to fulfill the demands imposed by the benchmarking and test suite. One of the key features is the MicroBlaze's ability to cut down on the features, such as the floating point operation. This in turn cuts down on the resources utilized by a single core and allows for a greater amount of total cores on the chip, or more advanced features requiring more logic.

3.3.1 Wrapping

To be able to easily add new cores to the design, a thin proxy layer has been wrapped around the MicroBlaze core. The wrapper layer translates the signal from the CoreConnect and the OPB to a scaled down version following the specifications described in section 3.2 and table 3.1. By designing a proxy between the MicroBlaze core and the rest of the system it is possible to exchange the core themselves without having to affect the rest of the system. I.e. by abstracting the memory interface, if the system were to exchange its cores with a NanoBlaze core the only part having to be rewritten is the Off-chip Memory Access (OMA) interface to the core.

Internally the MicroBlaze has two main busses through which it can communicate with different peripheral devices such as external memory, SDRAM, different media access controllers and more traditional units found on a processor such as on chip ram blocks. The MicroBlaze core divides its internal buses between those peripherals that are able to deliver data with one cycle latency, the LMB, and multiple cycles, the OPB [45]. One cycle latency is possible if the data is stored on chip, and hence it is meant to be the bus that supports on-chip memory access. However, since the memory subsystem cannot guarantee a one cycle access to memory it has to be connected to the OPB bus.

Although the MicroBlaze core itself is a Harvard Architecture[15] core, differentiating the data space from the instruction space, these two spaces are both joined in the OPB and thus minimizing the logic needed to implement different address spaces. Although by doing this the complexity of implementing an instruction cache is greatly reduced. This opposed to exposing both a data and instruction bus outside the core itself.

A single core is conceptually made up of mainly 4 components as shown in figure 3.3.1, communicating to the rest of the system through the OMA interface. As its kernel the

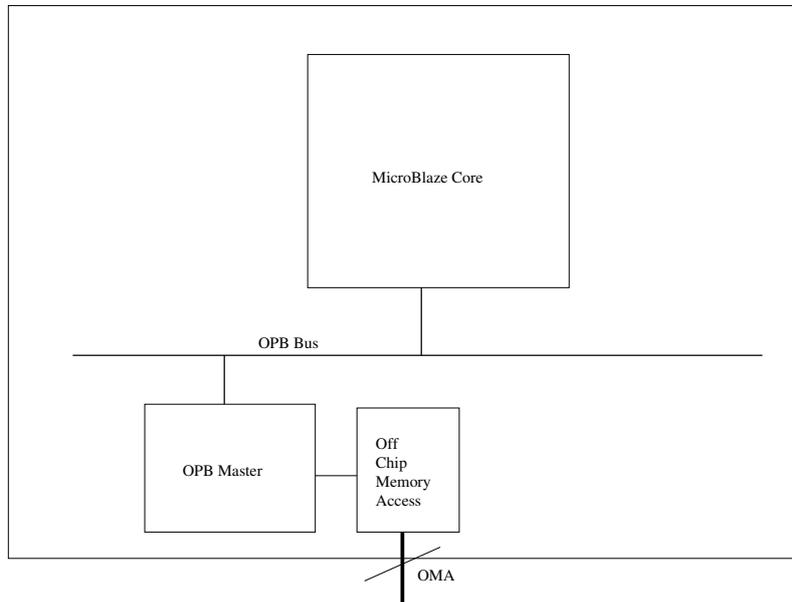


Figure 3.2: Core

core is build up around one MicroBlaze soft core. The soft core is connected to a OPB bus which conforms to the CoreConnect standard as provided by IBM[17], and on the other side a OPB Master which connects to the Off-Chip Memory Access component.

3.3.2 CPU Identification Component

To be of use in a multicore environment, a method of separating one core from another is required. This could be solved by using a special CPU version register containing the ID of the core itself, or split the address space of each core, each prepending the first bits of the address to identify the core's ID. In this architecture, that would have a serious impact on the available memory space as soon as the number of cores reaches more than 4 cores. Splitting the address space would also involve a lot of complexity in any shared level cache. However, the MicroBlaze core version 4 does not provide any user defined registers, which is a function in the newer versions of the microcontroller.

This led to the development of a mechanism using MicroBlaze's Fast Simplex Link bus. As described in section 2.4, each MicroBlaze core can be connected to up to 8 FSL-interfaces. One of those is a unit hardwired to deliver the core identification upon a request from the core itself. This gives the software developer the possibility to differentiate the different cores. This is done by issuing a GET instruction on port 0, storing the result in a register. Afterwards it is possible to check the given register and act accordingly as seen in example code 1.

```

// Get the CUID from port 0 and store it in register A.
GET REG_A, PORT0;
switch REG_A do
  case 0
    // Code for core 0
  end
  case 1
    // Code for core 1
  end
end
end

```

Algorithm 1: CPU Identification

3.3.3 Implementation Challenges

One of the major challenges by utilizing the MicroBlaze microcontroller is the proprietary solution which is bundled with the core itself. The MicroBlaze code is distributed to the end user using an encrypted VHDL format to protect the IP. Hence to be able to instantiate the core in a design, the encrypted VHDL must first be configured through 3rd party applications which outputs a netlist. The major problem with the netlist is that specific parts of the components are bound to specific components on the FPGA, relying on knowledge of the supporting FPGA. This is done by using the LOC-constraint, and the purpose is to deliver optimized hardware which fully utilizes the FPGA.

The problem arises when, through the VHDL code, several cores are being duplicated. The best way to do this would be to use the generate statement which is a key function in VHDL, and which is heavily used generating structures which are repeating in nature. Algorithm 3 shows an easy and clean way to do this. Here the *numCores* variable could be a globally set variable.

```

Data: numCores
cores : for I in 0 to numCores -1 generate do
  instance : microBlaze PORT (
    signals => signals_core(I)
  )
end

```

Algorithm 2: Generating several cores

However, due to the nature of the MicroBlaze core, this method led to several complications. First, the Xilinx tools generates files which depends on the correct name of the MicroBlaze instance. E.g. if the core is configured with the name “System”, the above mentioned example would fail. The generate statement prepends the name of the entity with a composite of the generate label and the counter. In this example a specific core would be named “cores(I).instance” in the netlist. The core being configured with the name “instance” would fail, having resources depending on the “instance” resource. Second is

the LOC-constraints, each instance of the core will try to occupy the same resources on the FPGA, and thus will fail.

```
instance : microBlaze PORT (
  signals_core0 => signals_core(0)
  signals_core1 => signals_core(1)
);
```

Algorithm 3: Generating several cores without generate

To solve this, another strategy was adopted. Instead of having XPS generate a single core per design, the design had to have several cores. The downside to this method is that there is no easy way of configuring the number of cores in the form of VHDL, and it is dependant of configuring XPS and MicroBlaze files instead of having a variable in the VHDL code. This means that for each core in the system, a whole new set of MicroBlaze core with its assorted accessory such as the OPB-bus and OMA interface manually has to be configured.

3.4 Cache

The cache block implementation is a fully parametrisable component which can be extended with different evict strategies. One of the key points is to be able to simulate different cache organisations, different line sizes, and different cache line sizes³. One of the key points when cache structures are concerned is to identify fully and directly associative caches as sub-set of set-associative caches. I.e. directly associative caches can be described as 1-set associative cache, and that fully associative caches in fact are 1-set associative caches where the tag length is the full size of the address width, leaving no bits left for the index. When that is identified, two key parameters that decides the structure of a cache block is clear, namely (1) set-associativity and (2) tag-size. The bits of the address used for indexing is given by the tag-size since the index concatenated with the tag must form the full address of the cache line. I.e. if the tag size is 25, and address width is 32, width of the index must be 32-25, which is 7.

One important parameter which is not included in this design is the cache line size. The cache line size describes the amount of bytes stored per cache lines. This is omitted in the design due to several reasons. Having a cache line size greater than the data width is favorable in an environment such as the Intel Pentium 4 where the main memory is a dual port RAM which delivers 64 bytes, and the data width of the CPU is 32 bit. If the cache would have been 32 bit, the cache block would discard 60 bytes per memory access. In this case a cache line size which is greater than the data width would make sense. However, the BenERA motherboard provides a 32 bit channel off the FPGA through which the chip communicates. This means that FPGA can request 32 bits of data per off-chip memory access, and for each time it requests data it needs to multiplex the address, direction and

³The amount of data stored per cache line

data transferred over the bus, which is overly expensive in terms of latency. If the cache line size is to be varied from a minimum of 32 bits to a multiple of 32, it would have cost that multiple of 32 times one memory transfer to data from to and fro the main memory. Another important factor that speaks for the omit of cache line size is the heavy limitation of the SelectBlock Ram available on the VirtexE FPGA. The cache line size would mean that the number of BlockRAM used per line would increase with the with $\frac{cachelinesize}{32}$ per 127 possible index times the degree of set-associative. Seeing how the number of available blocks on the FPGA is 94 this would lead to a reduction of the address space and sets available.

3.4.1 Storage Unit

To be able efficiently store data on the FPGA itself the on-board SelectRAM blocks have been utilized. This in comparison to an implementation written in VHDL which, if written incorrectly, would failed to recognised as RAM structure and instead use CLB. The first attempt to write a cache block without SelectRAM used 103% of the total FPGA resources which obviously was unrealisable, however it was less restricted by the limitations of the SelectRAM. A single memory block, which is the component wrapping the on-board block ram is shown in figur 3.4.1 and can hold up to 127 lines of cache, one set per line. Each SelectRam block can be configured as a dual port ram with an address width of 8 bits, and a data width of 16. Seeing how the data width of the CPU is 32 bit, an encoding scheme is needed. This is solved connecting the most significant bit to either 1 or 0, as shown in figure 3.4.1, leaving 7 bits to represent the address. This decoding scheme is also used at the SelectRAM block containing the meta data such as tag, LRU count and various amounts of flags such as valid bit and dirty bit used in the evict strategy. This makes the design use 2 SelectRAM blocks per 127 lines of data stored, which is a large part of the 94 available on the VirtexE FPGA, but it frees up CLB's that can be used to implement several cores and logic.

Seeing how an address width of 7 bits, i.e. 127 different indexes, can be a little restrictive, the design can expand the address space. This is done by stacking several memory blocks, into an array. When a given address is requested, the right memory block is addressed by using the most significant bits of the address⁴ to identify the right memory block. E.g., if the address 1000000_2 is requested, it will address the memory block with address 1, and in turn its memory cell 0. An expansion of the address space means that additional 2 to the power of extra bits of Single Memory blocks would be used per set, rapidly exhuming the available SelectRam blocks available. The indexes available is a direct result of the tag length parameter.

Besides the number of lines available, one of the parameters available is the number of sets per line, or the set associativity. The naive implementation of set associativity in the given FPGA environment would have been to store any number of sets in the same memory block. Although possible, that would have led to an increase in clock cycles

⁴address(7 to length)

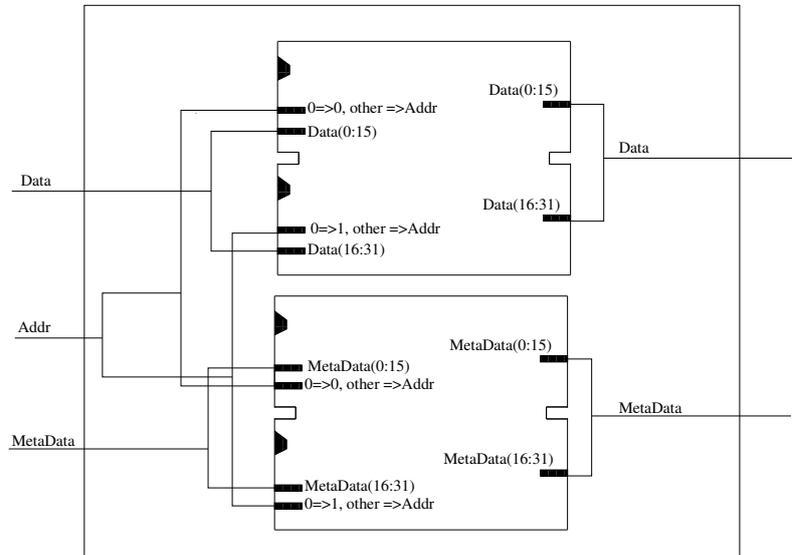


Figure 3.3: Single Memory block

wasted searching for the cache with the right tag. Done wrong the benefits of having a set associative cache would have diminished, and one would in terms of latency and logic overhead reach the scale of fully associative cache. Another more efficient way to store sets in the same SelectRAM would be to say that address 0 and 1 keeps the first set and so on. This would have required 2 extra clock cycles, or one extra clock cycle per set associativity, seeing how the chip must wait one clock cycle per memory access to the embedded SelectRAM block. Besides the waste of clock cycles, it would reduce the available indexes by the factor set associativity. Although a mutt point, since it is a direct result of all available strategies. Figure 3.4.1 shows how this implementation have done it, using a constant amount of clock cycles independent of the set associativity. For each new set added to the design, a new duplicate of the memory block gets added. Seeing how the address in the SelectRAM block is depending on only the most significant bits, the same index will be hit in all of the sets, and all sets will send both their meta data and data to a multiplexer as shown in the figure. The multiplexer will match each tag with the requested address, and output the data from the right memory set.

3.5 PCI Communication

The PCI Communication device, as seen in figure 3.5, acts as a gateway between the User FPGA and the PCI FPGA. Its main task is to provide to the system a way of communicating out of the User FPGA confirming to the OMA interface. Its role is to translate the internal signals into a format that suits the PCI FPGA. When the data are sent to the PCI FPGA, it alone will pass the data along to the host computer through DMA, and finally it will reach the controlling software. The communication between the

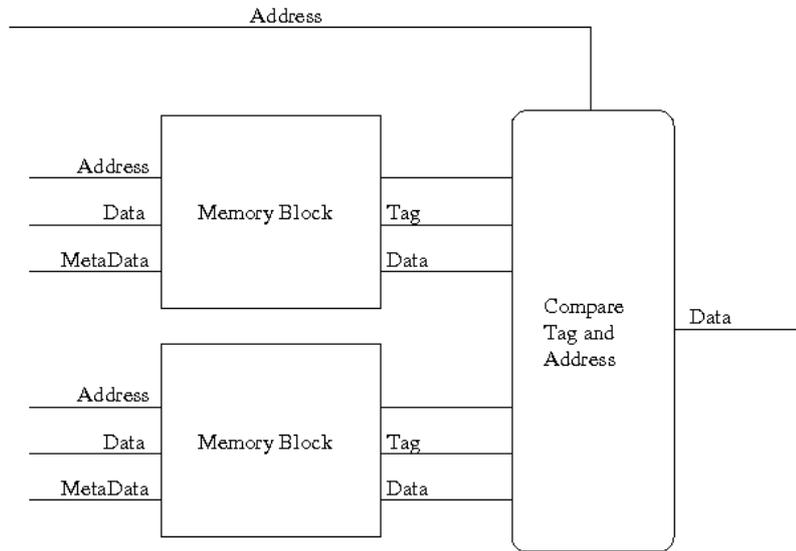


Figure 3.4: 2-way set associative block

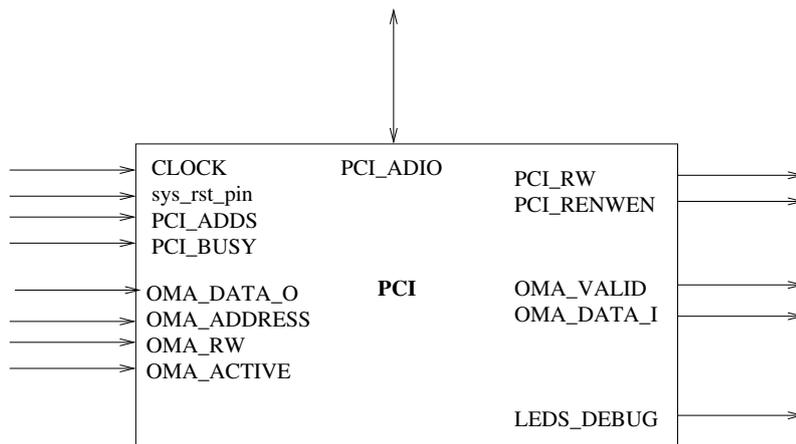


Figure 3.5: PCI Unit

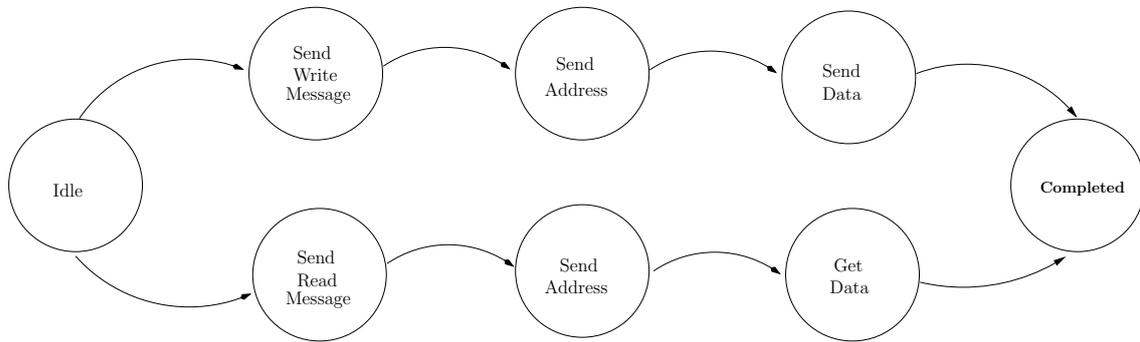


Figure 3.6: Com States

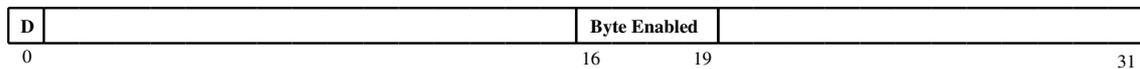


Figure 3.7: Control Word

User FPGA and the PCI FPGA happens through a dedicated 32 bit wide FIFO buffer on the PCI FPGA to which the User FPGA both reads and writes. However, seeing how the aggregated width of the data to be transfered to and fro the PCI Communication module exceeds 32 bit, a multiplexing scheme has been implemented.

To be able to access the memory banks implemented on the host computer in software, 3 important bits of informations is needed. First the direction of the transfer needs to be relayed, next the address to be operated on and finally the transaction of data which are to be stored or read from memory. With a 32-bit address and data bus, this would exceed the 32 bits available for communication through the PCI FPGA. Hence the communication will happen in the stages shown in figure 3.5. First the communication module will signal the direction of the transfer, either a write or read message. Then it will transfer the address of the cell in question. Pending on the direction of transfer, it will either get data from the PCI FPGA, or it will drive signals onto the bus itself.

A transfer gets initiated when the Active flag is asserted. When the active flag is asserted, the communication module stores the address, and in the case of a write transfer in an internal register to be able to meet the timing constraints put forth by the PCI FPGA. To initiate the communication with the controlling software on the host computer, the communication module puts a control word on the bus. The control word, shown in figure 3.5 contains information about the direction(D) of memory transfer and byte enabled bits(BE). The direction, if asserted indicates that the data should be transfered from the host computer to the User FPGA. Otherwise the transfer is a write action, and the data will flow from the User FPGA to the PCI FPGA. The Byte Enabled masks the bytes, and is ignored in the case of a read statement. When a read transaction happens, the software

based RAM will fill up the entire word from memory. However, if the transaction is a write action, the software based memory will only write to the part of memory signaled by the BE-flag.

3.6 Software

To be able to communicate with the processor, and to provide a opportunity for the processor to store its memory, dedicated software has been written. The controlling application loads the bitfile containing the processor onto the FPGA, then loads the requested application into memory before it resets the design and starts its role as memory.

When the application starts, it will initialize the memory. By hosting the main memory on the host computer, the design can allocate the amount of memory available on the host computer instead of storing data in the BlockRam on the FPGA. This effectively increases the amount of storage from 96 BlockRams à 4096 bits⁵ to the amount of memory available on the host computer to the maximum of 4 Gigabyte⁶. Depending on the user-input, the application will either load a given program or all zeroes into the memory. The latter intended for debugging purposes only. When the memory is loaded, the FPGA is initialized and the communication channels are set up. When the bitfile is loaded onto the FPGA, all buffers and reset signals are cleared. When everything is configured, the control is transfered to the FPGA.

```

Input: data
if state == control then
    state = getReadOrWrite ( data )
else
    if state == read then
        state = handleRead ( data )
    else
        state = handleWrite ( data )
    end
end

Function handle(data)

```

When the software has transfered the control to the hardware, it goes into the main loop shown in algorithm 5. The goal of the loop is to retrieve a word from the PCI Commu-

⁵376 Kb

⁶The address space of the processor being 32 bits

```
while 1 do  
    data = getWord32(from FPGA);  
    handle(data);  
end
```

Algorithm 5: Main loop

nication module⁷ and pass it on to the correct function. This is done by implementing a basic decode-and-dispatch loop[36]. For each turn the main loop fetches a word from the FIFO buffer on the PCI FPGA, and dispatches this to the handler function, as seen in algorithm 4. Depending on the internal state of the software itself, the handler function will further dispatch the data to the correct method. Due to the tight coupling between the internal states of the hardware and software, it is important that the software adheres to the states defined in section 3.5. If the software fails to send a message when the hardware is locked in the “Get Data” state, both the hardware and the software will go into a deadlock. Although a bad idea in real-world application, it was an acceptable trade-off due to the decreased complexity compared to a fault tolerant communication model.

3.7 Benchmark

To help benchmark the platform, several benchmarks has been developed. Several commercial benchmarking applications, such as Spec 2006[19] are available to hardware developers. However in the case of the Spec benchmark suit, it would require an underlying operating system and a ported version of the standard C library. Although a version of Linux exists for the MicroBlaze core⁸, it has not been ported to support a multicore architecture. This combined with fact that such benchmarks are not only testing the raw processor performance, but also the 3rd party libraries deemed the traditional benchmark suites unnecessary.

The suite developed presents two different kernels, each trying to benchmark different properties in the design. The first presented is testing the raw performance available to the processor by iterating through a sequence of number, accessing the memory as little as possible. The second benchmark does the opposite, loading and storing as much to the memory as possible.

3.7.1 BogoMIPS kernel

The BogoMIPS kernel, appendix A.1, is implemented as simple for-loop, iterating a various amount of times as seen in algorithm 6. What the BogoMIPS kernel tries to to benchmark is the raw performance of the overall system just doing enough calculations to increment

⁷See section 3.5

⁸µc Linux

```
i = 0
for i < MAX ITERATIONS do
  i += 1
end
```

Algorithm 6: bogoMIPS

the loop counter. By benchmarking this, it forms the basis for comparison for the rest of benchmark kernels.

3.7.2 Load-Store kernel

The load-store kernel, found in appendix A.4, is a kernel designed to stress test the bandwidth between the processors and the memory. It does so by first storing a certain value to the Nth memory addresses after the text segment, as seen in the code below.

```
storeLoop:
  sw   R11, R2, R0      ;store R11 to address R2+R0
  addi R1, R1, 0x4      ;add 0x4 to R1
  add  R2, R1, R2       ;add R1 to R2
  cmp  R4, R3, R1       ;R4 = R1 - R3
  blei R4, storeLoop   ;if R4 <= 0
```

When the store loop is done, the kernel starts over loading the data back in to the processor. By doing so the kernel is a memory bound kernel. Besides the benchmarking function, the load-store kernel also works as a memory boundary checker.

Chapter 4

Results

This chapter will present the findings of the benchmarks presented in section 3.7. This will provide foundation for the discussion in chapter 5 where a more final analysis will be presented.

4.1 Introduction

In a modern processor design, several factors decides if the processor is regarded as an high performance CPU. Even though the goal of the thesis was not to deliver a high performance multicore architecture, it presents several results varying the architectural parameters to show a proof of concept, and that the architecture described in chapter 3 allowed for rapid prototyping.

Being a prototype processor, the performance cannot match the one of a Intel or IBM. The clock frequency of the core operated at 42MHz, and thus it had been relentless to match it to an ASIC processor running at 4GHz. The timing comparisons done in this thesis was matched to a single core processor involving no cache on the same FPGA. That allowed the benchmarks to show if a configuration was better, but can in no way be compared to another configuration running on another system solely based on the timing information.

The timing model used is based on the Pentium instruction rdtsc which loads the number of cycles into a register specified by the developer. This was done in order to get a more accurate view than the coarse system clock could deliver. However, the draw-back to using the host computer as a timing device was that other processes acquiring resources affected the end results. To help diminish the effect of resource sharing, an average of several runs is presented as the results.

The final set of data extracted from the results was operations done at the main memory

Name	Number of Cores	Cache Size	Set Associativity	Index size
sCore	1	0	0	0
mCore	2	0	0	0
mCache	2	2Kb	2	7

Table 4.1: Processor Configurations

on the host computer. For each memory read or memory write to the host computer, an internal counter was increased which kept track of the total number of memory accesses.

The tested configurations are shown in table 4.1. The first configuration, sCore, is a single core without no cache, which will form the basis for comparison. The second configuration, mCore is a dual core CMP without cache. The final configuration, the mCache, is a dual core with a two-set associative cache, using 7 bits as the index, giving it 2Kb of available storage.

4.2 Benchmark results

4.2.1 BogoMIPS Kernel

The BogoMIPS kernel was designed to investigate how long the overall system used to perform a single performance bound problem. It did this by returning the amount of time spent performing a single loop.

Figure 4.1 shows the amount of time each configuration used for a given problem set. Both the single core and the multicore configurations uses between 59301819.8 and 11619442516 cycles on the host computer to finish. Without any cache the configurations are forced to request the same instructions and data on the host computer, and thus a lot of cycles will be wasted due to the processor design that has to wait for the memory transfer to complete. The mCache configuration managed to load the data requested into cache, and always had a cache hit when the data was re-requested. This lead the mCache to finish faster than the other cores, and had a approximately time of 150'0000 cycles independent of the problem size. This showed that the at problems of this size, the results are memory bound.

As presented by figure 4.2, both of the configurations without cache, the sCore and mCore designs has between 1223 and 240046 reads from the memory itself. This shows that for each instruction, the processor has gone outside the chip and requested data from the host computer. The mCache configuration, with a cache size of 2Kb, read a constant number of memory locations independent of the problem size. It also shows that the amount of writes varies with the mCache configuration.

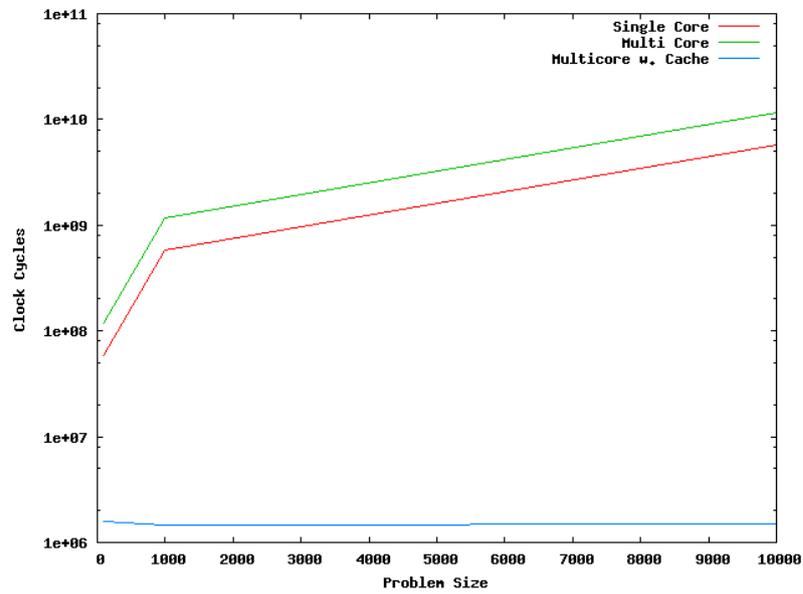


Figure 4.1: Different Configuration performance using the BogoMIPS kernel.
Logarithmic Plot

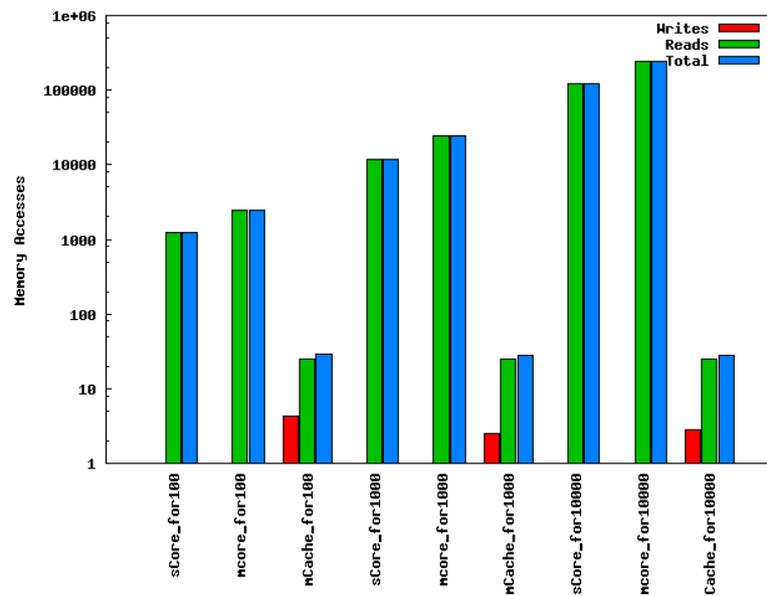


Figure 4.2: Different Memory Access patterns used by different configurations with the BogoMIPS kernel.

Logarithmic Plot

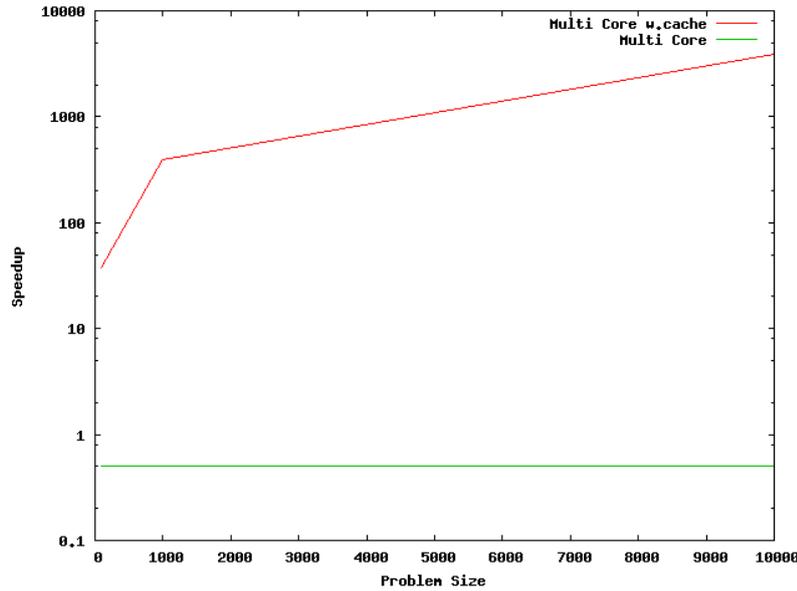


Figure 4.3: BogoMIPS Speedup
Logarithmic Plot

Figure 4.3 shows the amount of speedup gained by enabling cache on the multicore processor, which is in the magnitude of 3000. The speedup graph supports the finding presented in the above graphs, that the problem without cache is memory bound. The speedup graph also shows that with arbiter, in the case of the multi core design, that the round robin algorithm works as described in appendix B.5. By alternating between the cores on the chip, it should take approximately twice the time loading the same amount of instructions into each of the cores as the time used by a single core design. This is confirmed by graph 4.3 as it shows speedup of 0.5, compared to the single core processor.

4.2.2 Load/Store kernel

The load store kernel was designed to stress test the memory system of the processor design. The kernel tests is memory bound by first storing to the Nth first addresses, before it tries to re-read the same locations into memory.

Figure 4.4 shows the amount of cycles from used by the host computer waiting for the test to finish. The mCache configuration failed the verification phase of the test, and thus it has been omitted from the results. When the mCache kernel ran it would go in to a stall loop without sending the termination signal to the host computer to let it know that the application kernel was done. This might be due to a flaw in the internal memory structure of the cache itself, or somewhere else.

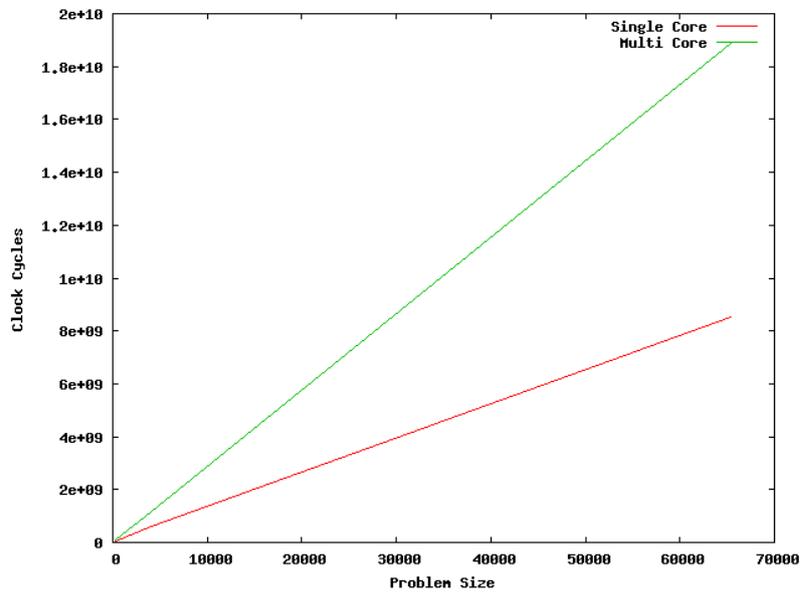


Figure 4.4: Different Configuration performance using the MemTest kernel.
Logarithmic Plot.

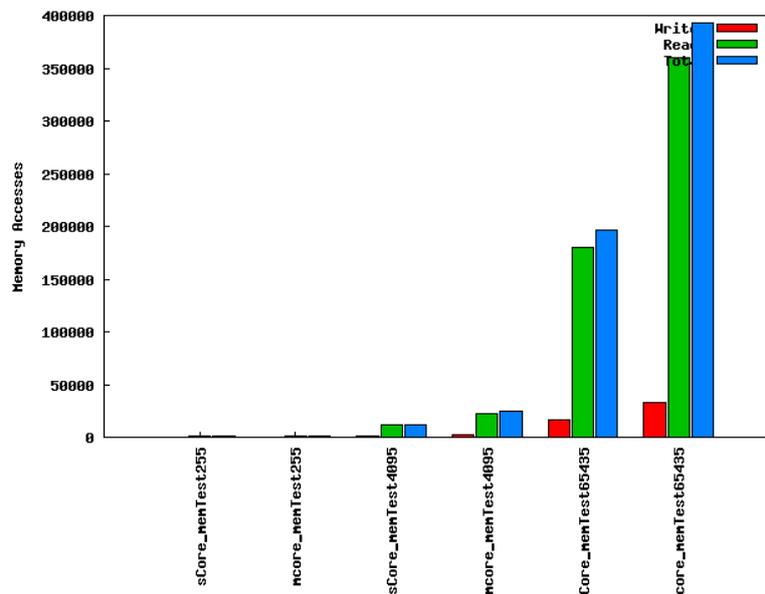


Figure 4.5: Different Memory Access patterns used by different configurations using the MemTest kernel.
Logarithmic Plot.

Figure 4.5 shows the number of memory accesses that the different configurations that cleared the validation phase had. Shown in the same figure is that the multicore architecture approximately has the double amount of memory requests as the single core configuration. This validated against the source code found in appendix A.4, where the kernel is ignorant of the number of cores available, leaving each core to request the same amount of memory as a single core found on the sCore configuration.

Chapter 5

Discussion

This chapter will present the analysis of the various topics touched in the thesis. It will start with a discussion of the core used, before it analyses the cache solution. It will then discuss the timing motivations before it continues with a presentation of the challenges involving 3rd party products. It will finish with a discussion of the results found in chapter 4.

5.1 Cores

The cores available at the beginning of the project were the MicroBlaze microcontroller and the PicoBlaze statemachine. Later in the project the Virtex-II Pro and its two PowerPC cores were made available. Of the cores available at the start of the project, Micro- and PicoBlaze, only the MicroBlaze was the one closest to a modern processor in design. The PicoBlaze controller, at 8 bit, didn't implement the most fundamental instructions found in a modern processor, hence rendered it useless in a Chip Multiprocessor design. The PowerPC hard cores on the Virtex-II Pro would have brought better compiler support to the project, and the bug presented in section 5.4.1 would most likely have been avoided. The PowerPC cores would also have done the porting of several operating systems easier, being a full processor supporting advanced features such as memory management by having a Memory Management Unit. The PPC would also have freed up a lot of resources such as LUT s, already embedded on the FPGA. However, the hard cores would break with one of the fundamental ideas in this project. By opting for a hard core, the system would have been locked down to the two cores delivered with the FPGA itself. One of the ideas with this project is that is should be possible to extend and test new architectures. Although two cores could be enough in testing cache and pre-fetching strategies, it would fall short when testing new interconnects where the amount of cores are vital to the outcome. Also, IBM predicting that they will see processors with 60 cores in the soon future, a locked down design depending solely on the two cores provided by the Virtex-II Pro would be

futile. Thus the argument that embedded PPC would have freed up resources falls on the ground that with two embedded cores, there is nothing left than cache and interconnect to use resources for. This, unless an array of extra components, not including cores, are to be tested.

By utilizing the MicroBlaze core, the design gained the ability to grow in both the number of cores, and still remained flexible enough to test several solutions. One of the key features of the MicroBlaze core was its ability to turn of unneeded features. This included the division unit and the floating point unit of the processor. However, should a benchmark or an application require those two features it is possible to enable them. This would lead to the cores requiring more areal on the FPGA, but the option has been retained. By having this option the developer can decide upon which feature set being the more important one. A researcher developing an intrinsic interconnect would most likely increase the overall feature size of the core itself in order to be able to have more cores on the FPGA. A person testing new multicore algorithms on the FPGA, might be more interested in having floating point operations instead of having a large amount of cores available.

5.2 Cache

One of the features in this project is the extendible cache solution. In a modern day processor, several levels of cache exists to ensure a reduced latency between the processor core and the memory. Another possibility to separate the different cache banks, allowing each core to have a private cache. This is done both to ensure a faster access time, and to “protect” important data from being evicted by another processor core.

Since each of the components implemented in the design enforces the OMA signal interface as described in section 3.2, it is possible to connect one cache solution to another one, without having to rewrite the system. Since all the sinks¹ have the same interface, the cache component can easily be connected to the arbiter, PCI Communication module or another cache module. By connecting the cache to another cache, one will in effect get a hierarchical memory. By connecting two caches with their correspondent cores to an arbiter, which in turn connects to an new cache bank, one will have created an two level cache with level 1 private and level 2 shared.

On the Virtex FPGA, having a hierarchical memory structure will solve the evict problem, since one core would not have the possibility to request memory from another core’s cache bank, and hence force forth an evict. However, to decrease the latency from having the signal traveling great distances demands great planning. On board the Virtex-E FPGA, each block RAM is evenly placed in columns as shown in figure 5.2, with the configurable logic blocks in between, all blocks guaranteeing a one cycle lookup time[42]. Hence to reduce the latency, either the logic of the cache unit must be reduced, or the distance the signal has to travel must be decreased.

¹Components handling the incoming request

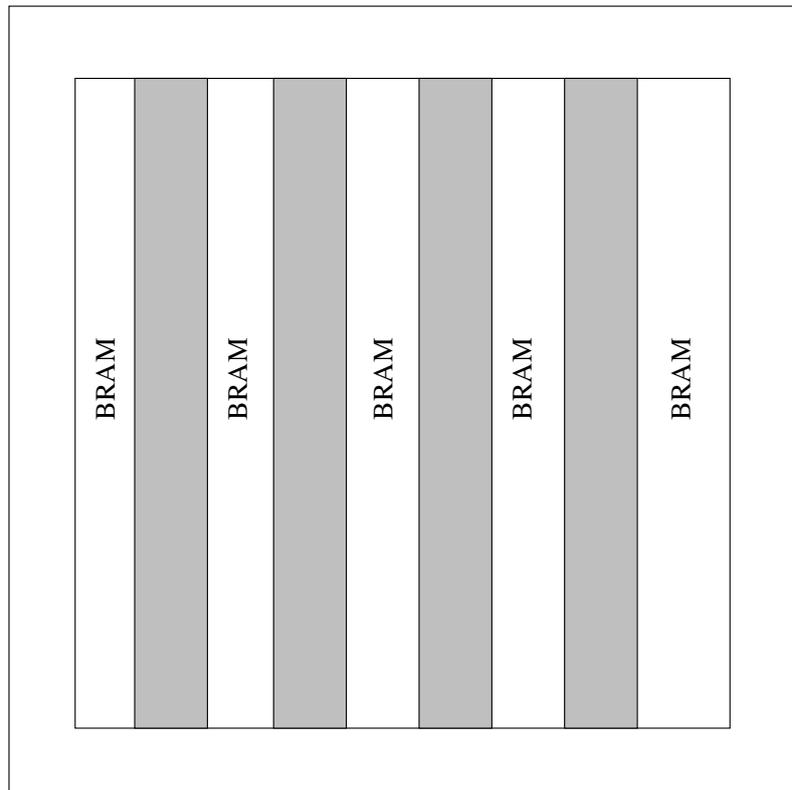


Figure 5.1: BRAM Location on the FPGA.

One of the motivations of the thesis' is to create an extendible architecture, thus the overhead in logic by making the cache solution extendible is required by the project's nature. This leaves the option of a better placement for level one cache compared to the placement of level 2 cache.

5.3 Time

One of the motivations behind the thesis was to reduce the amount of time spent waiting for the results, and thus be able to produce results in a shorter amount of time. The normal work flow for using the simulator at the NCAR group is do compile the new design, and let the simulators run over the weekend. When implementing the design in a FPGA, there are a significantly reduce in the time spent stalling waiting for the results to arrive.

Of the time spent in development, most of the time went into designing and verifying the hardware developed. A lot of time went into debugging errors that came during run-time at the FPGA itself, but which failed to appear during the simulations. This was due to the hardware giving "random" feedback, such as the FIFO buffer being full, the software might

be stalling waiting for execution time hence not requesting the available data. Situations which was impossible to emulate in the simulator, and hence hard to detect even with the proper use of LEDs. Although the cost of debugging the hardware could have been significantly reduced by access to the right equipment. Xilinx delivers a product named ChipScope, which through JTAG allows the developer to embed probes in their hardware design, extracting information at run-time. At the time this thesis was written the sole way of communicating to the JTAG interface on the FPGA was through a specialized cable not available to the writer.

However, with the behavioral code working, it took at a maximum of 5 minutes integrating it to the existing code base. A lot of time were saved defining a strict signal interface through which the different subsystems could communicate. From the code were done until the bitfile came out it took at a maximum 20 minutes. This is attributed the different processes required for creating a bitfile. I.e. Translation, Placement and Routing and the Mapping phase. From the bitfile was done to all of the benchmarks and validation kernels were done an additional of 10 minutes were spent waiting.

This leaves an overall time spent waiting for results at a maximum of 30 minutes depending on the problem size of the benchmark suite. This time would differ if the complexity of the design itself changes, i.e. designing hardware which would force the placement and routing tool to spend more time.

5.4 3rd Party Challenges

This section will discuss some of the challenges introduced by the various 3rd party software and hardware solutions. It will start off with introducing a bug found in the MicroBlaze branch of the GNU Compiler Collection, before describing some of the challenges with the glue code provided by Xilinx to connect to the OPB.

5.4.1 GNU Compiler Collection for MicroBlaze

GCC as described in section 2.6.1 has been ported to suit the MicroBlaze processor core. This allows the software developer to write code in a high level language such as C or C++ to help speed up development. However in the branch of GCC developed by Xilinx, a particularly bug was found. This had to with the way that the compiler handled 32 bits immediate values in the case of a branch and loading memory addresses. The payload of an instruction can at max be 16 bits, seeing how some bits are reserved the Op-code and registers. To solve this Xilinx have introduced a special atomic Immediate instruction. The payload of a Immediate instruction will be placed in a special register on the CPU. The following instruction, which must be a compatible Immediate instruction or else the register will be flushed, will send an additional 16 bits. The 16 bits from the immediate

```
12: 30600024      brlid   r3, r0, 0x24
```

Figure 5.2: Correct assembly code

```
12: b0000024      imm     0x24
16: 30600000      brlid   r3, r0, 0
```

Figure 5.3: Output from GCC

instruction will form the most significant bits, leaving the 16 bits from the following instruction to fill the least significant bits.

32 bit values are especially important when handling addresses. This might be the case when the developer either wants to load a word from main memory, or when the developer wants to do a jump. Figure 5.4.1 shows the correct code in case of a jump to current position + 0x24. However the output from the MicroBlaze port of GCC gave the output shown in figure 5.4.1. The code shown in figure 5.4.1 loads 0x24 into the immediate register, before it performs a brlid instruction, brlid being the branch immediate instruction. Internally the MicroBlaze core will concatenate the value of the immediate register to the payload of the brlid, the immediate register forming the high part of the word. This will have the effect of jumping to the address IP + 0x24000000, instead of the wanted IP + 0x24. The bug is due to having the software developers not being aware of the importance of order in which the immediate value has to be sent.

5.4.2 Xilinx IP Glue for OPB

To allow for rapid development of external peripheral devices, Xilinx have bundled in their Embedded Development Kit a customizable IP which glues the user developed hardware unit to the more advanced OPB. The main goal of the IP is to hide the complexity of the CoreConnect architecture. It does so by handling address decoding, making the user developed Property(IP) respond only to the given address, i.e. make certain that the component only respond to the configured address range. It also have a lot of other functionality to used by the developer, as hiding unused signals.

To help implement the wrapper layer between each MicroBlaze core and the rest of the design, the design utilized the IP glue code. This was utilized in the core wrappers described in section 3.3.1. Doing so led forth a new set of challenges. Each IP attached to the OPB or LMB has a set of parameters automatically configured by the tools provided by Xilinx. Amongst others this included the addresses space that the implemented IP should respond to. I.e. a IP can be configured to only respond to memory request between address X and Y . However, in the Xilinx IP there were two places where this address range is set, and the tools provided by Xilinx only updated one of the spots. This left implemented IPs to respond to address at $baseAddress + 0xFF$, which in effect gave the processor a 255

byte address space. The fix to this problem was to hard code in a address range which suited the project.

The second bug found in the IP glue code was regarding the address buses from the MicroBlaze core itself. From the master that the proxy connected to, two sets of address buses were provided. In the IP glue code it was asserted that the implemented IP only needed the signal of one the buses. The MicroBlaze core on the other hand alternated between the two buses when it requested addresses. This left the processor design only responding to the requested addresses which were 8-byte aligned, e.g. 0x0, 0x8, 0x16, instead of wanted 4-byte aligned addresses. This left the wrapper layer only responding to half of the addresses requested by the processor core. The fix to this problem was to multiplex the signals from each bus into a single bus in the wrapper layer, by checking which of the buses currently active.

5.5 Benchmarks

The benchmarks done in chapter 4, were designed to validate the different configurations, and to give a rough estimate of the performance of the design itself. The first benchmark was a simple for loop which tested the performance of the design itself. As presented by the graphs 4.1 and 4.2 most of the time spent by the different designs went into stalling for the memory transactions to complete. This was also the dominating part in the application designed to be memory bound. Seen in the first benchmark, the execution time of the multicore was roughly twice the time of a single core processor. This coincide with the hardware designed. As seen in appendix B.8 the multicore architecture has a arbiter between the PCI communication and the cores themselves. This means that for each request, the multicore architecture must wait for the core utilizing the memory bus to finish before it can request memory from the subsystem. This arbitration is not featured on the single core processor as shown in appendix B.9. Seeing how the first kernel is negligent of the number of cores available it forces both cores on the multicore architecture to fetch all the memory location. Since the arbitration mechanism is based upon a round robin algorithm the multicore design will have to wait twice the amount of time as a single core processor for the memory transfer to complete.

The multicore design with the cache ran much faster than both the single- and multicore architectures without cache available. This was due to most of the time was spent waiting for the memory transfers. The cache managed to store all of the requested data lines, and thus it only had to request the each data once. As shown in figure 4.2, the multicore cache design only requested 25 memory reads independent of the kernel problem size, which coincide with the size of the application itself.

In the load store benchmark, see section 4.2.2, the multicore with cache designed failed to verify during runtime. During the benchmark the core with the cache managed to go into a infinite loop without sending the termination signal. This probably means that

there is a bug in the write branch of the cache solution. Thoroughly tests were done in the simulation bench on the cache sub-system forcing the microblaze core to do a number of writes without managing to produce the same results, the cache subsystem managed to pass all of the test vectors. The test vectors in the simulation bench also tried to investigate if the possibility that bug had to do with the amount of memory read into the cache. One possibility could have been that the tag system was flawed, an address with the same index had the possibility to falsely trigger a cache hit. However, this did not turn out to be the case in the simulator. That test vector also passed.

Besides that the mCache design did not pass its verification, the same tendency found in the for-kernel appeared in the load/store kernel. With all cache disabled, the designs were more bound by the time it takes to access memory, than the actual computations. Each memory request has to travel out on the host computer's PCI bus, through the PCI FPGA. Then the data would have to be copied from kernel to userspace before it reaches the controlling software before a acknowledging signal is sent back the same route.

This also behavior concurs with the tendencies in computer architecture, and is why a software based memory was opted for. Having the memory on the FPGA itself would have significantly decreased the run time of a single kernel. By putting the memory on the host computer, it put forth a artificial memory access delay found in most modern computers. It also presented that including cache in the design, the for loop kernel was virtually depending on the time it used waiting for the compulsive memory transactions. This was seen on the constant number of cycles spent, independent of the iterations spent.

Chapter 6

Conclusion

This chapter presents some final remarks about the project, before it presents some ideas for further works.

6.1 Conclusion

To be able to do research on the field of Computer Architecture in a period where there are a lot of parallel research projects, a platform which brings results within the hour is to a great advantage. There are a lot of methodologies for speeding up the traditional simulators, such as using clusters and a various amount of techniques to decrease the time spent in simulation[16]. However in terms of both accuracy and time, hardware outperforms software in simulation of hardware architectures.

During the period of the thesis the author has implemented a framework for simulating and experimenting with chip multiprocessors in hardware. A core has been decided upon and integrated into the chip. A parametrizable cache solution has been designed and implemented. Benchmarks which tested and probed the implemented design has been written. The benchmarks showed that both the multicore and single core architecture worked. However, the benchmarks produced found a bug in the design of the cache when accessing memory above a given size. Of software, an application which hosts the main memory, and controls the application has been written. Also, by testing the different configurations the thesis has shown that it has made it possible to test different processor configurations in a rapid manor.

This thesis has shown that it has significantly reduced the time waiting for results. The thesis has also laid the groundwork for further development of a hardware platform which can be used to prototype novel computer architectures. It has implemented the most common structures found on a chip multicore, and abstracted most of the hardware specific

challenges put forth by the 3rd part vendors. Due to the strict interfaces between the different modules of the design, future developers can easily configure the platform in virtual any desired configuration. By enforcing a strict, but simplified interface to the different hardware modules the next research project would have to focus on the inherent challenges of his own design. This frees the next researcher from the implementation challenges put forth by the 3rd party vendors, leaving him to concentrate on his project alone.

Although none of the techniques involved are technologically groundbreaking by themselves, this thesis has put them together allowing the next project to test new configurations and parameters in a rapid and more accurate fashion. This will help the future researchers to increase their productivity, producing more results, one of which might be the next breakthrough in Chip Multiprocessors.

6.2 Future Work

To be able to better benchmark the design, a benchmark suite should be investigated. However, most benchmark suites requires a underlying port of the standard C library. To be able to get the standard C library working in the design a suitable operating system should be ported, this in order to provide the library with the required system calls. Some operating systems already has been ported to the MicroBlaze platform, but none of them have support for MicroBlaze in a symmetric multiprocessing environment. One idea could have been to port the embedded version of Linux to support a multicore MicroBlaze solution. Having a operating system would give the project access to a wider array of benchmarks and applications.

One of the more challenging aspects of designing the hardware was that situations impossible to predict in the simulator arose runtime on the FPGA. Without any possibility to analyze the internal state of FPGA itself, some bugs took a lot of resources locating. Xilinx has an logic analyzer, ChipScope, which can be integrated into a design, probing and returning the information found. However, it was not possible to connect to the JTAG interface at BenERA motherboard using only software, a specialized cable was needed. To help ease the development of future hardware, an solution where one connects to the JTAG interface using ChipScope would lead to an decrease in time spent in verifying stage.

One of the major changes would be to upgrade the system to use a more modern version of the Virtex FPGA used in this thesis. A newer FPGA will give the system more available BlockRAM, which is strongly advisable if the design is to incorporate more than two cores and a small shared level two cache. An modern FPGA such as the Virtex-II Pro would open for the possibility of using a newer version of the MicroBlaze core than the MicroBlaze version 4. The newer versions includes user defined registers which would have allowed the design to skip the CPU ID coprocessor designed, its sole task responding on requests for each core's id.

An future expansion would be to integrate the design into an already existing software based simulator such as the SimpleScalar or M5 simulator. To be able to do so, the design would have to be automatically generated from the configuration files. The current design allows for easy customization, configuring only needed at the top level. Thus it would be possible to automate the creation of the hardware itself. However, the challenge would lie in gathering statistics from the system. There exists no method of extracting statistics in the current design. It is, however, two approaches for information gathering from the hardware. The first would be to use the earlier mentioned ChipScope to probe the wanted signals, using an external application to interpret the data gathered. This would lead to a fully customizable method, letting the designer worry about only designing the hardware and afterwards attaching the test probes. The latter option of gathering statistics would be to design a subsystem connected to either the peripheral bus or the fast simplex link. This module could have had an shared bus which through the different components on the chip could send statistics. This is a model loosely based on the workings of the JTAG interface.

Bibliography

- [1] Ben Ames. NASA backs quantum computing claim, September 2007. <http://www.itworld.com/Tech/3494/070309nasaquantum/index.html>.
- [2] Arvind, Krste Asanovic, Derek Chiou, James C. Hoe, Christoforos Kozyrakis, Shih-Lien Lu, Mark Oskin, David Patterson, Jan Rabaey, and John Wawrzynek. Project summary for nsf solicitation 04-588 cri: Ramp: Research accelerator for multiple processors - a community vision for a shared experimental parallel hw/sw platform. Technical report, 2005.
- [3] Nathan L. Binkert, Ronald G. Dreslinski, Lisa R. Hsu, Kevin T. Lim, Ali G. Saidi, and Steven K. Reinhardt. The m5 simulator: Modeling networked systems. *IEEE Micro*, 26(4):52–60, 2006.
- [4] Doug Burger and Todd M. Austin. The simplescalar tool set, version 2.0. *SIGARCH Comput. Archit. News*, 25(3):13–25, 1997.
- [5] J. Adam Butts and Gurindar S. Sohi. A static power model for architects. In *International Symposium on Microarchitecture*, pages 191–201, 2000.
- [6] Jichuan Chang and Gurindar S. Sohi. Cooperative Caching for Chip Multiprocessors. *SIGARCH Comput. Archit. News*, 34(2):264–276, 2006.
- [7] Mikael Collin, Raimo Haukilahti, Mladen Nikitovic, and Joakim Adomat. SoCrates - A Multiprocessor SoC in 40 days. In *Conference on Design, Automation and Test in Europe 2001, Designer*, Munich, Germany, March 2001.
- [8] Katherine Compton and Scott Hauck. Reconfigurable computing: a survey of systems and software. *ACM Comput. Surv.*, 34(2):171–210, 2002.
- [9] International Business Machines Corp. IBM Unleashes World’s Fastest Chip in Powerful New Computer. <http://www-03.ibm.com/press/us/en/pressrelease/21580.wss>, May 2007.
- [10] Intel Corporation. Intel First to Demonstrate Working 45nm Chips. <http://www.intel.com/pressroom/archive/releases/20060125comp.htm>, January 2006.
- [11] Haakon Dybdahl, Per Stenström, and Lasse Natvig. A cache-partitioning aware replacement policy for chip multiprocessors. In *HiPC*, pages 22–34, 2006.

-
- [12] S. Ghiasi and D. Grunwald. Aide de camp: Asymmetric dual core design for power and energy reduction. *Technical Report CU-CS-964-03*, May 2003. Department of Computer Science, University of Colorado, Boulder.
- [13] Simcha Gochman, Ronny Ronen, Ittai Anati, Ariel Berkovits, Tsvika Kurts, Alon Naveh, Ali Saeed, Zeev Sperber, and Robert C. Valentine. The Intel® Pentium® M Processor: Microarchitecture and Performance. *Intel Technology Journal*, 02, 2003.
- [14] Lance Hammond, Benedict A. Hubbert, Michael Siu, Manohar K. Prabhu, Michael Chen, and Kunle Olukotun. The Stanford Hydra CMP. *IEEE Micro*, 20(2):71–84, March 2000.
- [15] John L. Hennessy and David A. Patterson. *Computer Architecture A Quantitative Approach, 3rd Edition*. Morgan Kaufman, San Mateo, California, 2003.
- [16] Ken Hines and Gaetano Borriello. Dynamic communication models in embedded system co-simulation. In *DAC '97: Proceedings of the 34th annual conference on Design automation*, pages 395–400, New York, NY, USA, 1997. ACM Press.
- [17] IBM. The CoreConnect Bus Architecture White Paper. "<http://www-306.ibm.com/chips/techlib/techlib.nsf/techdocs/852569B20050FF77852569910050C0FB>, January 1999.
- [18] Ahmed Jeraya and Wayne Wolf. *Multiprocessor Systems-on-Chips*. Morgan Kaufmann, 2004.
- [19] John L. Henning and SPEC CPU Subcommittee. SPEC CPU2006 Benchmark Descriptions. *ACM SIGARCH newsletter, Computer Architecture News*, 34, 2006.
- [20] Changkyu Kim, Doug Burger, and Stephen W. Keckler. Nonuniform cache architectures for wire-delay dominated on-chip caches. *IEEE Micro*, 23(6):99–107, 2003.
- [21] Rakesh Kumar, Victor Zyuban, and Dean M. Tullsen. Interconnections in multi-core architectures: Understanding mechanisms, overheads and scaling. In *ISCA '05: Proceedings of the 32nd Annual International Symposium on Computer Architecture*, pages 408–419, Washington, DC, USA, 2005. IEEE Computer Society.
- [22] Prathap Muthana, Madhavan Swaminathan, Rao Tummala, Venkatesh Sundaram, Lixi Wan, S.K Bhattacharya, and P.M. Raj. Packaging of multi-core microprocessors: Tradeoffs and potential solutions. In *Electronic Components and Technology Conference, 2005. Proceedings. 55th*, volume 2, pages 1895 – 1903. IEEE, 2005.
- [23] Nallatech. FUSE C/C++ API Developers Guide. www.es.ele.tue.nl/mininoc/doc/fuse_api.pdf, July 2002. Document Number: NT107-0068.
- [24] Nallatech. Compact PCI DIME-II Motherboard. <http://www.nallatech.com/mediaLibrary/images/english/1865.pdf>, April 2007.
- [25] Lasse Nattvig. NCAR NTNU Computer Architecture Research Group...improving multicore memory systems. <http://ncar.idi.ntnu.no>.
-

-
- [26] NetBSD. NetBSD/macppc. <http://www.netbsd.org/Ports/macppc/>. Last visited: Mars, 2007.
- [27] Kunle Olukotun and Lance Hammond. The future of microprocessors. *Queue*, 3(7):26–29, 2005.
- [28] Kunle Olukotun, Mark Horowitz, and Monica Lam. Flexible architecture for simulation and testing (fast), 2003. Last Visited: 5. June 2007.
- [29] OpenBSD. OpenBSD/macppc. <http://www.openbsd.org/macppc/>. Last visited: Mars, 2007.
- [30] David A. Patterson. Computer Science Education in the 21st Century. *Commun. ACM*, 49(3):27–30, 2006.
- [31] penguinppc.org. Introduction to PowerPC Linux. <http://penguinppc.org/about/intro.php#hardware>. Last visited: Mars, 2007.
- [32] Jr. Philip Enslow. Multiprocessor Organization - a Survey. *ACM Comput. Surv.*, 9(1):103–129, 1977.
- [33] Gordon D. Robinson. Why 1149.1 (jtag) really works. *Electro/94 International Conference Proceedings. Combined Volumes*, pages 749–754, 1994.
- [34] Robert R. Schaller. Moore’s law: past, present, and future. *IEEE Spectr.*, 34(6):52–59, 1997.
- [35] Alan Jay Smith. Cache memories. *ACM Comput. Surv.*, 14(3):473–530, 1982.
- [36] Jim Smith and Ravi Nair. *Virtual Machines: Versatile Platforms for Systems and Processes*. Morgan Kaufmann, June 2005.
- [37] S. W. Song, J. D. Zheng, and W. B. Gardner. Prototyping a Residential Gateway Using Xilinx ISE. In *IEEE International Workshop on Rapid System Prototyping*, pages 267–269. IEEE Computer Society, 2005.
- [38] Lawrence Spracklen and Santosh G. Abraham. Chip multithreading: Opportunities and challenges. In *HPCA '05: Proceedings of the 11th International Symposium on High-Performance Computer Architecture*, pages 248–252, Washington, DC, USA, 2005. IEEE Computer Society.
- [39] Richard M. Stallman. *Using and Porting the GNU Compiler Collection, For GCC Version 2.95*. Free Software Foundation, 1999.
- [40] Zhenlin Wang, Kathryn S. McKinley, Arnold L. Rosenberg, and Charles C. Weems. Using the compiler to improve cache replacement decisions. In *PACT '02: Proceedings of the 2002 International Conference on Parallel Architectures and Compilation Techniques*, page 199, Washington, DC, USA, 2002. IEEE Computer Society.

-
- [41] John Wawrzynek, Mark Oskin, Christoforos Kozyrakis, Derek Chiou, David A. Patterson, Shih-Lien Lu, James C. Hoe, and Krste Asanovic. Ramp: A research accelerator for multiple processors. Technical Report UCB/EECS-2006-158, EECS Department, University of California, Berkeley, November 24 2006.
- [42] Xilinx. Xilinx Virtex-E 1.8V Field Programmable Gate Arrays Data Sheet. <http://inst.eecs.berkeley.edu/~cs150/Documents/virtexE-datasheet.pdf#search=%22VirtexE%22>, July 2002. Revision: 2.4.
- [43] Xilinx. PicoBlaze 8-bit Embedded Microcontroller User Guide. <http://www.xilinx.com/bvdocs/userguides/ug129.pdf>, November 2005. Revision: 1.1.1.
- [44] Xilinx. Embedded System Tools Reference Manual, Embedded Development Kit, EDK 8.2i. http://www.xilinx.com/ise/embedded/est_rm.pdf, June 2006. Revision: v6.0.
- [45] Xilinx. MicroBlaze Processor Reference Guide. http://www.xilinx.com/ise/embedded/mb_ref_guide.pdf, September 2006. Revision: 7.0.
- [46] Xilinx. Virtex-5 Family Overview - LX and LXT Platforms. <http://direct.xilinx.com/bvdocs/publications/ds100.pdf>, May 2007. Revision: 3.1.

Appendix A

Benchmarks

A.1 BogoMIPS kernel

```
1 #include "utils.h"
2
3 int main(){
4     int i = 0 ;
5     int x;
6
7     for ( i = 0 ; i < 2; i++ ){
8         x++;
9     }
10    halt();
11    return 0;
12 }
13
14 void halt () {
15     asm("addi_r1,_R0,_0x70000000\n\t\
16     __addi_r11,_R0,_0x7FFFFFFF\n\t\
17     __sw__R1,_R0,_R11\n\t");
18 }
```

A.2 Sample Linker script

```
1 ENTRY(main)
2
3 MEMORY
4 {
5     mm : ORIGIN = 0x00000000 , LENGTH = 0x00001000
6 }
7
8 SECTIONS {
9     . = 0x0 ;
10    .text : { *(.text) } > mm
11    .rodata : { *(.text) } > mm
12    .bss : { *(.bss) } > mm
13    .data : { *(.data) } > mm
14 }
```

A.3 Sample Makefile

```
1 APPNAME = for
2 GCC    = mb-gcc
3 LD     = mb-ld
4 LSCRIPT = linker.ld
5
6 all: ${APPNAME}
7
8 ${APPNAME}: ${APPNAME}.o
9     ${LD} -p --oformat=binary -o $@ ${LSCRIPT}  $^ /opt/EDK/gnu/microblaze/lin
        /lib/gcc/microblaze/3.4.1/libgcc.a
10
11 %.o: %.s
12     ${GCC} -c $^
13
14 %.s: %.c
15     ${GCC} -S $^
16
17
18 clean:
19     rm *.o *.s ${APPNAME}
```

A.4 LoadStore kernel

```
1  .text
2
3  _start: .global _start
4  .global _main
5
6  addi R11, R11, 0xFFFFFFFF
7
8  addi R2, R0, fnord // Address lookup.
9  addi R1, R0, 0     // Counter variable.
10
11         // Change here to set the iterations.
12  addi R7, R0, 0x1  // Check the 255 * 4 bytes after fnord.
13
14  add  R3, R0, R7   // Saves one imm instruction during the reset op.
15
16  storeLoop:
17  sw   R11, R2, R0
18  addi R1, R1, 0x4
19  add  R2, R1, R2
20  cmp  R4, R3, R1
21  blei R4, storeLoop
22
23
24  // Reset...
25  addi R2, R0, fnord // Address lookup.
26  addi R1, R0, 0     // Counter variable.
27  add  R3, R0, R7   // ^-- See explanation above.
28
29  loadLoop:
30  lw   R11, R2, R0
31  addi R1, R1, 0x4
32  add  R2, R1, R2
33  cmp  R4, R3, R1
34  blei R4, loadLoop
35
36  terminate:
37  addi R11, R0, 0x7FFFFFFF
38  sw  R11, R0, R11
39
40  fnord:
```

Appendix B

Hardware

B.1 Core OMA Interface

1 —

2 — *user_logic.vhd - entity/architecture pair*

3 —

4 —

5 —

6 — *** Copyright (c) 1995-2006 Xilinx, Inc. All rights reserved.*

7 — ****

8 — *** Xilinx, Inc.*

9 — *** XILINX IS PROVIDING THIS DESIGN, CODE, OR INFORMATION "AS IS"*

10 — *** AS A COURTESY TO YOU, SOLELY FOR USE IN DEVELOPING PROGRAMS AND*

11 — *** SOLUTIONS FOR XILINX DEVICES. BY PROVIDING THIS DESIGN, CODE,*

12 — *** OR INFORMATION AS ONE POSSIBLE IMPLEMENTATION OF THIS FEATURE,*

13 — *** APPLICATION OR STANDARD, XILINX IS MAKING NO REPRESENTATION*

14 — *** THAT THIS IMPLEMENTATION IS FREE FROM ANY CLAIMS OF INFRINGEMENT,*

15 — *** AND YOU ARE RESPONSIBLE FOR OBTAINING ANY RIGHTS YOU MAY REQUIRE*

16 — *** FOR YOUR IMPLEMENTATION. XILINX EXPRESSLY DISCLAIMS ANY*

17 — *** WARRANTY WHATSOEVER WITH RESPECT TO THE ADEQUACY OF THE*

```

18  —  **
19  —  ** IMPLEMENTATION, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OR
20  —  ** REPRESENTATIONS THAT THIS IMPLEMENTATION IS FREE FROM CLAIMS OF
21  —  ** INFRINGEMENT, IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS
22  —  ** FOR A PARTICULAR PURPOSE.
23  —  **
24  —  **
25  —  **
26  —  *****
27  —
28  —
29  —
30  —
31  —
32  —  Naming Conventions:
33  —  active low signals:          "_n"
34  —  clock signals:              "clk", "clk_div#", "clk_#x"
35  —  reset signals:             "rst", "rst_n"
36  —  generics:                   "C_*"
37  —  user defined types:         "*_TYPE"
38  —  state machine next state:   "*_ns"
39  —  state machine current state: "*_cs"
40  —  combinatorial signals:      "*_com"
41  —  pipelined or register delay signals: "*_d#"
42  —  counter signals:           "*_cnt*"
43  —  clock enable signals:       "*_ce"
44  —  internal version of output port: "*_i"
45  —  device pins:               "*_pin"
46  —  ports:                      "- Names begin with Uppercase"
47  —  processes:                  "*_PROCESS"
48  —  component instantiations:  "<ENTITY>I_<#|FUNC>"
49  —
50
51  —  DO NOT EDIT BELOW THIS LINE —————
52  library ieee;
53  use ieee.std_logic_1164.all;
54  use ieee.std_logic_arith.all;
55  use ieee.std_logic_unsigned.all;
56
57  library proc_common_v2_00_a;
58  use proc_common_v2_00_a.proc_common_pkg.all;

```

```

59 -- DO NOT EDIT ABOVE THIS LINE -----
60
61 --USER libraries added here
62
63 --
-----

64 -- Entity section
65 --
-----

66 -- Definition of Generics:
67 --   CAWIDTH           -- User logic address bus width
68 --   CDWIDTH          -- User logic data bus width
69 --   CNUMLCE          -- User logic chip enable bus width
70 --
71 -- Definition of Ports:
72 --   Bus2IP_Clk        -- Bus to IP clock
73 --   Bus2IP_Reset      -- Bus to IP reset
74 --   Bus2IP_Addr       -- Bus to IP address bus
75 --   Bus2IP_Data       -- Bus to IP data bus for user logic
76 --   Bus2IP_BE         -- Bus to IP byte enables for user logic
77 --   Bus2IP_RNW        -- Bus to IP read/not write
78 --   Bus2IP_RdCE       -- Bus to IP read chip enable for user
   logic
79 --   Bus2IP_WrCE       -- Bus to IP write chip enable for user
   logic
80 --   IP2Bus_Data       -- IP to Bus data bus for user logic
81 --   IP2Bus_Ack        -- IP to Bus acknowledgement
82 --   IP2Bus_Retry      -- IP to Bus retry response
83 --   IP2Bus_Error      -- IP to Bus error response
84 --   IP2Bus_ToutSup    -- IP to Bus timeout suppress
85 --
-----

86
87 entity user_logic is
88   generic
89   (
90     -- ADD USER GENERICS BELOW THIS LINE -----
91     --USER generics added here
92     -- ADD USER GENERICS ABOVE THIS LINE -----
93
94     -- DO NOT EDIT BELOW THIS LINE -----
95     -- Bus protocol parameters, do not add to or delete
96     CAWIDTH           : integer           := 32;
97     CDWIDTH          : integer           := 32;
98     CNUMLCE          : integer           := 2
99     -- DO NOT EDIT ABOVE THIS LINE -----
100  );
101  port
102  (
103    -- ADD USER PORTS BELOW THIS LINE -----
104    --USER ports added here
105    -- The data I send out off the chip
106    OMA_DATA_O        : out std_logic_vector(0 to CDWIDTH-1);

```

```

107  — The data I get in from outside the chip.
108  OMA_DATA_I      : in std_logic_vector(0 to C_DWIDTH-1) := (others =>
      '0');
109  — If the data on OMA_DATA_* is the data you requested.
110  OMA_VALID      : in std_logic := '0';
111  — Which addr to operate on.
112  OMA_ADDRESS    : out std_logic_vector(0 to C_AWIDTH-1);
113  — Read or write. On asserted we read.
114  OMA_RW        : out std_logic;
115  — If there is a request pending.
116  OMA_ACTIVE     : out std_logic;
117  — Leds debug
118  LEDS_DEBUG     : out std_logic_vector( 0 to 15 );
119  OMA_BE        : out std_logic_vector(0 to C_DWIDTH/8-1);
120
121  — ADD USER PORTS ABOVE THIS LINE —————
122
123  — DO NOT EDIT BELOW THIS LINE —————
124  — Bus protocol ports, do not add to or delete
125  Bus2IP_Clk     : in  std_logic;
126  Bus2IP_Reset  : in  std_logic;
127  Bus2IP_Addr   : in  std_logic_vector(0 to C_AWIDTH-1);
128  Bus2IP_Data   : in  std_logic_vector(0 to C_DWIDTH-1);
129  Bus2IP_BE     : in  std_logic_vector(0 to C_DWIDTH/8-1)
      ;
130  Bus2IP_RNW    : in  std_logic;
131  Bus2IP_RdCE   : in  std_logic_vector(0 to C_NUMLCE-1);
132  Bus2IP_WrCE   : in  std_logic_vector(0 to C_NUMLCE-1);
133  IP2Bus_Data   : out std_logic_vector(0 to C_DWIDTH-1);
134  IP2Bus_Ack    : out std_logic;
135  IP2Bus_Retry  : out std_logic;
136  IP2Bus_Error  : out std_logic;
137  IP2Bus_ToutSup : out std_logic
138  — DO NOT EDIT ABOVE THIS LINE —————
139  );
140 end entity user_logic;
141
142  —

```

```

143 — Architecture section
144 —

```

```

145
146 architecture IMP of user_logic is
147   constant zeros      : std_logic_vector(0 to C_NUMLCE-1) := (others => '0');
148   type STATE is ( IDLE, READING, WRITING, RCOMP, WCOMP );
149
150   signal writeEnabled : std_logic := '0';
151   signal readEnabled  : std_logic := '0';
152   signal clock        : std_logic;
153   signal iState       : STATE := IDLE;
154   signal leds         : std_logic_vector(0 to 15) := (others => '1');
155
156 begin

```

```

157  — My own signals.
158  writeEnabled <= (Bus2IP_WrCE(0) or Bus2IP_WrCE(1) ) and not Bus2IP_RNW ;
159  readEnabled  <= (Bus2IP_RdCE(0) or Bus2IP_RdCE(1) ) and Bus2IP_RNW ;
160
161  clock        <= Bus2IP_Clk;
162
163  — Let some signals out independently of anything else.
164  OMA_ADDRESS  <= Bus2IP_Addr;
165  OMA_DATA_O   <= Bus2IP_Data;
166  OMA_BE       <= Bus2IP_BE;
167
168  LEDES_DEBUG  <= leds;
169
170  — Everything on time.
171  INN: process(clock)
172  begin
173    if rising_edge(clock) then
174      — Default sending out 0x00 on the databus
175      — to the core.
176      IP2Bus_Data <= (others => '0');
177      IP2Bus_Ack <= '0';
178
179      case iState is
180      when IDLE =>
181        if readEnabled = '1' then
182          OMA_ACTIVE <= '1';
183          OMA_RW <= '1';
184          IP2Bus_Data <= OMA_DATA_I;
185          iState <= READING;
186        elsif writeEnabled = '1' then
187          OMA_ACTIVE <= '1';
188          OMA_RW <= '0';
189          iState <= WRITING;
190        end if;
191
192      when READING =>
193        — We will use more than 16 cycles.
194        IP2Bus_ToutSup <= '1';
195        IP2Bus_Data <= OMA_DATA_I;
196
197        — Check for doneness.
198        if OMA_VALID = '1' then
199          iState <= RCOMP;
200          OMA_ACTIVE <= '0';
201          IP2Bus_ToutSup <= '0';
202          IP2Bus_Ack <= '1';
203          leds <= OMA_DATA_I(0 to 15);
204        end if;
205
206      when WRITING =>
207        — We will use more than 16 cycles.
208        IP2Bus_ToutSup <= '1';
209        — Check for doneness.
210        if OMA_VALID = '1' then
211          iState <= WCOMP;
212          OMA_ACTIVE <= '0';

```

```
213         IP2Bus_ToutSup    <= '0';
214         IP2Bus_Ack        <= '1';
215     end if;
216
217     —         when RCOMP =>
218     —         when WCOMP =>
219     when others =>
220         OMA_ACTIVE <= '0';
221         IP2Bus_Ack <= '0';
222         iState <= IDLE;
223     end case;
224
225 end if;
226 end process;
227
228 IP2Bus_Error    <= '0';
229 IP2Bus_Retry    <= '0';
230
231 end IMP;
```

B.2 CPU Identifier

```

1  ---
2  --- cpu_identifier - entity/architecture pair
3  ---
4  ---
5  ---
6  --- *****
7  --- ** Copyright (c) 1995-2006 Xilinx, Inc. All rights reserved.
8  --- **
9  --- ** Xilinx, Inc.
10 --- ** XILINX IS PROVIDING THIS DESIGN, CODE, OR INFORMATION "AS IS"
11 --- ** AS A COURTESY TO YOU, SOLELY FOR USE IN DEVELOPING PROGRAMS AND
12 --- ** SOLUTIONS FOR XILINX DEVICES. BY PROVIDING THIS DESIGN, CODE,
13 --- ** OR INFORMATION AS ONE POSSIBLE IMPLEMENTATION OF THIS FEATURE,
14 --- ** APPLICATION OR STANDARD, XILINX IS MAKING NO REPRESENTATION
15 --- ** THAT THIS IMPLEMENTATION IS FREE FROM ANY CLAIMS OF INFRINGEMENT,
16 --- ** AND YOU ARE RESPONSIBLE FOR OBTAINING ANY RIGHTS YOU MAY REQUIRE
17 --- ** FOR YOUR IMPLEMENTATION. XILINX EXPRESSLY DISCLAIMS ANY
18 --- ** WARRANTY WHATSOEVER WITH RESPECT TO THE ADEQUACY OF THE
19 --- ** IMPLEMENTATION, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OR
20 --- ** REPRESENTATIONS THAT THIS IMPLEMENTATION IS FREE FROM CLAIMS OF
21 --- ** INFRINGEMENT, IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS
22 --- ** FOR A PARTICULAR PURPOSE.
23 --- **
24 --- *****
25 ---
26 --- Filename:          cpu_identifier
27 --- Version:           1.00.a

```

```

28 — Description:           Example FSL core (VHDL).
29 — Date:                Wed May 2 13:52:15 2007 (by Create and Import
   Peripheral Wizard)
30 — VHDL Standard:      VHDL'93
31 —

```

```

32 — Naming Conventions:
33 —   active low signals:           "*_n"
34 —   clock signals:              "clk", "clk_div#", "clk_#x"
35 —   reset signals:              "rst", "rst_n"
36 —   generics:                   "C_*"
37 —   user defined types:         "*_TYPE"
38 —   state machine next state:    "*_ns"
39 —   state machine current state: "*_cs"
40 —   combinatorial signals:      "*_com"
41 —   pipelined or register delay signals: "*_d#"
42 —   counter signals:           "*_cnt*"
43 —   clock enable signals:       "*_ce"
44 —   internal version of output port: "*_i"
45 —   device pins:                "*_pin"
46 —   ports:                     "- Names begin with Uppercase"
47 —   processes:                  "*_PROCESS"
48 —   component instantiations:   "<ENTITY>I_<#|FUNC>"
49 —

```

```

50
51 library ieee;
52 use IEEE.STD_LOGIC_1164.ALL;
53 use IEEE.STD_LOGIC_ARITH.ALL;
54 use IEEE.STD_LOGIC_UNSIGNED.ALL;
55 —

```

```

56 —
57 —
58 — Definition of Ports
59 — FSL_Clk           : Synchronous clock
60 — FSL_Rst           : System reset, should always come from FSL bus
61 — FSL_S_Clk        : Slave asynchronous clock
62 — FSL_S_Read       : Read signal, requiring next available input to be read
63 — FSL_S_Data       : Input data
64 — FSL_S_CONTROL    : Control Bit, indicating the input data are control word
65 — FSL_S_Exists     : Data Exist Bit, indicating data exist in the input FSL
   bus
66 — FSL_M_Clk        : Master asynchronous clock
67 — FSL_M_Write      : Write signal, enabling writing to output FSL bus
68 — FSL_M_Data       : Output data
69 — FSL_M_Control    : Control Bit, indicating the output data are control word
70 — FSL_M_Full       : Full Bit, indicating output FSL bus is full
71 —
72 —

```

```

73

```

```

74  --
75  -- Entity Section
76  --
77
78  entity cpu_identifier is
79    generic (
80      CPU_ID : integer := 32
81    );
82    port
83    (
84      -- DO NOT EDIT BELOW THIS LINE -----
85      -- Bus protocol ports, do not add or delete.
86      FSL_Clk : in  std_logic;
87      FSL_Rst : in  std_logic;
88      FSL_S_Clk : out std_logic;
89      FSL_S_Read : out std_logic;
90      FSL_S_Data : in  std_logic_vector(0 to 31);
91      FSL_S_Control : in  std_logic;
92      FSL_S_Exists : in  std_logic;
93      FSL_M_Clk : out std_logic;
94      FSL_M_Write : out std_logic;
95      FSL_M_Data : out std_logic_vector(0 to 31);
96      FSL_M_Control : out std_logic;
97      FSL_M_Full : in  std_logic
98      -- DO NOT EDIT ABOVE THIS LINE -----
99    );
100
101  attribute SIGIS : string;
102  attribute SIGIS of FSL_Clk : signal is "Clk";
103  attribute SIGIS of FSL_S_Clk : signal is "Clk";
104  attribute SIGIS of FSL_M_Clk : signal is "Clk";
105
106  end cpu_identifier;
107
108  --
109  -- Architecture Section
110  --
111
112  architecture EXAMPLE of cpu_identifier is
113
114    -- Total number of input data.
115    constant NUMBER_OF_INPUT_WORDS : natural := 1;
116
117    -- Total number of output data
118    constant NUMBER_OF_OUTPUT_WORDS : natural := 1;
119
120  begin
121    FSL_M_Write <= not FSL_M_Full;

```

```
122
123     FSL_M_Data <= CONV_STD_LOGIC_VECTOR(CPU_ID, FSL_M_Data'length);
124
125 end architecture EXAMPLE;
```

B.3 RAM Block

```

1  --
2  -- Company:
3  -- Engineer:
4  --
5  -- Create Date:      10:39:13 03/22/2007
6  -- Design Name:
7  -- Module Name:      myRam - Behavioral
8  -- Project Name:
9  -- Target Devices:
10 -- Tool versions:
11 -- Description:
12 --
13 -- Dependencies:
14 --
15 -- Revision:
16 -- Revision 0.01 - File Created
17 -- Additional Comments:
18 --
19 --
20 library IEEE;
21 use IEEE.STD_LOGIC_1164.ALL;
22 use IEEE.STD_LOGIC_ARITH.ALL;
23 use IEEE.STD_LOGIC_UNSIGNED.ALL;
24
25 -- Uncomment the following library declaration if instantiating
26 -- any Xilinx primitives in this code.
27 library UNISIM;
28 use UNISIM.VComponents.all;
29
30 entity myRam is
31   GENERIC(
32     DATA_WIDTH      : natural := 32;
33     ADDRLENGTH       : natural := 8
34   );
35   PORT(
36     clock            : IN std_logic := '0';
37     reset            : IN std_logic := '0';
38     enable           : IN std_logic := '0';
39     writeEnabled     : IN std_logic := '0';
40     address          : IN std_logic_vector(0 to ADDRLENGTH-1) := (others =>
41       '0');
42     dataIn           : IN std_logic_vector(0 to DATA_WIDTH-1) := (others =>
43       '0');
44     metaIn           : IN std_logic_vector(0 to DATA_WIDTH-1) := (others =>
45       '0');
46     dataOut          : OUT std_logic_vector(0 to DATA_WIDTH-1);
47     metaOut          : OUT std_logic_vector(0 to DATA_WIDTH-1)
48   );
49 end myRam;

```

```

48 architecture Behavioral of myRam is
49   constant numBlocks : integer := 2 ** (ADDRLENGTH - 7);
50
51   subtype Data is std_logic_vector(0 to DATA_WIDTH-1);
52
53   type DOut      is array(0 to numBlocks-1) of std_logic_vector(0 to 15);
54
55   signal enabled    : std_logic_vector(0 to numBlocks-1);
56   signal WEnabled  : std_logic_vector(0 to numBlocks-1);
57
58   signal DoutA     : DOut;
59   signal DoutB     : DOut;
60   signal MoutA     : DOut;
61   signal MoutB     : DOut;
62   signal addrA     : std_logic_vector(0 to 7);
63   signal addrB     : std_logic_vector(0 to 7);
64 begin
65
66   mem: for I in 0 to numBlocks - 1 generate
67     begin
68       MetaData: RAMB4_S16_S16
69       generic map (
70         INIT_00 => X"
71           0000000000000000000000000000000000000000000000000000000000000000",
72         INIT_01 => X"
73           0000000000000000000000000000000000000000000000000000000000000000",
74         INIT_02 => X"
75           0000000000000000000000000000000000000000000000000000000000000000",
76         INIT_03 => X"
77           0000000000000000000000000000000000000000000000000000000000000000",
78         INIT_04 => X"
79           0000000000000000000000000000000000000000000000000000000000000000",
80         INIT_05 => X"
81           0000000000000000000000000000000000000000000000000000000000000000",
82         INIT_06 => X"
83           0000000000000000000000000000000000000000000000000000000000000000",
84         INIT_07 => X"
85           0000000000000000000000000000000000000000000000000000000000000000",
86         INIT_08 => X"
87           0000000000000000000000000000000000000000000000000000000000000000",
88         INIT_09 => X"
89           0000000000000000000000000000000000000000000000000000000000000000",
90         INIT_0A => X"
91           0000000000000000000000000000000000000000000000000000000000000000",
92         INIT_0B => X"
93           0000000000000000000000000000000000000000000000000000000000000000",
94         INIT_0C => X"
95           0000000000000000000000000000000000000000000000000000000000000000",
96         INIT_0D => X"
97           0000000000000000000000000000000000000000000000000000000000000000",
98         INIT_0E => X"
99           0000000000000000000000000000000000000000000000000000000000000000",
100        INIT_0F => X"
101          0000000000000000000000000000000000000000000000000000000000000000")
102       port map (
103         DOA => MoutA(I),           — Port A 16-bit data output

```

```

88     DOB => MoutB(I),           — Port B 16-bit data output
89     ADDRA => addrA,           — Port A 8-bit address input
90     ADDR8 => addrB,           — Port B 8-bit address input
91     CLKA => clock,            — Port A clock input
92     CLKB => clock,            — Port B clock input
93     DIA => metaIn(0 to 15),    — Port A 16-bit data input
94     DIB => metaIn(16 to 31),   — Port B 16-bit data input
95     ENA => enabled(I),        — Port A RAM enable input
96     ENB => enabled(I),        — Port B RAM enable input
97     RSTA => reset,            — Port A Synchronous reset input
98     RSTB => reset,            — Port B Synchronous reset input
99     WEA => WEnabled(I),       — Port A RAM write enable input
100    WEB => WEnabled(I)        — Port B RAM write enable input
101  );
102  DataMem: RAMB4_S16_S16
103  generic map (
104    INIT_00 => X"
105      0000000000000000000000000000000000000000000000000000000000000000",
106    INIT_01 => X"
107      0000000000000000000000000000000000000000000000000000000000000000",
108    INIT_02 => X"
109      0000000000000000000000000000000000000000000000000000000000000000",
110    INIT_03 => X"
111      0000000000000000000000000000000000000000000000000000000000000000",
112    INIT_04 => X"
113      0000000000000000000000000000000000000000000000000000000000000000",
114    INIT_05 => X"
115      0000000000000000000000000000000000000000000000000000000000000000",
116    INIT_06 => X"
117      0000000000000000000000000000000000000000000000000000000000000000",
118    INIT_07 => X"
119      0000000000000000000000000000000000000000000000000000000000000000",
120    INIT_08 => X"
121      0000000000000000000000000000000000000000000000000000000000000000",
122    INIT_09 => X"
123      0000000000000000000000000000000000000000000000000000000000000000",
124    INIT_0A => X"
125      0000000000000000000000000000000000000000000000000000000000000000",
126    INIT_0B => X"
127      0000000000000000000000000000000000000000000000000000000000000000",
128    INIT_0C => X"
129      0000000000000000000000000000000000000000000000000000000000000000",
130    INIT_0D => X"
131      0000000000000000000000000000000000000000000000000000000000000000",
132    INIT_0E => X"
133      0000000000000000000000000000000000000000000000000000000000000000",
134    INIT_0F => X"
135      0000000000000000000000000000000000000000000000000000000000000000")
136  port map (
137    DOA => DoutA(I),           — Port A 16-bit data output
138    DOB => DoutB(I),           — Port B 16-bit data output
139    ADDRA => addrA,            — Port A 8-bit address input
140    ADDR8 => addrB,            — Port B 8-bit address input
141    CLKA => clock,            — Port A clock input
142    CLKB => clock,            — Port B clock input
143    DIA => dataIn(0 to 15),    — Port A 16-bit data input

```

```

128     DIB => dataIn(16 to 31),    — Port B 16-bit data input
129     ENA => enabled(I),        — Port A RAM enable input
130     ENB => enabled(I),        — Port B RAM enable input
131     RSTA => reset,            — Port A Synchronous reset input
132     RSTB => reset,            — Port B Synchronous reset input
133     WEA => WEnabled(I),        — Port A RAM write enable input
134     WEB => WEnabled(I)        — Port B RAM write enable input
135 );
136 end generate;
137
138 lines: for I in 0 to numBlocks -1 generate
139     WEnabled(I) <= writeEnabled and enabled(I);
140 end generate;
141
142 clocked: process( clock, enabled, writeEnabled, address, dataIn ) is
143 begin
144     addrA(0) <= '0';
145     addrB(0) <= '1';
146     addrA(1 to 7) <= address(0 to 6);
147     addrB(1 to 7) <= address(0 to 6);
148     — decodedAddr
149     if rising_edge(clock) then
150         enabled <= (others => '0');
151
152         if enable = '1' then
153             if numBlocks > 1 then
154                 enabled(conv_integer(address(7 to 7+(ADDR_LENGTH-8)))) <= '1';
155             else
156                 enabled(0) <= '1';
157             end if;
158         end if;
159     end if;
160
161     — Multiplex the data out.
162     if numBlocks > 1 then
163         metaOut(0 to 15) <= MoutA(conv_integer(address(7 to 7+(ADDR_LENGTH
164             -8))));
165         metaOut(16 to 31) <= MoutB(conv_integer(address(7 to 7+(
166             ADDR_LENGTH-8))));
167         dataOut(0 to 15) <= DoutA(conv_integer(address(7 to 7+(ADDR_LENGTH
168             -8))));
169         dataOut(16 to 31) <= DoutB(conv_integer(address(7 to 7+(
170             ADDR_LENGTH-8))));
171     else
172         metaOut(0 to 15) <= MoutA(0);
173         metaOut(16 to 31) <= MoutB(0);
174         dataOut(0 to 15) <= DoutA(0);
175         dataOut(16 to 31) <= DoutB(0);
176     end if;
177 end if;
178 end process;
179 end Behavioral;

```

B.4 Cache Block

```

1  --
2  --
3  -- Create Date:      12:02:44 03/29/2007
4  -- Design Name:
5  -- Module Name:     cacheBlock - Behavioral
6  -- Project Name:
7  -- Target Devices:
8  -- Tool versions:
9  -- Description:
10 --
11 -- Dependencies:
12 --
13 -- Revision:
14 -- Revision 0.01 - File Created
15 -- Additional Comments:
16 -- Important to notice that all data coming out of the microBlaze core
17 -- is bit reversed. Hence the index and tag of the address is switched
18 -- compared to the normal way of doing it.
19 --
20 --
21 library IEEE;
22 use IEEE.STD_LOGIC_1164.ALL;
23 use IEEE.STD_LOGIC_ARITH.ALL;
24 use IEEE.STD_LOGIC_UNSIGNED.ALL;
25
26 library UNISIM;
27 use UNISIM.VComponents.all;
28
29 entity cacheBlock is
30   generic(
31     DATA_WIDTH      : integer := 32;
32     ADDRESS_WIDTH    : integer := 32; --Tag length is independent of the
33     SET_ASSOCIATIVITY : integer := 2; -- be able to create different caches,
34     TAG_LENGTH       : integer := 25; -- just set associative ones.
35     ADDR_LENGTH      : integer := 7+1;
36
37     LRU_SIZE         : integer := 2
38   );
39
40 port (
41   CLOCK      : IN std_logic;
42   RESET      : IN std_logic;
43   IN_OMA_VALID  : OUT std_logic;
44   IN_OMA_DATA_I : OUT std_logic_vector(0 to 31);
45   IN_OMA_DATA_O : IN std_logic_vector(0 to 31);
46   IN_OMA_ADDRESS : IN std_logic_vector(0 to 31);
47   IN_OMA_BE     : IN std_logic_vector(0 to 3);
48   IN_OMA_RW     : IN std_logic;

```

```

49     IN_OMA_ACTIVE    : IN std_logic;
50
51     OUT_OMA_VALID    : IN std_logic;
52     OUT_OMA_DATA_I   : IN std_logic_vector(0 to 31);
53     OUT_OMA_DATA_O   : OUT std_logic_vector(0 to 31);
54     OUT_OMA_ADDRESS  : OUT std_logic_vector(0 to 31);
55     OUT_OMA_BE       : OUT std_logic_vector(0 to 3);
56     OUT_OMA_RW       : OUT std_logic;
57     OUT_OMA_ACTIVE   : OUT std_logic
58 );
59 end cacheBlock;
60
61 architecture Behavioral of cacheBlock is
62     type dlines is array (0 to SET_ASSOCIATIVITY - 1) of
63         std_logic_vector(0 to DATA_WIDTH - 1);
64
65     type STATE is ( IDLE, FIND_VICTIM, EVICT, GET_DATA, OVERRIDE, DONE );
66
67     constant tagNotFound : std_logic_vector(0 to SET_ASSOCIATIVITY - 1) := (
68         others => '0');
69
70     — Which parts shall we address, and which are left for tags.
71     constant setAddrBase : integer := ADDRESS_WIDTH - ADDR_LENGTH;
72     constant lruBase     : integer := 0; — Base addr of LRU;
73     constant lruSize     : integer := 1; — Log2(SET_ASSOCIATIVITY);
74     constant dirtyBase   : integer := lruBase + lruSize;
75     constant validBase   : integer := dirtyBase + 1;
76     constant tagBase     : integer := validBase + 1;
77     constant tagSize     : integer := setAddrBase;
78
79     — Control signals.
80     signal enabled       : std_logic_vector(0 to SET_ASSOCIATIVITY - 1);
81     signal writeEnabled  : std_logic_vector(0 to SET_ASSOCIATIVITY - 1);
82     signal tagFound      : std_logic_vector(0 to SET_ASSOCIATIVITY - 1);
83
84     — Cache state
85     signal readState     : STATE := IDLE;
86     signal writeState    : STATE := IDLE;
87
88     — Set up the signals to be multiplexed.
89     signal dout          : dlines := (others => (others => '0'));
90     signal mout          : dlines := (others => (others => '0'));
91     signal din           : dlines := (others => (others => '0'));
92     signal min           : dlines := (others => (others => '0'));
93
94     — register the signal to meet timing constraints.
95     signal dataIn        : std_logic_vector(0 to IN_OMA_DATA_O'length - 1);
96
97     — BlockRAM needs delay.
98     signal delay         : std_logic_vector(0 to 1) := "00";
99     signal inValid       : std_logic := '0';
100
101     impure function findVictim return integer is
102         variable minValue : integer := 90; — Change this if you have set
103             associative > 90
104         variable minPos   : integer := 0;

```

```

103     variable current : integer := 0;
104     begin
105     for I in 0 to SET_ASSOCIATIVITY - 1 loop
106         if mout(I)(validBase) = '0' then
107             return I;
108         else
109             current := conv_integer(mout(I)(lruBase to lruBase + lruSize - 1));
110             if current < minValue then
111                 minValue := current;
112                 minPos := I;
113             end if;
114         end if;
115     end loop;
116
117     return minPos;
118 end function;
119
120 function lruReduction(A: std_logic_vector) return integer is
121     variable x : integer ;
122     begin
123     x := conv_integer(A);
124     if x >= 0 then
125         return x;
126     else return 0;
127     end if;
128 end function;
129 begin
130     — Generate the storage blocks.
131     setBlock: for I in 0 to SET_ASSOCIATIVITY - 1 generate
132     begin
133         cacheBlock: entity work.myRam
134             generic map(
135                 DATA_WIDTH => DATA_WIDTH,
136                 ADDR_LENGTH => ADDR_LENGTH
137             )
138             port map(
139                 clock      => CLOCK,
140                 reset      => RESET,
141                 enable     => enabled(I),
142                 writeEnabled => writeEnabled(I),
143                 address    => IN_OMA_ADDRESS(IN_OMA_ADDRESS'length - ADDR_LENGTH to
144                     IN_OMA_ADDRESS'length - 1 ),
145                 dataIn     => din(I),
146                 metaIn     => min(I),
147                 dataOut    => dout(I),
148                 metaOut    => mout(I)
149             );
150         end generate;
151     — Check if we've found something.
152     — update the tagFound signal.
153     fnd: process(CLOCK, IN_OMA_ADDRESS, mout) is
154     begin
155     for I in 0 to SET_ASSOCIATIVITY - 1 loop
156         if (mout(I)(tagBase to tagBase + tagSize - 1) = IN_OMA_ADDRESS(0 to
157             IN_OMA_ADDRESS'length - ADDR_LENGTH - 1))

```

```

157     and mout(I)(validBase) = '1'
158   then
159     tagFound(I) <= '1';
160   else
161     tagFound(I) <= '0';
162   end if;
163 end loop;
164 end process;
165
166 — The main function.
167 main: process(CLOCK, IN_OMA_ACTIVE) is
168   variable victim : natural := 0;
169
170 — Function decl.
171 procedure updateLRU is
172 begin
173   for I in 0 to SET_ASSOCIATIVITY - 1 loop
174     — Copy back, rude and elementary. From the output to the input...
175     min(I) <= mout(I);
176
177     — Update only the LRU.
178     min(I)(lruBase to lruBase + lruSize - 1) <= CONV_STDLOGIC_VECTOR(
179       lruReduction(mout(I)(lruBase to lruBase + lruSize - 1)),
180       lruSize);
181   end loop;
182 end updateLRU;
183
184 procedure evictLine is
185 begin
186   — If victim dirty
187   — dump it to memory.
188   OUT_OMA_BE      <= (others => '1');
189   OUT_OMA_RW      <= '0';
190   OUT_OMA_ACTIVE  <= '1';
191   OUT_OMA_DATA_O  <= dout(victim);
192   — The address..
193
194   OUT_OMA_ADDRESS(IN_OMA_ADDRESS'length - ADDR_LENGTH to IN_OMA_ADDRESS'
195     length - 1) <=
196     IN_OMA_ADDRESS( IN_OMA_ADDRESS'length - ADDR_LENGTH to
197       IN_OMA_ADDRESS'length - 1);
198   — Get the tag from the metadata.
199   OUT_OMA_ADDRESS(0 to OUT_OMA_ADDRESS'length - ADDR_LENGTH - 1) <= mout
200     (victim)(tagBase to tagBase + tagSize - 1);
201 end evictLine;
202
203 procedure doWriteHit is
204 begin
205   updateLRU;
206   min(victim)(validBase) <= '1';
207   min(victim)(dirtyBase) <= '1';
208   min(victim)(lruBase to lruSize - 1) <= (others => '1');
209   min(victim)(tagBase to tagBase + tagSize - 1) <= mout(victim)(tagBase
210     to tagBase+tagSize-1);
211   din(victim)      <= dataIn;
212   writeEnabled(victim) <= '1';

```

```

209
210     inValid <= '1';
211 end doWriteHit;
212
213 procedure doReadHit is
214 begin
215     updateLRU;
216     for I in 0 to SET_ASSOCIATIVITY -1 loop
217         if ( tagFound(I) = '1' ) then
218             IN_OMA_DATA_I <= dout(I);
219             exit;
220         end if;
221     end loop;
222
223     inValid <= '1';
224 end doReadHit;
225
226 procedure doReadMiss is
227 begin
228     — 1) Find victim
229     — 2) Evict
230     — 3) Get data from memory
231     — 4) Override
232     case readState is
233     when IDLE =>
234         readState <= FIND_VICTIM;
235     when FIND_VICTIM =>
236         victim := findVictim;
237         readState <= EVICT;
238     when EVICT =>
239         if OUT_OMA_VALID = '1' then
240             OUT_OMA_ADDRESS <= IN_OMA_ADDRESS;
241             OUT_OMA_RW <= '1';
242             OUT_OMA_ACTIVE <= '1';
243             — Send out a new request.
244             readState <= GET_DATA;
245         else
246             if mout(victim)(dirtyBase) = '1' and
247                mout(victim)(validBase) = '1' then
248
249                 — If it's dirty evict it.
250                 evictLine;
251             else
252                 OUT_OMA_ADDRESS <= IN_OMA_ADDRESS;
253                 OUT_OMA_RW <= '1';
254                 OUT_OMA_ACTIVE <= '1';
255                 readState <= GET_DATA;
256             end if;
257         end if;
258
259     when GET_DATA =>
260         OUT_OMA_ACTIVE <= '1';
261         if ( OUT_OMA_VALID = '1' ) then
262
263             — Update metadata.
264             min(victim)(validBase) <= '1';

```

```

265     min(victim)(dirtyBase)    <= '0';
266     min(victim)(lruBase to lruSize -1) <= (others => '1');
267
268     min(victim)(tagBase to tagBase + tagSize -1) <= IN_OMA_ADDRESS
      (0 to IN_OMA_ADDRESS'length - ADDR_LENGTH -1);
269     din(victim)                <= OUT_OMA_DATA_I;
270
271     — Send stuff out.
272     IN_OMA_DATA_I    <= OUT_OMA_DATA_I;
273
274     — writeEnabled asserted for two cycles.
275     readState <= DONE;
276     writeEnabled(victim)    <= '1';
277     end if;
278     when DONE =>
279         inValid <= '1';
280         writeEnabled(victim)    <= '1';
281         readState <= IDLE;
282
283     when others =>
284         readState <= IDLE;
285     end case;
286 end doReadMiss;
287
288
289 procedure doWriteMiss is
290 begin
291     — 1) Find victim.
292     — 2) Evict. E.g., die bastard cache line
293     — 3) Override.
294
295     case writeState is
296     when IDLE =>
297         writeState <= FIND_VICTIM;
298     when FIND_VICTIM =>
299         victim := findVictim;
300         writeState <= EVICT;
301     when EVICT =>
302         if OUT_OMA_VALID = '1' then
303             writeState <= OVERRIDE;
304         else
305             if mout(victim)(dirtyBase) = '1' and
306                mout(victim)(validBase) = '1' then
307                 — If it's dirty and valid evict it.
308                 evictLine;
309             else
310                 writeState <= OVERRIDE;
311             end if;
312         end if;
313     when OVERRIDE =>
314         min(victim)(validBase) <= '1';
315         min(victim)(dirtyBase) <= '1';
316         min(victim)(lruBase to lruSize -1) <= (others => '1');
317         min(victim)(tagBase to tagBase + tagSize -1) <= IN_OMA_ADDRESS(0
      to IN_OMA_ADDRESS'length - ADDR_LENGTH -1);
318         din(victim)                <= dataIn;

```

```

319         writeState <= DONE;
320
321         -- writeEnabled asserted for two clock cycles.
322         writeEnabled(victim) <= '1';
323         inValid <= '1';
324     when DONE =>
325         writeEnabled(victim) <= '1';
326         writeState <= IDLE;
327
328     when others =>
329         writeState <= IDLE;
330 end case;
331
332 end doWriteMiss;
333 -- decl. ends.
334
335 begin
336     dataIn <= IN_OMA_DATA_O;
337     IN_OMA_VALID <= inValid;
338
339     if rising_edge(CLOCK) then
340         inValid <= '0';
341         OUT_OMA_ACTIVE <= '0';
342
343         writeEnabled <= (others => '0');
344         enabled <= (others => '0');
345
346         if ( inValid = '0' ) then
347             enabled <= (others => '1');
348         else
349             enabled <= (others => '0');
350         end if;
351
352         if IN_OMA_ACTIVE = '1' and inValid = '0' then
353
354             if IN_OMA_ADDRESS = x"7FFFFFFF" and IN_OMA_RW = '0' then
355                 -- DO NOT cache the termination signal.
356                 -- We're done anyways, so lets do this the
357                 -- dirty way.
358                 OUT_OMA_ADDRESS <= IN_OMA_ADDRESS;
359                 OUT_OMA_RW <= '0';
360                 OUT_OMA_ACTIVE <= '1';
361             else
362                 -- We need to let the signals propagate from the block RAM to us.
363                 if delay = "11" or ( writeState /= IDLE or readState /= IDLE )
364                     then
365                     if IN_OMA_RW = '1' then
366                         if tagFound /= tagNotFound then
367                             -- Ladies and Gentlemen, we have hit.
368                             doReadHit;
369                         else
370                             doReadMiss;
371                         end if;
372                     elsif IN_OMA_RW = '0' then
373                         if tagFound /= tagNotFound then
374                             doWriteHit;

```

```
374         else
375             doWriteMiss;
376         end if;
377     end if;
378
379         delay <= "00";
380     elsif delay = "00" then
381         — Block RAM need time to propagate through...
382         delay <= "01";
383     else
384         delay <= "11";
385     end if;
386
387     end if;
388 end if;
389 end if;
390 end process;
391
392 end Behavioral;
```

B.5 Arbiter

```

1  --
2  -- Company:
3  -- Engineer:
4  --
5  -- Create Date:      10:28:34 04/13/2007
6  -- Design Name:
7  -- Module Name:      arbiter - Behavioral
8  -- Project Name:
9  -- Target Devices:
10 -- Tool versions:
11 -- Description:
12 --
13 -- Dependencies:
14 --
15 -- Revision:
16 -- Revision 0.01 - File Created
17 -- Additional Comments:
18 --
19 --
20 library IEEE;
21 use IEEE.STD_LOGIC_1164.ALL;
22 use IEEE.STD_LOGIC_ARITH.ALL;
23 use IEEE.STD_LOGIC_UNSIGNED.ALL;
24
25 --- Uncomment the following library declaration if instantiating
26 --- any Xilinx primitives in this code.
27 ---library UNISIM;
28 ---use UNISIM.VComponents.all;
29
30 library WORK;
31 use WORK.coreStuff.ALL;
32
33 entity arbiter is
34   port (
35     CLOCK      : IN std_logic;
36     RESET      : IN std_logic;
37
38     IN_OMA_VALID   : OUT coreBit;
39     IN_OMA_DATA_I  : OUT coreVector;
40     IN_OMA_DATA_O  : IN coreVector;
41     IN_OMA_ADDRESS : IN coreVector;
42     IN_OMA_BE      : IN coreBe;
43     IN_OMA_RW      : IN coreBit;
44     IN_OMA_ACTIVE  : IN coreBit;
45
46     OUT_OMA_VALID  : IN std_logic;
47     OUT_OMA_DATA_I : IN std_logic_vector(0 to 31);
48     OUT_OMA_DATA_O : OUT std_logic_vector(0 to 31);
49     OUT_OMA_ADDRESS : OUT std_logic_vector(0 to 31);
50     OUT_OMA_BE     : OUT std_logic_vector(0 to 3);

```

```

51     OUT_OMARW      : OUT std_logic;
52     OUT_OMA_ACTIVE : OUT std_logic
53   );
54 end arbiter;
55
56 architecture Behavioral of arbiter is
57   signal active      : std_logic := '0';
58   shared variable currentActive : natural := 0;
59 begin
60
61   f: process is
62     begin
63       IN_OMA_VALID <= (others => '0');
64       IN_OMA_VALID(currentActive) <= OUT_OMA_VALID;
65       IN_OMA_DATA_I(currentActive) <= OUT_OMA_DATA_I;
66       OUT_OMA_DATA_O <= IN_OMA_DATA_O(currentActive);
67       OUT_OMA_ADDRESS <= IN_OMA_ADDRESS(currentActive);
68       OUT_OMA_BE <= IN_OMA_BE(currentActive);
69       OUT_OMARW <= IN_OMARW(currentActive);
70       OUT_OMA_ACTIVE <= IN_OMA_ACTIVE(currentActive);
71     end process;
72
73   switcher: process(CLOCK, IN_OMA_ACTIVE, OUT_OMA_VALID) is
74     begin
75       if rising_edge(CLOCK) then
76         — If someone already are sending and receiving
77         — data, we'll wait for the ack to get back..
78         if active = '1' then
79           — Set no active on VALID data.
80           if OUT_OMA_VALID = '1' then
81             active <= '0';
82           end if;
83
84
85         else
86           for I in 0 to numCores - 1 loop
87             if IN_OMA_ACTIVE((currentActive + I) mod numCores) = '1' then
88               — Find the active one.
89               currentActive := (currentActive + I) mod numCores;
90               active <= '1';
91             exit;
92             end if;
93           end loop;
94         end if;
95       end if;
96     end process;
97
98 end Behavioral;

```

B.6 PCI COM

```

1  --
2  -- Company:
3  -- Engineer:
4  --
5  -- Create Date:      17:14:16 11/10/2006
6  -- Design Name:
7  -- Module Name:     com - Behavioral
8  -- Project Name:
9  -- Target Devices:
10 -- Tool versions:
11 -- Description:
12 --
13 -- Dependencies:
14 --
15 -- Revision:
16 -- Revision 0.01 - File Created
17 -- Additional Comments:
18 --
19 --
20 library IEEE;
21 use IEEE.STD_LOGIC_1164.ALL;
22 use IEEE.STD_LOGIC_ARITH.ALL;
23 use IEEE.STD_LOGIC_UNSIGNED.ALL;
24
25 --- Uncomment the following library declaration if instantiating
26 --- any Xilinx primitives in this code.
27 ---library UNISIM;
28 ---use UNISIM.VComponents.all;
29
30 entity pciCom is
31   Port (
32     --- PCI FPGA Spesific.
33     CLOCK      : in std_logic;
34     RESET      : in std_logic;
35
36     PCLADDS    : in STD_LOGIC; --- ADIO is DATA when high, else Address.
37     PCLEMPY   : in STD_LOGIC; --- Is empty
38     PCLBUSY   : in STD_LOGIC; --- Can't write quite yet.
39     PCLRW     : out STD_LOGIC; --- Will write to the PCI FPGA on high.
40     PCLRENWEN : out STD_LOGIC; --- High disables communication.
41     PCLADIO   : inout STD_LOGIC_VECTOR (0 to 31);
42
43     --- LED DEBUG..
44     LED_DEBUG : OUT std_logic_vector(0 to 15);
45
46     --- Internal chat with the memory bus
47     OMA_VALID : OUT std_logic;
48     OMA_DATA_I : OUT std_logic_vector(0 to 31);
49     OMA_DATA_O : IN std_logic_vector(0 to 31);
50     OMA_ADDRESS : IN std_logic_vector(0 to 31);

```

```

51     OMA_BE      : IN std_logic_vector(0 to 3);
52     OMA_RW      : IN std_logic;
53     OMA_ACTIVE  : IN std_logic
54     );
55 end pciCom;
56
57 architecture Behavioral of pciCom is
58     — State machine variables.
59     type States is ( IDLE,
60                     READ_START, READ_ADDR, READ_DATA, READ_IDLE,
61                     WRITE_START, WRITE_ADDR, WRITE_DATA,
62                     COMPLETE
63                 );
64
65     — constant signals to write out to the databus.
66     constant writeSignal : std_logic_vector(31 downto 0) := (1 => '1', others
67     => '0');
68     constant readSignal  : std_logic_vector(31 downto 0) := (0 => '1', others
69     => '0');
70
71     — state signals
72     constant sIDLE      : std_logic_vector(0 to 15) := (0 => '0', others => '1')
73     ;
74     constant sREAD_START : std_logic_vector(0 to 15) := (1 => '0', others =>
75     '1');
76     constant sREAD_ADDR  : std_logic_vector(0 to 15) := (2 => '0', others =>
77     '1');
78     constant sREAD_DATA   : std_logic_vector(0 to 15) := (3 => '0', others =>
79     '1');
80     constant sREAD_IDLE   : std_logic_vector(0 to 15) := (4 => '0', others =>
81     '1');
82     constant sWRITE_START : std_logic_vector(0 to 15) := (5 => '0', others =>
83     '1');
84     constant sWRITE_ADDR  : std_logic_vector(0 to 15) := (6 => '0', others =>
85     '1');
86     constant sWRITE_DATA  : std_logic_vector(0 to 15) := (7 => '0', others =>
87     '1');
88     constant sCOMPLETE   : std_logic_vector(0 to 15) := (8 => '0', others =>
89     '1');
90
91     signal state : States;
92     signal data_o : std_logic_vector(0 to 31) := (others => '1');
93     signal data_i : std_logic_vector(0 to 31) := (others => '1');
94     signal address : std_logic_vector(0 to 31) := (others => '1');
95
96     signal clockCount : std_logic_vector(0 to 31) := (others => '0');
97
98 begin
99     countClock: process(CLOCK, RESET, state) is
100     begin
101         if rising_edge(CLOCK) then
102             if RESET = '1' or state = WRITE_START or state = READ_START then
103                 clockCount <= (others => '0');
104             else
105                 clockCount <= clockCount + 1;

```

```

96     end if;
97     end if;
98     end process;
99
100    fnord: process(CLOCK, RESET) is
101    begin
102        if RESET = '1' then
103            PCLRENWEN <= '1' ;
104            state <= IDLE;
105
106        elsif rising_edge(CLOCK) then
107            -- PCLRENWEN <= '0'; -- Default to enable com.
108            -- OMA_VALID <= '0'; -- NOT valid data per default.
109
110            if PCLEMPY = '0' and state = READ.DATA then
111                -- Gather data from input.
112                data_i <= PCLADIO;
113                -- Debug fnord.
114                LED_DEBUG <= PCLADIO(0 to 15);
115                OMA_VALID <= '1';
116                state <= COMPLETE;
117
118            elsif PCLBUSY = '0' then
119                -- else
120                OMA_VALID <= '0'; -- NOT valid data per default.
121                case state is
122                    when IDLE =>
123                        PCLRENWEN <= '1'; -- We're not enabling com when idle.
124                        PCLADIO <= ( others => 'Z' );
125
126                        -- Write on low!
127                        if OMARW = '0' and OMA_ACTIVE='1' then
128                            state <= WRITE.START;
129                            PCLRENWEN <= '0';
130                        elsif OMARW = '1' and OMA_ACTIVE = '1' then
131                            state <= READ.START;
132                            PCLRENWEN <= '0';
133                        else
134                            state <= IDLE;
135                        end if;
136
137                        -- Writing to the bus.
138                        when WRITE.START =>
139                            PCLADIO <= writeSignal;
140                            PCLADIO(16 to 19) <= OMA_BE;
141                            state <= WRITE.ADDR;
142                        when WRITE.ADDR =>
143                            PCLADIO <= address;
144                            state <= WRITE.DATA;
145                        when WRITE.DATA =>
146                            PCLADIO <= data_o;
147                            -- Debug fnord
148                            OMA_VALID <= '1';
149                            state <= COMPLETE;
150
151                        -- Reading from the bus.

```

```

152     when READ.START =>
153         PCLADIO <= readSignal;
154         state <= READ_ADDR;
155     when READ.ADDR =>
156         PCLADIO <= address;
157         state <= READ_IDLE;
158     when READ.IDLE =>
159         PCLADIO <= (others => 'Z');
160         state <= READ_DATA;
161
162
163     -- The requested transaction is done!
164     when COMPLETE =>
165         PCLADIO <= data_i;
166         state <= IDLE;
167     when others =>
168         null;
169     end case;
170 end if;
171 end if;
172
173 end process;
174
175 -- -- show state on led
176 -- lSate: process(state) is
177 -- begin
178 --     case state is
179 --         when IDLE =>
180 --             LED_DEBUG <= sIDLE;
181 --         when READ.START =>
182 --             LED_DEBUG <= sREAD_START;
183 --         when READ.ADDR =>
184 --             LED_DEBUG <= sREAD_ADDR ;
185 --         when READ.DATA =>
186 --             LED_DEBUG <= sREAD_DATA;
187 --         when READ.IDLE =>
188 --             LED_DEBUG <= sREAD_IDLE ;
189 --         when WRITE.START =>
190 --             LED_DEBUG <= sWRITE_START;
191 --         when WRITE.ADDR=>
192 --             LED_DEBUG <= sWRITE_ADDR;
193 --         when WRITE.DATA=>
194 --             LED_DEBUG <= sWRITE_DATA;
195 --         when COMPLETE=>
196 --             LED_DEBUG <= sCOMPLETE;
197 --     end case;
198 -- end process;
199
200 comb: process(state, OMA_DATA_O, OMA_ADDRESS, RESET, data_i) is
201     begin
202
203
204     if RESET = '1' then
205         data_o <= (others => 'Z');
206         address <= (others => 'Z');
207     else

```

```
208     PCLRW    <= '1'; -- We're writing as default
209     --OMA_VALID <= '0';
210
211     OMA_DATA_I <= data_i;
212     address <= OMA_ADDRESS;
213     data_o <= OMA_DATA_O;
214
215     case state is
216     when READ_DATA =>
217         PCLRW <= '0'; -- READ !!
218     when COMPLETE =>
219         -- OMA_VALID <= '1';
220         null;
221     when others =>
222         null;
223     end case;
224 end if;
225 end process;
226
227 end Behavioral;
```

B.7 Toplevel - mCache

```

1  —


---


2  — Create Date:      13:58:29 11/20/2006
3  — Design Name:
4  — Module Name:    toplevel - Behavioral
5  — Project Name:
6  — Target Devices:
7  — Tool versions:
8  — Description:
9  —
10 — Dependencies:
11 —
12 — Revision:
13 — Revision 0.01 - File Created
14 — Additional Comments:
15 —
16 —


---


17 library IEEE;
18 use IEEE.STD_LOGIC_1164.ALL;
19 use IEEE.STD_LOGIC_ARITH.ALL;
20 use IEEE.STD_LOGIC_UNSIGNED.ALL;
21
22 ——— Uncomment the following library declaration if instantiating
23 ——— any Xilinx primitives in this code.
24 library UNISIM;
25 use UNISIM.VComponents.all;
26
27 use work.coreStuff.all;
28
29 entity toplevel is
30   PORT(
31     sys_clk_pin      : IN std_logic;
32     reset            : IN std_logic;
33     fpga_0_LEDS_GPIO_d_out_pin : OUT std_logic_vector(0 to 15);
34     PCLADDS          : in  STD_LOGIC;
35     PCLEMPY          : in  STD_LOGIC;
36     PCLBUSY          : in  STD_LOGIC;
37     PCLRW            : out STD_LOGIC;
38     PCLRENWEN        : out STD_LOGIC;
39     PCLADIO          : inout STD_LOGIC_VECTOR (0 to 31));
40 end toplevel;
41
42 architecture Behavioral of toplevel is
43   signal OMA_VALID : std_logic;
44   signal OMA_DATA_I : std_logic_vector(0 to 31) := (others => '0');
45   signal OMA_DATA_O : std_logic_vector(0 to 31) := (others => '0');
46   signal OMA_ADDRESS : std_logic_vector(0 to 31) := x"DEADBEEF";
47   signal OMA_RW      : std_logic;
48   signal OMA_ACTIVE : std_logic;
49   signal OMA_BE     : std_logic_vector (0 to 3) := (others => '0');
50

```

```

51  signal CACHEVALID  : std_logic;
52  signal CACHE_DATA_I : std_logic_vector(0 to 31) := (others =>'0');
53  signal CACHE_DATA_O : std_logic_vector(0 to 31) := (others =>'0');
54  signal CACHE_ADDRESS : std_logic_vector(0 to 31) := ( others => '0');
55  signal CACHERW       : std_logic;
56  signal CACHE_ACTIVE : std_logic;
57  signal CACHE_BE     : std_logic_vector (0 to 3) := (others => '0');
58
59  signal CORE_VALID   : corebit;
60  signal CORE_DATA_I  : corevector := (others => (others =>'0'));
61  signal CORE_DATA_O  : corevector := (others => (others =>'0'));
62  signal CORE_ADDRESS : corevector := (others => (others =>'0'));
63  signal CORE_RW      : corebit;
64  signal CORE_ACTIVE  : corebit;
65  signal CORE_BE      : corebe;
66
67
68  signal sPCLADDS     : std_logic := '0';
69  signal sPCLEEMPTY   : std_logic := '0';
70  signal sPCLBUSY     : std_logic := '0';
71  signal sPCLRW       : std_logic := '0';
72  signal sPCLRENWEN   : std_logic := '1';
73
74  signal rst_inv      : std_logic;
75  signal rst_internal : std_logic;
76
77
78  COMPONENT core
79  PORT(
80    sys_clk_pin : IN std_logic;
81    sys_rst_pin : IN std_logic;
82    opb_external_memory_0_OMA_DATA_I_pin : IN std_logic_vector(0 to 31);
83    opb_external_memory_0_OMA_VALID_pin  : IN std_logic;
84    opb_external_memory_1_OMA_DATA_I_pin : IN std_logic_vector(0 to 31);
85    opb_external_memory_1_OMA_VALID_pin  : IN std_logic;
86    opb_external_memory_0_OMA_DATA_O_pin : OUT std_logic_vector(0 to 31);
87    opb_external_memory_0_OMA_ADDRESS_pin : OUT std_logic_vector(0 to 31);
88    opb_external_memory_0_OMA_RW_pin     : OUT std_logic;
89    opb_external_memory_0_OMA_ACTIVE_pin : OUT std_logic;
90    opb_external_memory_0_LEDS_DEBUG_pin : OUT std_logic_vector(0 to 15);
91    opb_external_memory_0_OMA_BE_pin     : OUT std_logic_vector(0 to 3);
92    opb_external_memory_1_OMA_DATA_O_pin : OUT std_logic_vector(0 to 31);
93    opb_external_memory_1_OMA_ADDRESS_pin : OUT std_logic_vector(0 to 31);
94    opb_external_memory_1_OMA_RW_pin     : OUT std_logic;
95    opb_external_memory_1_OMA_ACTIVE_pin : OUT std_logic;
96    opb_external_memory_1_LEDS_DEBUG_pin : OUT std_logic_vector(0 to 15);
97    opb_external_memory_1_OMA_BE_pin     : OUT std_logic_vector(0 to 3)
98  );
99  END COMPONENT;
100
101
102
103  COMPONENT pciCom
104  Port (
105    CLOCK      : IN  STD_LOGIC;
106    RESET      : IN  STD_LOGIC;

```

```

107     PCLADDS      : IN    STD_LOGIC;
108     PCLEEMPTY   : IN    STD_LOGIC;
109     PCLBUSY     : IN    STD_LOGIC;
110     PCLRW       : OUT   STD_LOGIC;
111     PCLRENWEN   : OUT   STD_LOGIC;
112     PCLADIO     : INOUT STD_LOGIC_VECTOR (0 to 31);
113 —   LED_DEBUG   : OUT   std_logic_vector(0 to 15);
114     OMA_VALID    : OUT   STD_LOGIC;
115     OMA_DATA_I   : OUT   STD_LOGIC_VECTOR(0 to 31);
116     OMA_BE       : IN    STD_LOGIC_VECTOR(0 to 3);
117     OMA_DATA_O   : IN    STD_LOGIC_VECTOR(0 to 31);
118     OMA_ADDRESS  : IN    STD_LOGIC_VECTOR(0 to 31);
119     OMA_RW       : IN    STD_LOGIC;
120     OMA_ACTIVE   : IN    STD_LOGIC
121   );
122   END COMPONENT;
123
124   begin
125
126     Inst_system: core PORT MAP(
127       opb_external_memory_0_LEDS_DEBUG_pin => fpga_0_LEDS_GPIO_d_out_pin ,
128       sys_clk_pin                          => sys_clk_pin ,
129       sys_rst_pin                          => rst_inv ,
130       opb_external_memory_0_OMA_DATA_O_pin => CORE_DATA_O (0) ,
131       opb_external_memory_0_OMA_DATA_I_pin => CORE_DATA_I (0) ,
132       opb_external_memory_0_OMA_VALID_pin  => CORE_VALID(0) ,
133       opb_external_memory_0_OMA_ADDRESS_pin => CORE_ADDRESS(0) ,
134       opb_external_memory_0_OMA_RW_pin     => CORE_RW (0) ,
135       opb_external_memory_0_OMA_ACTIVE_pin => CORE_ACTIVE(0) ,
136       opb_external_memory_0_OMA_BE_pin     => CORE_BE(0) ,
137       opb_external_memory_1_OMA_DATA_O_pin => CORE_DATA_O (1) ,
138       opb_external_memory_1_OMA_DATA_I_pin => CORE_DATA_I (1) ,
139       opb_external_memory_1_OMA_VALID_pin  => CORE_VALID(1) ,
140       opb_external_memory_1_OMA_ADDRESS_pin => CORE_ADDRESS(1) ,
141       opb_external_memory_1_OMA_RW_pin     => CORE_RW (1) ,
142       opb_external_memory_1_OMA_ACTIVE_pin => CORE_ACTIVE(1) ,
143       opb_external_memory_1_OMA_BE_pin     => CORE_BE(1)
144   );
145
146     arbiter: entity work.arbiter PORT MAP(
147       CLOCK      => sys_clk_pin ,
148       RESET      => rst_inv ,
149
150       IN_OMA_DATA_O   => CORE_DATA_O,
151       IN_OMA_DATA_I   => CORE_DATA_I,
152       IN_OMA_VALID    => CORE_VALID,
153       IN_OMA_ADDRESS  => CORE_ADDRESS,
154       IN_OMA_RW       => CORE_RW,
155       IN_OMA_ACTIVE   => CORE_ACTIVE,
156       IN_OMA_BE       => CORE_BE,
157
158       OUT_OMA_VALID   => CACHE_VALID,
159       OUT_OMA_DATA_I   => CACHE_DATA_I,
160       OUT_OMA_DATA_O   => CACHE_DATA_O,
161       OUT_OMA_ADDRESS  => CACHE_ADDRESS,
162       OUT_OMA_BE       => CACHE_BE ,

```

```

163     OUT_OMARW      => CACHERW,
164     OUT_OMA_ACTIVE => CACHE_ACTIVE
165
166 );
167
168 cache : entity work.cacheBlock PORT MAP (
169     CLOCK          => sys_clk_pin ,
170     RESET          => rst_inv ,
171     -- From the arbiter to Cache
172     IN_OMA_VALID   => CACHE_VALID,
173     IN_OMA_DATA_I  => CACHE_DATA_I,
174     IN_OMA_BE      => CACHE_BE,
175     IN_OMA_DATA_O  => CACHE_DATA_O,
176     IN_OMA_ADDRESS => CACHE_ADDRESS,
177     IN_OMARW       => CACHERW,
178     IN_OMA_ACTIVE  => CACHE_ACTIVE,
179
180     -- From the cache to comm.
181     OUT_OMA_VALID  => OMA_VALID,
182     OUT_OMA_DATA_I => OMA_DATA_I,
183     OUT_OMA_BE     => OMA_BE,
184     OUT_OMA_DATA_O => OMA_DATA_O,
185     OUT_OMA_ADDRESS => OMA_ADDRESS,
186     OUT_OMARW     => OMARW,
187     OUT_OMA_ACTIVE => OMA_ACTIVE
188 );
189
190 communication : pciCom PORT MAP(
191     CLOCK    => sys_clk_pin ,
192     RESET    => rst_inv ,
193     PCLADDS  => sPCLADDS,
194     PCLEMPY  => sPCLEMPY,
195     PCLBUSY  => sPCLBUSY,
196     PCLRW    => sPCLRW,
197     PCLRENWEN => sPCLRENWEN,
198     PCLADIO  => PCLADIO,
199
200     OMA_BE   => OMA_BE,
201     OMA_VALID => OMA_VALID,
202     OMA_DATA_I => OMA_DATA_I,
203     OMA_DATA_O => OMA_DATA_O,
204     OMA_ADDRESS => OMA_ADDRESS,
205     OMARW     => OMARW,
206     OMA_ACTIVE => OMA_ACTIVE
207 );
208
209 IBUF_inst : IBUF
210 generic map (
211     IBUF_DELAY_VALUE => "0",
212     IFD_DELAY_VALUE => "AUTO",
213     IOSTANDARD => "DEFAULT")
214 port map (
215     O => rst_internal ,
216     I => reset
217 );
218

```

```
219 f: process(PCLADDS, PCLEMPY, PCLBUSY, sPCLRW, sPCLRENWEN,  
    rst_internal) is  
220 begin  
221     sPCLADDS <= PCLADDS ;  
222     sPCLEMPY <= PCLEMPY;  
223     sPCLBUSY <= PCLBUSY;  
224     PCLRW <= sPCLRW;  
225     PCLRENWEN <= sPCLRENWEN ;  
226     rst_inv <= not rst_internal;  
227 end process;  
228  
229 end Behavioral;
```

B.8 Toplevel - mCore

```

1  --
2  -- Create Date:      13:58:29 11/20/2006
3  -- Design Name:
4  -- Module Name:     toplevel - Behavioral
5  -- Project Name:
6  -- Target Devices:
7  -- Tool versions:
8  -- Description:
9  --
10 -- Dependencies:
11 --
12 -- Revision:
13 -- Revision 0.01 - File Created
14 -- Additional Comments:
15 --
16 --

```

```

17 library IEEE;
18 use IEEE.STD_LOGIC_1164.ALL;
19 use IEEE.STD_LOGIC_ARITH.ALL;
20 use IEEE.STD_LOGIC_UNSIGNED.ALL;
21
22 --- Uncomment the following library declaration if instantiating
23 --- any Xilinx primitives in this code.
24 library UNISIM;
25 use UNISIM.VComponents.all;
26
27 use work.coreStuff.all;
28
29 entity toplevel is
30     PORT(
31         sys_clk_pin      : IN std_logic;
32         reset             : IN std_logic;
33         fpga_0_LEDS_GPIO_d_out_pin : OUT std_logic_vector(0 to 15);
34         PCLADDS          : in  STD_LOGIC;
35         PCLEMPY          : in  STD_LOGIC;
36         PCLBUSY          : in  STD_LOGIC;
37         PCLRW            : out  STD_LOGIC;
38         PCLRENWEN        : out  STD_LOGIC;
39         PCLADIO          : inout STD_LOGIC_VECTOR (0 to 31));
40 end toplevel;
41
42 architecture Behavioral of toplevel is
43     signal OMA_VALID : std_logic;
44     signal OMA_DATA_I : std_logic_vector(0 to 31) := (others => '0');
45     signal OMA_DATA_O : std_logic_vector(0 to 31) := (others => '0');
46     signal OMA_ADDRESS : std_logic_vector(0 to 31) := x"DEADBEEF";
47     signal OMA_RW      : std_logic;
48     signal OMA_ACTIVE : std_logic;
49     signal OMA_BE      : std_logic_vector (0 to 3) := (others => '0');
50

```

```

51  signal CORE_VALID    : corebit;
52  signal CORE_DATA_I  : corevector := (others => (others => '0'));
53  signal CORE_DATA_O  : corevector := (others => (others => '0'));
54  signal CORE_ADDRESS : corevector := (others => (others => '0'));
55  signal CORE_RW      : corebit;
56  signal CORE_ACTIVE  : corebit;
57  signal CORE_BE      : corebe;
58
59
60  signal sPCLADDS      : std_logic := '0';
61  signal sPCLEMPY     : std_logic := '0';
62  signal sPCLBUSY     : std_logic := '0';
63  signal sPCLRW       : std_logic := '0';
64  signal sPCLRENWEN   : std_logic := '1';
65
66  signal rst_inv      : std_logic;
67  signal rst_internal : std_logic;
68
69
70  COMPONENT core
71  PORT(
72    sys_clk_pin : IN std_logic;
73    sys_rst_pin : IN std_logic;
74    opb_external_memory_0_OMA_DATA_I_pin : IN std_logic_vector(0 to 31);
75    opb_external_memory_0_OMA_VALID_pin  : IN std_logic;
76    opb_external_memory_1_OMA_DATA_I_pin : IN std_logic_vector(0 to 31);
77    opb_external_memory_1_OMA_VALID_pin  : IN std_logic;
78    opb_external_memory_0_OMA_DATA_O_pin : OUT std_logic_vector(0 to 31);
79    opb_external_memory_0_OMA_ADDRESS_pin : OUT std_logic_vector(0 to 31);
80    opb_external_memory_0_OMA_RW_pin     : OUT std_logic;
81    opb_external_memory_0_OMA_ACTIVE_pin : OUT std_logic;
82    opb_external_memory_0_LEDS_DEBUG_pin : OUT std_logic_vector(0 to 15);
83    opb_external_memory_0_OMA_BE_pin     : OUT std_logic_vector(0 to 3);
84    opb_external_memory_1_OMA_DATA_O_pin : OUT std_logic_vector(0 to 31);
85    opb_external_memory_1_OMA_ADDRESS_pin : OUT std_logic_vector(0 to 31);
86    opb_external_memory_1_OMA_RW_pin     : OUT std_logic;
87    opb_external_memory_1_OMA_ACTIVE_pin : OUT std_logic;
88    opb_external_memory_1_LEDS_DEBUG_pin : OUT std_logic_vector(0 to 15);
89    opb_external_memory_1_OMA_BE_pin     : OUT std_logic_vector(0 to 3)
90  );
91  END COMPONENT;
92
93
94
95  COMPONENT pciCom
96  Port (
97    CLOCK      : IN  STD_LOGIC;
98    RESET      : IN  STD_LOGIC;
99    PCLADDS    : IN   STD_LOGIC;
100   PCLEMPY    : IN   STD_LOGIC;
101   PCLBUSY    : IN   STD_LOGIC;
102   PCLRW      : OUT  STD_LOGIC;
103   PCLRENWEN  : OUT  STD_LOGIC;
104   PCLADIO    : INOUT STD_LOGIC_VECTOR (0 to 31);
105   — LED_DEBUG : OUT std_logic_vector(0 to 15);
106   OMA_VALID  : OUT  STD_LOGIC;

```

```

107     OMA_DATA_I    : OUT  STD_LOGIC_VECTOR(0 to 31);
108     OMA_BE       : IN   STD_LOGIC_VECTOR(0 to 3);
109     OMA_DATA_O    : IN   STD_LOGIC_VECTOR(0 to 31);
110     OMA_ADDRESS   : IN   STD_LOGIC_VECTOR(0 to 31);
111     OMA_RW        : IN   STD_LOGIC;
112     OMA_ACTIVE    : IN   STD_LOGIC
113 );
114 END COMPONENT;
115
116 begin
117
118     Inst_system : core PORT MAP(
119         opb_external_memory_0_LEDS_DEBUG_pin => fpga_0_LEDS_GPIO_d_out_pin ,
120         sys_clk_pin                        => sys_clk_pin ,
121         sys_rst_pin                        => rst_inv ,
122         opb_external_memory_0_OMA_DATA_O_pin => CORE_DATA_O (0) ,
123         opb_external_memory_0_OMA_DATA_I_pin => CORE_DATA_I (0) ,
124         opb_external_memory_0_OMA_VALID_pin  => CORE_VALID(0) ,
125         opb_external_memory_0_OMA_ADDRESS_pin => CORE_ADDRESS(0) ,
126         opb_external_memory_0_OMA_RW_pin     => CORE_RW (0) ,
127         opb_external_memory_0_OMA_ACTIVE_pin => CORE_ACTIVE(0) ,
128         opb_external_memory_0_OMA_BE_pin     => CORE_BE(0) ,
129         opb_external_memory_1_OMA_DATA_O_pin => CORE_DATA_O (1) ,
130         opb_external_memory_1_OMA_DATA_I_pin => CORE_DATA_I (1) ,
131         opb_external_memory_1_OMA_VALID_pin  => CORE_VALID(1) ,
132         opb_external_memory_1_OMA_ADDRESS_pin => CORE_ADDRESS(1) ,
133         opb_external_memory_1_OMA_RW_pin     => CORE_RW (1) ,
134         opb_external_memory_1_OMA_ACTIVE_pin => CORE_ACTIVE(1) ,
135         opb_external_memory_1_OMA_BE_pin     => CORE_BE(1)
136 );
137
138     arbiter : entity work.arbiter PORT MAP(
139         CLOCK      => sys_clk_pin ,
140         RESET      => rst_inv ,
141
142         IN_OMA_DATA_O  => CORE_DATA_O,
143         IN_OMA_DATA_I  => CORE_DATA_I,
144         IN_OMA_VALID   => CORE_VALID,
145         IN_OMA_ADDRESS => CORE_ADDRESS,
146         IN_OMA_RW      => CORE_RW,
147         IN_OMA_ACTIVE  => CORE_ACTIVE,
148         IN_OMA_BE      => CORE_BE,
149
150         OUT_OMA_VALID  => OMA_VALID,
151         OUT_OMA_DATA_I => OMA_DATA_I,
152         OUT_OMA_DATA_O => OMA_DATA_O,
153         OUT_OMA_ADDRESS => OMA_ADDRESS,
154         OUT_OMA_BE     => OMA_BE ,
155         OUT_OMA_RW     => OMA_RW,
156         OUT_OMA_ACTIVE => OMA_ACTIVE
157 );
158 );
159
160
161     communication : pciCom PORT MAP(
162         CLOCK      => sys_clk_pin ,

```

```

163     RESET    => rst_inv ,
164     PCLADDS => sPCLADDS,
165     PCLEMPY => sPCLEMPY,
166     PCLBUSY => sPCLBUSY,
167     PCLRW   => sPCLRW ,
168     PCLRENWEN => sPCLRENWEN,
169     PCLADIO => PCLADIO,
170
171     OMA_BE   => OMA_BE,
172     OMA_VALID => OMA_VALID,
173     OMA_DATA_I => OMA_DATA_I,
174     OMA_DATA_O => OMA_DATA_O,
175     OMA_ADDRESS => OMA_ADDRESS,
176     OMA_RW    => OMA_RW,
177     OMA_ACTIVE => OMA_ACTIVE
178     );
179
180     IBUF_inst : IBUF
181     generic map (
182         IBUF_DELAY_VALUE => "0",
183         IFD_DELAY_VALUE => "AUTO",
184         IOSTANDARD => "DEFAULT")
185     port map (
186         O => rst_internal ,
187         I => reset
188     );
189
190     f: process(PCLADDS, PCLEMPY, PCLBUSY, sPCLRW, sPCLRENWEN,
191               rst_internal) is
192     begin
193         sPCLADDS    <= PCLADDS ;
194         sPCLEMPY    <= PCLEMPY;
195         sPCLBUSY    <= PCLBUSY;
196         PCLRW       <= sPCLRW ;
197         PCLRENWEN   <= sPCLRENWEN ;
198         rst_inv     <= not rst_internal;
199     end process;
200 end Behavioral;

```

B.9 Toplevel - sCore

```

1  --


---


2  -- Create Date:      13:58:29 11/20/2006
3  -- Design Name:
4  -- Module Name:     toplevel - Behavioral
5  -- Project Name:
6  -- Target Devices:
7  -- Tool versions:
8  -- Description:
9  --
10 -- Dependencies:
11 --
12 -- Revision:
13 -- Revision 0.01 - File Created
14 -- Additional Comments:
15 --
16 --


---


17 library IEEE;
18 use IEEE.STD_LOGIC_1164.ALL;
19 use IEEE.STD_LOGIC_ARITH.ALL;
20 use IEEE.STD_LOGIC_UNSIGNED.ALL;
21
22 --- Uncomment the following library declaration if instantiating
23 --- any Xilinx primitives in this code.
24 library UNISIM;
25 use UNISIM.VComponents.all;
26
27 entity toplevel is
28     PORT(
29         sys_clk_pin : IN std_logic;
30         reset       : IN std_logic;
31         fpga_0_LEDS_GPIO_d_out_pin : OUT std_logic_vector(0 to 15);
32         PCLADDS     : in  STD_LOGIC;
33         PCLEMPY     : in  STD_LOGIC;
34         PCLBUSY     : in  STD_LOGIC;
35         PCLRW       : out STD_LOGIC;
36         PCLRENWEN   : out STD_LOGIC;
37         PCLADIO     : inout STD_LOGIC_VECTOR (0 to 31));
38 end toplevel;
39
40 architecture Behavioral of toplevel is
41     signal OMA_VALID : std_logic;
42     signal OMA_DATA_I : std_logic_vector(0 to 31) := (others =>'0');
43     signal OMA_DATA_O : std_logic_vector(0 to 31) := (others =>'0');
44     signal OMA_ADDRESS : std_logic_vector(0 to 31) := x"DEADBEEF";
45     signal OMA_RW      : std_logic;
46     signal OMA_ACTIVE : std_logic;
47     signal OMA_BE      : std_logic_vector (0 to 3) := (others => '0');
48
49
50     signal sPCLADDS : std_logic := '0';

```

```

51     signal sPCLEMPY    : std_logic := '0';
52     signal sPCLBUSY   : std_logic := '0';
53     signal sPCLRW     : std_logic := '0';
54     signal sPCLRENWEN : std_logic := '1';
55
56     signal rst_inv     : std_logic;
57     signal rst_internal : std_logic;
58
59
60     COMPONENT system
61     PORT(
62         opb_external_memory_0_LEDS_DEBUG_pin : OUT std_logic_vector(0 to 15);
63         sys_clk_pin : IN std_logic;
64         sys_rst_pin : IN std_logic;
65         opb_external_memory_0_OMA_DATA_I_pin : IN std_logic_vector(0 to 31);
66         opb_external_memory_0_OMA_VALID_pin : IN std_logic;
67         opb_external_memory_0_OMA_DATA_O_pin : OUT std_logic_vector(0 to 31);
68         opb_external_memory_0_OMA_ADDRESS_pin : OUT std_logic_vector(0 to 31);
69         opb_external_memory_0_OMA_RW_pin : OUT std_logic;
70         opb_external_memory_0_OMA_BE_pin : OUT std_logic_vector(0 to 3);
71         opb_external_memory_0_OMA_ACTIVE_pin : OUT std_logic
72     );
73     END COMPONENT;
74
75     COMPONENT pciCom
76     Port (
77         CLOCK      : in std_logic;
78         RESET      : in std_logic;
79         PCLADDS    : in STD_LOGIC;
80         PCLEMPY    : in STD_LOGIC;
81         PCLBUSY    : in STD_LOGIC;
82         PCLRW      : out STD_LOGIC;
83         PCLRENWEN  : out STD_LOGIC;
84         PCLADIO    : inout STD_LOGIC_VECTOR (0 to 31);
85 —     LED_DEBUG : OUT std_logic_vector(0 to 15);
86         OMA_VALID  : OUT std_logic;
87         OMA_DATA_I : OUT std_logic_vector(0 to 31);
88         OMA_BE     : IN std_logic_vector(0 to 3);
89         OMA_DATA_O : IN std_logic_vector(0 to 31);
90         OMA_ADDRESS : IN std_logic_vector(0 to 31);
91         OMARW     : IN std_logic;
92         OMA_ACTIVE : IN std_logic
93     );
94     END COMPONENT;
95
96     begin
97
98     Inst_system: system PORT MAP(
99         opb_external_memory_0_LEDS_DEBUG_pin => fpga_0_LEDS_GPIO_d_out_pin ,
100        sys_clk_pin           => sys_clk_pin ,
101        sys_rst_pin           => rst_inv ,
102        opb_external_memory_0_OMA_DATA_O_pin => OMA_DATA_O ,
103        opb_external_memory_0_OMA_DATA_I_pin => OMA_DATA_I ,
104        opb_external_memory_0_OMA_VALID_pin  => OMA_VALID,
105        opb_external_memory_0_OMA_ADDRESS_pin => OMA_ADDRESS,
106        opb_external_memory_0_OMA_RW_pin    => OMARW ,

```

```

107     opb_external_memory_0_OMA_ACTIVE_pin => OMA_ACTIVE,
108     opb_external_memory_0_OMA_BE_pin    => OMA_BE
109 );
110
111 communication: pciCom PORT MAP(
112     CLOCK      => sys_clk_pin ,
113     RESET      => rst_inv ,
114     PCLADDS    => sPCLADDS,
115     PCLEMPY    => sPCLEMPY,
116     PCLBUSY    => sPCLBUSY,
117     PCLRW      => sPCLRW ,
118     PCLRENWEN => sPCLRENWEN,
119     PCLADIO    => PCLADIO ,
120     OMA_BE     => OMA_BE,
121     OMA_VALID  => OMA_VALID,
122     OMA_DATA_I => OMA_DATA_I,
123     OMA_DATA_O => OMA_DATA_O,
124     OMA_ADDRESS => OMA_ADDRESS,
125     OMA_RW     => OMA_RW,
126     OMA_ACTIVE => OMA_ACTIVE
127 );
128
129 IBUF_inst : IBUF
130     generic map (
131         IBUF_DELAY_VALUE => "0" ,
132         IFD_DELAY_VALUE => "AUTO" ,
133         IO_STANDARD => "DEFAULT")
134     port map (
135         O => rst_internal ,
136         I => reset
137     );
138
139 f: process(PCLADDS, PCLEMPY, PCLBUSY, sPCLRW, sPCLRENWEN,
140           rst_internal) is
141     begin
142         sPCLADDS <= PCLADDS ;
143         sPCLEMPY <= PCLEMPY;
144         sPCLBUSY <= PCLBUSY;
145         PCLRW <= sPCLRW;
146         PCLRENWEN <= sPCLRENWEN ;
147         rst_inv <= not rst_internal;
148     end process;
149 end Behavioral;

```


Appendix C

Software

C.1 Controller.c

```
1 #include <stdio.h>
2 #include <stdlib.h>
3 #include <signal.h>
4
5 #include "controller.h"
6 #include "cpu.h"
7 #include "memory.h"
8 #include "util.h"
9
10
11 // #define DEBUG
12
13 static Memory *m;
14 static STATE state;
15
16 int run = 1;
17
18 unsigned long long int memRead = 0, memWrite = 0;
19
20 unsigned char byteEnabled = 0;
21
22 __inline__ unsigned long long int rdtsc() {
23     unsigned long long int x;
24     __asm__ volatile (".byte_0x0f, _0x31" : "=A" (x));
25     return x;
26 }
27
28
29 void termi(int code){
30     printf("\nNow exiting with code: %d\n", code);
31
32 #ifdef DEBUG
33     printf("Want to display a memory map?\n");
34     if ( getchar() == 'y' ){
```

```

35     printf("Memory_Map\n");
36     mem_print(m);
37 }
38 #endif
39
40     printf("\n");
41     exit(code);
42 }
43
44 void trap_seg(int signal){
45     printf("Seg.fault\n");
46     termi(5);
47 }
48
49 int handle(unsigned int command){
50     unsigned int reversed = reverseInt(command);
51
52 #ifdef DEBUG
53     printf("Doing_state:_0x%08x,r:_0x%08x\n",command,reversed);
54 #endif
55
56     if( state == IDLE ) {
57 #ifdef DEBUG
58         printf("Reverse:_0x%08X\n", reversed);
59 #endif
60         if( reversed == READ.SIGNAL ){
61             state = R_ADDR;
62         } else if ( (reversed & 0xFF) == 0x02 ){
63             state = W_ADDR;
64             // Byte Enabled Hack.
65             byteEnabled = (reversed >> 4*3) & 0xF;
66 #ifdef DEBUG
67             printf("Byte_Lines:_0x%08X\n", byteEnabled);
68 #endif
69         }
70     } else if( WRITING <= state && state <= W_DATA){
71         state = handleWrite(command, state);
72     } else if( READING <= state && state <= R_DATA) {
73         state = handleRead(command, state);
74     } else if ( state == COMPLETE ){
75 #ifdef DEBUG
76         printf("Operation_took:_0x%08X_operations\n",reversed);
77 #endif
78         state = IDLE;
79     }
80
81 }
82
83 int handleRead(unsigned int command,const STATE state){
84     /* The FPGA tries to access a memory location */
85     unsigned int reversed = reverseInt(command);
86     unsigned int r;
87
88 #ifdef DEBUG
89     printf("Now_handling_read_to_memory:_0x%X\n",reversed);
90 #endif

```

```

91     int i = mem_read(reversed ,NULL,m);
92     r = reverseInt(i);
93
94 #ifndef DEBUG
95     printf("Sending value: 0x%08x, r=0x%08x\n", i, r);
96 #endif
97     writeWord32(i);
98     memRead++;
99
100    // IDLE STAGE HACK!
101    getWord32();
102
103    return COMPLETE;
104 }
105
106 int handleWrite(unsigned int command, const STATE state){
107     unsigned int reversed = reverseInt(command);
108     /* The FPGA tries to write to a memory location */
109
110     unsigned int i = 0;
111
112     switch( state ) {
113         case W_ADDR:
114             i = getWord32();
115 #ifndef DEBUG
116             printf("Storing to mem: 0x%08x, data: 0x%08x\n", reversed, i);
117 #endif
118             if ( reversed == 0x7FFFFFFF ) {
119                 printf("\nTerminate signal from CPU\n");
120                 run = 0 ;
121                 // termi(2);
122             }
123             memWrite++;
124
125             if ( run ) mem_write(reversed, i, byteEnabled, m);
126         default: return COMPLETE;
127     }
128 }
129 }
130
131 print_state() {
132     switch(state) {
133         case IDLE:
134             printf("State: IDLE\n");
135             break;
136         case WRITING:
137             printf("State: WRITING\n");
138             break;
139         case W_ADDR:
140             printf("State: W_ADDR\n");
141             break;
142         case W_DATA:
143             printf("State: W_DATA\n");
144             break;
145         case READING:
146             printf("State: READING\n");

```

```

147     break;
148     case R_ADDR:
149         printf("State:_R_ADDR\n");
150         break;
151     case R_DATA:
152         printf("State:_R_DATA\n");
153         break;
154     case COMPLETE:
155         printf("State:_COMPLETE\n");
156         break;
157     default: printf("State:_UNKNOWN\n");
158 }
159 }
160
161 int main(int argc, char **argv) {
162     MemBlock b;
163     int f = 0;
164
165     if( argc < 2 ) {
166         printf("Usage:_%s_<bitfile>_[application]\n", argv[0]);
167         exit(1);
168     }
169
170     // signal(SIGSEGV, trap_seg);
171
172     // Init the memory
173     m = mem_init();
174     int mm;
175     if ( argc > 2 ){
176         if ( (mm = mem_load(m, argv[2] ) < 0 ) ){
177             printf("Error_opening_file:_%s_,_error:_%d\n", argv[2], mm );
178             exit(1);
179         }
180     }
181
182     state = IDLE;
183
184     #ifdef DEBUG
185         printf("Now_opening_the_card\n");
186     #endif
187     f = openCard(1, argv[1]);
188     #ifdef DEBUG
189         printf("OpenCard_returned:_%d\n\n", f);
190     #endif
191
192     //===== Configuration done.
193     unsigned long long int start, end, runtime;
194     start = end = 0 ;
195
196     f = '\0';
197
198     start = rdtsc();
199     do {
200     #ifdef DEBUG
201         print_state();
202     #endif

```

```

203
204     switch (f) {
205         case 's':
206 #ifdef DEBUG
207             printf("Now_writing\n");
208 #endif
209             writeWord32(0xFF);
210             break;
211         case 'r':
212 #ifdef DEBUG
213             printf("Reset\n");
214 #endif
215             break;
216         default:
217             break;
218     }
219
220     handle(getWord32());
221 #ifdef DEBUG
222     printf("Press_q_to_quit ,_everything_else_to_continue\n");
223     getchar();
224 #endif
225     }while ( run );
226     end = rdtsc();
227
228     printf("Stats:\n");
229     printf("\tRunning_time:_%llu\n", end - start);
230     printf("\tTotal_mem_access:_%llu\n", memRead + memWrite);
231     printf("\tTotal_mem_reads:_%llu\n", memRead);
232     printf("\tTotal_mem_writes:_%llu\n", memWrite);
233
234     printf("Did_you_find_anything?\nNow_exiting\n");
235     f = closeCard();
236     printf("Close_card_returned:_%d\n", f);
237     printf("Now_printing_memory\n");
238
239     for (f=0; f<m->mem_used; f++){
240         b = m->data[f];
241         printf("Cell_%d_contains_%02x\n", b.addr, b.data);
242     }
243
244     termi(3);
245
246     return 1;
247 }

```

C.2 Controller.h

```
1 #ifndef _CONTROLLER_H
2 #define _CONTROLLER_H
3
4 #define READ_SIGNAL    0x0000001
5 #define WRITE_SIGNAL   0x0000002
6
7
8 typedef enum {
9     IDLE = 0,
10    WRITING,
11    W_ADDR,
12    W_DATA,
13    READING,
14    R_ADDR,
15    R_DATA,
16    COMPLETE
17 }STATE;
18
19 #endif
```

C.3 memory.h

```
1
2 #ifndef _MEMORY_H
3 #define _MEMORY_H
4
5 #define START_MEM 20
6 #define GROW_RATE 20
7
8 // #define MEM_SIZE 0x4FC000C
9 #define MEM_SIZE 0x5000000
10
11 #include "rb.h"
12
13 struct _MemBlock {
14     unsigned int addr;
15     unsigned char data;
16     struct _MemBlock *prev, *next;
17 };
18
19 typedef struct _MemBlock MemBlock;
20
21 typedef struct {
22     unsigned int max_addr;
23     unsigned int mem_used;
24     unsigned int mem_size;
25     MemBlock *data;
26     struct rb_table *table;
27     unsigned char *bytes;
28 }Memory;
29
30
31 Memory *mem_init();
32 int mem_load(Memory *m, char *filename);
33 unsigned int mem_read(unsigned int addr, int *error, Memory *m);
34 unsigned int mem_write(unsigned int addr, unsigned int data, unsigned char
    byteEnabled, Memory *m);
35 unsigned int mem_writeByte(unsigned int addr, unsigned char data, Memory *m)
    ;
36 int mem_clean(Memory *m);
37 int mem_print(Memory *m);
38 int mem_zero(Memory *m, unsigned int offset);
39
40 #endif
```

C.4 memory.c

```

1  /**
2  * Simple memory holder..
3  * Implementing a array,
4  * A hash isn't worth the implementation
5  * time, and I don't want to allocate 4Gb
6  * of data for a traditional array.
7  *
8  * Kenneth Oestby <kenneo@idi.ntnu.no>
9  */
10
11 #include <stdio.h>
12 #include <stdlib.h>
13 #include <string.h>
14 #include <unistd.h>
15
16 #include "memory.h"
17 #include "util.h"
18 #include "rb.h"
19
20 int mem_load(Memory *m, char *filename){
21     FILE *f = NULL;
22     int memLoc = 0;
23     unsigned int *data = 0;
24     data =(unsigned int*) malloc(sizeof(unsigned int));
25     *data = 0;
26
27     if ( !m || !filename ) return -1;
28
29     if( ! ( f= fopen(filename,"r") ) ) {
30         return -2;
31     }
32
33     while( fread((void*)data, sizeof(unsigned char),
34         1, f) ) {
35         mem_writeByte(memLoc++, reverseInt8(*data),m);
36     }
37
38     free(data);
39
40     return memLoc;
41 }
42
43 int mem_print(Memory *m){
44     if ( ! m || !m->data) return -1;
45
46     int i = 0;
47     int x ;
48     for( i = 0 ; i < MEMSIZE; i++ ) {
49         if ( !(i % 32) )
50             printf("\n_%03X_", i );
51
52         printf("%02hhX_", mem_read(i, &x, m) );
53     }
54 }

```

```

55
56 int mem_clean(Memory *m){
57     return 1;
58 }
59
60 /*
61  * Zero out the memory from the given offset.
62  */
63 int mem_zero(Memory *m, unsigned int offset){
64     if ( !m ) return -1;
65
66     bzero(m->data + offset , (m->mem_size - offset) * sizeof(MemBlock));
67     return 1;
68 }
69
70 Memory *mem_init(){
71     Memory *m = NULL;
72     m = (Memory*) malloc(sizeof(Memory));
73     m->data = NULL;
74     m->mem_size = 0;
75     m->mem_used = 0;
76     m->max_addr = 0;
77
78     /*
79      * Bite of some bits of memory..
80      *
81      * If you're going use the CPU,
82      * you're probably going to want
83      * some memory to go along in the
84      * first place. Malloc is _slow_!
85      * Default is 20.. 20 * 32 should be enough
86      * for everybody.
87      */
88
89     if ( ! ( m->data = (MemBlock*) malloc(sizeof(MemBlock) * STARTMEM) ) ) {
90         mem_clean(m);
91         return NULL;
92     }
93
94     // m->table = rb_create(compare_ints, NULL, NULL);
95
96     m->bytes = (unsigned char*) malloc(MEM_SIZE * sizeof(char));
97
98     m->mem_size = STARTMEM;
99     mem_zero(m, 0);
100    return m;
101 }
102
103 unsigned int mem_read(unsigned int addr, int *error, Memory *m){
104     if ( !m ){
105         if (error) *error = -1;
106     }/* else if( !m->mem_used ){
107         if (error) *error = -2;
108     } else if( addr > m->max_addr ) {
109         if (error) *error = -3;
110     }*/

```

```

111
112     if ( error && *error < 0 ) return 0;
113
114     /*
115     int i = 0, j = 0;
116     MemBlock *tmp = NULL;
117     unsigned int returnMe = 0;
118
119     for ( i = 0; i < m->mem_used; i++ ) {
120         if ( m->data[i].addr == addr ){
121             tmp = m->data + i;
122             do {
123                 printf("Fnord:%p\n",tmp);
124                 printf("Next:%p\n",tmp->next);
125                 returnMe |= tmp->data << j * 8;
126             }while( (tmp = tmp->next) && ++j < 4 ) ;
127         }
128     }
129
130     return returnMe;
131     */
132
133     return m->bytes[addr+3] << 24 | m->bytes[addr+2] << 16 | m->bytes[addr+1]
134         << 8 | m->bytes[addr];
135 // return m->bytes[addr];
136 }
137
138 int updatePointers(Memory *m, unsigned int addr){
139     if ( !m ) return -1;
140
141     // Slow implementation on write to help speed up read.
142     // Update the next and prev-pointer.
143     int i = 0 ;
144     for(i = 0; i < m->mem_used; i++){
145         if( m->data[i].addr == addr - 1 )
146             m->data[i].next = m->data + m->mem_used - 1;
147         else if( m->data[i].addr == addr + 1 )
148             m->data[i].prev = m->data + m->mem_used - 1;
149     }
150
151     return 1;
152 }
153
154 unsigned int mem_writeByte(unsigned int addr, unsigned char data, Memory *m)
155 {
156     if( !m ) return 0;
157     static int i = 0;
158
159     /*
160     // Go through the array of memory to find the memory block.
161     // Future implementation would need a hash or something
162     // faster..
163     for ( i = 0; i < m->mem_used; i++ ) {
164         if ( m->data[i].addr == addr ){
165             m->data[i].data = data;
166             return 1;
167         }
168     }
169     return 0;
170 }

```

```

165     }
166 }
167
168 if ( m->max_addr < addr )
169     m->max_addr = addr;
170
171 // If not found in the array.. Expand it..
172 if ( m->mem_used < m->mem_size ) {
173     m->data[m->mem_used].addr = addr;
174     m->data[m->mem_used++].data = data;
175 } else {
176     // We need to grow the actual data array.
177     m->mem_size += GROWRATE;
178     printf("Now growing to %d\n",m->mem_size);
179     m->data = (MemBlock*)realloc((void*)m->data,m->mem_size*sizeof(MemBlock)
180         );
181     m->data[m->mem_used].addr = addr;
182     m->data[m->mem_used++].data = data;
183 }
184
185 updatePointers(m, addr);
186 */
187
188 m->bytes[addr] = data;
189
190 return 1;
191 }
192
193 unsigned int mem_write(unsigned int addr, unsigned int data, unsigned char
194     byteEnabled, Memory *m){
195     int i = 0 ;
196     char d;
197     // Split up the data in based upon byteEnabled field!
198     for(i = 0; i<3; i++){
199         if ( byteEnabled & ( 1 << i ) ){
200             d = (data >> sizeof(char) * 4) & 0xFF;
201             mem_writeByte(addr+i, d, m);
202         }else break;
203     }
204
205     return 1;
206 }

```

C.5 cpu.c

```

1 #include "cpu.h"
2 #include "util.h"
3
4 #include <stdio.h>
5
6 int openCard(int cardNum, char *bitfile){
7     char error[1024];
8     DWORD errorNum;
9
10    cpu.locate = NULL;
11    cpu.card = NULL;
12    cpu.dma = NULL;
13    cpu.send = NULL;
14    cpu.recv = NULL;
15
16    cpu.locate = DIME_LocateCard(dIPCI,
17                                mbtALL,
18                                NULL,
19                                dldrDEFAULT,
20                                dlDEFAULT);
21
22    if( !cpu.locate ){
23        DIME_GetError(NULL,&errorNum, error);
24        printf("\nLocate_Error_#%d\n%s\n",errorNum, error);
25        closeCard();
26        return -1;
27    }
28
29    if(!( cpu.card = DIME_OpenCard(cpu.locate,
30                                  cardNum,dccOPEN_DEFAULT)))
31    {
32        DIME_GetError(NULL,&errorNum, error);
33        printf("\nCard_Error_#%d\n%s\n",errorNum, error);
34        closeCard();
35        return -2;
36    }
37
38    // Setup chat.
39    if( setupDMA() < 1 ) {
40        printf("DMA_failed\n");
41        closeCard();
42        return -3;
43    }
44
45    if( DIME_JTAGControl(cpu.card, djtagCONFIGSPEED,djtagMAXSPEED100) ) {
46        printf("JTAG_died\n");
47        DIME_GetError(NULL,&errorNum, error);
48        printf("\nError_#%d\n%s\n",errorNum, error);
49        closeCard();
50        return -4;
51    }
52
53    // Finally set the clock before configuring the card:
54    DIME_SetOscillatorFrequency(cpu.card, CLOCKNUM, CLOCK_FREQ, NULL);

```

```

55
56  if(configureCard(bitfile) < 1){
57
58      printf("Could_not_configure_FPGA\n");
59      closeCard();
60      return -5;
61  }
62
63  if( resetCard() < 1 ){
64      printf("Error_resetting_card.\n");
65      closeCard();
66      return -6;
67  };
68
69  return 1;
70 }
71
72 int setupDMA(){
73     char error[1024];
74     DWORD errorNum;
75
76     bzero(cpu.recvBuffer, sizeof(int)*BUFFER_SIZE);
77     bzero(cpu.sendBuffer, sizeof(int)*BUFFER_SIZE);
78
79     if( !cpu.locate || !cpu.card )
80         return -1;
81
82     if( !( cpu.send = DIME_LockMemory(cpu.card,
83                                     (DWORD*)cpu.sendBuffer,
84                                     sizeof(cpu.sendBuffer)))
85     {
86         printf("Could_not_setup_send_buffer\n");
87         return -2;
88     }
89
90     if( !( cpu.recv = DIME_LockMemory(cpu.card,
91                                     (DWORD*)cpu.recvBuffer,
92                                     sizeof(cpu.recvBuffer)))
93     {
94         if( cpu.send ) DIME_UnLockMemory(cpu.card, cpu.send);
95         printf("Could_not_setup_send_buffer\n");
96         return -3;
97     }
98
99     if( ! ( cpu.dma = DIME_DMAOpen(cpu.card, 1,0) ) ) {
100         if( cpu.send ) DIME_UnLockMemory(cpu.card, cpu.send);
101         if( cpu.recv ) DIME_UnLockMemory(cpu.card, cpu.recv);
102
103         printf("Error_opening_DMA_Channel_numero_uno\n");
104
105         return -4;
106     }
107
108     // DO NOT INCREASE THE LOCAL MEMORY ADDR AFTER USE!!
109     DIME_DMAControl(cpu.card, cpu.dma, ddmaLOCALNOINC, 0);
110

```

```

111     return 1;
112 }
113
114 int configureCard(char *bitfile){
115     char error[1024];
116     DWORD errorNum;
117
118     if( !bitfile || !cpu.card )
119         return -1;
120
121     if( !DIME_BootVirtexSingle(cpu.card, bitfile) ) {
122         DIME_GetError(NULL,&errorNum, error);
123         printf("Configure_Error_###d\n%s\n",errorNum, error);
124     }
125
126     return 1;
127 }
128
129 int resetCard(){
130     // ENABLE = 0, DISABLE = 1
131     DIME_CardResetControl(cpu.card, drONBOARDFPGA, drENABLE, 0);
132     DIME_CardResetControl(cpu.card, drINTERFACE, drTOGGLE, 0);
133     DIME_CardResetControl(cpu.card, drONBOARDFPGA, drDISABLE, 0);
134
135     return 1;
136 }
137
138 int closeCard(){
139     if( cpu.card ) {
140         if( cpu.send ) DIME_UnLockMemory(cpu.card, cpu.send);
141         if( cpu.recv ) DIME_UnLockMemory(cpu.card, cpu.recv);
142         if( cpu.dma ) DIME_DMAClose(cpu.card, cpu.dma, ddmaCLOSETERMINATE);
143
144         DIME_CloseCard(cpu.card); }
145
146     // Close down the card
147     if( cpu.locate ) DIME_CloseLocate(cpu.locate);
148
149     cpu.card = (cpu.locate = NULL);
150
151     return 1;
152 }
153
154 int writeCard(int size) {
155     if( !cpu.card ) return -1;
156     if( !cpu.dma ) return -2;
157     if( !cpu.send ) return -3;
158
159
160     if( ddmaOK !=
161         DIME_DMAWriteFromLockedMem(cpu.card, cpu.dma,
162             cpu.send, 0, 1, ddmaBLOCKING) )
163     {
164         printf("Error_writing_stuff..!");
165     }
166

```

```
167     return 1;
168 }
169
170 int readCard(int size){
171     if( !cpu.card ) return -1;
172     if( !cpu.dma ) return -2;
173     if( !cpu.recv ) return -3;
174
175     bzero(cpu.recvBuffer ,BUFFER_SIZE);
176
177     if ( ddmaOK != DIME_DMAReadToLockedMem(cpu.card ,
178         cpu.dma, cpu.recv ,
179         0, size , ddmaBLOCKING ) ) {
180         printf("Error_reading_stuff\n");
181     }
182
183     return size;
184 }
185
186 // User functions.
187 unsigned int getWord32(){
188     readCard(1);
189     return cpu.recvBuffer[0];
190 }
191
192 unsigned int writeWord32(unsigned int word){
193     cpu.sendBuffer[0] = HostToCoreI(word);
194     writeCard(1);
195     return 1;
196 }
```

C.6 cpu.h

```
1  /*
2  * CPU Controller for Master Thesis Project
3  * 2006 - Kenneth Oestby <kenneo@idi.ntnu.no>
4  *
5  *
6  */
7
8  #include <dimesdl.h>
9
10 #ifndef _CPU_H
11
12 #define BUFFER_SIZE 2048
13 #define CLOCK_NUM 2
14 #define CLOCK_FREQ 40
15
16 typedef struct {
17     DIME_HANDLE card;
18     LOCATE_HANDLE locate;
19     DWORD leds;
20     unsigned int sendBuffer[BUFFER_SIZE];
21     unsigned int recvBuffer[BUFFER_SIZE];
22     DIME_MEMHANDLE send;
23     DIME_MEMHANDLE recv;
24     DIME_DMAHANDLE dma;
25 }CPU;
26
27 static CPU cpu;
28
29 int openCard(int cardNum, char *bitfile);
30 int setupDMA();
31 int closeCard();
32 int configureCard(char *filename);
33 int writeCard(int size);
34 int readCard(int size);
35 int resetCard();
36
37
38 unsigned int getWord32();
39 unsigned int writeWord32(unsigned int word);
40
41
42 #endif
```

C.7 util.h

```
1  /**
2   * Several utility functions
3   * to make the life easier
4   *
5   * 2007 – Kenneth Oestby <kenneo@idi.ntnu.no>
6   *
7   */
8
9  unsigned int CoreToHostI(unsigned int i);
10 unsigned int HostToCoreI(unsigned int i);
11 unsigned long ByteSwap2 (unsigned long nLongNumber);
12 unsigned int reverseInt(unsigned int i);
13 unsigned char reverseInt8(unsigned char i);
```

C.8 util.c

```

1  /**
2   * Several utility functions
3   * to make the life easier
4   *
5   * 2007 - Kenneth Oestby <kenneo@idi.ntnu.no>
6   *
7   */
8
9  static char isBigEndian = 1;
10
11 unsigned long ByteSwap2 (unsigned long nLongNumber)
12 {
13     return (((nLongNumber&0x000000FF)<<24)+((nLongNumber&0x0000FF00)<<8)+
14             ((nLongNumber&0x00FF0000)>>8)+((nLongNumber&0xFF000000)>>24));
15 }
16
17 unsigned int CoreToHostI(unsigned int i){
18     if ( isBigEndian )
19         return i;
20
21     return ByteSwap2(i);
22 }
23 unsigned int HostToCoreI(unsigned int i){
24     if ( isBigEndian )
25         return i;
26
27     return ByteSwap2(i);
28 }
29
30 /*
31  * Reverses the integer..
32  */
33 unsigned int reverseInt(unsigned int number){
34     int i = 0, j = 0;
35     unsigned int tmp = 0;
36
37     for(i=sizeof(int)*8-1,j=0;i>=0;i--,j++)
38         tmp |= ((number >> i) & 1) << j;
39
40     return tmp;
41 }
42
43 unsigned char reverseInt8(unsigned char number){
44     int i = 0, j = 0;
45     unsigned char tmp = 0;
46
47     for(i=sizeof(unsigned char)*8-1,j=0;i>=0;i--,j++)
48         tmp |= ((number >> i) & 1) << j;
49
50     return tmp;
51 }

```