Super Scalar Processor using Chip Level Optical Interconnections

D. M. Chiarulli, S. P. Levitan, R. P. Menon, N. Wattanapongsakorn
University of Pittsburgh, Pittsburgh, PA 15260

ABSTRACT

In this paper we present the design of a super-scalar processor constructed using optoelectronic components interconnected via high-speed free-space optical buses.

Keywords: Free Space Optical Interconnections, Chip Level Optical Interconnect, Super Scalar Processing, Instruction Level Parallelism.

Super Scalar Processor using Chip Level Optical Interconnections

D. M. Chiarulli, S. P. Levitan, R. P. Menon, N. Wattanapongsakorn
University of Pittsburgh, Pittsburgh, PA 15260*

1. INTRODUCTION

With the prospect of a "billion transistor" microprocessor chip becoming a reality in the next decade, VLSI photonics will be an important technology for high bandwidth, chip level, I/O. In this paper we describe a system which demonstrates the use of free space optical channels to connect the functional units in a super scalar microprocessor. This approach enables architectures in which the number of functional units available for parallel instruction execution is significantly larger than can be implemented in a purely electronic design. The design is implemented with SEED devices, flip-chip bonded on a 0.5μm CMOS silicon chip[1] and is currently being fabricated as part of the 1997 CMOS-SEED Coop program[2].

In a super scalar microprocessor, high performance is achieved by executing multiple instructions in parallel. The architecture consists of multiple functional units, each capable of independent execution, with source and result operands delivered via local interconnection busses. During program execution, a control unit works on a buffer filled with instructions that are eligible for execution and selects those instructions that can be executed without a conflict for resources or data. For example, two instructions may be in conflict over a specific functional unit, a bus, or a port to memory or a register file. These types of conflicts are collectively referred to as *structural hazards*. There may also be a conflict caused by data dependence within a sequence of instructions. In other words, the operand of one instruction depends on the result of another. These are called *data hazards.* Other conflicts can be caused by uncertainty over the outcome of a conditional branch instruction. In this case it may be unknown whether or not a particular instruction will be executed at all. This is called a *control hazard*. These conflicts place a limit on the performance of super scalar machine by limiting the number of instructions that can be executed in parallel.

In contemporary designs, architects have attempted to circumvent this limit by building additional functional units. This has an obvious impact on structural hazards but can also be effective on data and control hazards when speculative or redundant execution techniques are used[3]. For example, if a control hazard introduces uncertainty about the outcome of a conditional branch, both execution threads are allowed to proceed until the uncertainty is resolved. At that time, the computation from the untaken branch is simply discarded. Similarly data dependencies can be resolved by speculating as to the result of a dependent computation and discarding an execution thread if the guess was wrong.

In general, the more speculative instruction execution that is possible, the greater the effective level of instruction parallelism. However, the number of functional units that can be built and connected on a single chip limits electronic designs. As an alternative, we are suggesting a design which implements free space optical channels as the interconnection busses in a multi-chip super scalar system. These high-speed inter-chip busses allow us to create systems where the number of functional units is significantly larger than can be implemented in a purely electronic design. In this paper we describe a prototype system which implements six integer functional units and three registers files in a three-chip super scalar ALU design. The control unit implements a subset of the MIPS RS-2000 instruction set architecture and is capable of full dynamic (runtime) scheduling of ALU resources.

The rest of this paper is organized as follows. We begin with a logical description of the optical bus structure between the chips and the optoelectronic interface. This is followed by a description of the internal organization of the ALU chips. The optical system used to implement the optical interconnect is presented followed by simulation data showing the performance of the optical interconnect. Finally, we give a brief outline of our future research.

- Correspondence: D. M. Chiarulli, Computer Science Dept, Univ. of Pittsburgh, Pittsburgh, PA 15260, e-mail: don@cs.pitt.edu. voice: 412-624-8849 fax: 412-624-5429.
- Partially supported by USAF/Rome Laboratory under contract # AFRL F30602-96-C-0206

2. SYSTEM DESCRIPTION

We begin our system description with an outline of the logical structure of the ALU and the topology of the optical buses that connect the three ALU chips. The optoelectronic interface is implemented in a 10x20 array of seed devices, which is partitioned into three regions as shown on the right side of figure 1. Each partition consists of 52 devices leaving 44 devices in the array unused. The two outer partitions are biased such that the SEED devices operate as detectors and the center partition is biased such that the devices operate as modulators. The modulators in the center partition constantly reflect the contents of three internal 16-bit buses and a 4-bit slice of the data on a global load/store bus. As shown in the center of figure 1, this data is optically split in two and directed to one of the outer partitions of each of the other two chips. Thus any of the three chips has access to the contents of the internal buses of the other two. Any or all of the 16-bit buses within a chip can be switched to connect the incoming optical data onto the local bus, thus creating a common bus connection. On the left side of the figure we show the Control Unit, which generates all the control signals for all three chips, and the memory unit, which holds both the program and data.

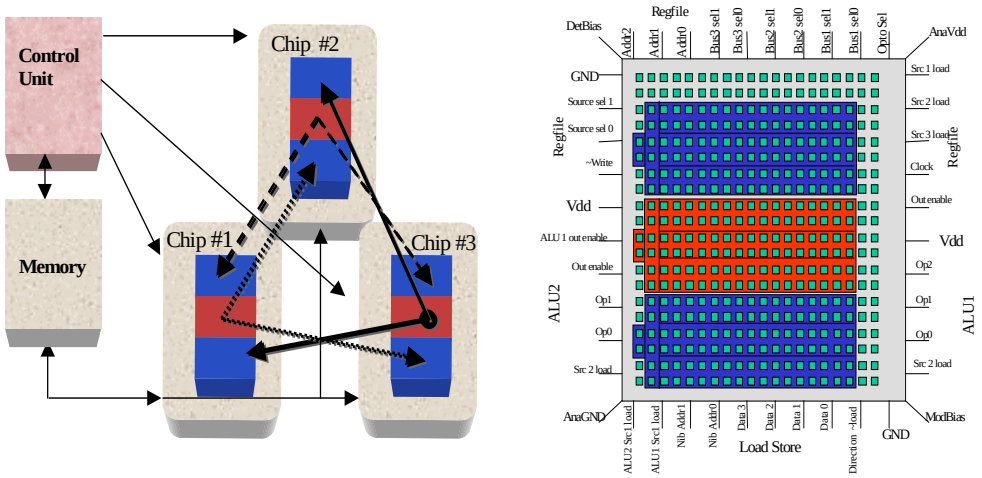

Figure 1: Optical Bus Layout and Optoelectronic Interface

Turning now to the internal organization, figure 2 shows a block diagram of the functional units and bus structure within a single ALU chip. All of the ALU chips in the architecture are identical. Thus the three chip ensemble contains six integer ALUs, three register files, and common, bit sliced, load store unit. Three 16-bit buses connect the functional units within a chip and pipeline registers have been placed at the inputs and outputs all functional units. During any clock cycle, each of these buses can transfer data originating in the output pipeline register of any functional unit or it can be driven by the optical data channels from the corresponding bus on either of the other two chips. The four main functional units on each chip are:

- **Load/store unit**: This provides the path for electrical data, i.e., memory references, to move in and out of the system. Due to pin limitations imposed by the pad frame supplied for the CMOS-SEED run, load/store operations are limited to twelve bits with each of the three chip contributing 4 bits in a bit sliced organization. Bus 4 shown in figure 2 is dedicated to load and store operations. Unlike the other buses, only a four-bit portion, carrying the locally input nibble, is transmitted optically to the other chips.
- **Register file:** This unit contains eight 16 bit registers. Register 0 always contains zero. A three-bit address line is used to select which register is to be written or read from.
- **ALU 1:** The primary ALU can perform one of eight operations; add, subtract, negate, less than, less than or equal to, equal to, start multiply, and continue multiply. Three control lines determine which operation is to be executed.

- **ALU 2:** The logical ALU can perform one of four functions: logical and, logical or, left-shift one bit, and right-shift one bit. Two control lines determine which operation is to be executed.

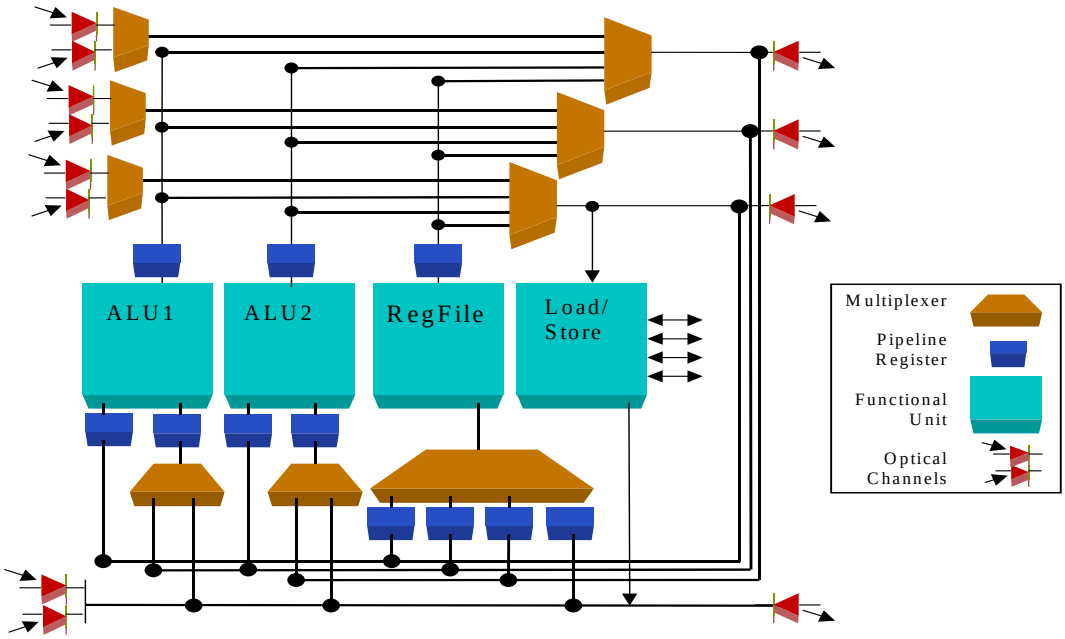

Figure 2: Single Chip Architecture

The final components of the chip are the bus mutiplexers and pipeline registers. The shared external control unit also controls these.

3. OPTICAL SETUP

Figure 3 shows the layout of the optical system[4]. All three of the ALU chips and the controller chip are mounted on a common ceramic substrate with the seed arrays aligned to tolerances of a few microns. A single laser source shown to the left of the figure acts as the optical power supply and generate an array of spots which illuminate the center group of modulators on each chip. The reflective beams after a polarization shift induced by the ¼ wave plate is directed by the beam splitter to a reflective CGH interconnect element. This element induces a 4-way beam split of which two of the split beams are discarded and two are directed back through the polarizing beam splitter and focused on the receiver of the appropriate chip.

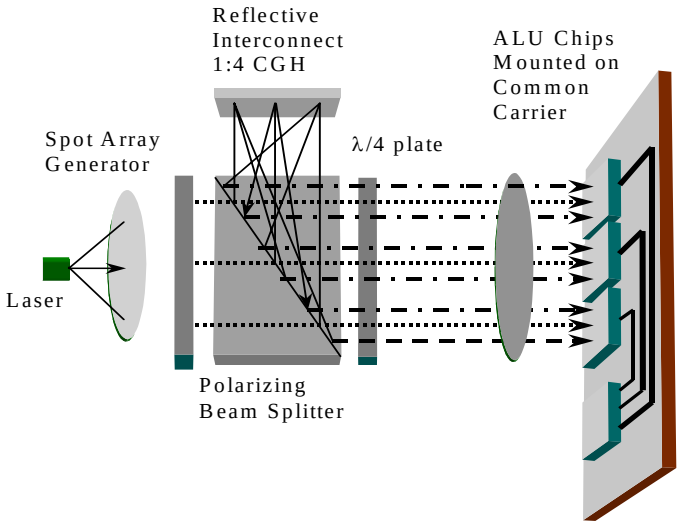

Figure 3: Optical Setup

4. PERFORMANCE DATA

Prior to fabrication, the chip design was tested using both functional and electrical simulators. For testing purposes we developed a micro-assembler which translates RS 2000 assembly language input to single clock cycle, datapath level operations on the architecture. The output is the form of a sequence of control

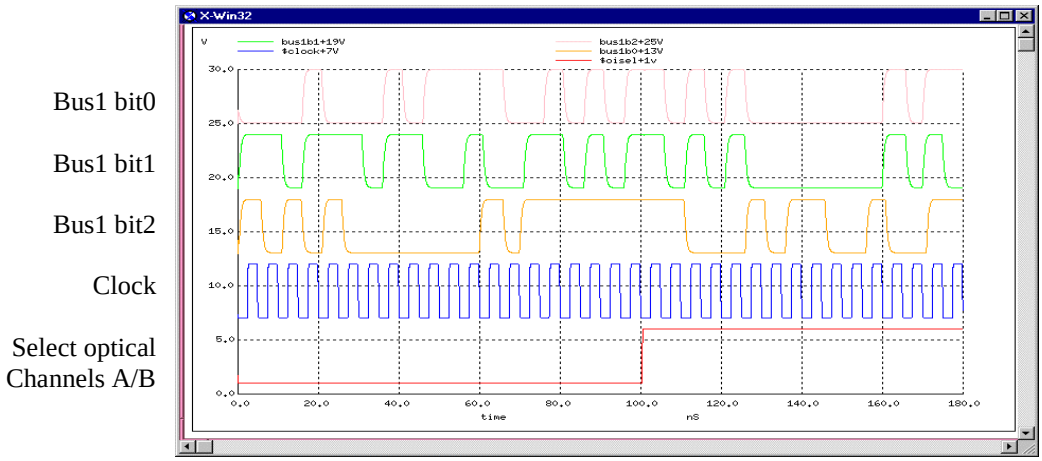

Figure 4: Electrical Simulation, optics in to optics out at 200Mhz

words coded in a format compatible with IRSIM for functional simulations or as piece-wise linear waveforms for SPICE simulator input. Using this tool, control sequences were generated and tested that functionally tested each instruction and electrically tested each datapath in the hardware. Performance data from spice was particularly encouraging. The output shown below is a waveform trace for the system clock running at 200MHz and the three low order bits of bus 1 stimulated using a random data sequence applied to the input pad of the receiver SEED devices. Bus 1 is sampled at the pads of the modulator seeds.

5. FUTURE WORK

The goal of this research is two-fold. First, it represents a significant application of chip-to-chip optical interconnections in a well-understood and non-trivial application. Second, it is an application in which the use of optics can potentially remove some of the limits of on electronic super scalar architectures. The extent to which speculative execution techniques can take advantage of the large number of resources made available in an optoelectronic architecture is an open question. Using this prototype as well as simulation studies of larger systems we hope to gain insight into this question.

6. ACKNOWLEDGEMENTS

This work has been supported in part by the United States Air Force/Rome Laboratory under contract number on AFRL F30602-96-C-0206. The authors also wish to acknowledge the support of the Consortium for Optical and Optoelectronic Technologies in Computing and Lucent Corporation for the CMOS-SEED foundry services used in this research.

Finally, we would like to acknowledge the contribution of Joe Ford of Lucent Corporation for the elegant optical system design which he sketched on the back of a napkin and which now appears in this paper.

7. REFERENCES

1. K.W. Goosen, et. al. , GaAs MQW Modulators Integrated With Silicon CMOS, IEEE Photonics Technology Letters, 7(4), pp. 360-362, April 1995.
2. http://co-op.gmu.edu
3. M. H. Lipasti and J. P. Shen, “Superspeculative Microarchitecture for Beyond AD 2000”, *IEEE Computer*, 30(8), pp. 59-66, September 1997.
4. Joe Ford, *Personal Communication*, Lucent Corporation