Fred: An Architecture for a Self-Timed Decoupled Computer

William F. Richardson and Erik Brunvand

UUCS-95-008

Department of Computer Science
University of Utah
Salt Lake City, UT 84112 USA

May 8, 1995

Abstract

Decoupled computer architectures provide an effective means of exploiting instruction level parallelism. Self-timed micropipeline systems are inherently decoupled due to the elastic nature of the basic FIFO structure, and may be ideally suited for constructing decoupled computer architectures. Fred is a self-timed decoupled, pipelined computer architecture based on micropipelines. We present the architecture of Fred, with specific details on a micropipelined implementation that includes support for multiple functional units and out-of-order instruction completion due to the self-timed decoupling.

Asynchronous Systems Research Group
University of Utah, Department of Computer Science

Fred: An Architecture for a Self-Timed Decoupled Computer

WILLIAM F. RICHARDSON (*willrich@cs.utah.edu*)
ERIK L. BRUNVAND (*brunvand@cs.utah.edu*)

Computer Science Department
3190 Merrill Engineering Building
University of Utah
Salt Lake City, Utah 84112

Keywords: processor design, computer architecture, micropipelines, FIFO, asynchronous systems

Abstract. Decoupled computer architectures provide an effective means of exploiting instruction level parallelism. Self-timed micropipeline systems are inherently decoupled due to the elastic nature of the basic FIFO structure, and may be ideally suited for constructing decoupled computer architectures. Fred is a self-timed decoupled, pipelined computer architecture based on micropipelines. We present the architecture of Fred, with specific details on a micropipelined implementation that includes support for multiple functional units and out-of-order instruction completion due to the self-timed decoupling.

1 Introduction

As computer systems have grown in size and complexity, the difficulty in synchronizing the system components has also grown. For example, simply distributing the clock signal throughout a large synchronous system can be a major source of complication. Clock skew is a serious concern in a large system, and is becoming significant even within a single chip. At the chip level, more and more of the power budget is being used to distribute the clock signal, while designing the clock distribution network can take a significant portion of the design time.

These symptoms have led to an increased interest in asynchronous designs. General asynchronous circuits do not use a global clock for synchronization, but instead rely on the behavior and arrangement of the circuit elements to keep the signals proceeding in the correct sequence. However, these circuits can be very difficult to design and debug without some additional structure to help the designer deal with the complexity. While there are many different asynchronous methodologies, one of the simplest to design, test, and debug is the self-timed micropipeline approach described by Sutherland [18], which avoids clock-related timing problems by enforcing a simple communication protocol between circuit elements. This is quite different from traditional synchronous signaling conventions where signal events occur at specific times and must remain asserted for specific time intervals. In self-timed systems it is important only that the correct *sequence* of signals be maintained. The timing of these signals is an issue of performance that can be handled separately.

Experience has shown the difficulty of writing parallel programs, yet most sequential programs have an (arguably) significant amount of instruction-level parallelism [13, 20][1]. One way of exploiting this parallelism is by decoupling the memory access portion of an instruction stream from the execution portion [7, 21, 5]. By performing the two operations independently, peaks and valleys in

[1]Nicolau claims there is lots of parallelism available. Wall claims there's some, but not much.

each may be smoothed, resulting in an overall performance gain.

Although decoupled architectures have been proposed and built using a traditional synchronous design style, a self-timed approach seems to offer many advantages. Typically the independent components of the machine are decoupled through a FIFO queue of some sort. As long as the machine components are all subject to the same system clock, connecting the components through the FIFOs is subject to only the usual problems of clock skew and distribution. If, however, the components are running at different rates or on separate clocks the FIFO must serve as a synchronizing element and thus presents even more serious problems.

The micropipeline approach is based on simple, self-timed, elastic, FIFO queues, which suggests that decoupled computer architectures may be implemented much more easily in a self-timed micropipeline form than with a clocked design. Because the FIFOs are self-timed, synchronization of the decoupled elements is handled naturally as a part of the FIFO communication. The elastic nature of a micropipeline FIFO allows the decoupled units to run at data-dependent speeds; producing or consuming data as fast as possible for the given program and data. Because the data are passed around in self-timed FIFO queues, and the decoupled processing elements are running at their own rate, the degree of decoupling is increased in this type of system organization, without the overhead of a global controller keeping track of the state of the decoupled components. This should allow increased performance due to the increased decoupling and potentially faster local control of the components, however it also means that exception handling must be considered carefully. Because each of the elements is running at its own rate, and data are possibly being transmitted through FIFO queues when the exception is signaled, care must be taken to make sure that the machine can process an exception in a functionally precise way without losing state that might be in the process of being modified by a different component.

Fred is a self-timed decoupled, pipelined processor architecture based on micropipelines. We present the architecture of Fred, with specific details on a micropipelined implementation that includes support for out-of-order instruction completion due to the decoupling, and a model for functionally precise exception processing.

2 Asynchronous Processors

In spite of the possible advantages, there have been very few asynchronous processors reported in the literature. Early work in asynchronous computer architecture includes the Macromodule project during the early 70's at Washington University [3] and the self-timed dataflow machines built at the University of Utah in the late 70's [4].

Although these projects were successful in many ways, asynchronous processor design did not progress much, perhaps because the circuit concepts were a little too far ahead of the available technology. With the advent of easily available custom ASIC technology, either as VLSI or FPGAs, asynchronous processor design is beginning to attract renewed attention. Some recent processor projects include the following:

The CalTech Asynchronous Microprocessor The first asynchronous VLSI processor was built by Alain Martin's group at CalTech [11]. It is completely asynchronous, using (mostly) delay-insensitive circuits and dual-rail data encoding. The processor as implemented has a small 16-bit instruction set, uses a simple two-stage fetch-execute pipeline, is not decoupled, and does not handle exceptions. It has been fabricated both in CMOS and GaAs.

The NSR The NSR (Non-Synchronous RISC) processor [2, 15] is structured as a five-stage pipeline where each pipe stage operates concurrently and communicates over self-timed data channels in the style of micropipelines. Branches, jumps, and memory accesses are also decoupled through the use of additional FIFO queues which can hide the execution latency of these instructions. The NSR was built using FPGAs. It is pipelined and decoupled, but doesn't handle exceptions. It is a simple 16-bit processor with only sixteen instructions, since it was built partially as an exercise in using FPGAs for rapid prototyping of self-timed circuits [1].

The Amulet A group at Manchester has built a self-timed micropipelined VLSI implementation of the ARM processor [6] which is an extremely power-efficient commercial microprocessor. The Amulet is a real processor in the sense that it mimics the behavior of an existing commercial processor and it handles simple exceptions. It is more deeply pipelined than the synchronous ARM, but it is not decoupled (although it does allow instruction prefetching), and its precise exception model is a simple one. The Amulet has been designed and fabricated. The performance of the first-generation design is within a factor of two of the commercial version [14]. Future versions of Amulet are expected to close this gap.

The Counterflow Pipeline Processor The Counterflow Pipeline Processor (CFPP) Architecture is an innovative architecture proposed by a group at Sun Microsystems Labs [17]. It derives its name from its fundamental feature, that instructions and results flow in opposite directions in a pipeline and interact as they pass. The nature of the Counterflow Pipeline is such that it supports in a very natural way a form of hardware register renaming, extensive data forwarding, and speculative execution across control flow changes. It should also be able to support exception processing.

A self-timed micropipeline-style implementation of the CFPP has been proposed. The CFPP is deeply pipelined and partially decoupled, with memory accesses launched and completed at different stages in the pipeline. It can handle exceptions, and a self-timed implementation which mimics a commercial RISC processor's instruction set has been simulated. The potential of this architecture is intriguing, but still unknown.

3 Micropipelines

Micropipelines are self-timed, event driven, elastic pipelines that may or may not contain processing between the pipe stages [18]. If no processing is done between the pipe stages, the micropipeline reduces to a simple first-in first-out (FIFO) buffer. A block diagram of a generic micropipeline is shown in Figure 1. It consists of three parts: a control network consisting of one C-element per micropipeline stage, a latch in each stage, and possibly some processing logic between the stages.

The Fred processor is implemented in a micropipeline style where concurrent processes cooperate using a request/acknowledge handshake, and connections between the processes may be pipelined to any desired degree by adding more micropipeline stages to the path. The pipelines and processes involved in the Fred processor use two-phase transition signaling and bundled data paths. Two-phase signaling involves a protocol whereby a process is requested to perform some action by receiving an event on its *Req* input, and will signal that it has completed the action by producing

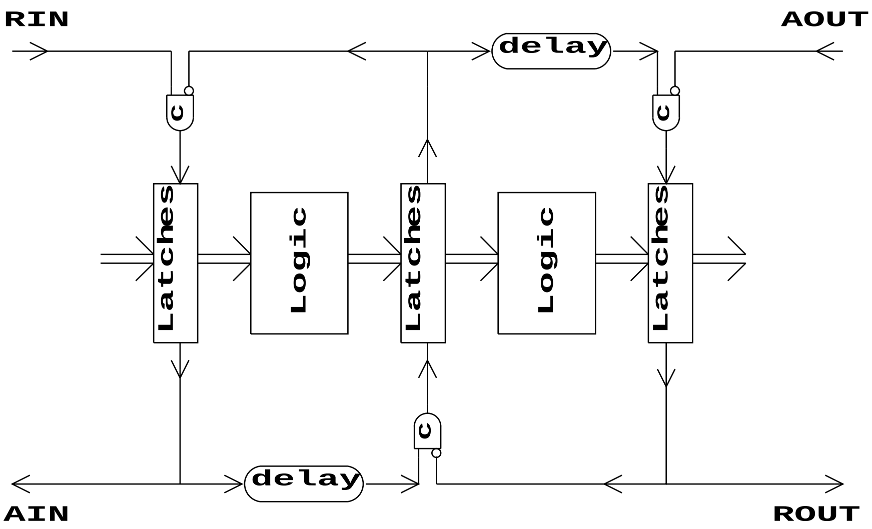

Figure 1: A Generic Micropipeline

an event on its *Ack* output. The fact that each self-timed process signals its completion allows for simple sequencing of operations in a larger self-timed system. "Transition signaling" means that the events that are communicated on the control lines are encoded in signal transitions rather than signal levels. A transition on a control line from low to high or from high to low is considered as an event and no distinction is made between two type of transitions.

"Bundled data path" means that data must be stable at a latch before the latching control signal arrives. The transition control signal indicates that the bundled data path is valid, and that the latch may update its contents. The delay elements shown in Figure 1 model the delay required to satisfy this constraint as the data move through the logic at each stage. This condition is a compromise to complete self-timed, but in practice the bundling constraint is a local constraint that is not hard to meet, and it allows the use of standard datapath circuits in the design. The logic could of course be self-timed and generate a completion signal, which would eliminate the need for the delay at the expense of more complicated logic for detecting completion of the processing. Latches used in such circuits are transition latches which latch new data upon receipt of a request event and produce an acknowledgment event when done.

The C-elements control the action of the micropipeline by acting as protocol preserving AND gates for the transition control signals of the micropipeline. These gates, drawn as an AND gate with a C inside, will drive their output low when both inputs are low, and high when both inputs are high. When the inputs are at different states, the output is held at its previous level. Note that one input of each C-element used in Figure 1 is inverted. Thus, assuming that all the control signals start low, the leftmost C-element will produce a transition to the leftmost latch when the incoming request (RIN) line first makes a transition from low to high. The acknowledge from the latch will produce a similar request through the next C-element to the right. Meanwhile, the leftmost C-element will not produce another request to the leftmost latch until there are transitions both on RIN (signaling that there are more data to be accepted) and the Ack from the next latch to the right (signaling that the next stage has finished with the current data). Each pipe stage acts as a concurrent process that will accept new data when the previous stage has data to give, and the next stage is finished with the data currently held. More details on building systems using a two-phase micropipeline circuit style can be found elsewhere [18, 15, 14].

4 The Fred Architecture

The Fred architecture is based roughly on the NSR architecture developed at the University of Utah [2, 15]. As such it consists of several decoupled independent processes connected by FIFO queues of various lengths, an approach which we believe offers a number of advantages over a clocked synchronous organization. The Fred architecture specifies the instruction set and the general layout and behavior of the processor. Other extensions to the Fred architecture may be made. New instructions may be added, and additional functional units may be incorporated. The existing functional units may be rearranged, combined, or replaced. The details of the exception handling mechanism is not specified by the architecture, but some means must be provided.

A prototype of Fred has been implemented in a detailed VHDL model. Figure 2 shows the overall organization. Each box in the figure is a self-timed process communicating via dedicated data paths rather than buses. Each of these data paths, shown as wires in Figure 2, may be pipelined to any desired depth without affecting the results of the computation. Because Fred uses self-timed micropipelines [18] in which pipeline stages communicate locally only with neighboring stages in order to pass data, there is no extra control circuitry involved in adding additional pipeline stages. Because buses are not used, the corresponding resource contention is avoided.

The VHDL version chooses particular implementations for each of the main pieces of Fred. In particular, the Dispatch unit is organized so as to issue instructions in order, but to allow out-of-order completion. This is of particular interest in a self-timed processor where the multiple functional units might take varying amounts of time to compute a result. An individual functional unit might even take different amounts of time to compute a result based on the data which will lead naturally to out of order instruction completion. The VHDL prototype is fully operational, including out-of-order instruction completion and a functionally precise exception model. The timing and configuration parameters can be adjusted for each component of the design.

Multiple independent functional units allow several instructions to be in progress at a given time. Because the machine organization is self-timed, the functional units may take as long or short a time as necessary to complete their function. One of the performance advantages of a self-timed organization is directly related to this ability to finish an instruction as soon as possible, without waiting for the next discrete clock cycle. It also allows the machine to be upgraded incrementally by replacing functional units with higher performance circuits after the machine is built with no global consequences or retiming. The performance benefits of the improved circuits are realized by having the acknowledgment produced more quickly and thus the instruction that uses that circuit finishes faster.

There are 32 general registers in the Fred architecture. Registers `r2` through `r31` are normal general-purpose registers, but `r0` and `r1` have special meaning. Register `r0` may be used as the destination of an instruction, but will always contain zero. Register `r1` is not really a register at all but provides read access to a data memory pipeline similar to that used in the WM machine [21]. Specifying `r1` as the destination of an instruction inserts the result into the pipeline. Each use of `r1` as a source for an instruction retrieves one word from the R1 Queue. For example, the instruction

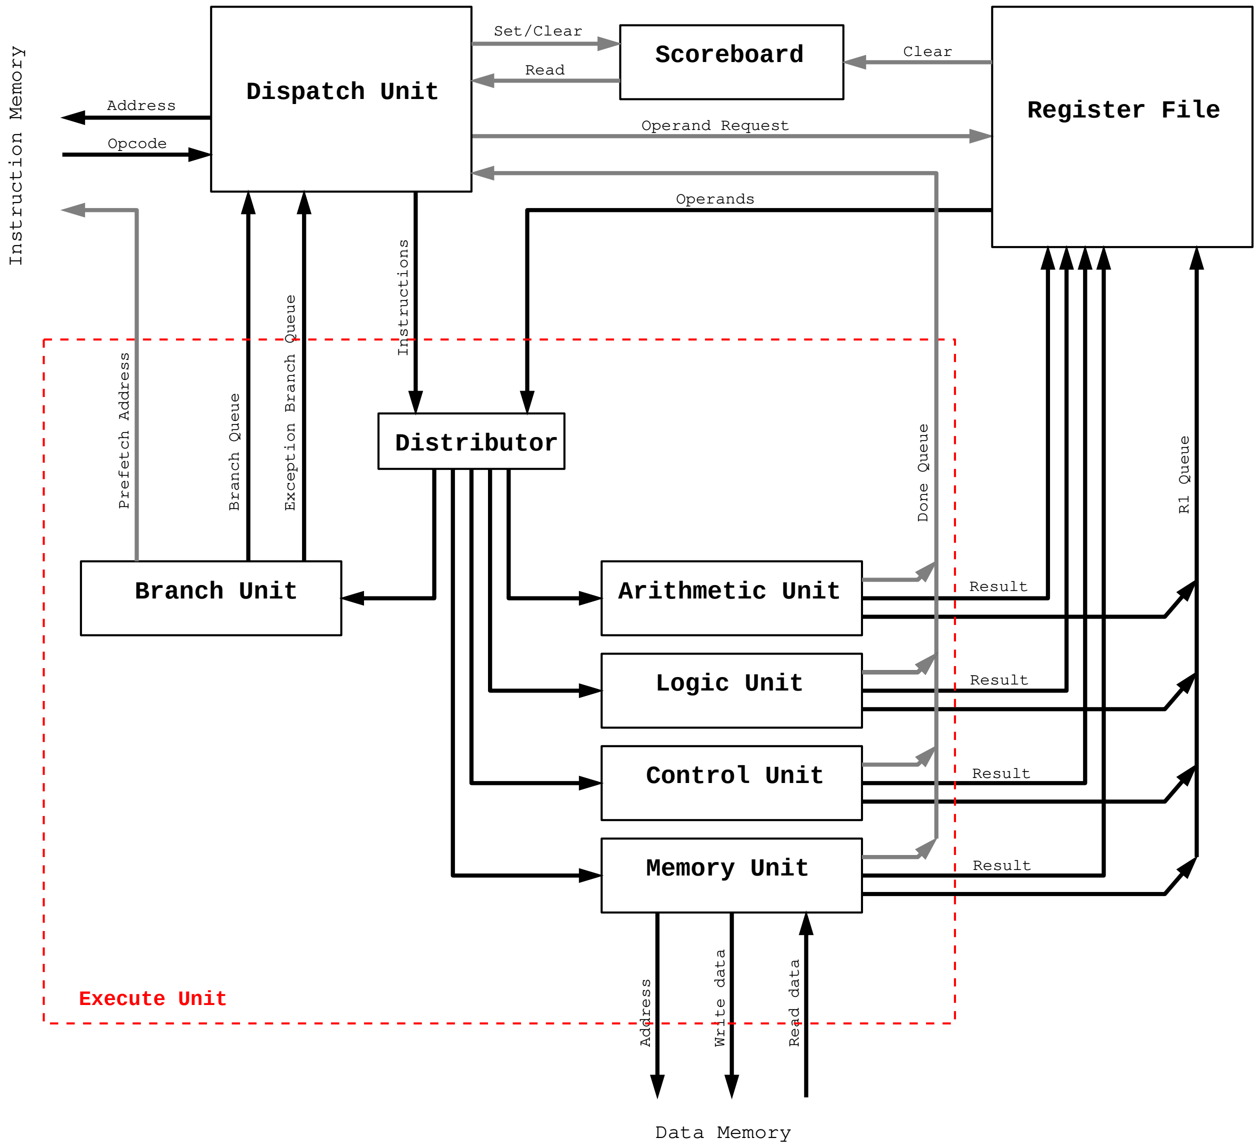

Figure 2: Fred Block Diagram. Solid black lines are primary data paths, grey lines are control paths. All data and control paths may be pipelined.

"`add r2,r1,r1`" would fetch two words from the R1 Queue, add them together, and place the sum in register `r2`. Likewise, assuming that sequential access to register `r1` would result in values A, B, and C, the instruction `st r1,r1,r1` would write the value C into memory location $A + B$. Data from any of the functional units may be queued into the R1 Queue, and loads from memory can also be queued. It may be possible to subsume some of the memory latency by queuing loaded data in the R1 Queue in advance of its use. Note that the program receives different information each time it performs a read access on register `r1`, thus achieving a form of register renaming directly in the R1 Queue. Instructions which write to the R1 Queue are forced to complete in-order, to provide deterministic behavior.

Deadlocking the processor is theoretically possible. Because both the R1 Queue and Branch Queue (section 4.2.2) are filled and emptied via two separate instructions, it is possible to issue an incorrect number of these instructions so that the producer/consumer relationship of the queues is violated. Fred's dispatch logic will detect these cases, and take an exception before an instruction sequence is issued that would result in deadlock.

4.1 Instruction Set

Choosing an instruction set for a RISC processor can be a complex task [9, 8, 10]. Rather than attempt to design a new instruction set from scratch, an instruction set from an existing commercial RISC processor was adapted. Much of the Fred instruction set is taken directly from the Motorola 88100 instruction set [12]. However, Fred does not implement all the 88100 instructions, and several of Fred's instructions do not correspond to any instructions of the 88100. The instructions, and the functional units that execute them, are shown in Figure 3.

Logic & Bitfield		Arithmetic	Memory	Branch		Control
and	clr	add	ld	blt	bb0	getcr
mask	ext	addu	st	ble	bb1	putcr
or	extu	cmp	xmem	bne	br	rte
xor	ff0	div		beq	doit	sync
	ff1	divu		bge	mvpc	trap
	mak	mul		bgt		
	rot	sub				
	set	subu				

Figure 3: Fred Instructions

4.2 Instruction Dispatch

Instruction Dispatch is, in some sense, the main control unit for the Fred processor. It is responsible for keeping track of the Program Counter, fetching new instructions, issuing instructions to the rest of the processor, and monitoring the instruction stream to watch for data hazards. Instructions are fetched and issued in program order to the rest of the machine as quickly as possible. Because different functional units may take different amounts of time to complete, individual instructions may complete in a different order than which they were issued.

4.2.1 The Instruction Window

An Instruction Window (IW) is used to buffer incoming instructions and to track the status of issued instructions [19]. A register scoreboard is used to avoid all data hazards. The IW is a set of internal registers located in the Dispatch unit which tracks the state of all current instructions. Each *slot* in the IW contains information about each instruction such as its opcode, address, current status, and various other parameters. As each instruction is fetched, it is placed into the IW. New instructions may continue to be added to the IW independently, as long as there is room for them. The scoreboard is also maintained in the Dispatch unit, and is cleared when results arrive at the Register File.

Instructions are issued from the IW in program order when all their data dependencies are satisfied (including WAW dependencies). Issuing an instruction does not remove it from the IW. Instead, instructions are removed from the IW only after they have completed successfully. Each issued instruction is assigned a tag which uniquely distinguishes it from all other current instructions. When an instruction completes, it uses this tag to report its status to back to the IW. The status is usually an indication that the instruction completed successfully, but is also used to report exceptions. Instructions signal completion as soon as the functional unit which processes them has generated a valid result, even though that result may not yet have reached its final destination. When an instruction is unsuccessful, it returns an exception status to the IW, which then begins exception processing. Instructions which can never cause exceptions do not have to report their status, and can be removed from the IW when they are dispatched. Because instructions may complete out-of-order, recoverable exceptions can cause unforseen WAW hazards. The Instruction Window contains enough information to resolve these issues.

The Dispatch unit uses the Instruction Window and scoreboard to determine when to issue new instructions to the rest of the machine. When instruction issue occurs, the required operands are requested from the Register File (possibly through a FIFO), and the instruction is issued to the EX unit (also possibly through a FIFO).

4.2.2 Branch Instructions

Flow control instructions are significantly affected by the degree of decoupling in Fred. By decoupling the branch instructions into an *address generating* part and a *sequence change* part, we gain the ability to prefetch instructions effectively. Fred does not require any special external memory system, but it can provide prefetching information which may be used by an intelligent cache or prefetch unit. This information is generated by the Branch unit when branch target addresses are computed, and is always correct.

The instructions for both absolute and relative branches compute a 32-bit value which will replace the program counter if the branch is taken, but the branch is not taken immediately. Instead, the branch target value is computed by the Branch unit and passed back to the Dispatch unit, along with a condition bit indicating whether the branch should be taken or not. These data are consumed by the Dispatch unit when a subsequent "`doit`" instruction is encountered, and the branch is either taken or not taken at that time. Although this action is similar to the synchronous concept of *squashing* instructions, Fred does not convert the `doit` instructions into NO-OPs, but instead removes them completely from the main processor pipeline.

Any number of instructions (including zero) may be placed between the branch target computation and the `doit` instruction. From the programmer's view, these instructions do not have to be common to both branches, nor must they be undone if the branch goes in an unexpected way. The only requirement for these instructions is that they not be needed to determine the direction of the branch. The branch instruction can be placed in the current block as soon as it is possible to compute the direction. The `doit` instruction should come only when the branch must be taken, allowing maximum time for instruction prefetching, as shown in Figure 4. Because the `doit` is consumed entirely within the Dispatch Unit, it will take effect as soon as the branch target data is available, allowing instructions past the branch point to be loaded into the IW before the prior instructions have completed (or even issued). This lets the IW act as an instruction prefetch buffer, but it is always correct, never speculative. Figure 5 shows an example, based on the code in Figure 4B.

This two-phase branch model allows for a variable number of "delay slots" by allowing an arbitrary number of instructions to be executed between the computation of the branch target and its use. It also allows other interesting behaviors such as achieving the effect of loop unrolling without increasing code size. This can be accomplished by computing several branch targets at one time and putting them in the branch queue before executing the loop code[2]. To avoid extra instruction fetches, the `doit` instruction can be implicitly inserted into the instruction stream by setting a bit available in the opcode of any other instruction.

```
Loop:
        addu  r3,r3,3
        mul   r9,r2,r3
        addu  r2,r9,2
        subu  r8,r8,1
        bcnd  gt,r8,Loop
        doit
```

A. Simple ordering.

```
Loop:
        subu    r8,r8,1
        bcnd    gt,r8,Loop
        addu    r3,r3,3
        mul     r9,r2,r3
        addu.d  r2,r9,2
```

B. Reordered, with implicit `doit`.

Figure 4: Two ways of ordering the same program segment.

4.3 Independent Functional Units

The Distributor is responsible for routing instructions to their proper functional unit. It takes incoming instructions and operands, matches them up where needed, and routes instructions to appropriate functional units. There are five independent functional units in the prototype implementation of Fred: Logic, Arithmetic, Memory, Branch, Control. Each functional unit is responsible for a particular type of instruction shown in Figure 3. Instructions arrive in program order but may complete in any order because the pipelines are self-timed, and the functional units themselves may take more or less time to execute a given instruction. The Distributor and its associated functional units collectively make up the Execution unit (EX).

Each of the functional units may produce results that are written back to the register file directly,

[2] This trick may not be of much use, but it *sounds* interesting.

Tag	Status	Instruction	Loop#
1	Issued	subu r8,r8,1	1
2	—	bcnd gt,r8,Loop	1
3	—	addu r3,r3,3	1
4	—	mul r9,r2,r3	1
5	—	addu r2,r9,2	1
6	—	doit	1

A. Branch target not yet available.

Tag	Status	Instruction	Loop#
4	Issued	mul r9,r2,r3	1
5	—	addu r2,r9,2	1
7	—	subu r8,r8,1	2
8	—	bcnd gt,r8,Loop	2
9	—	addu r3,r3,3	2
10	—	mul r9,r2,r3	2
11	—	addu r2,r9,2	2
12	—	doit	2

B. Branch target consumed.

Figure 5: Prefetching by the Instruction Window. A) Prefetching must wait until the branch target is available. B) When the target is available, the `doit` is consumed and prefetching continues with the next iteration of the loop.

or reenter the register file through the R1 Queue. In addition, forwarding may take place in each functional unit in a manner similar to that found in synchronous processors. The only difference is that in a synchronous processor the forwarded data will stay in the forwarding register only until the following clock tick. In a self-timed processor, data could stay in the forwarding register until the next instruction that wanted that data removes it. Although this is not implemented in the current version of Fred, this could be managed easily either by hardware or software (compiler).

The Memory subunit is treated as just another functional unit. The only difference is that the Memory unit sometimes produces data that is written to the data memory rather than the Register File.

4.4 Register File

The Register File responds to requests from the Dispatch unit for operands which it delivers through a FIFO to the EX unit. These operands are paired with instructions and passed to the appropriate functional unit. Because the instructions are issued in program order, there is no matching required to determine which operands should be paired with which instructions. They emerge from the FIFO queues in the correct sequence.

On the incoming side, the Register File accepts results from each functional unit that produces data. These results are accepted independently from each functional unit and are not multiplexed onto a common bus. Data hazards are prevented by the scoreboard and the Dispatch unit, which will not issue an instruction until all its data dependencies are satisfied, so there will never be conflicts for a register destination. The Register File clears the associated scoreboard bit when results arrive at a particular register. Instruction results may also be written into the R1 Queue as described earlier, but there is no actual register associated with the R1 Queue. Instead, the *Dispatch unit* clears the scoreboard bit for register `r1` when the producing instruction completes successfully.

5 Exceptions

Fred uses an Instruction Window [19] in the Dispatch unit to maintain the status of all current instructions. Exceptions are functionally precise. The exception model seen by the programmer is not that of a single point where the exception occurred. Instead, there is a *set* of instructions which were in progress. The hardware guarantees that this set (unless empty) will consist only of instructions which either faulted or which have been fetched but not yet issued when the exception occurred. The instructions in this set are a *subset* of a sequential portion of the dynamic program instructions. The missing elements are those instructions which completed successfully out of order, and so should not be re-issued. Because the total state of the processor is not available at one known time (such as on a clock tick), the details of the exception handling are somewhat complicated, but no more so than for a synchronous processor that is deeply pipelined and may issue or complete instructions out of order. This is described in more detail elsewhere [16].

6 Conclusions

Self-timed implementation seems to be a natural match for decoupled computer architectures. The ability to allow different parts of the machine to proceed at their own rate, and the natural use of self-timed FIFO queues, enhances the decoupling due to the architecture. The current prototype of Fred is in the form of a detailed VHDL model. This model is completely functional including the out-of-order instruction completion and functionally precise exceptions. We have a Fred assembler and a translator to convert 88100 assembly language into Fred's instruction set, so we can run compiled C programs through the VHDL simulation. We are in the process of investigating tradeoffs involved in queue depth, decoupled branches, functional unit performance, out-of-order instruction completion, exception handling, and other features of the architecture.

References

1. Erik Brunvand. Using FPGAs to prototype a self-timed computer. In *International Workshop on Field Programmable Logic and Applications*, Vienna University of Technology, September 1992.

2. Erik Brunvand. The NSR processor. In *Proceedings of the 26th Annual Hawaii International Conference on System Sciences*, pages 428–435, Maui, Hawaii, January 1993.

3. Wesley A. Clark and Charles A. Molnar. Macromodular system design. Technical Report 23, Computer Systems Laboratory, Washington University, April 1973.

4. A.L. Davis. The architecture and system method for DDM1: A recursively structured data-driven machine. In *5th Annual Symp. on Computer Architecture*, April 1978.

5. Matthew Farrens, Pius Ng, and Phil Nico. A comparison of superscalar and decoupled access/execute architectures. In *Proceedings of the 26th Annual ACM/IEEE International Symposium on Microarchitecture*, Austin, Texas, December 1993. IEEE,ACM.

6. S. B. Furber, P. Day, J. D. Garside, N. C. Paver, and J. V. Woods. A micropipelined ARM. In *Proceedings of the VII Banff Workshop: Asynchronous Hardware Design*, Banff, Canada,

August 1993.

7. J. R. Goodman, J. Hsieh, K. Liou, A. R. Pleszkun, P. B. Schechter, and H. C. Young. PIPE: A VLSI decoupled architecture. In *12th Annual International Symposium on Computer Architecture*, pages 20–27. IEEE Computer Society, June 1985.

8. Thomas R. Gross, John L. Hennessy, Stephen A. Przybylski, and Christopher Rowen. Measurement and evaluation of the MIPS architecture and processor. *ACM Transactions on Computer Systems*, 6(3):229–257, August 1988.

9. John Hennessy, Norman Jouppi, Forest Baskett, Thomas Gross, and John Gill. Hardware/software tradeoffs for increased performance. In *Proceedings of the Symposium on Architectural Support for Programming Languages and Operating Systems*, pages 2–11. ACM, April 1982.

10. Manolis G. H. Katevenis. *Reduced Instruction Set Computer Architectures for VLSI.* MIT Press, 1985.

11. Alain Martin, Steven Burns, T.K. Lee, Drazen Borkovic, and Pieter Hazewindus. The design of an asynchronous microprocessor. In *Proc. CalTech Conference on VLSI*, 1989.

12. Motorola. *MC88100 RISC Microprocessor User's Manual.* Prentice Hall, Englewood Cliffs, New Jersey 07632, second edition, 1990.

13. Alexandru Nicolau and Joseph A. Fisher. Measuring the parallelism available for very long instruction word architectures. *IEEE Transactions on Computers*, C-33(11):110–118, November 1984.

14. Nigel Charles Paver. *The Design and Implementation of an Asynchronous Microprocessor.* PhD thesis, Unversity of Manchester, 1994.
`http://www.cs.man.ac.uk/amulet/publications/thesis/paver94_phd.html`

15. William F. Richardson and Erik Brunvand. The NSR processor prototype. Technical Report UUCS–92–029, University of Utah, August 1992.
`ftp://ftp.cs.utah.edu/techreports/1992/UUCS-92-029.ps.Z`

16. William F. Richardson and Erik Brunvand. Precise exception handling for a self-timed processor. To appear in *1995 International Conference on Computer Design: VLSI in Computers & Processors*, October 1995.

17. Robert F. Sproull and Ivan E. Sutherland. Counterflow pipeline processor architecture. Technical Report SMLI TR-94-25, Sun Microsystems Laboratories, Inc., M/S 29-01, 2550 Garcia Avenue, Mountain View, CA 94043, April 1994.
`http://www.sun.com/smli/technical-reports/1994/smli_tr-94-25.ps`

18. Ivan Sutherland. Micropipelines. *Communications of the ACM*, 32(6):720–738, 1989.

19. H. C. Torng and Martin Day. Interrupt handling for out-of-order execution processors. *IEEE Transactions on Computers*, 42(1):122–127, January 1993.

20. David W. Wall. Limits of instruction-level parallelism. WRL Technical Note TN-15, Digital Western Research Laboratory, 100 Hamilton Avenue, Palo Alto, CA 94301, December 1990.

21. Wm. A. Wulf. The WM computer architecture. *Computer Architecture News*, 16(1), March 1988.