

A Technology-Scalable Architecture for Fast Clocks and High ILP

Karthikeyan Sankaralingam Ramadass Nagarajan Doug Burger Stephen W. Keckler

Computer Architecture and Technology Laboratory
Department of Computer Sciences
The University of Texas at Austin
cart@cs.utexas.edu — www.cs.utexas.edu/users/cart

Abstract

CMOS technology scaling poses challenges in designing dynamically scheduled cores that can sustain both high instruction-level parallelism and aggressive clock frequencies. In this paper, we present a new architecture that maps compiler-scheduled blocks onto a two-dimensional grid of ALUs. For the mapped window of execution, instructions execute in a dataflow-like manner, with each ALU forwarding its result along short wires to the consumers of the result. We describe our studies of program behavior and a preliminary evaluation that show that this architecture has the potential for both high clock speeds and high ILP, and may offer the best of both the VLIW and dynamic superscalar architectures.

1 Introduction

Conventional microarchitectures have been improving in performance by approximately 50-60% per year, improving the instructions per cycle (IPC) using more transistors on a chip and increasing the clock speed. However both strategies will fail for future technologies (50nm and below), with clock speed growth slowing down because of fundamental pipelining limits and wire delays making architectures communication bound [1]. Thus, today's architectures will not scale, showing diminishing returns in IPC even with increasing chip transistor budgets. New designs must address these issues, efficiently utilizing the increasing transistor budget while overcoming communication bottlenecks.

One approach for extracting ILP is through conventional superscalar cores that detect parallelism at run-time. The amount of ILP that can be detected is limited by the issue window, whose logic complexity grows as the square of the number of entries [12]. Conventional architectures also rely on many frequently accessed global structures, such as register files, re-order buffers and issue windows, which become bottlenecks limiting clock speed or pipeline depths.

Another approach for extracting parallelism is taken by VLIW machines, in which ILP analysis is performed at compile time. Instruction scheduling is performed by the compiler, orchestrating the flow of execution statically. This approach performs well only for regular workloads and suffers from the drawback that dynamic events are not handled well — a stall in one functional unit forces the entire ma-

chine to stall, since all functional units must be synchronized.

In this paper, we describe a new architecture called the Grid Processor that takes into consideration the technology constraints of wire delays and pipelining limits. The compiler is used to detect parallelism and statically schedule instructions on a computation substrate, but instructions are issued dynamically. We propose an execution substrate that consists of a set of *named distributed computing elements* to which the compiler statically assigns individual instructions.

The architecture does not suffer from VLIW issue restrictions, as instructions are issued dynamically and executed in a dataflow fashion. Instructions from a compiler-generated basic block or hyperblock are mapped statically to nodes in the computation array, with each node being assigned one or more instructions. The nodes issue instructions dynamically when the input operands are available. Temporary values produced and consumed inside a block are not visible to the architectural state, and are instead forwarded directly from the producers to their consumers.

We propose a fine-grained partitioning of the issue window and associated functional units (FU). The computation array includes both a grid of issue window-FU pairs (nodes) and a dedicated communication network for passing data. Data produced at a node are routed dynamically through intermediate nodes to their eventual destinations. The architecture is a hybrid between conventional superscalar and conventional VLIW architectures, issuing instructions dynamically with static scheduling. The hardware substrate is designed to extract high ILP and run at aggressive clock rates.

In the Grid Processor, the available transistor budget is used to build an array of computation elements aimed at overcoming several challenges of communication overhead in future systems. First, by forwarding values directly between producers and consumers, the reliance on centralized structures is reduced. Second, compiler controlled physical layout ensures that the critical path is scheduled along the shortest physical path. Finally, instruction blocks are mapped onto the grid as single units of computation amortizing scheduling and decode overhead over a large number of instructions. The reduced reliance on centralized structures allows the computation substrate to be clocked at high speeds.

The remainder of this paper is organized as follows. Sec-




Figure 1: Grid block diagram. Express channels connect the last row with the first row

tion 2 describes the key features of the Grid Processor and demonstrates how programs are mapped onto it. Section 3 characterizes certain aspects of program behavior indicating that existing applications are amenable for execution on the Grid Processor. Section 4 describes related work pertaining to wide-issue and dataflow oriented machines. Section 5 concludes with a discussion on some of the secondary advantages including power reduction and speculation control as well as the remaining issues to be solved.

2 Architecture

The Grid Processor consists of a computation substrate that is configured as a two-dimensional grid of fine-grained computation nodes connected by an interconnection network. The compiler partitions the program into a sequence of blocks (basic blocks or hyperblocks [16]), performs renaming of temporaries, and schedules instructions in a block to nodes of the grid. Instruction traces generated at run-time could be used instead of blocks generated by the compiler. Blocks are fetched one at a time and their instructions are mapped to the computation nodes *en masse* as assigned by the compiler. Execution proceeds in a dataflow fashion with each instruction sending its results directly to other instructions that use them. A set of interfaces are used by the computation substrate to access external data.

2.1 Computation Nodes

Figure 1 shows a high level overview of the grid with some of the associated interfaces. Nearby neighbors in the grid are connected by short wires that have small communica-

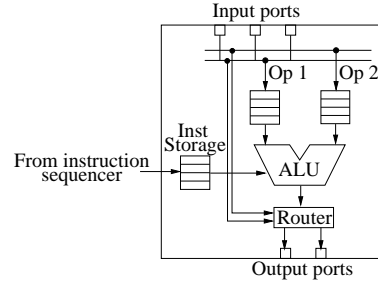


Figure 2: Organization of a computation node

tion delays ¹. Fast express channels connect nodes that are physically far apart in the grid. The instruction sequencer fetches blocks of instructions from the instruction memory and places those instructions on the nodes as scheduled by the compiler. The block termination control interfaces with the register file and the memory interface, detects when a block completes execution, and commits architecturally visible data to the register file and memory. The memory interface is used to communicate with the load/store queue, caches, and main memory.

The computation nodes are lightweight units that perform the function of execution, temporary storage, and data forwarding. Each computation node consists of a set of functional units, storage structures, an instruction wakeup unit, a router, and read/write ports for communication. Figure 2 shows the layout of a computation node. The functional units consist of an integer unit and optionally a floating point unit that perform the actual execution. The storage structures include a set of queues and buffers for storing the instructions, their input operands, and data tokens that need to be forwarded to other nodes in the grid. The instruction wakeup unit matches instructions with their operands as they arrive and issues them to the functional units for execution. The router examines tokens in the storage structures and forwards them along one of the many paths out of this node to their eventual targets. Data tokens meant for other nodes bypass the ALU and are directly forwarded by the router to their destinations.

2.2 Execution Model

The compiler partitions the program into a sequence of *blocks*. Blocks are constructed such that there are no internal control flow changes, and all control transfers out of a block initiate instructions in other blocks. These blocks may be basic blocks, hyperblocks, or program traces generated at run-time. Figure 3 shows a stream of instructions that has been partitioned by the compiler into three different blocks

¹The figure shows one possible grid interconnect topology as an example only.

```

0x0000 add r1, r2, r3    // I1
0x0004 add r2, r2, r1    // I2
0x0008 ld r4, (r1)      // I3
0x000c add r5, r4, 1    // I4
0x0010 beqz r5, 0xdeac  // I5

// End of block B1

0x0014 add r10, r2, r3   // I6
0x0018 add r11, r2, r3  // I7
0x001c ld r4, (r10)     // I8
0x0020 ld r5, (r11)     // I9
0x0024 mul r31, r4, r5  // I10
0x0028 bne r31, 0xbee0  // I11

// End of block B2

0x002c xor r8, r5, 1    // I12
0x0030 sll r9, r4, r8   // I13
0x0034 add r13, r9, 8   // I14
0x0038 add r12, r9, r2  // I15
0x004c sw r13, r12     // I16

// End of block B3

0x0050 add r1, r6, r9
.....
0x0070 jmp 0x0050

```

Figure 3: A sample instruction stream

(basic blocks in this case) B1, B2, and B3. Explicit *move* instructions, separate from the computation instructions, are generated for the registers read by every block. The move instructions fetch block inputs from the register file and pass them as internal (temporary) values to the block. Figure 4 shows the Data Flow Graph (DFG) of the blocks B1, B2, and B3 in Figure 3 along with the *move* instructions. As shown in the figure, all instructions have been renamed with temporary registers for their operands and *move* instructions generated for every input register. For example, in block B1, two move instructions, *move t2, r2* and *move t3, r3* are generated by the compiler for input registers r2 and r3. Inside a block, all values are referenced using temporary names. The *move* instructions associates register inputs of the block and temporaries. Data values that must be passed to other blocks are written to the register file.

At run-time, the instruction sequencer fetches a block from the instruction memory and maps it onto the grid *en masse*; there is no serialization of fetch, decode and rename for the instructions within a block. Individual instructions are written to the storage structures of the nodes to which they have been assigned at compile time. Block execution is initiated by the move instructions which read register data and send them to their consumers. The instruction wakeup unit matches incoming data with an instruction and issues ready instructions to the functional unit for execution. The results of the computation are tagged and forwarded by the router through the interconnect to their eventual destinations.

2.2.1 Instruction Mapping

The compiler generates a mapping by physically laying out the data flow graph of each block on a grid. Every computation instruction in the block is assigned to a node in the grid, with the critical path scheduled along the shortest possible physical path. All output operands are renamed with the positions of consumer nodes. Move instructions serve the purpose of associating register data with positions of their consumer nodes. Figure 5 illustrates a layout of the grid with instructions mapped on the computation elements. Instructions I1 and I2 of block B1 in Figure 4 are mapped on the grid at positions (0,1) and (1,1) respectively. Correspondingly, the move instruction *move t2, r2* has (0,1) and (1,1) encoded in its destination fields and register name r2 in its input field.

2.2.2 Instruction Wakeup and Execution

As described in section 2.1, multiple instructions are mapped onto a single node and data are written to an operand buffer when they arrive. Upon arrival of a data token, the instruction and operand buffers are examined to wake-up and issue ready instructions. The wakeup delays will be considerably smaller than seen in conventional cores because of smaller issue windows. The computation node serially performs two operations whenever an operand arrives — *wakeup* and *execute*. Serializing wakeup and execute may increase the cycle time along the execute-execute path of dependent instructions. Wakeup-execute can be pipelined into two stages, if the instruction wakeup does slow the clock. The execute phase of the producer can be overlapped with the wakeup phase of its consumer.

Conventional superscalar cores have dedicated bypass paths to forward data which can be used to guarantee that following the execution of an instruction its dependents will have that data in the next cycle. In the Grid Processor, since data are routed dynamically, there is no dedicated path that is guaranteed to be free when an instruction completes execution to forward its data to its consumer. However, there are several mechanisms that can alleviate this problem. Special wakeup tokens could be generated during the issue stage of a producer instruction. They reach the consumer nodes at the end of the stage, reserving a channel for the data to follow in the next cycle. Alternately, speculative instruction issue could be used to hide the select latency with local rollback mechanisms in the event of incorrect issue.

2.2.3 Block Mapping

The instruction and operand storage structures at a node can be used to buffer multiple instructions and data, which are associated through tags. There are three reasons to have multiple instructions mapped on a node. First, graphs larger than the physical grid can be folded over and mapped on

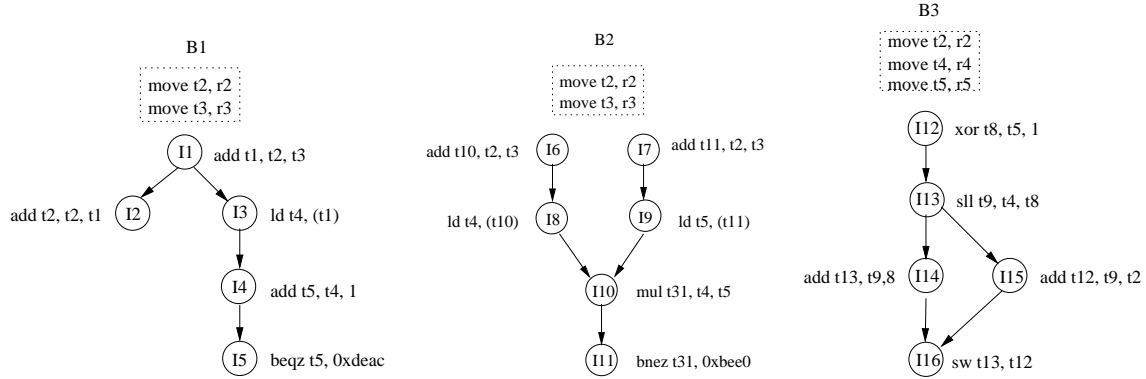


Figure 4: Basic blocks shown as dataflow graphs. Registers are marked with “r” and temporaries with “t”.




Figure 5: Basic blocks mapped on a grid of dimension 4x4, with the 3 nearest neighbors reachable directly. Instruction destinations are ordered pairs (x, y) , which identify a consumer with a relative position x nodes below and y nodes to the right of the producer.

the grid with more than one instruction at a node. Second, instructions from different blocks that are fetched speculatively (using control speculation or from speculative threads [19]) can be mapped at a node. Finally, blocks from different threads can also be mapped to support multithreading.

2.3 Instruction Encoding

The Grid Processor ISA is divided into data movement instructions and computation instructions. Data movement instructions include *move*, *split* and *repeat* instructions. The *move* instructions fetch block inputs from the register file and pass them as temporary values to the block. Encoding space limitations restrict the number of targets that can be specified in an instruction. The *split* instructions replicate data to reach additional targets. The range of each target (distance from the producer) that can be specified is finite. The *repeat* instructions are used to forward data to targets

outside the range. There is a trade-off between the instruction size, number of specifiable targets and the range of each target.

Every instruction is encoded with an opcode field, destination field, and in the case of *move* instructions, an input field. The destination field consists of multiple targets, with each target encoded with the position of the consumer expressed as an offset. The *move* instructions are encoded with an input register name and its destinations. Figure 5 shows a sample encoding of a *move* instruction and a computation instruction. The *move* instruction `move t2, r2` is encoded as `move (0, 1), (1, 1), r2` corresponding to the input register `r2` and consumer instructions `I1` and `I2` that are mapped at $(0, 1)$ and $(1, 1)$ respectively. Instruction `I1` is encoded with destinations corresponding to consumers `I2` and `I3` of the temporary `t1`. `I2` is mapped at the node directly below `I1` and `I3` is mapped at the node one to the right and one below `I1`. An extra bit (not shown in the figure) is necessary to specify the order of the input operands.

2.4 Role of the Compiler

The compiler plays an important role in the Grid Processor. Apart from detecting ILP, the compiler constructs blocks that are scheduled on the grid, and defines mechanisms for intra and inter-block communication. The compiler must also generate data movement instructions to overcome encoding space limitations.

In the Grid Processor, blocks are fetched as a single unit, mapped on the grid, and executed in a dataflow fashion. Since the instructions are fetched at a block granularity, it is desirable to have large blocks and good *block utilization*. Block utilization is defined as the ratio of dynamically executed instructions to the static size of the block. One method of building large blocks is to build hyperblocks based on profiling information. Register file communication for data passed between successively executed blocks can be bypassed using the grid interconnect, thereby “stitching” these blocks as a single dataflow graph. The compiler must define interfaces and mechanisms to stitch together multiple blocks. Since data movement instructions add overhead, when scheduling the graph, the compiler should minimize the critical path and attempt to minimize the number of such instructions.

3 Preliminary Analysis

In this section, we investigate the amenability of existing applications to the Grid Processor and examine few aspects of program behavior that affect performance. Large blocks with a significant number of block temporaries and a few input and output registers are desirable because they have low register file bandwidth. It is also desirable to have large blocks with high utilization to amortize the cost of block fetch and map. The encoding space needed for representing temporaries is determined by the number of targets of an instruction. Fewer average targets per value produced permits a compact encoding. We examine these characteristics in existing applications to determine how well they map onto the Grid Processor.

In our experimental analysis, SPEC CPU2000 benchmarks were compiled using the Trimaran [20] tool set. Three floating point (*equake*, *ammp* and *art*) and three integer (*parser*, *gzip* and *mcf*) benchmarks were selected for analysis. Hyperblocks were generated for these benchmarks using Trimaran’s IMPACT compiler with the *train* input set for profiling. All of the benchmarks were simulated using the Trimaran simulator for 500 million instructions with the *ref* input set. We collected dynamic statistics using modifications made to the simulator to track block size profiles and register usage.

Benchmark	Average Instructions per block			
	Static Size	Dynamically Executed	Never executed due to early branches	NOPs
gzip	144	77	59	8
mcf	48	35	8	5
parser	29	27	1	1
art	129	125	2	2
equake	57	52	2	3
ammp	124	103	8	13

Table 1: Block utilization

3.1 Instruction Behavior

In this section, we examine some aspects of program behavior. Performance is affected by the block sizes in programs and grid configurations. A profile of the dynamic block size was obtained for all the benchmarks to analyze the trade-off of block size with respect to block utilization. Wide grids have better performance at the cost of increased area with fewer nodes having mapped instructions. We analyze this trade-off for three different grid widths.

3.1.1 Block Size

From our analysis of the SPEC CPU2000 benchmarks, we observed that large hyperblocks can be built. Figure 6 shows the dynamic block size profiles for the different benchmarks. For each of the benchmarks, the figure plots the percentage of execution time spent for each dynamic block size as a cumulative distribution function. Dynamic block size is the number of instructions in a block that are actually executed, excluding predicated instructions that are converted to NOPs.

Nearly 70% of the execution time is spent in blocks of size greater than 26 for the integer benchmarks and blocks of size greater than 65 for the floating point benchmarks. Across the benchmarks, the average number of dynamic instructions in a block ranges from 27 to 125.

High utilization percentages are desirable because blocks are fetched and mapped as a single unit. Blocks with poor utilization have a large number of instructions that are fetched and mapped without being executed. In a hyperblock, instructions may not be executed, either because of early exits from the block or because of predicated instructions being converted to NOPs. Table 1 shows the average number of instructions in each block that belong to the categories described above. For four of the six benchmarks, the average utilization is nearly 90% with average static block sizes ranging from 29 to 129. The benchmark *gzip* shows worse utilization than the other benchmarks, as nearly 40% of the instructions fetched are never executed. However, these preliminary results show a potential for building large




Figure 6: CDF of block size profiles. The X-axis represents the number of dynamically executed instructions in a block. The Y-axis represents the percentage execution time spent by blocks of corresponding sizes expressed as a cumulative distribution function. For a block, we approximate the execution time as its dynamic instruction count.

Width	Min. height	Nodes	Utilization
ammp — 218 insts			
4	55	220	99%
8	43	344	63%
16	39	624	35%
parser — 43 insts			
4	13	52	83%
8	9	72	72%
16	7	112	46%

Table 2: Grid utilization for different grid widths, for the most frequently executed blocks in *ammp* and *parser*.

blocks with good utilization.

3.1.2 Grid Utilization

The grid dimensions and connectivity determine its performance and utilization. We define *grid utilization* as the fraction of nodes in the grid that have instructions scheduled on them. This is affected by the richness of the interconnect: fewer connections will result in a sparser schedule, in turn resulting in a lower utilization. We conservatively chose a connectivity that is relatively restrictive, guaranteeing technological scalability. Future studies will explore the connectivity design space. For a preliminary analysis, we used a simulator that schedules the DFG of a block on a grid of finite width and infinite depth, using a greedy critical path scheduling strategy. This strategy schedules one instruction per node, with the longest path in the DFG on the shortest possible physical path in the grid. Sufficient encoding space was assumed to specify up to 3 targets per instruction, at any grid position. We selected the most frequently executed blocks in two benchmarks and examined the grid utilization for three grid configurations. These blocks in *ammp* and *parser* account for 44% and 15% of the dynamic instructions executed. Their static sizes were 218 and 43 instructions, respectively.

The blocks were scheduled on three grids with the same connectivity of 3, but different widths: 4, 8, and 16. Increasing the width of the grid increases the number of independent instructions that can be scheduled in a single row. A wider row thus reduces the height of the schedule, but wastes more nodes. For each schedule, we computed the grid utilization and the minimum grid height required.

Table 2 shows the grid utilization for the three cases. The table shows that at a width of 4, both the blocks exhibit a high utilization of 99% and 83%, requiring a grid height of 55 and 13 respectively. An increase in the grid width to 8 decreases the height by 8 in *ammp* and 4 in *parser*, improving performance but with poorer node utilization. As the width is further increased to 16, the drop in the height is less, but the utilization drops dramatically. The best balance of the grid height and node utilization is at 8. We note that

Input registers						
Number of regs	% of blocks					
	ammp	equake	art	gzip	parser	mcf
0-4	88.6	32.1	1.0	48.2	82.9	93.5
5-9	5.2	32.9	0.3	34.8	15.2	2.9
10-14	4.4	19.5	19.8	11.6	1.0	1.3
15-19	0.5	1.5	56.8	5.0	0.4	2.2
>= 20	0.4	13.5	22.0	0.0	0.3	0.0
Temporary registers						
0-9	39.8	47.4	20.6	37.6	85.8	32.9
10-19	14.9	16.6	38.1	19.1	3.5	3.1
20-29	5.1	10.3	18.8	18.4	10.0	59.6
30-39	13.0	1.7	0.0	0.4	0.2	2.9
40-49	1.3	4.5	0.8	14.3	0.0	0.0
50-59	0.8	0.3	0.0	0.0	0.0	0.0
60-69	2.3	5.4	1.6	0.5	0.0	7.1
70-79	0.0	0.0	3.6	2.2	0.0	0.0
80-89	0.0	0.0	0.6	0.5	0.0	0.0
90-99	22.3	1.0	0.5	0.1	0.0	0.0
>= 100	0.0	11.0	0.0	0.3	0.0	7.1
Output registers						
0-4	94.6	63.4	20.1	62.9	97.6	97.6
5-9	4.2	6.4	0.8	15.0	2.0	1.9
10-14	0.2	21.8	37.8	18.5	0.2	0.4
15-19	0.9	6.8	25.6	3.3	0.0	0.1
>= 20	0.0	1.2	15.6	0.0	0.0	0.0

Table 3: Input, Output and Temporary registers used by the blocks.

mapping multiple blocks from same or different threads may increase the utilization at a minimal cost.

3.2 Register Behavior

Block inputs and block outputs are data read from and written to the register file. The number of such inputs and outputs determine the bandwidth and number of ports required at the register file. Block temporaries are data created and used within a block. With a large number of temporaries, a significant amount of communication through the register file can be eliminated. We report the results of our simulations to estimate the number of input registers, output registers and block temporaries used during program execution.

3.2.1 Input, Output and Temporary Data

In a “well-behaved” block, the number of register inputs and outputs is small and the number of temporaries large. Table 3 shows the number of input, output and temporaries used on the six selected benchmarks. The percentage of executed blocks versus the number of registers for the three different types is shown. For example, in the benchmark *ammp*, fewer than 5 registers are written out by 94.6% of the blocks executed, and an additional 4.2% of the blocks output between 5 and 9 registers.

The integer and FP benchmarks show slightly different behavior. More than 75% of the blocks in the integer bench-

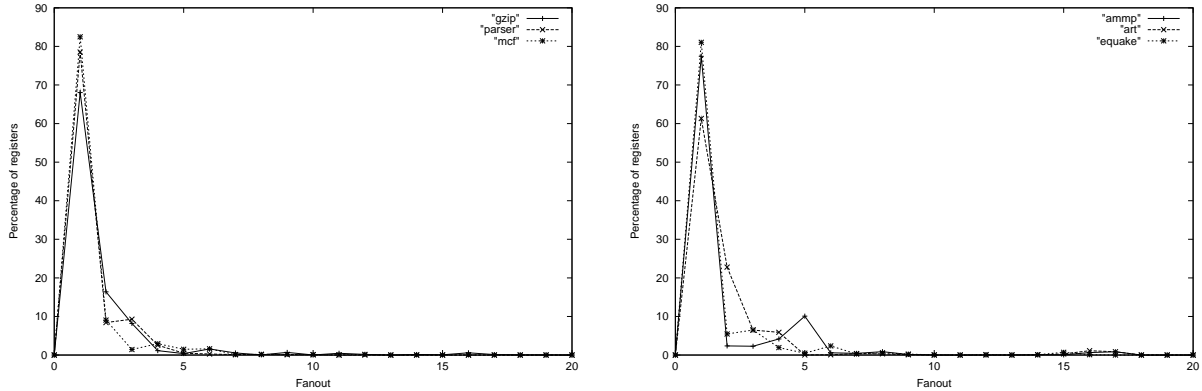


Figure 7: Register fanout — Number of targets for each produced value

marks read or write fewer than 10 registers. For the FP benchmarks, more than 75% of the blocks read or write less than 20 registers. The larger number is due to the larger block sizes for the FP benchmarks, for which the average block size is almost twice that of the integer benchmarks. Except for *parser*, all of the benchmarks have at least 10 temporaries in over 50% of the blocks, showing that a significant reduction in register bandwidth can be achieved by the internal renaming provided in this architecture.

3.2.2 Register fanout

Limited instruction encoding space constrains the number of consumers that can be specified. *Fanout* refers to the number of targets for which an instruction can produce data. Large-fanout instructions require extra *split* instructions to reach all consumers. Figure 7 shows the average fanout for each produced value in the various benchmarks. For all 6 benchmarks, the fanout is 1 for more than 60% of instructions and less than or equal to 2 for more than 80% of the instructions. This shows that, if we support at least two targets, only one-fifth of the producers require *split* instructions.

These results are consistent with those obtained by Franklin and Sohi [8]. Their work studied the liveness of registers in terms of number of instructions and the fanout for registers, but did not analyze register behavior at the block level.

3.3 Technology Evaluation

To evaluate the Grid Processor for technology scaling, we used the technology-independent area and delay estimates described by Gupta [9] and Agarwal [2]. We used minimum features, which include a 64-bit integer ALU and multiplier, 8-entry instruction and operand buffers, and an 8-bit ALU and buffers for the router at each node. We assumed that 50% of the chip area is reserved for the grid interconnect. We computed both the area occupied by a node at various

Technology (nm)	Clock Speed (GHz)	Grid Dimensions in 400 mm^2	Row Delay (cycles)	Grid Height Delay (cycles)
100	3.5	10x10	0.53	3.57
70	6.0	20x20	0.71	5.83
50	10.0	27x27	1.02	11.02
35	13.5	40x40	1.16	21.30

Table 4: Area and Delay estimates at various technologies

technologies and the number of nodes that can be accommodated on a 400 mm^2 chip.

Using node dimensions derived from the area estimates, we computed the wire delay between nodes in adjacent rows and the delay to traverse the entire height of the grid at SIA projected clock rates [17]. The results are summarized in Table 4. At all technologies, the wire delay between adjacent rows is close to a single cycle. This result shows that a grid with a fast local interconnect can be built at all technologies. As feature size shrinks, both the node density and delay across the height of the grid increases super-linearly. At technologies below 50 nm, the grid sizes are much larger than what is required to map a typical block and the express channels have prohibitively long delays. In such cases, the computation substrate could be further partitioned into multiple grids.

4 Related Work

There have been a number of related approaches preceding the Grid Processor. Dennis and Misunas proposed a static dataflow architecture with programs expressed in a Fortran-like dataflow language [6]. Arvind proposed the Tagged-Token Dataflow architecture with purely data-driven instruction scheduling for programs expressed in a dataflow language [3]. Culler proposed a hybrid dataflow execution model where programs are partitioned into code

blocks made up of instruction sequences called threads with dataflow execution between threads [5]. Our approach uses a conventional programming interface with dataflow execution for a limited window of instructions.

We take a hybrid approach between VLIW [7] and conventional superscalar architectures by statically scheduling the instructions using the compiler and dynamically issuing them. There have been other efforts to enhance dynamic execution in VLIW machines. Rau proposed a split-issue mechanism to separate register read and execute from write-back and a delay buffer to support dynamic scheduling for VLIW processors [14].

Others have looked at various naming mechanisms for values to reduce the register pressure and register file size. Smelyanskiy proposed Register Queues for allocating live values in software pipelined loops [18]. Llosa proposed register sacks, which are low bandwidth port-limited register files for allocating live values in pipelined loops [11]. Corporaal proposed an explicitly specified communication mechanism in Transport Triggered Architectures for general purpose computing [4].

Many researchers are exploring distributed or partitioned uniprocessor designs. Waingold proposed a distributed execution model with extensive compiler support in the RAW architecture [21]. The RAW architecture assumes a coarser-grain execution than does the Grid Processor, exploiting parallelism across multiple compiler-generated instruction streams. Ranganathan and Franklin described an empirical study of decentralized ILP execution models [13]. Sohi proposed Multiscalar processors where a single program is broken up into a collections of tasks. The tasks are distributed to a number of parallel processing units which reside within a processor complex [19]. Each of these units fetches and executes instructions belonging to its assigned task. Rotenberg proposed trace processors where several processing elements work on different traces of the program, passing data values using a common register bus [15]. Unlike the trace processor, the Grid Processor executes a trace in a fine-grain dataflow fashion and overlaps multiple traces on the same computation substrate.

Finally, Patt proposed a Block-Structured Instruction Set Architecture for increasing the fetch rate for wide issue machines where the atomic unit of execution is a block and not an instruction [10].

5 Conclusion

In this paper, we have proposed a technology-driven architecture that combines the advantages of compiler-scheduled instruction level parallelism with data-driven execution on a fast clocked grid of execution units. Instruction blocks are mapped onto the grid as single units of computation, amortizing the fetch and decode over a large number of instruc-

tions. Access to global storage elements such as register files is reduced by maintaining temporaries as transient values within the grid. Overheads can further be reduced by overlapping the execution of one instruction block with the fetch and mapping of the next.

Our initial evaluation indicates that existing programs are ripe to be mapped to this substrate. Typical block sizes range from 27 to 125 dynamically executed instructions, which we anticipate to be sufficiently large to amortize scheduling overheads. The number of input and output values required for a large fraction of the blocks is less than 10 in five of the six benchmarks, indicating that the amount of register file communication between blocks is small. The average number of temporary registers per block is larger, ranging from 10 to 30, depending on the benchmark. This range indicates that a substantial amount of communication to the centralized register file can be eliminated through the producer/consumer communication within the grid. Finally, the average number of consumers of a produced value is only 1.9, which shows that the network within the grid does not require large bandwidth for intra-block communication.

In addition to these direct performance advantages, the proposed Grid Processor provides several other benefits. It offers substantial power savings because much of the scheduling is performed at compile time and each execution unit is idle until all operands arrive. Furthermore, since particular blocks of the program, such as loops, can be mapped once and reused many times, the time and power required for block mapping can be reduced substantially. In conventional architectures, the instructions must be fetched repeatedly for each iteration of the loop. This *mapping reuse* may permit the Grid Processor to act as a high performance substrate for DSP codes as well. Finally, the data driven computation model on the Grid Processor is amenable to both polypath execution and selective re-execution upon mis-speculation. Conditionally executed instructions within the block can be started speculatively. If the speculation was incorrect, the block may be re-executed by loading correct values into the appropriate graph nodes and letting the values propagate to the data dependent downstream instructions. Thus no instructions need be refetched, and independent instructions do not need to be re-executed upon a mis-speculation rollback.

While this paper outlines the basic Grid Processor architecture and its potential for performance improvement, substantial challenges lie ahead in the complete design. The dataflow-style execution is amenable to both computation and memory parallelism. However, dynamic dependences between loads and stores must be detected to ensure proper ordering in the memory system. Dynamic dataflow execution with compiler-controlled static scheduling removes the need for synchronization in the execution substrate that conventional architectures enforce. The challenge lies in detecting when all of the instructions in the block have terminated

and architecturally visible storage can be committed. Finally, the traditional precise exception model in which an exception can occur at any point in the instruction stream is particularly challenging for the Grid Processor. Changing the granularity of rollback from an instruction to a block level may enable more efficient exception implementation.

Acknowledgements

Many thanks to the anonymous reviewers and the CART group members for their feedback on early versions of this paper. This work is supported by the National Science Foundation under CAREER awards CCR-9985109 and CCR-9984336, CISE Research Instrumentation grant EIA-9985991, University Partnership Awards from IBM, and a grant from the Intel Research Council.

References

- [1] V. Agarwal, M. S. Hrishikesh, S. W. Keckler, and D. Burger. Clock rate versus IPC: The end of the road for conventional microarchitectures. In *Proceedings of the 27th Annual International Symposium on Computer Architecture*, pages 248–259, June 2000.
- [2] V. Agarwal, S. W. Keckler, and D. Burger. Scaling of microarchitectural structures in future process technologies. Technical Report TR2000-02, Department of Computer Sciences, The University of Texas at Austin, Austin, TX, February 2000.
- [3] Arvind and R. S. Nikhil. Executing a program on the MIT Tagged-Token Dataflow Architecture. *IEEE Transactions on Computers*, 39(3):300–318, 1990.
- [4] H. Corporaal. *Transport Triggered Architectures*. PhD thesis, Delft University of Technology, September 1995.
- [5] D. E. Culler, A. Sah, K. E. Schauer, T. von Eicken, and J. Wawrzynek. Fine-grain parallelism with minimal hardware support: A compiler-controlled threaded abstract machine. In *Proceedings of the 4th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 164–175, April 1991.
- [6] J. Dennis and D. Misunas. A preliminary architecture for a basic data-flow processor. In *Proceedings of the 2nd Annual Symposium on Computer Architecture*, pages 126–132, January 1975.
- [7] J. Fisher. Very long instruction word architectures and the ELI-512. In *Proceedings of the Tenth Annual International Symposium on Computer Architecture*, pages 140–150, June 1983.
- [8] M. Franklin and G. Sohi. Register traffic analysis for streamlining inter-operation communication in fine grain parallel processors. In *Proceedings of the 25th Annual International Symposium on Microarchitecture*, pages 236–245, 1992.
- [9] S. Gupta, S. W. Keckler, and D. Burger. Technology independent area and delay estimates for microprocessors building blocks. Technical Report TR2000-01, Department of Computer Sciences, The University of Texas at Austin, Austin, TX, February 2000.
- [10] E. Hao, P. Chang, M. Evers, and Y. Patt. Increasing the instruction fetch rate via block-structured instruction set architectures. In *Proceedings of the 29th International Symposium on Microarchitecture*, pages 191–200, December 1996.
- [11] J. Llosa, M. Valero, J. Fortes, and E. Ayguade. Using sacks to organize register files in VLIW machines. In *CONPAR 94 - VAPP VI*, September 1994.
- [12] S. Palacharla, N. P. Jouppi, and J. E. Smith. Complexity-effective superscalar processors. In *Proceedings of the 24th Annual International Symposium on Computer Architecture*, pages 206–218, June 1997.
- [13] N. Ranganathan and M. Franklin. An empirical study of decentralized ILP execution models. In *8th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 272–281, October 1998.
- [14] B. Rau. Dynamically scheduled VLIW processors. In *Proceedings of the 26th Annual International Symposium on Microarchitecture*, pages 80–90, December 1993.
- [15] E. Rotenberg, Q. Jacobson, Y. Sazeides, and J. Smith. Trace processors. In *Proceedings of the 30th Annual International Symposium on Microarchitecture*, pages 138–148, December 1997.
- [16] S.A.Mahlke, D.C.Lin, W.Y.Chen, R.E.Hank, and R.A.Bringmann. Effective compiler support for predicated execution using the hyperblock. In *Proceedings of the 25th Annual International Symposium on Microarchitecture*, pages 45–54, June 1992.
- [17] The national technology roadmap for semiconductors. Semiconductor Industry Association, 1999.
- [18] M. Smelyanskiy, G. Tyson, and E. Davidson. Register queues: A new hardware/software approach to efficient software pipelining. In *International Conference on Parallel Architectures and Compilation Techniques (PACT 2000)*, October 2000.
- [19] G. Sohi, S. Breach, and T. Vijaykumar. Multiscalar processors. In *Proceedings of the 22nd International Symposium on Computer Architecture*, pages 414–425, June 1995.
- [20] Trimaran : An infrastructure for research in instruction-level parallelism. <http://www.trimaran.org>.
- [21] E. Waingold, M. Taylor, D. Srikrishna, V. Sarkar, W. Lee, V. Lee, J. Kim, M. Frank, P. Finch, R. Barua, J. Babb, S. Amarsinghe, and A. Agarwal. Baring it all to software: RAW machines. *IEEE Computer*, pages 86–93, September 1997.