

Balanced Scheduling:

Instruction scheduling
when memory latency is uncertain

Daniel R. Kerns

Department of Computer Science and Engineering
University of Washington

Technical Report 92-11-03
November 1992

Balanced Scheduling:

Instruction scheduling when memory latency is uncertain

Daniel R. Kerns
Department of Computer Science and Engineering, FR-35
University of Washington
Seattle, WA 98195

December 14, 1992

Abstract

Traditional list schedulers order instructions based on an optimistic estimate of the load delay imposed by the implementation. Therefore they cannot respond to variations in load latencies (due to cache hits or misses, congestion in the memory interconnect, etc.) and cannot easily be applied across different implementations. We have developed an alternative algorithm, known as balanced scheduling, that schedules instructions based on an estimate of the amount of instruction level parallelism in the program. Since scheduling decisions are program- rather than machine-based, balanced scheduling is unaffected by implementation changes. Since it is based on the amount of instruction level parallelism that a program can support, it can respond better to variations in load latencies. Performance improvements over a traditional list scheduler on a Fortran workload and simulating several different machine types (cache-based workstations, large parallel machines with a multipath interconnect and a combination, all with non-blocking processors) are quite good, averaging between 3% and 18%.

1 Introduction

Instruction schedulers for conventional machines generate code assuming a machine model in which load latencies are well-defined and fixed. Usually the latencies reflect the most optimistic execution situation, e.g., the time of a cache hit rather than a cache miss. Compiler optimizations intended to improve performance through instruction scheduling, such as reordering instructions to avoid pipeline stalls, insert independent instructions after loads to keep the CPU busy while memory references are in progress. The number of instructions inserted (in the best case) depends on this predefined latency value.

When a load reference exceeds the implementation-defined latency, the processor architecture generally stipulates that instruction execution be stalled. The advantage of this design (called *blocking* loads) is that it requires a simple and straightforward hardware implementation. The consequence for compiler technology is that the compiler does not have to consider multiple memory latencies during instruction scheduling.

Two architectural innovations make it worthwhile to reconsider how to schedule behind load instructions. The first is processor designs that do not stall on unsatisfied load references (called *nonblocking* loads)

through the use of lockup free caches[19, 20, 16, 13], multiple hardware contexts[2, 1] or an instruction lookahead scheme[2]. Nonblocking loads allow a processor to continue executing other instructions while a load is in progress. Although the design requires more complex hardware, more instruction level parallelism can be exploited, and therefore programs execute faster. The second innovation is machines that have a large variance in memory response time. These uncertain memory latencies may be due to congestion in a multipath interconnect or a hierarchy of memory, including both cache hierarchies and local and global memories.

Variable load instruction latencies, coupled with nonblocking loads, complicate scheduling, because the instruction scheduler does not know how many instructions to schedule after a load to maintain high processor utilization. If the memory reference is delayed beyond the scheduler's latency estimate, the processor will stall and processor utilization will drop. However, if the load latency is shorter than the estimate, the destination register of a load instruction will be tied up longer than necessary. This may increase register pressure enough to cause unnecessary spills to memory and a consequent increase in program execution time. In addition, an excessive number of instructions may migrate to the top of the schedule, leaving an insufficient number to hide load latencies near the bottom. In this case the CPU will also be needlessly idled.

In this report we present a code scheduling algorithm, called *balanced scheduling*, that has been specifically designed to tolerate a wide range of variance in load latency over the entire execution of a program. Balanced scheduling works within the context of a traditional list scheduler[9, 15, 23, 8, 6], but uses a new method for calculating load instruction weights. Rather than using weights that are determined by the implementation and therefore are fixed for all programs, the weight of each load is based on the amount of instruction level parallelism that is available to it. (We refer to this as *load level parallelism*.) This assignment is effective, since load instructions are scheduled for the maximum latency that can be sustained by the amount of load level parallelism in the code. In essence, our algorithm schedules for the code instead of scheduling for the machine. Looking at it another way, balanced scheduling amortizes the cost of incorrectly estimating actual load latencies over all load instructions in the program.

To validate the algorithm we compared the performance of several programs scheduled via balanced scheduling and a traditional list scheduler on a variety of processor and memory architectures. The processor models differed in their ability to exploit load level parallelism; each was coupled with three different memory systems, that exhibit dissimilar latency behavior. Both the balanced scheduler and the traditional scheduler were incorporated into the GCC[21] compiler and generated code for the Perfect Club benchmarks[4]. Performance improvements for balanced scheduling averaged 3% to 18% over the traditional list scheduler, for different processor and system model combinations.

The remainder of this report is organized as follows. Section 2 introduces balanced scheduling, and

section 3 describes the algorithm in more detail. Section 4 explains our experimental methodology; Section 5 presents the experimental results. Section 6 discusses extensions and other applications of the balanced scheduling algorithm. The conclusion follows in section 7.

2 Balanced Scheduling

The traditional approach to instruction scheduling that considers machine resource constraints is *list scheduling*[9, 15, 23, 8, 6]. The primary data structure used by list schedulers is the *code DAG*, in which nodes represent instructions and edges represent dependences between them. Each node is labeled with a weight reflecting the latency of the instruction.¹ At each iteration of its algorithm a list scheduler creates a ready list of instructions that are eligible for scheduling, i.e., those whose predecessors in the code DAG have been scheduled or have had their latencies met. A set of heuristics is then applied to decide which instruction from the ready list should be scheduled next; the heuristics used depend on the particular list scheduler. For example, Gibbons and Muchnick[8] first schedule the instruction with the greatest operation latency. If more than one instruction qualifies, their scheduler breaks the tie by choosing the instruction(s) with the greatest number of successors. The final heuristic picks the the instruction with the largest sum of the latencies along the longest path from the instruction node to a leaf node. Other styles of list schedulers include those that combine several levels of heuristics into a single weight and schedule in decreasing weight order[17, 24] and update scheduling weights dynamically[23]. Our heuristics are described in detail in Section 4.1.

If a processor exposes the variations in actual memory reference latency to the compiler through non-blocking load instructions, instruction scheduling becomes more complicated. Traditional list schedulers use a single constant for the weight of all load instructions, usually an implementation-defined latency (e.g., cache hit time). They then schedule instructions independent of that load until the load latency has been consumed. As expected, traditional schedulers work best when the actual latency of each load matches the predefined (and optimistic) value. When it does not, a longer latency (e.g., the time of a cache miss) penalizes the program by stalling the CPU. This fixed estimate of memory latency prevents the scheduler from hiding latencies larger than the nominal value. Therefore, when the optimistic execution scenario does not occur, performance suffers. The worst scheduling situation exists when the actual latencies change over time, for example, as congestion in the interconnect varies.

In contrast, the balanced scheduler computes load instruction weights based on a measure of instruction level parallelism in the code rather than on an implementation-defined value. This measure, which we call *load level parallelism*, defines the number of instructions that may execute in parallel with each load

¹Edges can also be labeled, allowing latencies to differ among successor nodes of a given node, as on the Intel i860.

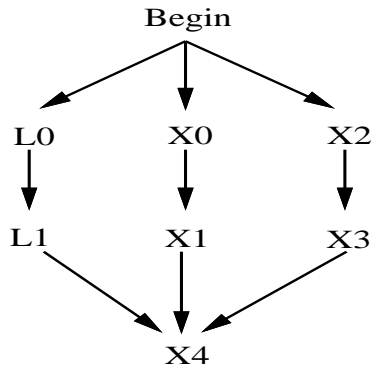


Figure 1: Code DAG of a hypothetical program.

Traditional Scheduling	Traditional Scheduling	Balanced Scheduling
W = 5	W = 1	
L0	L0	L0
X0	L1	X0
X1	X0	X1
X2	X1	L1
X3	X2	X2
L1	X3	X3
X4	X4	X4

(a) (b) (c)

Figure 2: Schedules generated from the code DAG in Figure 1, using the traditional and balanced schedulers. The traditional scheduler is illustrated with load instruction weights of 5 and 1, respectively.

instruction. The weight for each load is calculated separately, as a function of the number of instructions that may initiate execution during the load and the number of other loads that could also use them to hide latencies.

Both the balanced scheduling algorithm and the traditional scheduler operate on a basic block by basic block basis. The balanced scheduler simply incorporates the new method of computing weights for each load instruction into a traditional list scheduler.

Using the code DAG of Figure 1 as an example, Figure 2 illustrates the schedules generated by the traditional and the balanced schedulers. Nodes labeled L_n represent load instructions and nodes labeled X_n represent other non-load instructions of weight 1. The schedules in Figures 2a and 2b result from scheduling the graph of Figure 1 with a traditional scheduler, assuming load instruction weights of 5 and 1, respectively. These two schedules illustrate the effect of over- and under-estimating load instruction latency. In Figure 2a, if L1 incurs an actual latency greater than one, hardware interlocks will be inserted before X4. We say the scheduler is *greedy* in this case, because L0 captured all of the load level parallelism and left none for L1. The opposite situation occurs when load instruction weights are too small. Figure 2b illustrates the schedule produced when a weight of one is used. In this case we have not taken advantage of the load level parallelism with respect to L0. We say the scheduler was *lazy*, because it passed over opportunities for parallelism. Should the actual latency be greater than the scheduling assumption, the processor will needlessly stall. Figure 2c is the schedule that the balanced scheduler generates. The balanced scheduler has measured the load level parallelism in the DAG and determined that a weight of 3 assigned to each load instruction would generate an efficient schedule.

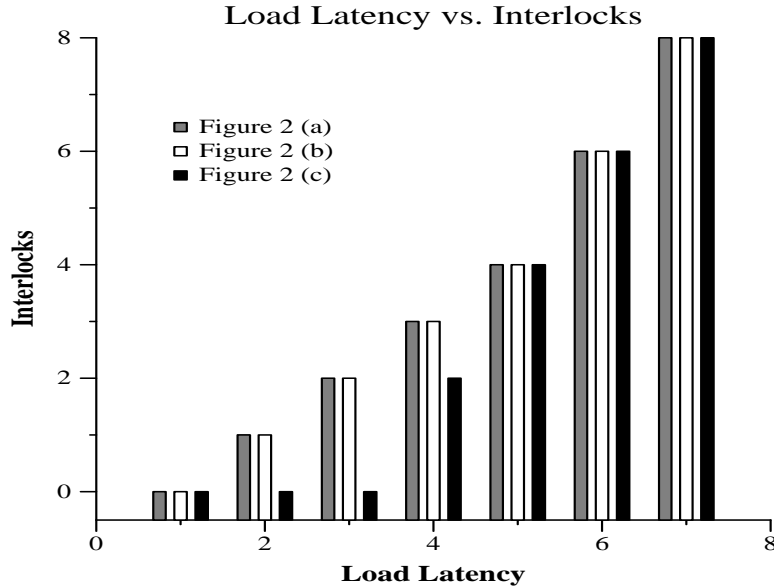


Figure 3: Interlocks generated from Figure 2 for various load latencies.

Figure 3 summarizes the number of interlocks that accrue when these schedules are executed with varying memory latencies. The chart shows that, for latencies in the range of 2–4, the balanced schedules are faster than both the greedy and lazy traditional schedules illustrated in Figure 2. Outside this range the balanced and traditional schedules perform equivalently.

In summary, balanced scheduling’s strength is its ability to look beyond fixed latencies, thereby exposing additional instruction level parallelism. Whereas traditional schedulers plan for the optimal latency, balanced schedulers make scheduling decisions based on the amount of load level parallelism the code can support. It therefore produces fewer interlocks when the optimal case doesn’t occur.

3 Balanced Scheduling Algorithm

This section presents the balanced scheduling algorithm. The algorithm is first illustrated through two simple examples. The examples depict the two relationships load instructions can have with each other, i.e., occurring in series and in parallel, and, therefore, the two cases the algorithm must handle.

When load instructions occur in series, the balanced scheduling algorithm equally distributes among them all instructions with which they can execute in parallel. Referring again to the code DAG of Figure 1, the two load instructions, L0 and L1, may execute independently of X0, X1, X2 and X3. Since L1 is dependent on L0, the obvious partitioning would schedule two instructions after L0 and two after L1. The weight on each load instruction is simply one (for the issue slot of the load), plus the number of instruction issue slots




Figure 4: Code DAG in which L0 and L1 are independent and execute in parallel with all other instructions.

Balanced Scheduling
L0
L1
X0
X1
X2
X3
X4

Figure 5: Balanced schedule generated from the code DAG in Figure 4.

that may be initiated independently of the load divided by the number of loads in series or, $1 + (4/2) = 3$. Issue slots are measured, because instruction weights represent the number of machine cycles that should pass before an instruction that uses the result of the load is initiated.

When load instructions are not dependent on each other, i.e., they occur in parallel, their latencies can be hidden, using instructions drawn from the same set. Referring to the code DAG in Figure 4, the balanced scheduling algorithm takes advantage of the fact that L0 and L1 can, and should, share the same set of padding instructions. In Figure 4 each load instruction may execute in parallel with five other instructions, so they are each assigned a weight of six ($1 + 5/1$). The final schedule is shown in Figure 5.

For a balanced scheduling algorithm to be successful, any combination of loads in series and loads in parallel must be accommodated.

A balanced scheduler operates by measuring load level parallelism and assigning weights accordingly. The algorithm, shown in Figure 6, examines each instruction i in the code DAG (G) and computes the set of instructions with which it may execute in parallel. It first eliminates from G those instructions that are predecessors or successors, recursively, producing G_{ind} (line 3). The resulting connected components of G_{ind} contain the sets of load instructions that may execute in parallel with i . Within each connected component, C , the path with the largest number of load instructions is located (lines 4–5). (We examine the longest load path, because loads on other paths can be overlapped with it.) Since the loads on this path execute in series, their sum (called *Chances*) represents the number of opportunities for scheduling i . Finally, the number of issue slots in the instruction execution pipeline that are required by i is divided by the number of loads in series ($IssueSlots(i)/Chances$), and is added to the accumulating weight of each load instruction in C (lines 6–7).

Term	Definition
G	the code DAG.
$\text{Pred}(i)$	the transitive closure of the predecessor function on node i .
$\text{Succ}(i)$	the transitive closure of the successor function on node i .
Chances	the maximum number of loads on any path in a connected component.
$\text{IssueSlots}(i)$	the number of issue slots in the instruction execution pipeline required by instruction i .

1. Initialize the latency of each load instruction to 1.
2. for each instruction i in G
3. $G_{ind} = G - (\text{Pred}(i) \cup \text{Succ}(i))$
4. for each connected component C in G_{ind}
5. Find the path with the maximum number of load instructions.
6. for each load instruction $l \in C$
7. add $\text{IssueSlots}(i)/\text{Chances}$ to the weight of l
8. end
9. end
10. end

Figure 6: Balanced scheduling algorithm

Figure 7a illustrates the balanced scheduling algorithm on a more challenging basic block. Using $i=X1$, step 4 generates the three connected components shown in Figure 7b. (L2 does not appear in a connected component because it is a predecessor of X1). The maximum path length in the component containing L1 is 1; therefore X1 contributes $1/1$ to L1's weight. The maximum path length in the second component is 3, and X1 contributes $1/3$ to the weights of each load instruction, L3, L4, L5 and L6. The third connected component has no load instructions. Table 1 shows the weight contributed by each instruction to each load at the completion of the algorithm. The latencies assigned to the five load instructions represent a distribution of load level parallelism that is representative of the load level parallelism in Figure 7.

If n is the number of nodes in the DAG, steps 4 and 5 together may be done in a worst case time of $O(n \alpha n)^2$, using the set union algorithm. First, each node in G_{ind} is labeled with its level from the farthest leaf. Next, it is combined with the nodes to which it is connected, using the set union function. Each time we perform set union, the set label is updated to reflect both the minimum and maximum level number that has been seen in that set. Therefore, the largest path length for each connected component is simply the maximum level number minus the minimum level number plus 1. Steps 6-7 are performed in $O(n)$ time and, therefore, do not impact the worst case time complexity. Connected component analysis is done for each instruction in the code DAG; therefore, the entire algorithm has a worst case time complexity of $O(n^2 \alpha n)$.

² α is the inverse Ackerman function. As a function of n , it increases very slowly and may be considered constant[22].

4 Experimental Methodology

We designed a series of experiments to compare balanced scheduling with a traditional scheduling approach. These experiments modeled the execution of real programs running on several different architectures. This section describes the methodology of these experiments. The integration of the balanced scheduler into the GCC compiler, the workload and the simulator we used for our measurements are described, in turn, in sections 4.1 through 4.3.

For our experiments we classify the target machine characteristics into two groups. The processor characteristics are those that control how the processor exploits parallelism with respect to load instructions. The system characteristics are the attributes of the memory system in a particular implementation. We used several alternatives for each model, to demonstrate that balanced scheduling works well on architectures that contribute to latency uncertainty in different ways. The processor and system models we used are described in sections 4.4 and 4.5.

4.1 Compiler

We modified the GNU GCC version 2.2.2 compiler[21] to perform balanced instruction scheduling. The default instruction scheduler within GCC was replaced by a new module that can schedule using either the traditional or balanced approaches. In addition, several changes were made to GCC to increase scheduling effectiveness and improve instruction level parallelism. The changes include alleviating the effect of dependences in spill code introduced by register allocation, our heuristics for picking instructions from the ready list (one of which helps control register pressure) and modifications to GCC's RTL intermediate language. Both schedulers take advantage of these modifications.

GCC performs instruction scheduling both before and after register allocation. Since register allocation may add spill code and/or copy instructions, the second scheduling pass serves to integrate these additional instructions into the final schedule. However, the effectiveness of the second scheduling pass is restricted because of dependences introduced by register allocation.

These false dependences negatively effect schedule performance in two ways. First, the final assignment of register numbers severely limits the code motion that a scheduler can perform. Second, when adding spill instructions, the GCC compiler always uses register numbers selected from a small pool of spill registers. The net effect is that spill code cannot be scheduled effectively with other instructions. We improve performance by increasing the size of GCC's spill register pool by two and implementing a FIFO queue-like ordering of the registers in the pool. An alternative approach would use software register renaming after register allocation to better integrate spill instructions.

As previously mentioned, both the balanced and traditional schedulers use the same list scheduler. Some list schedulers place instructions onto the ready list when all their predecessors in the code DAG have been scheduled. In contrast, our scheduler defers adding these instructions to the ready list until each predecessor has exhausted its expected latency. In the case of starvation the scheduler inserts virtual no-op's into the instruction stream. This delayed insertion of instructions into the ready list increases the accuracy of instruction placement within the schedule. Since our processors use the hardware interlock model of execution, the virtual no-ops are removed before actual code generation.

List schedulers select instructions from the ready list in priority order. In our case, the priority of an instruction is equal to its weight plus the maximum priority among its successors. In the event of ties we select instructions using alternate heuristics in the following order. The first selects the instruction that has the largest difference between consumed and defined registers; this heuristic helps control register pressure. The second ranks instructions based on the number of successors in the code DAG that would be exposed for scheduling if that instruction were to be selected; it gives the list scheduler more instructions from which to select. The final heuristic selects the instruction that was generated the earliest. Our list scheduler is a bottom-up scheduler, therefore we generate schedules in reverse order by scheduling from the leaves of the Code DAG toward the roots.

The compiler has been configured for the MIPS RISC processor[12]. GCC's intermediate language, RTL, is not sufficiently RISC-like for an instruction scheduler to get maximum benefit, since some primitive operations in RTL are actually multi-cycle macros. In the context of this work, memory-to-memory copies are the most notable, since it is load instructions that we are concentrating on scheduling. Our implementation extracts GCC's intermediate language after optimization but before register allocation and modifies it to replace certain non-RISC patterns, such as memory-to-memory copy, with their RISC equivalents. The modified RTL is at a lower level and therefore more suitable for instruction scheduling.

Loop unrolling is an optimization that increases instruction level parallelism. Due to a conflict with the way we use profiling information (section 4.3), GCC's unrolling capability is not usable for these experiments. Therefore, unrolling was performed manually.

4.2 Workload

The workload consisted of the Perfect Club suite of benchmarks[4]. (See Table 4.2) Since these programs are written in FORTRAN, they were converted to C using *f2c*[7]. The Fortran-to-C converter produces C programs that correctly represent the semantics of the original FORTRAN programs. However, these C programs are conservative translations: after being compiled by a C compiler, they will most likely execute more slowly than if they were compiled by a FORTRAN compiler. For example, since almost all data is referenced through

Name	Description
ADM	A 3D pollution concentration simulation evaluating systems of hydrodynamic equations.
ARC2D	Fluid flow modeling supersonic reentry.
BDNA	Uses the BIOMOL package to model the molecular dynamics of biomolecules in water.
FLO52Q	Analyses transonic inviscid flow past an airfoil by solving unsteady Euler equations.
MDG	Molecular dynamics calculation of water molecules in the liquid state at room temperature and pressure.
MG3D	Seismic migration code using FFT's to model below surface structures.
QCD2	Lattice gauge QCD simulation.
TRACK	Determines the course of a set of an unknown number of targets, such as rocket boosters, from observations.

Table 2: The benchmarks in the “Perfect Club.”

```

float func(a, b)
    float *a, *b;
    {
        a[1] = b[2];
        a[2] = b[3];
    }
float a[HUGE], b[HUGE];
float new_func(_a, _b)
    float *_a, *_b;
    {
        a[1] = b[2];
        a[2] = b[3];
    }

```

Figure 8: Example *f2c* program showing the disambiguation problem and our transformation. In `func` the load of `b[3]` must be considered dependent on the store of `a[1]`. Our transformation results in `new_func`. The resulting program produces incorrect results, but accurately models the code that would be generated by a FORTRAN compiler.

pointers in the C program, it is nearly impossible for a C compiler to do the memory reference disambiguation that might be obvious to a FORTRAN compiler. Instruction scheduling is effected, because load instructions are not free to move above stores. Since this problem severely restricts a scheduler’s ability to exploit load level parallelism, we apply a transformation which more correctly models the dependences in the FORTRAN program and increases the available parallelism.

The FORTRAN standard[3] specifically disallows aliasing among dummy arguments (formal parameters) if there will be any stores to the dummy arguments. If the function `func` in Figure 8 were produced by *f2c*, the FORTRAN standard would assume that array `a` and array `b` were disjoint; therefore the load for `b[3]` could be scheduled before the store of `a[1]`. However, the C semantics for `func` insert a true data dependence between the store of `a[1]` and the load of `b[3]`. This dependence is an artifact of the Fortran-to-C translation and does not exist in the original program.

Our compiler takes advantage of the FORTRAN semantics by performing a parallelism-exposing transformation on the input C programs. The transformation would replace `func` with `new_func`, as illustrated in Figure 8. New global variables are inserted with the same names as the original subroutine parameters. The

formal parameters are replaced with names that are never referenced. The program is no longer semantically correct, but the compiler is now able to correctly model the FORTRAN independence between references to array `a` and array `b`. The net effect is the generation of code that is comparable to that generated by a FORTRAN compiler. This transformation is a conservative representation of the data dependences that a FORTRAN compiler could discover, since FORTRAN is quite specific about when aliasing may occur.

4.3 Simulator

After the second scheduling pass, the machine instructions are extracted and run through an instruction level simulator. Given a particular model for load instruction latencies (explained in section 4.5), the simulator simulates instruction issue and completion for each basic block and computes its execution time in cycles.

As the simulator encounters load instructions, it draws latency samples from a random distribution that represents the system-level characteristics being modeled (see Section 4.5). The output of the simulator is one sample of the number of instruction and interlock cycles that comprise the execution time of the program on the modeled system. Because the results of the simulation are based on an independent and identically distributed random variable, we can take several steps to both reduce the execution time of the simulation and improve the quality of the results.

We have chosen to execute the full instruction-by-instruction simulation 30 times with new random numbers on each iteration. The number 30 represents an arbitrary choice which is large enough to avoid statistical noise.

Second, we measure the accuracy of our results by generating confidence intervals. Confidence intervals are computed for percentage improvement using a bootstrapping[5] procedure. From the 30 sample runtimes, we randomly draw 30 samples, with replacement, in order to generate a second sample mean. This process is repeated until we have 100 sample means for the block. These 100 sample mean runtimes are scaled by the profiled execution frequency to compute the actual runtime of the block. The sample means for each block are summed giving 100 sample runtimes for the entire program. The mean runtime reported is the mean of the 100 sample mean runtimes.

In order to report a percentage improvement for balanced scheduling, the 100 sample means from the balanced scheduler are paired with an equal number from the traditional scheduler, and the calculation is performed. After sorting, a 95% confidence interval is directly extracted.

4.4 Processor-level model

Processor-level attributes model a processor's ability to exploit load level parallelism. We model three different configurations. The first is unrealistically aggressive and serves as a best case reference. The

second two are restricted in ways that make them implementable. All of our processor models are assumed to maintain store/load consistency, i.e., if a load instruction follows a store, and they reference the same address, the load instruction receives the data that was written by the store instruction.

The first processor model (called UNLIMITED) can dispatch non-blocking load instructions with no limit on the number of loads outstanding. This model is similar to theoretical dataflow machines[10]. It is of interest because it exposes the maximum benefit that processor parallelism can achieve. The second (called MAX-8) allows a maximum of eight load instructions to be simultaneously executing. If a ninth load instruction is issued, the processor blocks until one of the eight outstanding loads completes. The third processor model (called LEN-8) restricts the maximum number of cycles a load instruction can take before blocking, as in the Tera Computer[2].³ In this model, if a load instruction has been outstanding for eight cycles, the processor blocks until the data is returned.

The balanced scheduler has not been specifically configured for any of the processor models. In particular, it may schedule more than eight load instructions before using loaded data (as is prohibited in MAX-8), and it might assign load instructions weights greater than eight (not effective in LEN-8). If this information were available to the compiler, the results for MAX-8 and LEN-8 would improve. We used a processor-independent version of balanced scheduling to demonstrate that a code scheduling approach that was not associated with a particular implementation, but instead was based solely on program characteristics, such as the amount of load level parallelism, would generate efficient code.

4.5 System-level model

Three memory systems are modeled and simulated, representing different latency behavior in both current and future architectures. The first has a data cache. A load instruction's data is returned after 2 cycles on a cache hit and either 5 or 10 cycles on a cache miss. The model represents a typical workstation-class RISC processor that implements nonblocking load instructions, such as the Motorola 88000 series[16]. It is simulated with cache hit rates of 80% and 95%, modeling first level caches of 4K and 32K bytes, respectively[11]. Four configurations are modeled, and are referred to as $Lhr(hl,ml)$, where Lhr stands for lockup-free caches with a hit rate of hr , and hl and ml are hit and miss latencies, respectively.

The second model has a memory interconnection network and no cache. The interconnection scheme uses a hashing function to assign addresses to memory modules, effectively randomizing memory access locations. In this architecture, memory latencies modeled by one of two zero-based probability mass functions, depicting normal distributions with standard deviations of 2 or 5. A standard deviation of 2 represents a machine

³The Tera restricts the number of instructions rather than cycles; since we assume that instructions other than loads execute in a single cycle, the two are equivalent.

in a relatively stable state (uniform network load, low to medium uncertainty). A standard deviation of 5 represents one with unpredictable memory latencies (changing network load, high uncertainty). The network machine is modeled in seven different configurations. Each distribution is combined with a mean of 2, 3 or 5, representing different base load levels. In a multithreaded processor such as the Tera, the different means are related to the number of active threads; the more threads, the lower the mean memory access time. We refer to these models as $N(\mu, \sigma)$ where μ is the mean of the distribution and σ is the standard deviation. All six configurations are reasonable design points for the machine. A seventh configuration models an unbalanced system, with a mean access time of 30 cycles and a standard deviation of 5 ($N(30, 5)$). Although we recognize that a compiler would not likely generate code specifically for such an unbalanced configuration, we include it in order to gauge balanced scheduling’s ability to handle a workload that has too little load level parallelism to hide the average latency.

The third machine has both a data cache and a Tera-style memory interconnection network. A cache hit occurs 80% of the time and takes two cycles. A cache miss is represented by a normal distribution with a mean of 30 and a standard deviation of 5. This configuration is referred to as L80- $N(30, 5)$ and has a mean latency of 7.6. In this case the 30 cycle latency is a reasonable design point, since the cache satisfies most requests. The model is intended to be representative of Alewife-like systems[1], where a commodity processor might be incorporated into a shared memory machine.

5 Experimental Results

The first set of results is the percentage improvement in execution time of the balanced scheduler over the traditional scheduler (positive values indicate an improvement due to balanced scheduling). These results appear in Tables 3–5, one for each processor model. For these experiments, the traditional scheduler uses load latencies equal to the cache hit time or effective access time for models with caches and the mean of the normal distribution for models without caches (labeled Optimistic Latency in the table). The percentage improvement of balanced scheduling over traditional scheduling is quite good. The average decrease in execution time for the UNLIMITED model varies from 3 to 18 percent for individual system models, with a mean improvement of 9.9%. The results for MAX-8 and LEN 8 are similar, with ranges of 7% to 16% and 3% to 16%, and means of 10.0% and 8.7%, respectively. These results demonstrate that balanced scheduling works well for several architectures, each of which contributes to latency uncertainty in a different way. It is important to emphasize that the balanced scheduler has not been customized for the restricted processors; these results represent the improvement from a machine-independent scheduler and would be better if the processor dependences were taken into account.

Processor model: UNLIMITED — Unlimited loads											
System	Optimistic Latency	Percentage improvement from balanced scheduling									
		ADM	ARC2D	BDNA	FLO52Q	MDG	MG3D	QCD2	TRACK	Mean	
Data cache; bus-based interconnection											
L80(2,5)	2	5.8	6.7	6.0	4.9	9.8	7.0	19.3	7.2	8.3	
	2.6	4.0	6.2	5.2	3.6	8.7	6.2	18.6	2.6	6.9	
L80(2,10)	2	9.9	13.1	10.6	8.7	14.4	11.9	27.8	6.7	12.9	
	3.6	7.5	11.7	8.1	6.7	11.6	10.7	25.8	2.2	10.5	
L95(2,5)	2	3.4	3.9	4.4	2.8	6.9	3.7	16.9	6.1	6.0	
	2.2	2.1	4.0	4.0	2.1	6.2	3.9	16.2	2.0	5.1	
L95(2,10)	2	4.6	5.8	5.8	3.9	8.0	5.4	19.9	4.9	7.3	
	2.4	3.2	6.1	5.6	3.8	7.2	5.8	19.0	1.7	6.6	
No cache; network interconnection											
N(2,2)	2	8.0	9.3	8.0	6.5	11.3	9.2	21.3	9.4	10.4	
N(3,2)	3	6.4	8.9	4.0	5.0	12.0	8.6	22.7	3.5	8.9	
N(5,2)	5	4.8	5.5	3.4	3.6	13.5	6.5	20.0	3.9	7.7	
N(2,5)	2	14.2	17.7	14.4	11.9	20.9	16.1	32.7	16.6	18.1	
N(3,5)	3	11.5	18.2	10.8	10.9	20.0	14.9	35.9	3.9	15.8	
N(5,5)	5	9.2	12.0	9.3	7.6	18.3	10.3	27.8	4.9	12.4	
N(30,5)	30	-3.5	-5.0	1.9	4.1	19.3	-0.9	7.1	0.6	3.0	
Mixed											
L80-N(30,5)	2	12.4	20.4	15.8	12.7	20.0	13.3	39.6	11.3	18.2	
	7.6	7.0	9.3	18.4	6.3	14.3	4.5	19.4	-2.5	9.6	

Table 3: Percent improvement in execution time from simulations using processor model UNLIMITED

Processor model: MAX 8 — Maximum of eight outstanding loads											
System	Optimistic Latency	Percentage improvement from balanced scheduling									
		ADM	ARC2D	BDNA	FLO52Q	MDG	MG3D	QCD2	TRACK	Mean	
Data cache; bus-based interconnection											
L80(2,5)	2	6.5	8.9	8.6	5.6	7.8	5.9	23.1	8.7	9.4	
	2.6	4.5	9.0	9.8	4.2	7.4	6.6	21.8	2.6	8.2	
L80(2,10)	2	9.8	13.4	11.2	9.0	10.8	7.8	26.2	7.2	11.9	
	3.6	7.7	12.5	6.8	6.6	8.9	8.3	25.2	2.3	9.8	
L95(2,5)	2	4.1	5.8	7.1	3.8	6.0	3.9	20.4	6.6	7.2	
	2.15	2.9	6.9	8.7	3.2	6.0	5.2	19.4	2.9	6.9	
L95(2,10)	2	5.4	7.3	8.0	5.0	6.4	4.8	21.6	7.5	8.3	
	2.4	4.2	9.0	10.2	4.9	6.8	5.9	20.6	3.2	8.1	
No cache; network interconnection											
N(2,2)	2	8.2	10.9	9.6	7.3	10.2	7.7	24.4	9.2	10.9	
N(3,2)	3	6.5	10.8	4.5	6.1	10.0	7.7	24.1	5.3	9.4	
N(5,2)	5	6.1	10.5	4.6	4.7	10.8	6.3	24.1	5.2	9.0	
N(2,5)	2	13.4	17.8	13.9	11.1	15.9	10.9	30.0	13.2	15.8	
N(3,5)	3	10.6	16.8	6.0	10.4	14.5	10.8	30.1	6.5	13.2	
N(5,5)	5	8.9	14.3	7.4	7.7	13.5	7.8	25.9	5.5	11.4	
N(30,5)	30	-1.7	-3.9	12.5	1.4	14.5	1.5	23.4	4.7	6.6	
Mixed											
L80-N(30,5)	2	10.6	17.4	12.3	11.2	13.6	9.4	30.4	11.9	14.6	
	7.6	5.2	11.3	8.8	7.4	11.5	5.1	19.8	1.5	8.9	

Table 4: Percent improvement in execution time from simulations using processor model MAX 8

Processor model: LEN 8 — Maximum of eight cycles for loads											
System	Optimistic Latency	Percentage improvement from balanced scheduling									
		ADM	ARC2D	BDNA	FLO52Q	MDG	MG3D	QCD2	TRACK	Mean	
Data cache; bus-based interconnection											
L80(2,5)	2	5.6	6.6	6.1	4.6	9.6	6.8	19.2	7.2	8.2	
	2.6	3.8	6.3	5.1	3.5	8.2	6.1	18.4	2.8	6.8	
L80(2,10)	2	8.5	11.2	8.5	7.7	14.6	10.8	24.9	10.8	12.1	
	3.6	6.3	9.5	5.9	5.2	11.6	9.6	24.4	4.6	9.6	
L95(2,5)	2	3.3	3.9	4.2	2.7	6.8	3.7	16.9	6.0	5.9	
	2.15	2.0	4.1	3.8	2.0	6.2	3.9	15.9	2.0	5.0	
L95(2,10)	2	4.5	5.7	5.0	3.5	8.0	4.8	18.6	7.6	7.2	
	2.4	3.5	5.7	4.6	2.7	7.2	5.5	18.2	4.1	6.4	
No cache; network interconnection											
N(2,2)	2	7.6	9.2	7.9	6.2	11.2	9.0	21.2	8.0	10.0	
N(3,2)	3	6.1	8.7	3.9	5.0	12.0	8.8	23.1	5.5	9.1	
N(5,2)	5	4.6	5.3	2.1	3.4	12.9	6.1	19.0	5.9	7.4	
N(2,5)	2	12.8	15.8	11.5	10.1	19.2	13.6	29.7	10.9	15.5	
N(3,5)	3	10.1	14.6	7.4	9.6	17.4	12.2	30.9	3.8	13.3	
N(5,5)	5	6.5	8.1	3.3	5.7	16.4	7.3	21.6	2.0	8.9	
N(30,5)	30	-2.7	-4.7	2.5	3.7	18.7	-3.8	6.3	2.3	2.8	
Mixed											
L80-N(30,5)	2	9.7	11.9	9.7	10.3	13.9	9.5	27.6	17.5	13.8	
	7.6	2.4	0.7	2.9	3.5	10.4	0.9	10.9	11.0	5.3	

Table 5: Percent improvement in execution time from simulations using processor model LEN 8

Label	Definition
TIns	The number of instructions executed, in millions, for the traditional scheduler.
BIns	The number of instructions executed, in millions, for the balanced scheduler.
TI%	The percentage of cycles which were interlock cycles for the traditional scheduler.
BI%	The percentage of cycles which were interlock cycles for the balanced scheduler.
Imp%	The percentage improvement of the balanced scheduler over the traditional scheduler.

Table 6: Index of labels used in figures.

The balanced scheduler does relatively better (over the traditional scheduler) as the uncertainty of the load instruction latencies increases. This can be seen in three different situations: when the cache hit rate is low (L80 vs. L95); when the cache miss penalty is high (L80(2,10) vs. L80(2,5) and L95(2,10) vs. L95(2,5)); and when the standard deviation of the normal is high (N(2,5) vs. N(2,2), etc.).

To better understand the reasons for the performance improvements, we did a component analysis of the execution times. All of our instructions execute in a single cycle; therefore the runtime of a program is the sum of the number of instructions executed and the number of interlocks incurred. Table 7 presents interlock information on the performance of one of the benchmarks, MDG. (Table 6 explains some of the column headings used in the remaining tables.) In this table, the percentage of the total number of cycles that were interlock cycles is reported for both the traditional and balanced schedulers. MDG’s performance gain with balanced scheduling (and also that of the other programs) is a result of both executing fewer instructions ($BIns < TIns$) and incurring fewer interlocks ($BI\% < TI\%$).

Balanced schedules often execute fewer instructions because their schedules contain less spill code. Table 8 presents data on the percentage of total instructions executed that was classified as spill code. (A spill instruction is defined to be any instruction that is inserted by the register allocator.) Balanced scheduling incurred fewer spills than the traditional scheduler for virtually all implementation-defined latencies on all programs. (The sole exceptions were ARC2D with an optimistic latency of 30 cycles and FLO52Q with 3.6 cycles.)

We hypothesize that the reduction in interlocks and spill code when using the balanced scheduler is a direct consequence of its always considering load level parallelism when calculating latency weights. It measures the parallelism, and, whether it is high or low, tries to use it to hide all load latencies in a basic block.

When there is significant load level parallelism, code DAGs tend to be bushy, causing all list schedulers to schedule independent instructions in parallel. The balanced scheduler manages this by assigning load instruction weights in such a way that load latencies are hidden by the other instructions. Register pressure

Program: MDG (BIns = 5,144 million)											
System	Optimistic Latency	TIns	UNLIMITED			MAX 8			LEN 8		
			Imp%	TI%	BI%	Imp%	TI%	BI%	Imp%	TI%	BI%
Data cache; bus-based interconnection											
L80(2,5)	2	5,358	9.8	10.4	5.6	7.8	13.9	10.9	9.6	10.4	5.7
	2.6	5,351	8.7	9.6		7.4	13.7		8.2	9.3	
L80(2,10)	2	5,358	14.4	21.6	13.6	10.8	25.2	20.6	14.6	22.5	14.7
	3.6	5,299	11.6	20.2		8.9	24.7		11.6	21.2	
L95(2,5)	2	5,358	6.9	5.9	3.4	6.0	8.8	7.1	6.8	5.8	3.4
	2.15	5,351	6.2	5.5		6.0	8.9		6.2	5.3	
L95(2,10)	2	5,358	8.0	9.4	6.1	6.4	12.1	10.2	8.0	9.9	6.6
	2.4	5,351	7.2	8.9		6.8	12.5		7.2	9.4	
No cache; network interconnection											
N(2,2)	2	5,358	11.3	12.8	6.9	10.2	16.6	11.8	11.2	13.1	7.2
N(3,2)	3	5,351	12.0	16.1	9.7	10.0	21.1	16.5	12.0	16.0	9.5
N(5,2)	5	5,297	13.5	24.4	16.8	10.8	32.1	27.0	12.9	24.3	17.0
N(2,5)	2	5,358	20.9	30.0	18.7	15.9	35.6	28.3	19.2	30.4	20.4
N(3,5)	3	5,351	20.0	31.8	21.3	14.5	37.2	30.9	17.4	31.9	23.0
N(5,5)	5	5,297	18.3	35.5	25.9	13.5	42.3	36.4	16.4	36.3	28.0
N(30,5)	30	5,393	19.3	71.8	67.9	14.5	79.4	77.5	18.7	73.1	69.6
Mixed											
L80-N(30,5)	2	5,358	20.0	49.9	42.3	13.6	52.3	47.9	13.9	49.5	44.7
	7.6	5,405	14.3	46.9		11.5	50.9		10.4	47.4	

Table 7: Detailed analysis of performance in MDG

Program	BIns	Percentage of Spill Instructions									
		Balanced Scheduler	Traditional Scheduler with Optimistic Latency of								
			2	2.15	2.4	2.6	3	3.6	5	7.6	30
ADM	2,494	7.43	9.59	9.15	9.15	9.15	9.22	9.42	9.50	8.70	7.49
ARC2D	11,149	10.47	13.52	13.74	13.74	13.68	13.27	13.46	13.89	12.25	10.11
BDNA	2,391	22.84	26.50	26.32	26.32	26.32	24.17	24.94	24.68	24.73	25.54
FLO52Q	3,323	4.61	7.14	6.82	6.82	6.82	6.97	3.91	6.55	5.89	4.90
MDG	5,144	7.49	7.86	8.04	8.04	8.04	8.04	8.13	8.00	8.86	9.21
MG3D	60,784	7.38	9.73	10.36	10.36	10.36	10.36	10.86	10.36	8.85	7.88
QCD2	1,176	19.91	29.30	28.92	28.92	28.92	28.92	28.78	28.02	26.89	28.02
TRACK	398	15.78	20.41	17.85	17.85	17.85	17.85	17.85	17.84	17.45	17.46

Table 8: Spill Instructions Executed

Program	TIns	BIns	UNLIMITED			MAX 8			LEN 8		
			Imp%	TI%	BI%	Imp%	TI%	BI%	Imp%	TI%	BI%
ADM	2,496	2,494	-3.5	67.8	69.0	-1.7	76.1	76.5	-2.7	69.9	70.7
ARC2D	11,108	11,149	-5.0	67.3	68.9	-3.9	78.4	79.1	-4.7	70.9	72.1
BDNA	2,478	2,391	1.9	65.0	65.6	12.5	85.3	84.0	2.5	71.1	71.4
FLO52Q	3,332	3,323	4.1	67.0	65.7	1.4	76.6	76.4	3.7	69.0	67.9
MDG	5,393	5,144	19.3	71.8	67.9	14.5	79.4	77.5	18.7	73.1	69.6
MG3D	61,116	60,784	-0.9	63.1	63.7	1.5	86.6	86.5	-3.8	67.4	68.8
QCD2	1,270	1,176	7.1	69.0	69.2	23.4	86.4	84.5	6.3	72.2	72.6
TRACK	406	398	0.6	81.6	81.9	4.7	85.6	85.2	2.3	82.3	82.3

Table 9: Analysis of N(30,5) results — the effect of spill code.

is unchanged, but interlocks go down. Traditional schedulers lack the guidance for efficient load placement. Therefore they incur similar register pressure, but also more interlocks.

When there is little load level parallelism, traditional schedulers greedily let independent instructions float to one end of the basic block. Therefore they incur spills at that end, and interlocks at the other. In contrast, the balanced scheduler spreads out the few independent instructions behind all loads. In all cases uses quickly follow definitions, and little or no spill code is generated. If the load level parallelism is less than the latency assumed by the traditional scheduler, balanced scheduling generates fewer spill instructions than the traditional technique.

In both situations (high and low load level parallelism) balanced scheduling contributes either little additional or less register pressure. When actual latencies differ from the optimistic latency, balanced scheduling incurs fewer interlocks; when both latencies are equal, the number of interlocks produced by the two schedulers is similar.

When load latencies are much larger than the amount of load level parallelism and therefore cannot be hidden via instruction scheduling, there is no guarantee the balanced scheduler will do better. In this case, register pressure can be a problem, and balanced scheduling can insert more spill code than the traditional scheduler. The situation is illustrated in Table 9, which summarizes the results for the N(30,5) model. This model assumes a mean latency much larger than the amount of load level parallelism of the programs in our workload. Two interrelated factors contribute to balanced scheduling’s poor performance with this model. First, as latencies get long, interlocks account for an increasingly large proportion of execution time. Both schedulers do poorly, and often equally poorly (for example, see TRACK). Second, a consequence of long load latencies is that each load instruction consumes more cycles relative to other instructions, and its contribution to execution time is greater. Therefore whichever scheduler generates more spill loads will have the poorer performance. Occasionally balanced scheduling chooses load instruction weights that cause higher than necessary register pressure and consequently issues more spill instructions (for example, see ARC2D).

The complete results generated from these experiments are reproduced in Appendix A.

In summary, these results show that balanced scheduling reduces execution time relative to traditional list scheduling in most cases. Because its schedules are based on the amount of load level parallelism that a program can support, they cause fewer interlocks during program execution and contain less spill code. The benefits are most apparent when memory latency uncertainty is high, as evidenced by greater miss rates and penalties, and larger standard deviations from mean latencies.

6 Future Work

Balanced scheduling has been presented in a specific form (weights calculated based on load level parallelism) to solve a specific problem (scheduling with uncertain load instruction latencies). The technique is applicable to a wider set of problems, such as other multi-cycle instructions (e.g., floating point operations coupled with asynchronous floating point units), disabling balanced scheduling when the latency is known (e.g., for the second access to a cache line), techniques that enlarge basic blocks (trace scheduling and software pipelining) and superscalar architectures. The remainder of this section discusses future improvements on this study and applications of balanced scheduling.

This study assumes unit latencies for all nonload instructions. A study should be done to measure the effect of multicycle instructions, such as floating point multiply, on our results. We expect this to reduce the usable load level parallelism (due to the possible introduction of no-ops) and therefore diminish the magnitude of our performance increases.

Balanced scheduling can be applied to other types of uncertain latencies, for example, those of an asynchronous divide unit. In this case, the algorithm would measure instruction level parallelism with respect to divides and schedule accordingly. As another example, consider machines with asynchronous floating point co-processors. In these machines the co-processor's instruction latencies might be known. However, if the latencies are long enough, and if the processor stalls while the co-processor is busy, the co-processor instructions might appear to have completed in less than the expected amount of time. Hence, from the point-of-view of the CPU, uncertain latency exists in the co-processor. Some modification to the algorithm would need to occur to better identify the instances where this could be profitable.

Another processor variation might have multiple classes of uncertain latency. Our algorithm could be modified for this type of machine by computing a *Chances* value for each type of instruction that exhibited uncertain latency and then repeatedly executing the loop in step 7, once for each value of *Chances*. This is possible because what counts is instruction initiation, not instruction latency. This modification does not change the worst case time complexity of the algorithm.

Some loads might not exhibit uncertain latencies. For example, if X and Y reside within in the same cache line, and X has been referenced, Y will hit in the cache. For programs with regular memory access patterns, a large fraction of the instructions might fall into this class. Balanced scheduling could use disambiguation information to recognize this case and then remove the load instructions that have certain latency from consideration in steps 6 and 7 of the algorithm.

Our algorithm does not take advantage of load instructions whose first use is outside the basic block containing the load instruction. Some work, namely Trace Scheduling[6], directly addresses load instructions whose first use is in future blocks by creating large basic blocks containing instructions from many blocks. Since trace scheduling uses list scheduling at the lowest level, our work is compatible with a trace scheduler. Another technique, software pipelining[18, 14], actively generates schedules where the first use of the register defined by a load instruction can be along cyclic edges in the flow graph. In particular, it can be in a future iteration of the same block with the load instruction. In that case, the first use precedes the load instruction in the code DAG. A technique such as Palem and Simons'[17] or Woodard's[24] might provide some of the benefits of balancing in the software pipelined environment.

Superscalar processors are fully supported by the balanced scheduling algorithm as presented here. The fact that a superscalar processor has the ability to initiate more than one instruction per cycle changes the number of issue slots an instruction requires by a constant factor. However, if the processor has some restriction on the types of instructions that can simultaneously be issued, an additional heuristic might improve code quality.

7 Conclusion

This paper describes an instruction scheduling algorithm, called balanced scheduling, that is appropriate for computers that expose uncertain memory latencies. Balanced scheduling is fundamentally different from previous list schedulers in two respects. First, it ignores the optimistic, implementation-determined memory latency when assigning scheduling priorities, basing them instead on the amount of parallel execution that is achievable in the program. Second, it computes individual scheduling weights for each load instruction separately, rather than using a single value for all loads in a basic block. Balanced scheduling thus insulates program execution from machine uncertainties by generating schedules that are optimized for the program rather than the machine.

To validate the algorithm we incorporated balanced scheduling into the GCC compiler and compared the performance of the Perfect Club benchmarks scheduled with both balanced scheduling and a traditional list scheduler. Three processors were modeled, representing machines with varying abilities to exploit instruction

level parallelism. Each of the processor models was coupled with several memory systems that exhibit dissimilar latency behavior. Execution time reductions of balanced scheduling over the traditional list scheduler averaged between 3% and 18%, depending on the processor model, system model and program. The results demonstrate that, if the capability to exploit uncertain memory latency is architected in future machines, balanced schedulers can effectively take advantage of the additional flexibility to generate faster schedules.

References

- [1] Anant Agarwal, Beng-Hong Lim, David Kranz, and John Kubiawicz. APRIL: A processor architecture for multiprocessing. In *Proceedings of the 17th Annual International Symposium on Computer Architecture*, pages 104–114. IEEE, May 1990.
- [2] Robert Alverson, David Callahan, Daniel Cummings, Brian Koblenz, Allan Porterfield, and Burton Smith. The Tera Computer System. In *1990 International Conference on Supercomputing*, pages 1–6. The ACM Press/SIGARCH, June 1990.
- [3] ANS X3.9-1978. *American National Standard Programming language FORTRAN*. American National Standards Institute, New York, 1978.
- [4] M. Berry, D. Chen, D. Kuck, S. Lo, Y. Pang, L. Pointer, R. Roloff, A. Samah, E. Clementi, S. Chin, D. Schneider, G. Fox, P. Messina, D. Walker, C. Hsiung, J. Schwarzmeier, K. Lue, S. Orszag, F. Seidl, O. Johnson, R. Goodrum, and J. Martin. The perfect club: Effective performance evaluation of supercomputers. *The International Journal of Supercomputer Applications*, 3(3), Fall 1989.
- [5] Bradley Efron. The jackknife, the bootstrap, and other resampling plans. SIAM/CBMS-NSF Regional conference series in applied mathematics volume 38, Philadelphia, PA, 1982.
- [6] John R. Ellis. *Bulldog: A Compiler for VLIW Architectures*. ACM doctoral dissertation award; 1985. The MIT Press, 1986.
- [7] S. I. Feldman, David M. Gay, Mark W. Maimone, and N. L. Schryer. A Fortran-to-C converter. Computer Science Technical Report 149, AT&T Bell Laboratories, Murray Hill, NJ 07974, April 1991.
- [8] Phillip B. Gibbons and Steven S. Muchnick. Efficient instruction scheduling for a pipelined architecture. *Proceedings of the SIGPLAN 1986 Symposium on Compiler Construction, SIGPLAN Notices*, 21(7), July 1986.
- [9] John L. Hennessy and Thomas R. Gross. Code generation and reorganization in the presence of pipeline constraints. In *Symposium on Principles of Programming Languages*, pages 120–127, January 1982.
- [10] John L. Hennessy and David A. Patterson. *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann, 1990.
- [11] Mark Donald Hill. *Aspects of Cache Memory and Instruction Buffer Performance*. PhD thesis, University of California, Berkeley, November 1987.
- [12] Gerry Kane. *mips RISC Architecture*. Prentice-Hall, 1988.
- [13] D. Kroft. Lockup-free instruction fetch/prefetch cache organization. In *8th Annual International Symposium on Computer Architecture*, pages 81–87, 1981.

- [14] Monica Lam. Software pipelining: An effective scheduling technique for vliw machines. In *Proceedings of the SIGPLAN '88 Conference on Programming Language Design and Implementation, SIGPLAN Notices*, pages 318–328, 1988.
- [15] E. Lawler, J. K. Lenstra, C. Martel, B. Simons, and L. Stockmeyer. Pipeline scheduling: A survey. Research Report RJ-5738, IBM, July 1987.
- [16] Motorola. *MC88100 RISC Microprocessor User's Manual*. Prentice Hall, 1990.
- [17] Krishna V. Palem and Barbara B. Simons. Scheduling time-critical instructions on RISC machines. In *ACM Symposium on Principles of Programming Languages*, January 1990.
- [18] B. R. Rau and C. D. Glaeser. Some scheduling techniques and an easily schedulable horizontal architecture for high performance scientific computing. In *Proc. 14th Annual Symposium on Microprogramming*, pages 183–198. SIGMICRO, IEEE, October 1981.
- [19] C. Scheurich and M. Dubois. Lockup-free caches in high-performance multiprocessors. *Journal of Parallel and Distributed Processing*, 11(1):25–36, January 1991.
- [20] G. S. Sohi and M. Franklin. High-bandwidth data memory systems for superscalar processor. In *Fourth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 53–62, April 1991.
- [21] Richard Stallman. *The GNU project optimizing C compiler*. Free Software Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.
- [22] Robert Endre Tarjan. *Data Structures and Network Algorithms*, volume 44 of *Regional Conference Series in Applied Mathematics*. Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania, 1983.
- [23] H. S. Warren, Jr. Instruction scheduling for the IBM RISC System/6000 processor. *IBM Journal of Research and Development*, 34(1), January 1990.
- [24] Michael J. Woodard. Personal communication. Scheduling techniques used in Sun SPARC compilers, September 1992.

A Raw Data

This appendix presents the raw data used in the evaluation of balanced scheduling. Each table contains the data for one benchmark executing on one processor model.

In addition to the column labels listed in Table 6, the headings listed in Table 10 are used in this appendix.

Table 10: Index of labels used in Appendix A.

Label	Definition
Tst	The number of static instructions in the program generated with the traditional scheduler
Bst	The number of static instructions in the program generated by the balanced scheduler.
Trt	The run time of the program generated with the traditional scheduler, in millions of cycles. The confidence interval is a percentage of the run time.
Brt	The run time of the program generated by the balanced scheduler, in millions of cycles. The confidence interval is a percentage of the run time.
Tsp%	The percentage of dynamic instructions classified as spill instructions in the program generated with the traditional scheduler.
Bsp%	The percentage of dynamic instructions classified as spill instructions in the program generated by the balanced scheduler.
Imp%	The percentage improvement, at the 95% confidence interval, of the balanced scheduler over the traditional scheduler.

Table 11: Simulation results for benchmark: ADM

system	OL	Tst	Bst	Tins	Bins	Trt	Brt	TI%	BI%	Imp %	Tsp%	Bsp%
Processor: UNLIMITED												
L80(2,5)	2	14553	14526	2554	2494	2760±0.15%	2609±0.11%	7.47	4.41	5.8±0.18	9.59	7.43
	2.6	14539	14526	2541	2494	2714±0.14%	2609±0.11%	6.35	4.41	4.0±0.20	9.15	7.43
L80(2,10)	2	14553	14526	2554	2494	3122±0.37%	2840±0.31%	18.19	12.18	9.9±0.53	9.59	7.43
	3.6	14538	14526	2549	2494	3053±0.34%	2840±0.31%	16.52	12.18	7.5±0.49	9.42	7.43
L95(2,5)	2	14553	14526	2554	2494	2648±0.09%	2562±0.06%	3.56	2.63	3.4±0.11	9.59	7.43
	2.15	14539	14526	2541	2494	2616±0.08%	2562±0.06%	2.84	2.63	2.1±0.12	9.15	7.43
L95(2,10)	2	14553	14526	2554	2494	2753±0.27%	2631±0.22%	7.23	5.21	4.6±0.39	9.59	7.43
	2.4	14539	14526	2541	2494	2714±0.23%	2631±0.22%	6.38	5.21	3.2±0.37	9.15	7.43
N(2,2)	2	14553	14526	2554	2494	2847±0.16%	2637±0.14%	10.30	5.40	8.0±0.24	9.59	7.43
N(3,2)	3	14539	14526	2543	2494	2878±0.17%	2707±0.16%	11.63	7.85	6.4±0.29	9.22	7.43
N(5,2)	5	14533	14526	2550	2494	3063±0.19%	2925±0.18%	16.77	14.74	4.8±0.26	9.50	7.43
N(2,5)	2	14553	14526	2554	2494	3455±0.33%	3025±0.33%	26.09	17.55	14.2±0.51	9.59	7.43
N(3,5)	3	14539	14526	2543	2494	3467±0.32%	3106±0.33%	26.65	19.71	11.5±0.50	9.22	7.43
N(5,5)	5	14533	14526	2550	2494	3598±0.35%	3297±0.35%	29.14	24.35	9.2±0.55	9.50	7.43
N(30,5)	30	14474	14526	2496	2494	7753±0.19%	8034±0.20%	67.81	68.95	-3.5±0.29	7.49	7.43
L80-	2	14553	14526	2554	2494	4769±0.87%	4252±0.90%	46.45	41.35	12.4±1.82	9.59	7.43
N(30,5)	7.6	14506	14526	2527	2494	4550±0.89%	4252±0.90%	44.46	41.35	7.0±1.17	8.70	7.43
Processor: MAX 8												
L80(2,5)	2	14553	14526	2554	2494	2879±0.16%	2703±0.13%	11.30	7.73	6.5±0.24	9.59	7.43
	2.6	14539	14526	2541	2494	2824±0.14%	2703±0.13%	10.02	7.73	4.5±0.22	9.15	7.43
L80(2,10)	2	14553	14526	2554	2494	3301±0.37%	3007±0.34%	22.63	17.06	9.8±0.52	9.59	7.43
	3.6	14538	14526	2549	2494	3239±0.35%	3007±0.34%	21.31	17.06	7.7±0.49	9.42	7.43
L95(2,5)	2	14553	14526	2554	2494	2729±0.09%	2621±0.08%	6.44	4.84	4.1±0.14	9.59	7.43
	2.15	14539	14526	2541	2494	2698±0.08%	2621±0.08%	5.83	4.84	2.9±0.10	9.15	7.43
L95(2,10)	2	14553	14526	2554	2494	2845±0.26%	2700±0.22%	10.24	7.63	5.4±0.42	9.59	7.43
	2.4	14539	14526	2541	2494	2815±0.25%	2700±0.22%	9.73	7.63	4.2±0.29	9.15	7.43
N(2,2)	2	14553	14526	2554	2494	2961±0.18%	2736±0.16%	13.75	8.85	8.2±0.29	9.59	7.43
N(3,2)	3	14539	14526	2543	2494	3042±0.19%	2856±0.19%	16.41	12.66	6.5±0.27	9.22	7.43
N(5,2)	5	14533	14526	2550	2494	3420±0.20%	3224±0.20%	25.45	22.63	6.1±0.34	9.50	7.43
N(2,5)	2	14553	14526	2554	2494	3740±0.34%	3299±0.34%	31.72	24.39	13.4±0.65	9.59	7.43
N(3,5)	3	14539	14526	2543	2494	3782±0.33%	3419±0.34%	32.75	27.04	10.6±0.49	9.22	7.43
N(5,5)	5	14533	14526	2550	2494	4043±0.36%	3709±0.35%	36.95	32.76	8.9±0.59	9.50	7.43
N(30,5)	30	14474	14526	2496	2494	10435±0.17%	10615±0.17%	76.08	76.50	-1.7±0.23	7.49	7.43
L80-	2	14553	14526	2554	2494	5135±0.89%	4652±0.92%	50.27	46.38	10.6±1.22	9.59	7.43
N(30,5)	7.6	14506	14526	2527	2494	4888±0.89%	4652±0.92%	48.29	46.38	5.2±1.23	8.70	7.43
Processor: LEN 8												
L80(2,5)	2	14553	14526	2554	2494	2756±0.15%	2610±0.12%	7.34	4.43	5.6±0.22	9.59	7.43
	2.6	14539	14526	2541	2494	2710±0.13%	2610±0.12%	6.24	4.43	3.8±0.19	9.15	7.43
L80(2,10)	2	14553	14526	2554	2494	3136±0.36%	2891±0.30%	18.56	13.73	8.5±0.49	9.59	7.43
	3.6	14538	14526	2549	2494	3071±0.33%	2891±0.30%	16.99	13.73	6.3±0.40	9.42	7.43
L95(2,5)	2	14553	14526	2554	2494	2646±0.09%	2562±0.06%	3.50	2.67	3.3±0.13	9.59	7.43
	2.15	14539	14526	2541	2494	2615±0.07%	2562±0.06%	2.82	2.67	2.0±0.12	9.15	7.43
L95(2,10)	2	14553	14526	2554	2494	2758±0.26%	2638±0.19%	7.40	5.47	4.5±0.33	9.59	7.43
	2.4	14539	14526	2541	2494	2729±0.25%	2638±0.19%	6.89	5.47	3.5±0.31	9.15	7.43
N(2,2)	2	14553	14526	2554	2494	2844±0.16%	2642±0.15%	10.19	5.59	7.6±0.24	9.59	7.43
N(3,2)	3	14539	14526	2543	2494	2876±0.17%	2710±0.16%	11.58	7.96	6.1±0.31	9.22	7.43
N(5,2)	5	14533	14526	2550	2494	3063±0.19%	2931±0.18%	16.76	14.89	4.6±0.30	9.50	7.43
N(2,5)	2	14553	14526	2554	2494	3469±0.32%	3073±0.32%	26.38	18.84	12.8±0.58	9.59	7.43
N(3,5)	3	14539	14526	2543	2494	3482±0.32%	3160±0.33%	26.97	21.06	10.1±0.59	9.22	7.43
N(5,5)	5	14533	14526	2550	2494	3612±0.34%	3391±0.36%	29.41	26.46	6.5±0.54	9.50	7.43
N(30,5)	30	14474	14526	2496	2494	8280±0.18%	8512±0.19%	69.86	70.70	-2.7±0.34	7.49	7.43
L80-	2	14553	14526	2554	2494	4924±0.86%	4489±0.85%	48.14	44.43	9.7±1.26	9.59	7.43
N(30,5)	7.6	14506	14526	2527	2494	4595±0.87%	4489±0.85%	44.99	44.43	2.4±1.17	8.70	7.43

Table 12: Simulation results for benchmark: ARC2D

system	OL	Tst	Bst	Tins	Bins	Trt	Brt	TI%	BI%	Imp %	Tsp%	Bsp%
Processor: UNLIMITED												
L80(2,5)	2	12491	12170	11548	11149	12428±0.15%	11644±0.10%	7.08	4.25	6.7±0.18	13.52	10.47
	2.6	12451	12170	11570	11149	12360±0.12%	11644±0.10%	6.39	4.25	6.2±0.17	13.68	10.47
L80(2,10)	2	12491	12170	11548	11149	14271±0.36%	12620±0.30%	19.08	11.66	13.1±0.55	13.52	10.47
	3.6	12475	12170	11537	11149	14088±0.33%	12620±0.30%	18.10	11.66	11.7±0.54	13.46	10.47
L95(2,5)	2	12491	12170	11548	11149	11871±0.09%	11424±0.06%	2.72	2.41	3.9±0.11	13.52	10.47
	2.15	12452	12170	11577	11149	11885±0.08%	11424±0.06%	2.59	2.41	4.0±0.09	13.74	10.47
L95(2,10)	2	12491	12170	11548	11149	12390±0.26%	11712±0.19%	6.80	4.81	5.8±0.37	13.52	10.47
	2.4	12452	12170	11577	11149	12429±0.25%	11712±0.19%	6.85	4.81	6.1±0.33	13.74	10.47
N(2,2)	2	12491	12170	11548	11149	12849±0.15%	11757±0.12%	10.13	5.18	9.3±0.25	13.52	10.47
N(3,2)	3	12450	12170	11517	11149	13106±0.16%	12038±0.14%	12.12	7.39	8.9±0.25	13.27	10.47
N(5,2)	5	12459	12170	11598	11149	13695±0.17%	12975±0.17%	15.31	14.08	5.5±0.26	13.89	10.47
N(2,5)	2	12491	12170	11548	11149	15769±0.32%	13397±0.31%	26.77	16.78	17.7±0.46	13.52	10.47
N(3,5)	3	12450	12170	11517	11149	16217±0.32%	13719±0.32%	28.98	18.74	18.2±0.45	13.27	10.47
N(5,5)	5	12459	12170	11598	11149	16326±0.32%	14587±0.34%	28.96	23.57	12.0±0.64	13.89	10.47
N(30,5)	30	12113	12170	11108	11149	33941±0.19%	35713±0.20%	67.27	68.78	-5.0±0.30	10.11	10.47
L80-	2	12491	12170	11548	11149	22857±0.90%	19009±0.88%	49.48	41.35	20.4±1.58	13.52	10.47
N(30,5)	7.6	12353	12170	11378	11149	20784±0.83%	19009±0.88%	45.25	41.35	9.3±1.50	12.25	10.47
Processor: MAX 8												
L80(2,5)	2	12491	12170	11548	11149	13111±0.17%	12045±0.12%	11.93	7.44	8.9±0.19	13.52	10.47
	2.6	12451	12170	11570	11149	13126±0.15%	12045±0.12%	11.86	7.44	9.0±0.19	13.68	10.47
L80(2,10)	2	12491	12170	11548	11149	15227±0.39%	13407±0.36%	24.16	16.84	13.4±0.69	13.52	10.47
	3.6	12475	12170	11537	11149	15095±0.36%	13407±0.36%	23.57	16.84	12.5±0.64	13.46	10.47
L95(2,5)	2	12491	12170	11548	11149	12351±0.10%	11677±0.07%	6.50	4.52	5.8±0.14	13.52	10.47
	2.15	12452	12170	11577	11149	12483±0.09%	11677±0.07%	7.26	4.52	6.9±0.13	13.74	10.47
L95(2,10)	2	12491	12170	11548	11149	12862±0.27%	11982±0.21%	10.22	6.96	7.3±0.41	13.52	10.47
	2.4	12452	12170	11577	11149	13066±0.26%	11982±0.21%	11.39	6.96	9.0±0.39	13.74	10.47
N(2,2)	2	12491	12170	11548	11149	13481±0.19%	12153±0.16%	14.34	8.26	10.9±0.27	13.52	10.47
N(3,2)	3	12450	12170	11517	11149	14156±0.20%	12780±0.20%	18.64	12.77	10.8±0.32	13.27	10.47
N(5,2)	5	12459	12170	11598	11149	16249±0.21%	14713±0.22%	28.62	24.23	10.5±0.42	13.89	10.47
N(2,5)	2	12491	12170	11548	11149	17550±0.37%	14903±0.37%	34.20	25.19	17.8±0.61	13.52	10.47
N(3,5)	3	12450	12170	11517	11149	18116±0.35%	15525±0.35%	36.43	28.19	16.8±0.54	13.27	10.47
N(5,5)	5	12459	12170	11598	11149	19496±0.38%	17052±0.38%	40.51	34.62	14.3±0.66	13.89	10.47
N(30,5)	30	12113	12170	11108	11149	51344±0.20%	53433±0.18%	78.37	79.14	-3.9±0.26	10.11	10.47
L80-	2	12491	12170	11548	11149	24331±0.94%	20739±0.92%	52.54	46.24	17.4±1.57	13.52	10.47
N(30,5)	7.6	12353	12170	11378	11149	23055±0.91%	20739±0.92%	50.65	46.24	11.3±1.43	12.25	10.47
Processor: LEN 8												
L80(2,5)	2	12491	12170	11548	11149	12408±0.15%	11640±0.10%	6.93	4.22	6.6±0.17	13.52	10.47
	2.6	12451	12170	11570	11149	12376±0.12%	11640±0.10%	6.52	4.22	6.3±0.22	13.68	10.47
L80(2,10)	2	12491	12170	11548	11149	14374±0.36%	12918±0.29%	19.66	13.70	11.2±0.62	13.52	10.47
	3.6	12475	12170	11537	11149	14147±0.32%	12918±0.29%	18.45	13.70	9.5±0.49	13.46	10.47
L95(2,5)	2	12491	12170	11548	11149	11873±0.09%	11429±0.06%	2.74	2.46	3.9±0.13	13.52	10.47
	2.15	12452	12170	11577	11149	11896±0.08%	11429±0.06%	2.68	2.46	4.1±0.11	13.74	10.47
L95(2,10)	2	12491	12170	11548	11149	12437±0.27%	11775±0.19%	7.15	5.32	5.7±0.36	13.52	10.47
	2.4	12452	12170	11577	11149	12451±0.25%	11775±0.19%	7.02	5.32	5.7±0.34	13.74	10.47
N(2,2)	2	12491	12170	11548	11149	12842±0.15%	11762±0.12%	10.08	5.21	9.2±0.24	13.52	10.47
N(3,2)	3	12450	12170	11517	11149	13125±0.17%	12072±0.15%	12.25	7.65	8.7±0.23	13.27	10.47
N(5,2)	5	12459	12170	11598	11149	13702±0.17%	13018±0.17%	15.36	14.36	5.3±0.21	13.89	10.47
N(2,5)	2	12491	12170	11548	11149	15913±0.32%	13751±0.31%	27.43	18.92	15.8±0.55	13.52	10.47
N(3,5)	3	12450	12170	11517	11149	16268±0.32%	14188±0.32%	29.20	21.42	14.6±0.59	13.27	10.47
N(5,5)	5	12459	12170	11598	11149	16448±0.32%	15216±0.33%	29.49	26.73	8.1±0.51	13.89	10.47
N(30,5)	30	12113	12170	11108	11149	38101±0.17%	40023±0.18%	70.85	72.14	-4.7±0.24	10.11	10.47
L80-	2	12491	12170	11548	11149	23467±0.88%	20944±0.88%	50.79	46.77	11.9±1.29	13.52	10.47
N(30,5)	7.6	12353	12170	11378	11149	21128±0.82%	20944±0.88%	46.15	46.77	0.7±1.17	12.25	10.47

Table 13: Simulation results for benchmark: BDNA

system	OL	Tst	Bst	Tins	Bins	Trt	Brt	TI%	BI%	Imp %	Tsp%	Bsp%
Processor: UNLIMITED												
L80(2,5)	2	10624	10460	2510	2391	2653±0.25%	2502±0.19%	5.39	4.45	6.0±0.37	26.50	22.84
	2.6	10602	10460	2504	2391	2633±0.20%	2502±0.19%	4.91	4.45	5.2±0.28	26.32	22.84
L80(2,10)	2	10624	10460	2510	2391	2930±0.63%	2649±0.52%	14.34	9.74	10.6±0.90	26.50	22.84
	3.6	10575	10460	2458	2391	2863±0.55%	2649±0.52%	14.17	9.74	8.1±0.80	24.94	22.84
L95(2,5)	2	10624	10460	2510	2391	2570±0.15%	2462±0.10%	2.34	2.89	4.4±0.20	26.50	22.84
	2.15	10602	10460	2504	2391	2560±0.12%	2462±0.10%	2.20	2.89	4.0±0.18	26.32	22.84
L95(2,10)	2	10624	10460	2510	2391	2642±0.39%	2495±0.28%	5.00	4.18	5.8±0.49	26.50	22.84
	2.4	10602	10460	2504	2391	2632±0.37%	2495±0.28%	4.87	4.18	5.6±0.48	26.32	22.84
N(2,2)	2	10624	10460	2510	2391	2719±0.25%	2518±0.22%	7.71	5.04	8.0±0.36	26.50	22.84
N(3,2)	3	10570	10460	2433	2391	2696±0.28%	2594±0.27%	9.77	7.83	4.0±0.46	24.17	22.84
N(5,2)	5	10547	10460	2449	2391	2849±0.31%	2753±0.31%	14.04	13.17	3.4±0.44	24.68	22.84
N(2,5)	2	10624	10460	2510	2391	3208±0.55%	2808±0.55%	21.77	14.86	14.4±1.02	26.50	22.84
N(3,5)	3	10570	10460	2433	2391	3178±0.58%	2865±0.54%	23.45	16.56	10.8±1.04	24.17	22.84
N(5,5)	5	10547	10460	2449	2391	3302±0.57%	3024±0.61%	25.84	20.94	9.3±0.87	24.68	22.84
N(30,5)	30	10483	10460	2478	2391	7086±0.35%	6955±0.35%	65.03	65.63	1.9±0.56	25.54	22.84
L80-	2	10624	10460	2510	2391	4288±1.48%	3712±1.55%	41.48	35.60	15.8±2.78	26.50	22.84
N(30,5)	7.6	10525	10460	2451	2391	4377±1.35%	3712±1.55%	43.99	35.60	18.4±2.54	24.73	22.84
Processor: MAX 8												
L80(2,5)	2	10624	10460	2510	2391	3110±0.32%	2863±0.31%	19.29	16.50	8.6±0.54	26.50	22.84
	2.6	10602	10460	2504	2391	3142±0.32%	2863±0.31%	20.31	16.50	9.8±0.59	26.32	22.84
L80(2,10)	2	10624	10460	2510	2391	3638±0.79%	3272±0.75%	31.02	26.94	11.2±1.30	26.50	22.84
	3.6	10575	10460	2458	2391	3493±0.73%	3272±0.75%	29.65	26.94	6.8±1.18	24.94	22.84
L95(2,5)	2	10624	10460	2510	2391	2899±0.20%	2706±0.18%	13.42	11.66	7.1±0.25	26.50	22.84
	2.15	10602	10460	2504	2391	2939±0.19%	2706±0.18%	14.83	11.66	8.7±0.33	26.32	22.84
L95(2,10)	2	10624	10460	2510	2391	3024±0.56%	2797±0.47%	17.00	14.53	8.0±0.71	26.50	22.84
	2.4	10602	10460	2504	2391	3078±0.54%	2797±0.47%	18.67	14.53	10.2±0.76	26.32	22.84
N(2,2)	2	10624	10460	2510	2391	3152±0.40%	2875±0.40%	20.39	16.85	9.6±0.63	26.50	22.84
N(3,2)	3	10570	10460	2433	2391	3251±0.42%	3112±0.45%	25.18	23.18	4.5±0.63	24.17	22.84
N(5,2)	5	10547	10460	2449	2391	3967±0.44%	3790±0.42%	38.27	36.92	4.6±0.65	24.68	22.84
N(2,5)	2	10624	10460	2510	2391	4282±0.70%	3757±0.69%	41.39	36.36	13.9±1.13	26.50	22.84
N(3,5)	3	10570	10460	2433	2391	4164±0.73%	3924±0.70%	41.58	39.08	6.0±1.14	24.17	22.84
N(5,5)	5	10547	10460	2449	2391	4713±0.71%	4386±0.77%	48.03	45.50	7.4±1.10	24.68	22.84
N(30,5)	30	10483	10460	2478	2391	16836±0.27%	14972±0.29%	85.28	84.03	12.5±0.48	25.54	22.84
L80-	2	10624	10460	2510	2391	5804±1.76%	5168±1.82%	56.76	53.74	12.3±3.10	26.50	22.84
N(30,5)	7.6	10525	10460	2451	2391	5618±1.70%	5168±1.82%	56.37	53.74	8.8±2.67	24.73	22.84
Processor: LEN 8												
L80(2,5)	2	10624	10460	2510	2391	2654±0.23%	2502±0.19%	5.44	4.44	6.1±0.27	26.50	22.84
	2.6	10602	10460	2504	2391	2629±0.20%	2502±0.19%	4.77	4.44	5.1±0.30	26.32	22.84
L80(2,10)	2	10624	10460	2510	2391	2989±0.57%	2754±0.51%	16.04	13.20	8.5±0.81	26.50	22.84
	3.6	10575	10460	2458	2391	2912±0.49%	2754±0.51%	15.61	13.20	5.9±0.77	24.94	22.84
L95(2,5)	2	10624	10460	2510	2391	2563±0.13%	2461±0.10%	2.10	2.85	4.2±0.20	26.50	22.84
	2.15	10602	10460	2504	2391	2555±0.11%	2461±0.10%	2.02	2.85	3.8±0.17	26.32	22.84
L95(2,10)	2	10624	10460	2510	2391	2665±0.37%	2539±0.32%	5.81	5.85	5.0±0.47	26.50	22.84
	2.4	10602	10460	2504	2391	2658±0.37%	2539±0.32%	5.80	5.85	4.6±0.55	26.32	22.84
N(2,2)	2	10624	10460	2510	2391	2723±0.25%	2524±0.23%	7.82	5.27	7.9±0.34	26.50	22.84
N(3,2)	3	10570	10460	2433	2391	2689±0.28%	2588±0.27%	9.54	7.61	3.9±0.37	24.17	22.84
N(5,2)	5	10547	10460	2449	2391	2849±0.31%	2791±0.32%	14.05	14.34	2.1±0.45	24.68	22.84
N(2,5)	2	10624	10460	2510	2391	3285±0.56%	2943±0.56%	23.59	18.77	11.5±0.98	26.50	22.84
N(3,5)	3	10570	10460	2433	2391	3265±0.55%	3042±0.56%	25.50	21.42	7.4±0.92	24.17	22.84
N(5,5)	5	10547	10460	2449	2391	3373±0.57%	3263±0.60%	27.38	26.74	3.3±0.97	24.68	22.84
N(30,5)	30	10483	10460	2478	2391	8578±0.34%	8363±0.31%	71.11	71.41	2.5±0.57	25.54	22.84
L80-	2	10624	10460	2510	2391	4952±1.38%	4524±1.44%	49.32	47.15	9.7±2.31	26.50	22.84
N(30,5)	7.6	10525	10460	2451	2391	4638±1.36%	4524±1.44%	47.15	47.15	2.9±2.11	24.73	22.84

Table 14: Simulation results for benchmark: FLO52Q

system	OL	Tst	Bst	Tins	Bins	Trt	BrT	TI%	BI%	Imp %	Tsp%	Bsp%
Processor: UNLIMITED												
L80(2,5)	2	10960	10785	3413	3323	3652±0.21%	3481±0.19%	6.55	4.56	4.9±0.28	7.14	4.61
	2.6	10938	10785	3401	3323	3607±0.18%	3481±0.19%	5.69	4.56	3.6±0.24	6.82	4.61
L80(2,10)	2	10960	10785	3413	3323	4144±0.53%	3810±0.50%	17.64	12.80	8.7±0.72	7.14	4.61
	3.6	10939	10785	3405	3323	4065±0.47%	3810±0.50%	16.24	12.80	6.7±0.71	6.91	4.61
L95(2,5)	2	10960	10785	3413	3323	3493±0.14%	3397±0.11%	2.30	2.18	2.8±0.19	7.14	4.61
	2.15	10938	10785	3401	3323	3470±0.12%	3397±0.11%	1.98	2.18	2.1±0.18	6.82	4.61
L95(2,10)	2	10960	10785	3413	3323	3625±0.38%	3492±0.34%	5.85	4.84	3.9±0.61	7.14	4.61
	2.4	10938	10785	3401	3323	3622±0.37%	3492±0.34%	6.09	4.84	3.8±0.50	6.82	4.61
N(2,2)	2	10960	10785	3413	3323	3762±0.21%	3532±0.20%	9.28	5.93	6.5±0.31	7.14	4.61
N(3,2)	3	10939	10785	3407	3323	3828±0.23%	3644±0.23%	11.01	8.83	5.0±0.34	6.97	4.61
N(5,2)	5	10886	10785	3391	3323	4081±0.23%	3937±0.25%	16.90	15.61	3.6±0.37	6.55	4.61
N(2,5)	2	10960	10785	3413	3323	4563±0.43%	4076±0.44%	25.20	18.49	11.9±0.79	7.14	4.61
N(3,5)	3	10939	10785	3407	3323	4647±0.43%	4194±0.45%	26.70	20.78	10.9±0.81	6.97	4.61
N(5,5)	5	10886	10785	3391	3323	4760±0.46%	4430±0.46%	28.75	24.99	7.6±0.84	6.55	4.61
N(30,5)	30	10756	10785	3332	3323	10096±0.24%	9695±0.26%	67.00	65.73	4.1±0.37	4.90	4.61
L80-	2	10960	10785	3413	3323	6441±1.32%	5699±1.33%	47.01	41.70	12.7±2.23	7.14	4.61
N(30,5)	7.6	10834	10785	3368	3323	6054±1.32%	5699±1.33%	44.37	41.70	6.3±2.17	5.89	4.61
Processor: MAX 8												
L80(2,5)	2	10960	10785	3413	3323	3858±0.23%	3652±0.21%	11.54	9.02	5.6±0.36	7.14	4.61
	2.6	10938	10785	3401	3323	3804±0.20%	3652±0.21%	10.59	9.02	4.2±0.28	6.82	4.61
L80(2,10)	2	10960	10785	3413	3323	4444±0.55%	4079±0.52%	23.20	18.55	9.0±0.80	7.14	4.61
	3.6	10939	10785	3405	3323	4344±0.51%	4079±0.52%	21.63	18.55	6.6±0.84	6.91	4.61
L95(2,5)	2	10960	10785	3413	3323	3637±0.15%	3504±0.13%	6.17	5.16	3.8±0.23	7.14	4.61
	2.15	10938	10785	3401	3323	3616±0.12%	3504±0.13%	5.95	5.16	3.2±0.21	6.82	4.61
L95(2,10)	2	10960	10785	3413	3323	3791±0.39%	3614±0.35%	9.98	8.07	5.0±0.55	7.14	4.61
	2.4	10938	10785	3401	3323	3797±0.41%	3614±0.35%	10.41	8.07	4.9±0.55	6.82	4.61
N(2,2)	2	10960	10785	3413	3323	3976±0.24%	3706±0.25%	14.16	10.34	7.3±0.35	7.14	4.61
N(3,2)	3	10939	10785	3407	3323	4149±0.26%	3909±0.27%	17.90	15.00	6.1±0.42	6.97	4.61
N(5,2)	5	10886	10785	3391	3323	4664±0.26%	4456±0.27%	27.30	25.43	4.7±0.40	6.55	4.61
N(2,5)	2	10960	10785	3413	3323	5066±0.46%	4557±0.48%	32.62	27.10	11.1±0.82	7.14	4.61
N(3,5)	3	10939	10785	3407	3323	5222±0.45%	4734±0.46%	34.77	29.81	10.4±0.73	6.97	4.61
N(5,5)	5	10886	10785	3391	3323	5524±0.48%	5130±0.49%	38.61	35.23	7.7±0.80	6.55	4.61
N(30,5)	30	10756	10785	3332	3323	14260±0.23%	14057±0.23%	76.63	76.36	1.4±0.31	4.90	4.61
L80-	2	10960	10785	3413	3323	6921±1.31%	6250±1.36%	50.69	46.84	11.2±1.98	7.14	4.61
N(30,5)	7.6	10834	10785	3368	3323	6718±1.34%	6250±1.36%	49.88	46.84	7.4±2.12	5.89	4.61
Processor: LEN 8												
L80(2,5)	2	10960	10785	3413	3323	3646±0.21%	3486±0.18%	6.37	4.68	4.6±0.29	7.14	4.61
	2.6	10938	10785	3401	3323	3608±0.18%	3486±0.18%	5.74	4.68	3.5±0.27	6.82	4.61
L80(2,10)	2	10960	10785	3413	3323	4176±0.53%	3879±0.49%	18.27	14.35	7.7±0.75	7.14	4.61
	3.6	10939	10785	3405	3323	4082±0.47%	3879±0.49%	16.59	14.35	5.2±0.74	6.91	4.61
L95(2,5)	2	10960	10785	3413	3323	3492±0.14%	3400±0.12%	2.25	2.27	2.7±0.21	7.14	4.61
	2.15	10938	10785	3401	3323	3469±0.12%	3400±0.12%	1.95	2.27	2.0±0.17	6.82	4.61
L95(2,10)	2	10960	10785	3413	3323	3641±0.38%	3518±0.36%	6.25	5.55	3.5±0.54	7.14	4.61
	2.4	10938	10785	3401	3323	3613±0.35%	3518±0.36%	5.85	5.55	2.7±0.66	6.82	4.61
N(2,2)	2	10960	10785	3413	3323	3757±0.21%	3538±0.20%	9.15	6.09	6.2±0.28	7.14	4.61
N(3,2)	3	10939	10785	3407	3323	3830±0.22%	3646±0.23%	11.07	8.87	5.0±0.36	6.97	4.61
N(5,2)	5	10886	10785	3391	3323	4080±0.24%	3945±0.24%	16.88	15.78	3.4±0.32	6.55	4.61
N(2,5)	2	10960	10785	3413	3323	4573±0.43%	4145±0.44%	25.36	19.84	10.1±0.68	7.14	4.61
N(3,5)	3	10939	10785	3407	3323	4650±0.42%	4248±0.42%	26.74	21.79	9.6±0.72	6.97	4.61
N(5,5)	5	10886	10785	3391	3323	4804±0.45%	4544±0.47%	29.41	26.87	5.7±0.71	6.55	4.61
N(30,5)	30	10756	10785	3332	3323	10739±0.25%	10355±0.25%	68.97	67.91	3.7±0.41	4.90	4.61
L80-	2	10960	10785	3413	3323	6525±1.27%	5927±1.31%	47.69	43.94	10.3±2.01	7.14	4.61
N(30,5)	7.6	10834	10785	3368	3323	6119±1.30%	5927±1.31%	44.96	43.94	3.5±2.28	5.89	4.61

Table 15: Simulation results for benchmark: MDG

system	OL	Tst	Bst	Tins	Bins	Trt	Brt	TI%	BI%	Imp %	Tsp%	Bsp%
Processor: UNLIMITED												
L80(2,5)	2	4259	4159	5358	5144	5979±0.27%	5447±0.22%	10.38	5.56	9.8±0.41	7.86	7.49
	2.6	4272	4159	5351	5144	5918±0.27%	5447±0.22%	9.57	5.56	8.7±0.37	8.04	7.49
L80(2,10)	2	4259	4159	5358	5144	6813±0.64%	5954±0.58%	21.35	13.61	14.4±0.99	7.86	7.49
	3.6	4257	4159	5299	5144	6644±0.63%	5954±0.58%	20.24	13.61	11.6±0.98	8.13	7.49
L95(2,5)	2	4259	4159	5358	5144	5694±0.16%	5326±0.13%	5.89	3.43	6.9±0.24	7.86	7.49
	2.15	4272	4159	5351	5144	5662±0.17%	5326±0.13%	5.48	3.43	6.2±0.26	8.04	7.49
L95(2,10)	2	4259	4159	5358	5144	5917±0.40%	5477±0.39%	9.43	6.08	8.0±0.77	7.86	7.49
	2.4	4272	4159	5351	5144	5874±0.43%	5477±0.39%	8.89	6.08	7.2±0.70	8.04	7.49
N(2,2)	2	4259	4159	5358	5144	6144±0.32%	5523±0.26%	12.78	6.86	11.3±0.49	7.86	7.49
N(3,2)	3	4271	4159	5351	5144	6379±0.33%	5696±0.31%	16.11	9.69	12.0±0.52	8.04	7.49
N(5,2)	5	4257	4159	5297	5144	7010±0.35%	6179±0.32%	24.44	16.75	13.5±0.55	8.00	7.49
N(2,5)	2	4259	4159	5358	5144	7651±0.57%	6330±0.52%	29.97	18.73	20.9±0.85	7.86	7.49
N(3,5)	3	4271	4159	5351	5144	7850±0.59%	6542±0.57%	31.83	21.38	20.0±1.03	8.04	7.49
N(5,5)	5	4257	4159	5297	5144	8213±0.63%	6942±0.59%	35.50	25.90	18.3±1.14	8.00	7.49
N(30,5)	30	4234	4159	5393	5144	19124±0.33%	16035±0.35%	71.80	67.92	19.3±0.57	9.21	7.49
L80-	2	4259	4159	5358	5144	10689±1.56%	8908±1.42%	49.87	42.26	20.0±2.72	7.86	7.49
N(30,5)	7.6	4237	4159	5405	5144	10185±1.60%	8908±1.42%	46.93	42.26	14.3±2.50	8.68	7.49
Processor: MAX 8												
L80(2,5)	2	4259	4159	5358	5144	6221±0.27%	5772±0.24%	13.87	10.88	7.8±0.37	7.86	7.49
	2.6	4272	4159	5351	5144	6200±0.26%	5772±0.24%	13.69	10.88	7.4±0.43	8.04	7.49
L80(2,10)	2	4259	4159	5358	5144	7165±0.63%	6474±0.58%	25.22	20.55	10.8±0.99	7.86	7.49
	3.6	4257	4159	5299	5144	7036±0.62%	6474±0.58%	24.68	20.55	8.9±1.01	8.13	7.49
L95(2,5)	2	4259	4159	5358	5144	5875±0.16%	5539±0.14%	8.79	7.14	6.0±0.22	7.86	7.49
	2.15	4272	4159	5351	5144	5872±0.15%	5539±0.14%	8.87	7.14	6.0±0.21	8.04	7.49
L95(2,10)	2	4259	4159	5358	5144	6094±0.39%	5726±0.36%	12.07	10.17	6.4±0.63	7.86	7.49
	2.4	4272	4159	5351	5144	6115±0.42%	5726±0.36%	12.49	10.17	6.8±0.59	8.04	7.49
N(2,2)	2	4259	4159	5358	5144	6424±0.30%	5831±0.28%	16.58	11.78	10.2±0.49	7.86	7.49
N(3,2)	3	4271	4159	5351	5144	6780±0.34%	6161±0.30%	21.08	16.51	10.0±0.53	8.04	7.49
N(5,2)	5	4257	4159	5297	5144	7800±0.34%	7042±0.30%	32.09	26.96	10.8±0.49	8.00	7.49
N(2,5)	2	4259	4159	5358	5144	8322±0.58%	7178±0.55%	35.61	28.34	15.9±1.06	7.86	7.49
N(3,5)	3	4271	4159	5351	5144	8518±0.57%	7440±0.55%	37.18	30.87	14.5±0.88	8.04	7.49
N(5,5)	5	4257	4159	5297	5144	9187±0.59%	8088±0.56%	42.34	36.40	13.5±0.86	8.00	7.49
N(30,5)	30	4234	4159	5393	5144	26158±0.28%	22848±0.27%	79.38	77.49	14.5±0.45	9.21	7.49
L80-	2	4259	4159	5358	5144	11230±1.48%	9868±1.40%	52.29	47.87	13.6±2.53	7.86	7.49
N(30,5)	7.6	4237	4159	5405	5144	11001±1.54%	9868±1.40%	50.87	47.87	11.5±2.53	8.68	7.49
Processor: LEN 8												
L80(2,5)	2	4259	4159	5358	5144	5978±0.27%	5456±0.22%	10.36	5.72	9.6±0.44	7.86	7.49
	2.6	4272	4159	5351	5144	5901±0.26%	5456±0.22%	9.31	5.72	8.2±0.37	8.04	7.49
L80(2,10)	2	4259	4159	5358	5144	6912±0.66%	6028±0.56%	22.47	14.67	14.6±1.02	7.86	7.49
	3.6	4257	4159	5299	5144	6727±0.65%	6028±0.56%	21.23	14.67	11.6±0.97	8.13	7.49
L95(2,5)	2	4259	4159	5358	5144	5687±0.15%	5323±0.12%	5.78	3.37	6.8±0.22	7.86	7.49
	2.15	4272	4159	5351	5144	5652±0.16%	5323±0.12%	5.32	3.37	6.2±0.20	8.04	7.49
L95(2,10)	2	4259	4159	5358	5144	5949±0.45%	5506±0.37%	9.93	6.58	8.0±0.64	7.86	7.49
	2.4	4272	4159	5351	5144	5907±0.44%	5506±0.37%	9.40	6.58	7.2±0.69	8.04	7.49
N(2,2)	2	4259	4159	5358	5144	6163±0.31%	5543±0.27%	13.06	7.21	11.2±0.53	7.86	7.49
N(3,2)	3	4271	4159	5351	5144	6368±0.34%	5681±0.29%	15.97	9.46	12.0±0.48	8.04	7.49
N(5,2)	5	4257	4159	5297	5144	6993±0.36%	6196±0.32%	24.25	16.99	12.9±0.61	8.00	7.49
N(2,5)	2	4259	4159	5358	5144	7699±0.59%	6465±0.56%	30.40	20.43	19.2±0.95	7.86	7.49
N(3,5)	3	4271	4159	5351	5144	7852±0.59%	6683±0.58%	31.85	23.04	17.4±1.01	8.04	7.49
N(5,5)	5	4257	4159	5297	5144	8318±0.62%	7145±0.58%	36.31	28.00	16.4±1.12	8.00	7.49
N(30,5)	30	4234	4159	5393	5144	20060±0.33%	16891±0.33%	73.11	69.55	18.7±0.59	9.21	7.49
L80-	2	4259	4159	5358	5144	10619±1.45%	9301±1.38%	49.54	44.70	13.9±2.66	7.86	7.49
N(30,5)	7.6	4237	4159	5405	5144	10280±1.56%	9301±1.38%	47.42	44.70	10.4±2.20	8.68	7.49

Table 16: Simulation results for benchmark: MG3D

system	OL	Tst	Bst	Tins	Bins	Trt	Brt	TI%	BI%	Imp %	Tsp%	Bsp%
Processor: UNLIMITED												
L80(2,5)	2	7695	7431	62371	60784	66646±0.15%	62305±0.09%	6.42	2.44	7.0±0.21	9.73	7.38
	2.6	7667	7431	62806	60784	66153±0.11%	62305±0.09%	5.06	2.44	6.2±0.16	10.36	7.38
L80(2,10)	2	7695	7431	62371	60784	75656±0.34%	67619±0.28%	17.56	10.11	11.9±0.56	9.73	7.38
	3.6	7677	7431	63154	60784	74814±0.28%	67619±0.28%	15.59	10.11	10.7±0.51	10.86	7.38
L95(2,5)	2	7695	7431	62371	60784	63782±0.09%	61509±0.05%	2.21	1.18	3.7±0.11	9.73	7.38
	2.15	7667	7431	62806	60784	63917±0.07%	61509±0.05%	1.74	1.18	3.9±0.09	10.36	7.38
L95(2,10)	2	7695	7431	62371	60784	66359±0.26%	62980±0.19%	6.01	3.49	5.4±0.34	9.73	7.38
	2.4	7667	7431	62806	60784	66669±0.25%	62980±0.19%	5.79	3.49	5.8±0.36	10.36	7.38
N(2,2)	2	7695	7431	62371	60784	68573±0.14%	62763±0.11%	9.05	3.15	9.2±0.20	9.73	7.38
N(3,2)	3	7667	7431	62806	60784	69552±0.14%	64038±0.13%	9.70	5.08	8.6±0.27	10.36	7.38
N(5,2)	5	7638	7431	62807	60784	72222±0.16%	67823±0.16%	13.04	10.38	6.5±0.26	10.36	7.38
N(2,5)	2	7695	7431	62371	60784	82426±0.31%	70992±0.29%	24.33	14.38	16.1±0.51	9.73	7.38
N(3,5)	3	7667	7431	62806	60784	83307±0.29%	72557±0.31%	24.61	16.23	14.9±0.47	10.36	7.38
N(5,5)	5	7638	7431	62807	60784	84721±0.31%	76773±0.32%	25.87	20.83	10.3±0.51	10.36	7.38
N(30,5)	30	7440	7431	61116	60784	165787±0.19%	167326±0.20%	63.14	63.67	-0.9±0.35	7.88	7.38
L80-	2	7695	7431	62371	60784	119690±0.82%	105574±0.80%	47.89	42.42	13.3±1.16	9.73	7.38
N(30,5)	7.6	7546	7431	61764	60784	110265±0.77%	105574±0.80%	43.99	42.42	4.5±1.13	8.85	7.38
Processor: MAX 8												
L80(2,5)	2	7695	7431	62371	60784	80961±0.17%	76427±0.15%	22.96	20.47	5.9±0.26	9.73	7.38
	2.6	7667	7431	62806	60784	81484±0.15%	76427±0.15%	22.92	20.47	6.6±0.27	10.36	7.38
L80(2,10)	2	7695	7431	62371	60784	97048±0.39%	89956±0.35%	35.73	32.43	7.8±0.62	9.73	7.38
	3.6	7677	7431	63154	60784	97357±0.36%	89956±0.35%	35.13	32.43	8.3±0.53	10.86	7.38
L95(2,5)	2	7695	7431	62371	60784	74354±0.10%	71577±0.08%	16.12	15.08	3.9±0.16	9.73	7.38
	2.15	7667	7431	62806	60784	75316±0.09%	71577±0.08%	16.61	15.08	5.2±0.12	10.36	7.38
L95(2,10)	2	7695	7431	62371	60784	78628±0.29%	75089±0.25%	20.68	19.05	4.8±0.39	9.73	7.38
	2.4	7667	7431	62806	60784	79539±0.27%	75089±0.25%	21.04	19.05	5.9±0.40	10.36	7.38
N(2,2)	2	7695	7431	62371	60784	84141±0.18%	78115±0.17%	25.87	22.19	7.7±0.26	9.73	7.38
N(3,2)	3	7667	7431	62806	60784	92046±0.21%	85446±0.20%	31.77	28.86	7.7±0.31	10.36	7.38
N(5,2)	5	7638	7431	62807	60784	112237±0.20%	105577±0.21%	44.04	42.43	6.3±0.30	10.36	7.38
N(2,5)	2	7695	7431	62371	60784	117900±0.36%	106270±0.34%	47.10	42.80	10.9±0.56	9.73	7.38
N(3,5)	3	7667	7431	62806	60784	124238±0.32%	112119±0.33%	49.45	45.79	10.8±0.46	10.36	7.38
N(5,5)	5	7638	7431	62807	60784	136204±0.37%	126311±0.36%	53.89	51.88	7.8±0.65	10.36	7.38
N(30,5)	30	7440	7431	61116	60784	456101±0.14%	449084±0.15%	86.60	86.46	1.5±0.24	7.88	7.38
L80-	2	7695	7431	62371	60784	163090±0.82%	149129±0.81%	61.76	59.24	9.4±1.48	9.73	7.38
N(30,5)	7.6	7546	7431	61764	60784	156650±0.77%	149129±0.81%	60.57	59.24	5.1±1.19	8.85	7.38
Processor: LEN 8												
L80(2,5)	2	7695	7431	62371	60784	66550±0.14%	62306±0.09%	6.28	2.44	6.8±0.20	9.73	7.38
	2.6	7667	7431	62806	60784	66095±0.11%	62306±0.09%	4.98	2.44	6.1±0.15	10.36	7.38
L80(2,10)	2	7695	7431	62371	60784	76662±0.34%	69193±0.27%	18.64	12.15	10.8±0.47	9.73	7.38
	3.6	7677	7431	63154	60784	75778±0.28%	69193±0.27%	16.66	12.15	9.6±0.45	10.86	7.38
L95(2,5)	2	7695	7431	62371	60784	63802±0.09%	61488±0.05%	2.24	1.14	3.7±0.11	9.73	7.38
	2.15	7667	7431	62806	60784	63881±0.07%	61488±0.05%	1.68	1.14	3.9±0.11	10.36	7.38
L95(2,10)	2	7695	7431	62371	60784	66715±0.25%	63688±0.18%	6.51	4.56	4.8±0.29	9.73	7.38
	2.4	7667	7431	62806	60784	67186±0.24%	63688±0.18%	6.52	4.56	5.5±0.36	10.36	7.38
N(2,2)	2	7695	7431	62371	60784	68584±0.15%	62906±0.11%	9.06	3.37	9.0±0.21	9.73	7.38
N(3,2)	3	7667	7431	62806	60784	69635±0.15%	64017±0.12%	9.81	5.05	8.8±0.19	10.36	7.38
N(5,2)	5	7638	7431	62807	60784	72470±0.17%	68304±0.17%	13.33	11.01	6.1±0.22	10.36	7.38
N(2,5)	2	7695	7431	62371	60784	83327±0.31%	73272±0.31%	25.15	17.04	13.6±0.51	9.73	7.38
N(3,5)	3	7667	7431	62806	60784	84446±0.30%	75229±0.30%	25.63	19.20	12.2±0.50	10.36	7.38
N(5,5)	5	7638	7431	62807	60784	86339±0.32%	80379±0.32%	27.26	24.38	7.3±0.38	10.36	7.38
N(30,5)	30	7440	7431	61116	60784	187416±0.18%	194718±0.20%	67.39	68.78	-3.8±0.24	7.88	7.38
L80-	2	7695	7431	62371	60784	127001±0.80%	116132±0.77%	50.89	47.66	9.5±1.34	9.73	7.38
N(30,5)	7.6	7546	7431	61764	60784	117195±0.73%	116132±0.77%	47.30	47.66	0.9±1.19	8.85	7.38

Table 17: Simulation results for benchmark: QCD2

system	OL	Tst	Bst	Tins	Bins	Trt	Brt	TI%	BI%	Imp %	Tsp%	Bsp%
Processor: UNLIMITED												
L80(2,5)	2	5813	5685	1366	1176	1472±0.20%	1234±0.19%	7.15	4.72	19.3±0.36	29.30	19.91
	2.6	5785	5685	1356	1176	1464±0.20%	1234±0.19%	7.34	4.72	18.6±0.36	28.92	19.91
L80(2,10)	2	5813	5685	1366	1176	1721±0.51%	1347±0.53%	20.59	12.73	27.8±1.02	29.30	19.91
	3.6	5798	5685	1351	1176	1695±0.46%	1347±0.53%	20.28	12.73	25.8±0.92	28.78	19.91
L95(2,5)	2	5813	5685	1366	1176	1412±0.12%	1208±0.10%	3.22	2.70	16.9±0.20	29.30	19.91
	2.15	5785	5685	1356	1176	1404±0.13%	1208±0.10%	3.42	2.70	16.2±0.22	28.92	19.91
L95(2,10)	2	5813	5685	1366	1176	1484±0.35%	1239±0.32%	7.92	5.11	19.9±0.59	29.30	19.91
	2.4	5785	5685	1356	1176	1474±0.35%	1239±0.32%	7.98	5.11	19.0±0.54	28.92	19.91
N(2,2)	2	5813	5685	1366	1176	1512±0.23%	1247±0.23%	9.61	5.71	21.3±0.44	29.30	19.91
N(3,2)	3	5785	5685	1356	1176	1571±0.26%	1280±0.25%	13.66	8.14	22.7±0.48	28.92	19.91
N(5,2)	5	5774	5685	1329	1176	1646±0.27%	1373±0.30%	19.25	14.39	20.0±0.52	28.02	19.91
N(2,5)	2	5813	5685	1366	1176	1892±0.47%	1426±0.51%	27.77	17.56	32.7±1.14	29.30	19.91
N(3,5)	3	5785	5685	1356	1176	1978±0.48%	1456±0.48%	31.44	19.29	35.9±0.92	28.92	19.91
N(5,5)	5	5774	5685	1329	1176	1981±0.47%	1551±0.54%	32.89	24.19	27.8±0.83	28.02	19.91
N(30,5)	30	5755	5685	1270	1176	4092±0.27%	3822±0.31%	68.95	69.24	7.1±0.50	28.02	19.91
L80-	2	5813	5685	1366	1176	2889±1.20%	2074±1.31%	52.70	43.32	39.6±2.38	29.30	19.91
N(30,5)	7.6	5751	5685	1297	1176	2477±1.17%	2074±1.31%	47.62	43.32	19.4±2.21	26.89	19.91
Processor: MAX 8												
L80(2,5)	2	5813	5685	1366	1176	1691±0.24%	1374±0.22%	19.18	14.43	23.1±0.46	29.30	19.91
	2.6	5785	5685	1356	1176	1673±0.23%	1374±0.22%	18.92	14.43	21.8±0.44	28.92	19.91
L80(2,10)	2	5813	5685	1366	1176	2021±0.54%	1599±0.54%	32.38	26.50	26.2±1.00	29.30	19.91
	3.6	5798	5685	1351	1176	2001±0.53%	1599±0.54%	32.47	26.50	25.2±0.96	28.78	19.91
L95(2,5)	2	5813	5685	1366	1176	1565±0.15%	1300±0.14%	12.69	9.54	20.4±0.26	29.30	19.91
	2.15	5785	5685	1356	1176	1551±0.14%	1300±0.14%	12.54	9.54	19.4±0.23	28.92	19.91
L95(2,10)	2	5813	5685	1366	1176	1645±0.37%	1353±0.35%	16.94	13.14	21.6±0.65	29.30	19.91
	2.4	5785	5685	1356	1176	1633±0.37%	1353±0.35%	16.95	13.14	20.6±0.58	28.92	19.91
N(2,2)	2	5813	5685	1366	1176	1747±0.27%	1405±0.27%	21.80	16.33	24.4±0.45	29.30	19.91
N(3,2)	3	5785	5685	1356	1176	1884±0.29%	1518±0.32%	28.03	22.58	24.1±0.51	28.92	19.91
N(5,2)	5	5774	5685	1329	1176	2280±0.29%	1838±0.30%	41.70	36.03	24.1±0.47	28.02	19.91
N(2,5)	2	5813	5685	1366	1176	2427±0.48%	1867±0.53%	43.69	37.02	30.0±0.84	29.30	19.91
N(3,5)	3	5785	5685	1356	1176	2546±0.44%	1956±0.51%	46.72	39.91	30.1±0.89	28.92	19.91
N(5,5)	5	5774	5685	1329	1176	2766±0.52%	2196±0.55%	51.95	46.47	25.9±1.07	28.02	19.91
N(30,5)	30	5755	5685	1270	1176	9343±0.19%	7574±0.22%	86.40	84.48	23.4±0.33	28.02	19.91
L80-	2	5813	5685	1366	1176	3391±1.19%	2603±1.28%	59.70	54.84	30.4±2.03	29.30	19.91
N(30,5)	7.6	5751	5685	1297	1176	3115±1.19%	2603±1.28%	58.35	54.84	19.8±2.07	26.89	19.91
Processor: LEN 8												
L80(2,5)	2	5813	5685	1366	1176	1472±0.20%	1234±0.18%	7.16	4.76	19.2±0.34	29.30	19.91
	2.6	5785	5685	1356	1176	1462±0.19%	1234±0.18%	7.22	4.76	18.4±0.32	28.92	19.91
L80(2,10)	2	5813	5685	1366	1176	1720±0.48%	1377±0.51%	20.56	14.61	24.9±0.84	29.30	19.91
	3.6	5798	5685	1351	1176	1712±0.48%	1377±0.51%	21.08	14.61	24.4±0.88	28.78	19.91
L95(2,5)	2	5813	5685	1366	1176	1413±0.11%	1208±0.11%	3.29	2.70	16.9±0.17	29.30	19.91
	2.15	5785	5685	1356	1176	1400±0.11%	1208±0.11%	3.15	2.70	15.9±0.19	28.92	19.91
L95(2,10)	2	5813	5685	1366	1176	1482±0.34%	1251±0.31%	7.81	6.00	18.6±0.60	29.30	19.91
	2.4	5785	5685	1356	1176	1479±0.36%	1251±0.31%	8.29	6.00	18.2±0.55	28.92	19.91
N(2,2)	2	5813	5685	1366	1176	1510±0.22%	1246±0.23%	9.49	5.63	21.2±0.40	29.30	19.91
N(3,2)	3	5785	5685	1356	1176	1575±0.26%	1279±0.24%	13.88	8.11	23.1±0.57	28.92	19.91
N(5,2)	5	5774	5685	1329	1176	1649±0.27%	1385±0.29%	19.37	15.15	19.0±0.55	28.02	19.91
N(2,5)	2	5813	5685	1366	1176	1909±0.47%	1471±0.50%	28.40	20.07	29.7±1.08	29.30	19.91
N(3,5)	3	5785	5685	1356	1176	1987±0.47%	1518±0.51%	31.74	22.57	30.9±1.07	28.92	19.91
N(5,5)	5	5774	5685	1329	1176	1997±0.48%	1642±0.55%	33.45	28.42	21.6±0.89	28.02	19.91
N(30,5)	30	5755	5685	1270	1176	4565±0.26%	4297±0.28%	72.17	72.64	6.3±0.40	28.02	19.91
L80-	2	5813	5685	1366	1176	2921±1.14%	2288±1.26%	53.22	48.62	27.6±2.26	29.30	19.91
N(30,5)	7.6	5751	5685	1297	1176	2530±1.12%	2288±1.26%	48.72	48.62	10.9±2.15	26.89	19.91

Table 18: Simulation results for benchmark: TRACK

system	OL	Tst	Bst	Tins	Bins	Trt	Br	TI%	BI%	Imp %	Tsp%	Bsp%
Processor: UNLIMITED												
L80(2,5)	2	6102	5853	421	398	517±1.16%	481±1.05%	18.52	17.28	7.2±1.52	20.41	15.78
	2.6	6083	5853	408	398	494±1.03%	481±1.05%	17.44	17.28	2.6±1.57	17.85	15.78
L80(2,10)	2	6102	5853	421	398	624±2.70%	585±2.68%	32.51	31.97	6.7±3.96	20.41	15.78
	3.6	6087	5853	408	398	600±2.55%	585±2.68%	31.96	31.97	2.2±3.92	17.85	15.78
L95(2,5)	2	6102	5853	421	398	482±0.75%	454±0.63%	12.61	12.43	6.1±1.00	20.41	15.78
	2.15	6083	5853	408	398	463±0.58%	454±0.63%	11.91	12.43	2.0±0.98	17.85	15.78
L95(2,10)	2	6102	5853	421	398	508±1.94%	483±2.02%	17.16	17.70	4.9±2.69	20.41	15.78
	2.4	6083	5853	408	398	492±1.85%	483±2.02%	17.07	17.70	1.7±2.99	17.85	15.78
N(2,2)	2	6102	5853	421	398	549±1.27%	502±1.29%	23.26	20.69	9.4±1.88	20.41	15.78
N(3,2)	3	6083	5853	408	398	551±1.30%	533±1.38%	25.97	25.33	3.5±2.02	17.84	15.78
N(5,2)	5	6082	5853	408	398	649±1.18%	625±1.27%	37.17	36.35	3.9±1.74	17.84	15.78
N(2,5)	2	6102	5853	421	398	754±1.98%	648±2.23%	44.10	38.57	16.6±3.32	20.41	15.78
N(3,5)	3	6083	5853	408	398	726±2.10%	696±2.25%	43.81	42.83	3.9±3.67	17.84	15.78
N(5,5)	5	6082	5853	408	398	789±2.09%	755±2.21%	48.34	47.31	4.9±3.64	17.84	15.78
N(30,5)	30	6063	5853	406	398	2206±0.90%	2194±0.88%	81.59	81.87	0.6±1.33	17.46	15.78
L80-	2	6102	5853	421	398	1064±5.74%	957±5.93%	60.41	58.44	11.3±9.64	20.41	15.78
N(30,5)	7.6	6053	5853	406	398	928±5.62%	957±5.93%	56.25	58.44	-2.5±9.33	17.45	15.78
Processor: MAX 8												
L80(2,5)	2	6102	5853	421	398	551±1.14%	508±1.09%	23.62	21.66	8.7±1.65	20.41	15.78
	2.6	6083	5853	408	398	521±1.01%	508±1.09%	21.79	21.66	2.6±1.44	17.85	15.78
L80(2,10)	2	6102	5853	421	398	671±2.51%	620±2.59%	37.26	35.83	7.2±3.63	20.41	15.78
	3.6	6087	5853	408	398	638±2.47%	620±2.59%	36.08	35.83	2.3±3.64	17.85	15.78
L95(2,5)	2	6102	5853	421	398	500±0.70%	469±0.66%	15.72	15.27	6.6±1.19	20.41	15.78
	2.15	6083	5853	408	398	483±0.65%	469±0.66%	15.58	15.27	2.9±0.91	17.85	15.78
L95(2,10)	2	6102	5853	421	398	537±2.02%	498±1.92%	21.50	20.17	7.5±3.05	20.41	15.78
	2.4	6083	5853	408	398	515±1.80%	498±1.92%	20.81	20.17	3.2±3.05	17.85	15.78
N(2,2)	2	6102	5853	421	398	567±1.21%	518±1.27%	25.73	23.25	9.2±2.03	20.41	15.78
N(3,2)	3	6083	5853	408	398	597±1.28%	567±1.35%	31.75	29.85	5.3±1.83	17.84	15.78
N(5,2)	5	6082	5853	408	398	728±1.07%	692±1.10%	44.00	42.52	5.2±1.75	17.84	15.78
N(2,5)	2	6102	5853	421	398	816±2.03%	717±2.05%	48.37	44.51	13.2±3.58	20.41	15.78
N(3,5)	3	6083	5853	408	398	790±1.87%	745±2.06%	48.40	46.58	6.5±3.22	17.84	15.78
N(5,5)	5	6082	5853	408	398	886±1.86%	837±1.89%	53.97	52.50	5.5±2.92	17.84	15.78
N(30,5)	30	6063	5853	406	398	2814±0.69%	2687±0.75%	85.57	85.20	4.7±1.21	17.46	15.78
L80-	2	6102	5853	421	398	1122±5.65%	999±5.67%	62.47	60.18	11.9±9.62	20.41	15.78
N(30,5)	7.6	6053	5853	406	398	1021±5.37%	999±5.67%	60.24	60.18	1.5±8.87	17.45	15.78
Processor: LEN 8												
L80(2,5)	2	6102	5853	421	398	522±1.14%	486±1.14%	19.25	18.08	7.2±1.94	20.41	15.78
	2.6	6083	5853	408	398	498±1.09%	486±1.14%	18.11	18.08	2.8±1.64	17.85	15.78
L80(2,10)	2	6102	5853	421	398	634±2.66%	571±2.65%	33.54	30.37	10.8±3.68	20.41	15.78
	3.6	6087	5853	408	398	596±2.58%	571±2.65%	31.54	30.37	4.6±3.07	17.85	15.78
L95(2,5)	2	6102	5853	421	398	483±0.79%	457±0.72%	12.83	12.91	6.0±1.37	20.41	15.78
	2.15	6083	5853	408	398	466±0.73%	457±0.72%	12.54	12.91	2.0±1.15	17.85	15.78
L95(2,10)	2	6102	5853	421	398	508±1.89%	473±1.64%	17.09	15.92	7.6±3.07	20.41	15.78
	2.4	6083	5853	408	398	490±1.81%	473±1.64%	16.84	15.92	4.1±2.43	17.85	15.78
N(2,2)	2	6102	5853	421	398	539±1.28%	501±1.25%	21.85	20.53	8.0±2.33	20.41	15.78
N(3,2)	3	6083	5853	408	398	559±1.24%	531±1.43%	27.09	25.11	5.5±2.12	17.84	15.78
N(5,2)	5	6082	5853	408	398	654±1.19%	616±1.20%	37.65	35.45	5.9±1.74	17.84	15.78
N(2,5)	2	6102	5853	421	398	737±2.46%	661±2.17%	42.87	39.82	10.9±3.78	20.41	15.78
N(3,5)	3	6083	5853	408	398	725±2.06%	695±2.32%	43.74	42.81	3.8±3.28	17.84	15.78
N(5,5)	5	6082	5853	408	398	792±1.94%	777±2.16%	48.54	48.78	2.0±2.91	17.84	15.78
N(30,5)	30	6063	5853	406	398	2294±0.87%	2242±0.86%	82.30	82.25	2.3±1.55	17.46	15.78
L80-	2	6102	5853	421	398	1094±5.89%	943±5.75%	61.51	57.82	17.5±10.77	20.41	15.78
N(30,5)	7.6	6053	5853	406	398	1049±5.77%	943±5.75%	61.32	57.82	11.0±10.05	17.45	15.78