

Evaluating the Performance of Multithreading and Prefetching in Multiprocessors*

Ricardo Bianchini

Beng-Hong Lim

COPPE Systems Engineering
Federal University of Rio de Janeiro
Rio de Janeiro, RJ 21945-970 Brazil

T.J. Watson Research Center
IBM Corporation
Yorktown Heights, NY 10598

`ricardo@cos.ufrj.br`

`bhlim@watson.ibm.com`

*This work was partially supported by the Alewife project that is funded in part by ARPA contract # N00014-94-1-09885, in part by NSF Experimental Systems grant # MIP-9012773, and in part by a NSF Presidential Young Investigator Award.

Proposed Running Head:

Multithreading and Prefetching in Multiprocessors

Contact Author:

Ricardo Bianchini

COPPE Sistemas e Computação

Universidade Federal do Rio de Janeiro

Ilha do Fundão, Centro de Tecnologia, Bloco H, Sala H-319

Rio de Janeiro, RJ 21945-970 Brazil

Tel/FAX: + 55 21 590 2552

ricardo@cos.ufrj.br

Abstract

This paper presents new analytical models of the performance benefits of multithreading and prefetching, and experimental measurements of parallel applications on the MIT Alewife multiprocessor. For the first time, both techniques are evaluated on a real machine as opposed to simulations. The models determine the region in the parameter space where the techniques are most effective, while the measurements determine the region where the applications lie. We find that these regions do not always overlap significantly.

The multithreading model shows that only 2–4 contexts are necessary to maximize this technique’s potential benefit in current multiprocessors. For these multiprocessors, multithreading improves execution time by less than 10% for most of the applications that we examined. The model also shows that multithreading can significantly improve the performance of the same applications in multiprocessors with longer latencies. Reducing context-switch overhead is not crucial.

The software prefetching model shows that allowing 4 outstanding prefetches is sufficient to achieve most of this technique’s potential benefit on current multiprocessors. Prefetching improves performance over a wide range of parameters, and improves execution time by as much as 20–50% even on current multiprocessors.

A comparison between the two models shows that prefetching has a significant advantage over multithreading for machines with low memory latencies and/or applications with high cache miss rates, because a prefetch instruction consumes less time than a context-switch.

List of Symbols

1 Introduction

The high latency of remote memory accesses is a major impediment to good application performance on scalable cache-coherent multiprocessors. Several techniques have been proposed for coping with this problem, including *block multithreading* and *software prefetching*. These techniques overlap communication with useful computation to reduce processor stalls. Block multithreading, or multithreading for short, switches between processor-resident threads on cache misses that involve remote communication. Software prefetching, or prefetching for short, requests in advance cache blocks that would otherwise lead to cache misses.

The performance benefits of multithreading and prefetching depend on parameters such as the multiprocessor's cache miss latency, the application's cache miss rate, and the amount of useful computation in between consecutive cache misses. We refer to the last parameter as the application's *run-length*.

This paper develops simple, but useful analytical models that use these parameters to predict the maximum performance improvement achievable through multithreading and prefetching. It introduces the notion of *gain* to characterize the benefit, where gain is defined as the ratio of program execution times without and with the technique. The multithreading model extends Agarwal's processor utilization model [2] to predict the maximum performance gain due to multithreading. The prefetching model is new, and predicts both utilization and gain due to prefetching. As an additional contribution, this paper also introduces a model of utilization and gain when both multithreading and prefetching are used at the same time.

The models are useful for investigating the design tradeoffs associated with multithreading and prefetching, to understand the effect of technological changes, and to predict the potential benefits of each latency-tolerating technique. They determine the range of architectural and application parameters, such as run-length, remote cache miss latency and context-switch overhead, where multithreading and prefetching are most likely to improve performance.

As a concrete example of the utility of the models, this paper applies them towards evaluating multithreading and prefetching on the MIT Alewife Machine [3], a distributed shared-memory multiprocessor that implements both techniques. For the first time, both multithreading and prefetching are evaluated on an actual machine as opposed to simulations. We measure the run-lengths of a number of parallel applications (including SPLASH [30] and NAS parallel benchmarks [6]) and use the models to

predict the potential benefits of multithreading and prefetching for these applications. To corroborate the models' predictions, we measure the execution times of the applications with the latency-tolerating techniques.

The multithreading model shows that 2–4 contexts are sufficient to maximize this technique's potential benefit in the current generation of shared-memory multiprocessors. For these multiprocessors, multithreading improves execution time by less than 10% for most of the applications in this paper. The model also shows that multithreading can significantly improve the performance of the same applications in future multiprocessors with much longer latencies, provided sufficient hardware contexts exist to tolerate the longer latencies. Reducing context-switch overhead will not be crucial in these multiprocessors.

The prefetching model shows that this technique can improve performance over a wide range of parameters. On current multiprocessors, allowing up to 4 outstanding prefetches is sufficient to achieve most of the potential benefit of this technique. Prefetching can improve the execution time of several of the applications by as much as 20–50%, even on current multiprocessors. The rest of the applications may also benefit in future multiprocessors.

A comparison of the two models shows that software prefetching has a significant advantage over multithreading for machines with low remote access latencies and/or applications with high remote cache miss rates because a prefetch instruction typically consumes less time than a context-switch. However, prefetching does not always realize its higher potential of improving performance over multithreading because it requires prefetch instructions to be inserted at the right places. Furthermore, long run-lengths and latencies reduce the negative impact of context-switch overhead in multithreading.

Even though most of the work on multithreading and prefetching has been for cache-coherent multiprocessors, our models show that these techniques can also be successfully applied *completely in software* to distributed shared-memory systems with software-based coherence. In fact, the potential benefits of multithreading and prefetching are enormous in these systems, since software implementations allow for greater flexibility, and the overheads associated with the techniques are negligible compared to page or diff fetch latencies.

This paper makes the following contributions.

- It presents models of multithreading, prefetching, and their combination that can be used to deter-

mine the usefulness of the latency-tolerating techniques in the hardware design and application space. The multithreading model exposes the limits of multithreading by considering performance gain as well as processor utilization, unlike previous models that focus only on processor utilization. As far as we know, this paper is the first to present a processor utilization and gain model for prefetching and for the combination of multithreading and prefetching.

- It presents an analytical and experimental comparison of multithreading and prefetching on a single multiprocessor platform. This allows us to isolate the differences between the two techniques while keeping other parameters constant.
- It introduces the concept of a “sweet spot,” a criterion for determining if a latency-tolerating technique is useful for a particular set of hardware parameters and application characteristics. A comparison of sweet spots against run-lengths of a representative set of applications on Alewife shows that some of the applications cannot benefit significantly from either technique in current multiprocessors.
- It describes, through an extensive analysis of the impact of different architectural parameters, the circumstances under which multithreading and prefetching are most beneficial.

The remainder of this paper is organized as follows. The next section presents some background on multithreading and prefetching. Section 3 presents our models of the two techniques. Section 4, describes the MIT Alewife Machine and our application workload, and presents the results of our analyses and experiments on Alewife. Section 5 considers the potential benefits of multithreading and prefetching for other architectures. Section 6 summarizes the conclusions drawn from this research. Finally, Appendix A presents our model that combines both multithreading and prefetching.

2 Background

Multithreading and prefetching have been the focus of work on tolerating communication latency for many years. Multithreaded processors have been proposed or used in several machines, such as the HEP [31], Monsoon [28], the Tera MTA [5], McGill MTA [19], and MIT Alewife [3]. A description of the history and status of multithreaded architectures and research appears in [9] and [20].

Early multithreaded architectures take an aggressive fine-grain approach that switches contexts at each instruction. This requires a large number of contexts and very low-overhead context switches. The main drawback of this approach is poor single-thread performance. To address this limitation, Alewife introduced the idea of block multithreading, where caches reduce the number of remote cache misses and context-switches occur only on such misses. This allows for a less aggressive implementation of multithreading [4].

Some recent architectures [21, 24] promise both single-cycle context-switches and high single-thread performance by allowing instructions from multiple threads to interleave arbitrarily in the processor pipeline. The Tera MTA [5] takes an extreme approach by providing 128 contexts on each processor. At each cycle, the processor executes an instruction from one of the contexts. Tera does not have caches, and requires a large number of concurrent threads and high network bandwidth to achieve high processor utilization.

Previous experimental research on multithreading performance shows that multithreading is effective at tolerating memory latencies for some applications [17, 33, 18, 32]. Previous analytical research [2, 29, 27, 14] focuses on modeling processor utilization and predicting the number of contexts needed for good processor utilization. In contrast, this paper combines analytical models and experimental measurements in a novel way. It defines a range of run-lengths where multithreading is profitable, and compares this range to the run-lengths in real applications.

Prefetching has been considered in many flavors, with hardware [12, 13], software [26, 10], and hybrid [7] approaches. Software prefetching, where a compiler inserts prefetch instructions for blocks ahead of actual use, is widely accepted as one of the most efficient prefetching techniques. It requires little hardware support and usually incurs the overhead of a single instruction per cache block prefetched. Several machines provide software prefetching, including Stanford DASH [25], the KSR1 [15], and MIT Alewife.

As far as we know, performance studies of software prefetching have all been experimental [26, 10, 7]. These studies indicate sizable performance improvements and consistently better performance than other prefetching strategies. We significantly extend these contributions by modeling software prefetching both in terms of processor utilization and gain, defining a range of run-lengths where prefetching is profitable, and comparing this range to experimentally observed run-lengths.

Boothe and Ranade [8] use a compiler-based approach that combines prefetching with multithreading. In their scheme, hardware provides prefetch and context-switch instructions. The compiler issues prefetches for a groups of shared-memory accesses before context switching. Grouping prefetches increases run-lengths and amortizes the overhead of context switches. To improve performance further, they suggest a “conditional-switch” scheme that switches contexts only upon cache misses. If prefetches are early enough, context-switches can be avoided. This is exactly what Alewife does when prefetching and multithreading are used together. In Appendix A, we present a combined model for prefetching and multithreading that applies to the conditional-switch scheme.

Empirical comparisons of multithreading and prefetching are extremely uncommon. Gupta *et al.* [17] find that prefetching leads to better performance when applied either in isolation or in combination with relaxed consistency. They show that multithreading with an aggressive 4-cycle context-switch overhead, combined with relaxed consistency achieves good performance. Our study uses both modeling and measurement to provide insight into such previous results by allowing parameters to be varied easily and their effects observed.

3 Performance Models of Multithreading and Prefetching

In this section, we develop analytical models of multithreading and prefetching. Appendix A presents an analytical model of the combination of the two techniques. The models predict processor utilization and performance gain, based on machine and application parameters such as remote cache miss latencies and application run-lengths. They allow us to evaluate and compare the performance of both latency-tolerating techniques.

The analysis makes the following assumptions:

- Multithreaded processors switch contexts only on remote cache misses. Software prefetches are issued only for blocks that would otherwise cause a remote cache miss.
- Remote cache miss latencies and run-lengths are constant within the section of an application that we wish to model.

- Multithreading and prefetching result in longer remote cache miss latencies and shorter run-lengths due to increased memory traffic and cache pollution.
- The models ignore the impact of multithreading and prefetching on program synchronization performance and scheduling.

3.1 A Performance Model for Multithreading

Previous research has presented sophisticated models of multithreading performance [2, 29, 27, 14] that focus on predicting processor utilization as accurately as possible. These models attempt to account for the effect of multithreading on other system parameters such as cache and network behavior.

In contrast, we use a simple model of processor utilization and extend it to predict the limits on the performance gain due to multithreading. This allows us to characterize precisely the regime where multithreading may significantly benefit an application. While more sophisticated utilization models may yield tighter upper bounds on the benefit of multithreading, the added accuracy is not worth the additional complexity for our purposes.

The main parameters that affect the performance of multithreaded processors are the number of contexts p , the context-switch overhead C , the latency of remote cache misses $T(p)$, and the average run-length $t(p)$. Our model uses these parameters to predict processor utilization, $U(p)$, and the performance gain due to multithreading, $G(p)$. Table I summarizes these parameters.

Figure 1 illustrates two cases to consider when modeling the utilization of a multithreaded processor. In the first case, enough contexts exist to overlap remote cache miss latency completely. This occurs when $T(p) < (p - 1)t(p) + pC$. In the second case, not enough contexts exist so that the processor suffers some idle time. Thus, the utilization of a multithreaded processor is

$$U(p) = \begin{cases} \frac{t(p)}{t(p)+C} & \text{if } T(p) < (p - 1)t(p) + pC \\ \frac{pt(p)}{t(p)+T(p)} & \text{otherwise} \end{cases} \quad (1)$$

This is essentially the same model for processor utilization presented in [2]. Although the model is useful for determining processor utilization, we carry the analysis a step further and derive an upper

bound on the gain achievable through multithreading, $G(p)$.

The processor utilization for a non-multithreaded processor is

$$U(1) = \frac{t(1)}{t(1) + T(1)}$$

As observed by Agarwal [2], multithreading may shorten run-lengths and increase remote memory latencies. That is, $t(p) \leq t(1)$ due to cache interference and finer computation grain sizes, and $T(p) \geq T(1)$ due to higher network and memory contention. With these inequalities, the maximum performance gain from using a multithreaded processor with p contexts, $G(p) = U(p)/U(1)$, is

$$G(p) = U(p)/U(1) \leq \begin{cases} \frac{t(p)+T(p)}{t(p)+C} & \text{if } T(p) < (p-1)t(p) + pC \\ p & \text{otherwise} \end{cases} \quad (2)$$

Equations (1) and (2) provide a useful tool for evaluating the potential benefits of multithreading for a given machine architecture. The machine’s hardware determines the values of $T(p)$ and C and allows us to determine the range of run-lengths, $t(p)$, such that multithreading results in *both* a significant gain *and* a reasonable level of processor utilization in order to justify hardware support for multithreading. We define that range of run-lengths to be the *sweet spot* of multithreading for that architecture.

Sections 4 and 5 use the model to evaluate multithreading in Alewife and other architectures. We now develop a corresponding model for processor utilization and gain due to prefetching.

3.2 A Performance Model for Prefetching

Unlike multithreading, software prefetching relies on a user or compiler to insert explicit prefetch instructions for data cache blocks that would miss otherwise. Often, such a task is program-dependent. Thus, we need to identify the type of programs we are interested in modeling.

In the following analysis, we focus on regular loops as a natural candidate for prefetching. We find that a number of the applications we consider in this paper contain loops that can be optimized through prefetching. Furthermore, such loops comprise a significant fraction of the total execution time of those applications.

The prototypical loop that we optimize with prefetching looks like:

```
for (i = 0; i < N; i++)  
    compute(i);
```

A prefetching compiler would transform such a loop to a piece of code similar to:

```
prefetch(0, p);                /* prefetch p blocks for first iteration */  
for (i = 0; i < N-step; i = i+step) { /* note increased step size */  
    prefetch(i+step, p);        /* prefetch p blocks for iter i+step to i+2*step-1 */  
    for (j = 0; j < step; j++)    /* compute iters i to i+step-1 */  
        compute(i*step + j);  
}  
for (j = 0; j < step; j++)      /* compute last 'step' iterations */  
    compute(i*step + j);
```

The parameters involved in modeling the behavior of the above prefetching loop are the average number of cache blocks prefetched at a time, p , the overhead of each prefetch instruction, c , the latency of remote cache misses, $T(p)$, and the average amount of computation per prefetched block, t . This last parameter is equivalent to the run-length $t(p)$ in the multithreading model. Table II summarizes these parameters.

Prefetching allows greater flexibility when trying to overlap communication and computation, since the compiler or user can schedule the prefetches according to the remote access latency of the machine and the amount of work per iteration of each loop in the program. This intelligent scheduling of prefetches adjusts p and the iteration step size appropriately (see code segments above).

The optimal number of blocks to prefetch at a time also depends on the available cache space; blocks prefetched into the cache may be replaced before being used. Another limitation is the number of prefetched blocks that can fit in a prefetch/transaction buffer. Under these constraints, the amount of useful work that can be overlapped with communication may be insufficient to hide memory latency

completely.

Figure 2 illustrates the behavior of the prefetching loop. The execution alternates between prefetch and computation intervals, with each prefetch accessing cache blocks to be used one computation block ahead. While it is possible to have loops that prefetch more than one computation block ahead, we can always transform such loops into an equivalent loop that prefetches a single computation block ahead. Again, there are two cases to consider. In the first case, enough blocks are prefetched to provide enough computation to overlap remote memory access latency completely. In the second case, the computation block is too short to avoid processor idle time. The resulting processor utilization is given by

$$U(p) = \begin{cases} \frac{t}{t+c} & \text{if } T(p) < pt + 2pc \\ \frac{2pt}{pt+T(p)} & \text{otherwise} \end{cases} \quad (3)$$

To derive an upper bound on the performance gain due to prefetching, we compute the ratio $U(p)/U(0)$. The processor utilization without prefetching is

$$U(0) = \frac{t}{t + T(0)}$$

Again, we assume $T(p) \geq T(0)$, so that the maximum performance gain achievable from prefetching, $G(p) = U(p)/U(0)$, is

$$G(p) = U(p)/U(0) \leq \begin{cases} \frac{t+T(p)}{t+c} & \text{if } T(p) < pt + 2pc \\ \frac{2p(t+T(p))}{pt+T(p)} & \text{otherwise} \end{cases} \quad (4)$$

In a similar way to the multithreading model, Equations (3) and (4) provide a useful tool for evaluating the potential benefits of prefetching. For a given machine architecture, we can determine the range of t such that prefetching results in *both* a significant gain *and* a reasonable level of processor utilization in order to justify hardware support for prefetching.

4 Evaluation of Multithreading and Prefetching on Alewife

This section applies our models towards an evaluation of multithreading and prefetching in the MIT Alewife Machine [3]. Alewife provides support for both multithreading and prefetching, and presents an ideal platform for evaluating both techniques.

The analysis proceeds as shown in Figure 3. First, we use the latency, context-switch and prefetch overheads, and the number of contexts parameters of Alewife to determine the sweet spots for each of the techniques. Next, we measure the run-lengths of a number of parallel applications executing on Alewife to find out whether they lie within the sweet spots. Last, to corroborate the models, we run a subset of the applications with multithreading and prefetching enabled, and measure the resulting performance gains.

4.1 The MIT Alewife Machine

Figure 4 presents an overview of the architecture and organization of Alewife. Each node consists of a Sparcle processor [4] and an associated floating point unit with up to 4 processor contexts, 64K bytes of direct-mapped cache, 8M bytes of DRAM, an Elko-series mesh routing chip (EMRC) from Caltech, and a custom-designed Communication and Memory Management Unit (CMMU). The EMRC network routers use wormhole routing and are connected in a 2-D mesh topology. The single-chip CMMU implements Alewife’s scalable LimitLESS cache coherence protocol [11], and provides Sparcle with a low-latency interface to the network. To support both multithreading and prefetching, the caches are lockup free [22].

Alewife supports multithreading via Sparcle’s register windows and fast trap handling. A cache miss involving remote communication generates a trap that performs an 18-cycle context-switch to another register window. (Sparcle actually supports a 14-cycle context-switch, but a minor bug in the current version of the CMMU requires a 4-cycle work-around.) Alewife supports both read and write prefetching by using special Sparcle memory loads and stores, and a transaction buffer in the CMMU to store prefetch data [23].

Alewife’s parameters for our models are its 18-cycle context-switch overhead, 2-cycle prefetch instruction overhead, and average remote cache miss latency of 53 cycles. (A two-party cache miss

takes 42 cycles and a three-party cache miss takes 63 cycles.) Thus, we assume that $C = 18$, $c = 2$, and $T(p) = 53$.

4.2 Model Predictions for Alewife

Figures 5 and 6 present the results of the multithreading and prefetching models with Alewife’s parameters. Each graph plots simultaneously both processor utilization and gain as a function of run-length. These graphs show that for both techniques, gains are significant for short run-lengths. However, processor utilization is extremely low for these same run-lengths. As expected, long run-lengths lead to high levels of processor utilization. However, gains due to the latency-tolerating techniques are very limited. These observations highlight the fact that the sweet spot for both techniques lies somewhere in between the extremes of run-lengths.

In our evaluation, we define the sweet spot to be between the shortest run-length that yields at least 50% processor utilization (left bound) and the longest run-length that yields at least 20% potential gain (right bound). Although we chose this range of values based on our own criteria of acceptable values for utilization and gain, others are free to decide on different criteria. The essential point here is that there is a bounded range of run-lengths where a latency-tolerating technique is profitable.

In contrast, previous research has focused on whether multithreading or prefetching can increase processor utilization without considering whether they are operating in a regime where the performance gain from latency tolerance is significant enough to care about.

Figure 5 shows that Alewife has a very narrow multithreading sweet spot of between about 18 and 105 cycles. This narrow range can be attributed to Alewife’s short remote cache miss latencies. Another interesting observation is that within the sweet spot, two contexts are sufficient to attain all of the potential gains of multithreading. In a sense, the Alewife hardware designers have aggressively minimized memory latencies to a point where multithreading has rather limited benefits.

Figure 6 shows that Alewife’s prefetching sweet spot ranges from 4 to 250 cycles. It also shows that the number of cache blocks that can be prefetched at a time, p , affects the performance of prefetching significantly. The reason for this is that the amount of computation that can be used to hide communication latency is expected to grow linearly with p . Within the sweet spot, allowing 4 outstanding prefetches is sufficient to achieve most of the potential performance gains of prefetching.

The figures show that the prefetching sweet spot is wider than the multithreading sweet spot mainly because of the lower overhead of prefetching: 2 versus 18 cycles. For run-lengths of less than 100 cycles, prefetching gains are significantly better than multithreading gains. For longer run-lengths however, the overhead of a context-switch becomes less significant and the two techniques perform comparably.

The analysis thus far poses an unanswered question: where do typical applications lie along the spectrum of run-lengths? The next section presents measurements of application run-lengths on the Alewife machine to determine if they fall within the sweet spot.

4.3 Application Performance on Alewife

In order to determine the range of run-lengths found in real applications on a real machine, we use a 32-processor Alewife multiprocessor to study a large number of parallel applications, including the SPLASH [30] and NAS parallel benchmarks [6] and a number of engineering kernels.

Table III lists the applications and their input parameters. Please refer to the SPLASH and NAS documents for descriptions of those benchmarks. Among the engineering kernels, EM3D simulates electromagnetic wave propagation through 3-D objects, FFT performs a Fast Fourier Transform, Gauss performs Gaussian elimination, BLU performs blocked LU decomposition, MergeSort sorts a list of integers, and SOR performs red-black Jacobi relaxation. SOR-Dyn is a dynamically scheduled version of SOR.

To measure multithreading run-lengths, $t(p)$, we use Alewife's built-in statistics hardware, which includes event and cycle counters, to count the number of remote cache misses and the execution time of each application. Assuming an average remote cache miss time of 53 cycles, we derive an average value for run-length.

To measure prefetching run-lengths, t , we hand-instrument each loop that has been optimized with prefetching. The instrumented loops measure the actual execution time for each prefetched cache block. Subtracting two cycles for each prefetch instruction from the execution time yields the value of t .

Figures 7 and 8 present the average multithreading and prefetching run-lengths in the applications on 32 Alewife processors and the corresponding potential gains. Results on 16 processors do not differ significantly from 32 processors. The multithreading experiments use 2 contexts, while the prefetching

experiments prefetch at most 4 blocks at a time. The applications are listed in order of increasing run-lengths. We should note that using the average run-length in our models yields a correct upper bound on the overall performance of an application even if the application consists of phases with different run lengths. However, we do not provide a proof of this here due to space limitations.

The prefetching run-lengths are generally shorter than the multithreading run-lengths, although they are comparable parameters in the models. One reason for this difference is that multithreading switches contexts only on remote cache misses while prefetches occur for every remote memory access that is *expected* to miss. This causes multithreading run-lengths to tend to be longer than prefetching run-lengths.

Another reason for this difference is that multithreading run-lengths are measured over the entire application, while prefetching run-lengths are measured only at loops optimized with prefetching. This strategy reflects the common practical use of these techniques: multithreading is implemented in hardware and is active throughout execution, while prefetching is implemented by the compiler or programmer to optimize loops. To evaluate the impact of this measurement strategy on our results, we measured multithreading run-lengths within the loops optimized with prefetching and found that both measurement strategies lead to minor differences in run-lengths, except for MP3D, BLU, and Gauss. In these applications, multithreading run-lengths within the loops are actually longer than the global multithreading run-lengths because cache miss rates outside the loops are higher.

Recall that run-lengths measure the time in between *remote* cache misses and prefetches. Thus, a high local cache miss rate can inflate the run-lengths since the local miss latency is included in the time. Our measurements find that the majority of the applications have local miss rates of well below 5% and that only EM3D has a significant local cache miss rate of 23% on 32 processors.

Most of the applications have long multithreading run-lengths of 200–2000 cycles that lie outside of the sweet spot. Only MP3D, CG, and Cholesky lie within the sweet spot and may benefit from multithreading. The prefetching run-lengths are shorter, and range from 50–500 cycles, with MP3D, CG, EM3D, Gauss, SOR-Dyn and MergeSort all falling within the sweet spot.

The most important observation here is that for Alewife’s parameters, most of the applications operate on a suboptimal part of the parameter space for multithreading. This implies that multithreading is not likely to yield much benefit for the applications we considered. The main reason for this is Alewife’s

short remote cache miss latencies relative to the application run-lengths. The prospect is better for prefetching, but only for applications where we can identify loops for optimization.

To validate our models’ predictions of the upper bound of application gains, we execute a subset of the applications on Alewife with the latency-tolerating techniques and measure the resulting execution times. For multithreading, we statically schedule a thread on each hardware context, with two contexts on each processor. (Experiments with three contexts show only minor performance differences.) This prevents scheduling overhead from diminishing any potential performance gain. We produced multithreading versions of only a subset of our applications, due to minor incompatibilities between Alewife’s thread-based model and the benchmarks’ process-based model. For prefetching, we identify loops in a subset of the applications that are amenable to optimization with prefetch instructions, and modified them in accordance with the description in Section 3.

Figures 9 and 10 present the results. We see that the applications respect the models’ upper bounds on the achievable gain for each technique. Some applications come close to achieving optimal performance, but many others do not as a result of factors such as cache interference and/or contention. We also find that the latency-tolerating techniques sometimes affect the synchronization and scheduling behavior of the application adversely. Characterizing or avoiding performance degradation from such factors is beyond the scope of this paper.

The models and experimental results for Alewife show that multithreading and prefetching yield rather limited benefits for the current generation of shared-memory machine architectures like Alewife and DASH, where remote cache miss latencies are relatively short. The models can be easily applied to other machine architectures, and the next section investigates the effect of different architectural parameters such as latency and context-switch/prefetch overhead on the performance of each technique.

5 Potential Benefits on Other Architectures

The latency of remote cache misses, $T(p)$, the overhead of context switching, C , and the overhead of prefetching, c , are significant factors on the performance of multithreading and prefetching. In order to investigate the potential of these techniques for current and future architectures, Sections 5.1 and 5.2 model the sensitivity of utilization and gain to each of these parameters. Whether a certain technique

can be beneficial to real applications depends on their associated run-lengths. Section 5.3 investigates the effect of the number of processors on the applications' run-lengths.

5.1 Effect of Remote Cache Miss Latency

The latency of remote cache misses determines the amount of communication that must be tolerated by multithreading and prefetching. Figures 11 and 12 show the effect of higher latency on gain and utilization as a function of p and t . These figures reproduce the results of Figures 5 and 6, except under a 200-cycle remote access latency as opposed to 53. The vertical lines in each graph delimit the sweet spot and the range of application run-lengths.

The figures show that longer latency tends to widen the sweet spots (as a result of increased gains for long run-lengths) and increase the overlap with the application run-lengths. A larger number of contexts or prefetch buffers is needed to achieve a high level of utilization for short run-lengths. Figure 11 also shows that, under higher latency than in Alewife, increasing the number of contexts beyond two becomes an effective way of improving processor utilization.

Our models show that longer latencies also increase the potential for improving performance via latency-tolerating techniques. Figures 13 and 14 model the effect of varying remote cache miss latencies, assuming that enough contexts or prefetch buffers are available to tolerate the latencies. Again, the vertical lines in each graph delimit the sweet spot and the range of application run-lengths.

Figure 13 shows the effect of varying remote cache miss latencies on multithreading, with $p = 10$ and $C = 15$. It shows that latency has a significant impact on the performance of multithreading: even though the processor utilization is virtually the same for all the latencies we considered, the potential gain due to multithreading is much greater with longer latencies, and the sweet spot extends to longer run-lengths.

Figure 14 shows the effect of varying remote cache miss latencies on prefetching, with $p = 10$ and $c = 2$. It shows that increasing latency has the effect of shifting the entire sweet spot towards longer run-lengths. For the run-lengths in the applications in this study, the effect of latency on utilization is negligible, while the effect on gain is much more pronounced.

These results imply that most of the SPLASH and NAS benchmarks would benefit from multithreading or prefetching on machines with remote latencies longer than 200 cycles. Recall that only three of

the applications have run-lengths longer than 1000 cycles. The rest of the applications fall within the sweet spots.

5.2 Effect of Multithreading and Prefetching Overhead

Another parameter that affects performance is the overhead associated with each technique: the context-switch overhead for multithreading and the prefetch instruction overhead for prefetching. Assuming enough contexts or prefetch buffers to cover the latency, lower overheads tend to extend the sweet spot towards shorter run-lengths. However, there is little or no impact on the sweet spot at longer run-lengths. Figures 15 and 16 illustrate this effect.

Figure 15 shows the effect of varying context-switch overhead on multithreading, with $p = 10$ and $T = 50$. The graph shows that context-switch overhead significantly impacts performance for short run-lengths. A lower overhead tends to extend the sweet spot towards shorter run-lengths. However, context-switch overhead does not have a significant impact within the range of application run-lengths we found.

With higher latencies, the models predict that context-switch overhead has even less impact on performance. This implies that Alewife's context-switch overhead is small enough for the applications we considered. A more aggressive implementation is necessary only for applications with run-lengths shorter than 100 cycles.

Figure 16 shows the effect of varying prefetch instruction overhead on prefetching, with $p = 10$ and $T = 50$. The graph shows that prefetch overhead significantly impacts performance for short run-lengths. We also see that prefetch overhead has very little effect for run-lengths that are longer than 50 cycles. As with multithreading, prefetch overhead is not a significant factor for the applications we considered.

The figure also shows the interesting result that even a two-cycle run-length is within the prefetching sweet spot for this set of parameters. This result means that even an application with no locality whatsoever can achieve reasonable performance with prefetching, provided that enough prefetch buffers and network bandwidth exist. A similar effect can be observed in Figure 15 but hardware support for a one-cycle context-switch is more expensive than for a one-cycle prefetch instruction overhead.

To summarize, lower overheads extend the sweet spot towards shorter run-lengths, in contrast with

longer latencies that extend the sweet spot towards longer run-lengths. Since the application run-lengths in our experiments either lie outside the sweet spots or only overlap with their rightmost region, overhead makes little difference on performance for the applications in our suite.

5.3 Effect of the Number of Processors on Run-Lengths

The applications' run-lengths determine whether a specific latency-tolerating technique can be effective in practice. This section investigates the impact of the number of processors and the input data size on multithreading run-lengths. Recall that we assume the prefetching run-length to be equal to the amount of work performed per shared cache block, and is thus constant across different machine sizes.

Figure 17 plots multithreading run-lengths as a function of the number of processors for a few representative applications. We find two basic types of applications: one where run-lengths decrease with increasing number of processors (as in SOR and Gauss), and another where run-lengths remain relatively constant beyond a small number of processors (as in MP3D and SOR-Dyn). In applications of the first type the performance benefit achievable through multithreading depends on the input size and number of processors. Run-lengths may be lengthened by increasing the input size so that communication latency becomes less of a problem. The second type tends to include fine-grained applications where latency tolerance is more crucial.

5.4 Implications for Other Architectures

We conclude this section by speculating on the implications of our models and analysis on the architecture of other computing platforms.

A recently promoted alternative to custom-designed multiprocessors for parallel computing is to use commodity workstations connected via high-speed networks. Multithreading and prefetching may have a larger impact here. Such a platform is likely to have remote memory access latencies in the order of $10 \mu s$. Assuming a $5 ns$ cycle time, this results in a latency of about 2000 cycles. At such latencies, our models predict that multithreading and prefetching are likely to improve performance significantly. Since latencies are so high, these techniques could even be implemented in software, handling memory pages as opposed to cache blocks. The hardware would only need to support non-blocking memory

loads and stores.

Figures 18 and 19 show the performance of all-software multithreading and prefetching, respectively, assuming a 5000-cycle latency to fetch pages or `diffs`, a 200-cycle software context-switch, and a 20-cycle prefetch (i.e. procedure call) overhead. We see from the figures that, again in software, the overhead of multithreading hurts the performance of this technique for all but the longest run-lengths. In practice however, the two techniques should perform comparably, as it seems unreasonable for applications to exhibit run-lengths shorter than 10^4 cycles given the remote latencies in these systems.

Multiprocessor latency-hiding techniques have also been suggested to improve performance in uniprocessor systems. As advances in processor speeds outstrip advances in memory speeds, local memory latencies become more significant. It is conceivable that cache misses will take upwards of 50 cycles in the next generation of high-speed processors, close to today's multiprocessors' remote memory latencies. However, uniprocessor local cache misses are likely to be much more frequent than multiprocessor remote cache misses. On such a platform, context-switch and prefetch overheads will be significant and will need to be reduced as much as possible.

6 Summary and Conclusions

This paper introduces analytical models of processor utilization and performance gain for two latency-tolerating techniques: multithreading and prefetching. The models, based on hardware and application characteristics such as memory latencies and application run-lengths, show the limits as well as the benefits of these techniques. The models define sweet spots: conditions under which the techniques may improve performance significantly. They provide a method for architects and programmers to evaluate if a particular latency-tolerating technique will be effective for a particular set of architectural parameters and applications.

A comparison of the sweet spots with run-lengths from a representative set of applications running on the MIT Alewife Machine shows that for current multiprocessors, very few of the applications can benefit significantly from multithreading, while some but not all of the applications can benefit from prefetching. The main reason behind this is the relatively short remote cache miss latencies (< 150 cycles) in these machines. With short latencies, prefetching has an advantage over multithreading

because a context-switch usually consumes more processor cycles than a prefetch instruction.

An evaluation of the effect of technological changes shows that remote cache miss latencies have a greater impact on the performance benefits of multithreading and prefetching than their associated overheads. A remote cache miss latency of greater than 200 cycles will allow multithreading and prefetching to improve application performance significantly. With longer latencies, the higher overhead of context-switching becomes less of a factor.

A comparison of the two models shows that prefetching is preferable over multithreading for machines with low remote access latencies and/or applications with poor locality and consequently short run-lengths. The performance of both techniques is comparable for applications with high remote access latencies and/or good locality. For the applications in our suite, prefetching is the technique of choice for current multiprocessors, while both multithreading and prefetching can be successful for future multiprocessors.

Although multithreading may not always be effective in improving the performance of parallel applications, there are other domains where multithreading may be profitable. An example is micro-kernel operating systems [1] where context-switching occurs frequently between client, kernel, and server threads. Here, context-switching occurs not as a method to tolerate latency, but as a protection mechanism for providing kernel services at the user-level.

Another domain where multithreading may be beneficial is distributed and interactive applications where multiple threads are used to handle multiple concurrent activities and to simplify the program structure. The Java [16] run-time system, for example, provides multithreading so that concurrent activities such as graphics animations, user input, and downloading data from the Internet can each be handled by separate threads.

For such domains, where context-switches occur naturally and frequently, the overhead of context-switching represents a significant fraction of the execution time and, thus, hardware support for multiple contexts can improve performance significantly.

A Modeling the Combination of Multithreading and Prefetching

This appendix models the combination of multithreading and prefetching techniques, where context-switches happen only when prefetching cannot fetch data early enough to avoid a processor stall. Table I summarizes the main parameters that affect the performance of the combined strategy. Our model uses these parameters to predict processor utilization, $U(p', p)$, and the performance gain, $G(p', p)$.

As Figure 20 illustrates, there are three cases to consider when modeling the processor utilization. In the first case, enough blocks are prefetched to provide enough computation to overlap remote memory access latency completely. This occurs when $T(p', p) < pt + 2pc$. In the second case, the computation block is too short for prefetches to be effective, but there are enough contexts to avoid processor idle time. This occurs when $pt + 2pc < T(p', p) < 2(p' - 1)(pc + pt) + p'C + pt + 2pc$. In the third case, not enough contexts exist so that the processor suffers some idle time. This occurs when $T(p', p) > 2(p' - 1)(pc + pt) + p'C + pt + 2pc$. Thus, the utilization entailed by the combination of multithreading and prefetching is given by

$$U(p', p) = \begin{cases} \frac{t}{t+c} & \text{if } T(p', p) < pt + 2pc \\ \frac{2pt}{2pt+2pc+C} & \text{if } pt + 2pc < T(p', p) < \\ & 2(p' - 1)(pc + pt) + p'C + pt + 2pc \\ \frac{2p'pt}{pt+T(p', p)} & \text{otherwise} \end{cases} \quad (5)$$

The processor utilization without multithreading or prefetching is

$$U(1, 0) = \frac{t}{t + T(1, 0)}$$

Again, we assume $T(p', p) \geq T(1, 0)$, so that the maximum performance gain achievable from the combination of multithreading and prefetching is $G(p', p) = U(p', p)/U(1, 0)$ or

$$G(p', p) \leq \begin{cases} \frac{t+T(p',p)}{t+c} & \text{if } T(p', p) < pt + 2pc \\ \frac{2p(t+T(p',p))}{2(pc+pt)+C} & \text{if } pt + 2pc < T(p', p) < \\ & 2(p' - 1)(pc + pt) + p'C + pt + 2pc \\ \frac{2p'p(t+T(p',p))}{pt+T(p',p)} & \text{otherwise} \end{cases} \quad (6)$$

Figure 21 shows the processor utilization and the performance gain predicted by this model for the architectural parameters of Alewife. The sweet spot is determined in the same way as for multithreading and prefetching in isolation.

The model shows that long application run-lengths lead to the same utilization and gain behavior as of prefetching alone. As we decrease run-lengths however, there is a point where run-lengths are too short for prefetching alone to hide all the latency, and a cache miss leads to a context switch. There is a sudden drop in utilization and gain at this point because the remaining latency is shorter than the context switch overhead. As we continue to decrease run-lengths, the benefits of multithreading start to show and performance is better than when either multithreading or prefetching is used in isolation. For very short run-lengths, both techniques fail to deal with the remote access latency. These results show that multithreading and prefetching together provide an effective combination, especially for applications exhibiting medium to short run-length.

Acknowledgments

We would like to thank Tom LeBlanc, Fredrik Dahlgren, Claudio Amorim, and the anonymous referees for comments that helped improve the quality of this paper. We would also like to thank the members of the Alewife group, especially John Kubiawicz, for help with the Alewife statistics hardware.

References

1. M. Acetta, R. Baron, W. Bolosky, D. Golub, R. Rashid, A. Tevanian, and M. Young. Mach: A New Kernel Foundation for UNIX Development. In *Proceedings of the Summer 1986 USENIX*

- Conference*, pages 93–112, July 1986.
2. A. Agarwal. Performance Tradeoffs in Multithreaded Processors. *IEEE Transactions on Parallel and Distributed Systems*, 3(5):525–539, September 1992.
 3. A. Agarwal, R. Bianchini, D. Chaiken, K. Johnson, D. Kranz, J. Kubiawicz, B.-H. Lim, K. Mackenzie, and D. Yeung. The MIT Alewife Machine: Architecture and Performance. In *Proceedings of the 22nd Annual International Symposium on Computer Architecture (ISCA)*, pages 2–13. ACM, June 1995.
 4. A. Agarwal, J. Kubiawicz, D. Kranz, B.-H. Lim, D. Yeung, G. D’Souza, and M. Parkin. Sparcle: An Evolutionary Processor Design for Multiprocessors. *IEEE Micro*, 13(3):48–61, June 1993.
 5. R. Alverson, D. Callahan, D. Cummings, B. Koblenz, A. Porterfield, and B. Smith. The Tera Computer System. In *Proceedings of the International Conference on Supercomputing*, pages 1–6, Amsterdam, June 1990. ACM.
 6. D. Bailey et al. The NAS Parallel Benchmarks. Technical Report RNR-94-007, NASA Ames Research Center, March 1994.
 7. R. Bianchini and T. J. LeBlanc. A Preliminary Evaluation of Cache-Miss-Initiated Prefetching Techniques in Scalable Multiprocessors. Technical Report TR 515, Department of Computer Science, University of Rochester, May 1994.
 8. B. Boothe and A. Ranade. Improved Multithreading Techniques for Hiding Communication Latency in Multiprocessors. In *Proceedings of the 19th International Annual Symposium on Computing Architecture (ISCA)*, pages 214–223, Gold Coast, Australia, June 1992. IEEE.
 9. G. T. Byrd and M. A. Holliday. Multithreaded Processor Architectures. *IEEE Spectrum*, pages 38–46, August 1995.
 10. D. Callahan, K. Kennedy, and A. Porterfield. Software Prefetching. In *Proceedings of the 4th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS IV)*, pages 40–52, Boston, MA, April 1991.

11. D. Chaiken, J. Kubiawicz, and A. Agarwal. LimitLESS Directories: A Scalable Cache Coherence Scheme. In *Proceedings of the 4th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS IV)*, pages 224–234. ACM, April 1991.
12. T.-F. Chen and J.-L. Baer. A Performance Study of Software and Hardware Prefetching Schemes. In *Proceedings of the 21st International Symposium on Computer Architecture (ISCA)*, Chicago, IL, April 1994. ACM.
13. F. Dahlgren, M. Dubois, and P. Stenstrom. Sequential Hardware Prefetching in Shared-Memory Multiprocessors. *IEEE Transactions on Parallel and Distributed Systems*, 6(7):733–746, July 1995.
14. P. K. Dubey, A. Krishna, and M. S. Squillante. Analytic Performance Modeling for a Spectrum of Multithreaded Processor Architectures. Computer Science RC 19661, IBM, July 1994.
15. S. Frank, H. B. III, and J. Rothnie. The KSR1: Bridging the Gap Between Shared Memory and MPPs. In *Proceedings of the 38th Annual IEEE Computer Society Computer Conference (COMPCON)*, pages 284–294, San Francisco, CA, 1993. IEEE.
16. J. Gosling and H. McGilton. *The Java Language Environment, A White Paper*. Sun Microsystems Computer Company, Mountain View, CA, October 1995.
17. A. Gupta, J. Hennessy, K. Gharachorloo, T. Mowry, and W.-D. Weber. Comparative Evaluation of Latency Reducing and Tolerating Techniques. In *Proceedings of the 18th International Symposium on Computer Architecture (ISCA)*, pages 254–263, Toronto, Canada, May 1991. ACM.
18. M. Haines and W. Bohm. An Evaluation of Software Multithreading in a Conventional Distributed Memory Multiprocessor. In *Proceedings of the 5th IEEE Symposium on Parallel and Distributed Processing*, December 1993.
19. H. Hum *et al.* The Multi-Threaded Architecture Multiprocessor. Technical Report ACAPS Technical Memo 88, McGill University School of Computer Science, December 1994.
20. R. A. Iannucci, editor. *Multithreaded Computer Architecture - A Summary of the State of the Art*. Kluwer Academic Publishers, 1994.

21. S. W. Keckler and W. J. Dally. Processor Coupling: Integrating Compile Time and Runtime Scheduling for Parallelism. In *Proceedings of the 19th Annual International Symposium on Computer Architecture (ISCA)*, pages 202–213, Gold Coast, Australia, June 1992. IEEE.
22. D. Kroft. Lockup-Free Instruction Fetch/Prefetch Cache Organization. In *Proceedings of the 8th Annual Symposium on Computer Architecture (ISCA)*, pages 81–87, June 1981.
23. J. Kubiawicz, D. Chaiken, and A. Agarwal. Closing the Window of Vulnerability in Multiphase Memory Transactions. In *Proceedings of the 5th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS V)*, pages 274–284, Boston, MA, October 1992. ACM.
24. J. Laudon, A. Gupta, and M. Horowitz. Interleaving: A Multithreading Technique Targeting Multiprocessors and Workstations. In *Proceedings of the Sixth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS VI)*, pages 308–318, San Jose, CA, October 1994. ACM.
25. D. Lenoski, J. Laudon, K. Gharachorloo, W. Weber, A. Gupta, J. Hennessy, M. Horowitz, and M. Lam. The Stanford DASH Multiprocessor. *IEEE Computer*, 25(3):63–79, March 1992.
26. T. Mowry and A. Gupta. Tolerating Latency Through Software-Controlled Prefetching in Shared-Memory Multiprocessors. *Journal of Parallel and Distributed Computing*, 12(2):87–106, June 1991.
27. S. Nemawarkar, R. Govindarajan, G. Gao, and V. Agarwal. Analysis of Multithreaded Architectures with Distributed Shared Memory. In *Proceedings of the 5th IEEE Symposium on Parallel and Distributed Processing*, pages 114–121, Dallas, TX, 1993. IEEE.
28. G. Papadopoulos and D. Culler. Monsoon: An Explicit Token-Store Architecture. In *Proceedings of the 17th Annual International Symposium on Computer Architecture (ISCA)*, pages 82–91, June 1990.
29. R. Saavedra-Barrera, D. Culler, and T. von Eicken. Analysis of Multithreaded Architectures for Parallel Computing. In *Proceedings of the 2nd Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 169–177, July 1990.

30. J. Singh, W.-D. Weber, and A. Gupta. SPLASH: Stanford Parallel Applications for Shared-Memory. Technical Report CSL-TR-92-526, Stanford University, June 1992.
31. B. Smith. Architecture and Applications of the HEP Multiprocessor Computer System. *Society of Photooptical Instrumentation Engineers*, 298:241–248, 1981.
32. R. Thekkath and S. Eggers. The Effectiveness of Multiple Hardware Contexts. In *Proceedings of the 6th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS VI)*, pages 328–337, San Jose, CA, October 1994. ACM.
33. C. A. Waldspurger and W. E. Weihl. Register Relocation: Flexible Contexts for Multithreading. In *Proceedings of the 20th Annual International Symposium on Computer Architecture (ISCA)*, pages 120–130, San Diego, CA, May 1993. IEEE.

Biographies

Ricardo Bianchini is a research associate at COPPE Systems Engineering/UFRJ, where he co-leads the NCP₂ project, a large research effort focused on building hardware support for software-coherent distributed shared-memory systems. He received his Ph.D. degree in Computer Science from the University of Rochester in 1995. His research interests include parallel and distributed computing, advanced operating systems, and parallel computer architecture. He is a member of the ACM and IEEE computer societies.

Beng-Hong Lim is a research staff member at the IBM T.J. Watson Research Center. He received his Ph.D. degree in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology in 1994. His current research interests lie in the design, implementation and performance evaluation of high-performance computer systems and applications, with focus on parallel and distributed environments. He is a member of the ACM, IEEE, and Sigma Xi.

List of Figures

1	<i>Time lines illustrating the overlap of communication latency $T(p)$, with computation $t(p)$, when multithreading with $p = 2$ contexts. In the second case, $T(p)$ is too long to be completely hidden and results in some idle processor cycles.</i>	32
2	<i>Time lines illustrating the overlap of communication latency $T(p)$, with computation pt, when prefetching p blocks at a time. pc_i prefetches cache blocks for computation block pt_i. In the second case, $T(p)$ is too long to be completely hidden and results in some idle processor cycles in every other computation block.</i>	33
3	<i>Applying the models and application parameters to estimate the potential benefits of latency-tolerating techniques.</i>	34
4	<i>Architecture of the MIT Alewife Machine.</i>	35
5	<i>Modeling the performance of multithreading on Alewife.</i>	36
6	<i>Modeling the performance of prefetching on Alewife.</i>	37
7	<i>Potential gain based on multithreading run-lengths on Alewife.</i>	38
8	<i>Potential gain based on prefetching run-lengths on Alewife.</i>	39
9	<i>Measured gain on Alewife due to multithreading.</i>	40
10	<i>Measured gain on Alewife due to prefetching.</i>	41
11	<i>Modeling the performance of multithreading under higher latency. Vertical lines delimit the sweet spot and the range of application run-lengths.</i>	42
12	<i>Modeling the performance of prefetching under higher latency. Vertical lines delimit the sweet spot and the range of application run-lengths.</i>	43
13	<i>Effect of remote cache miss latency on multithreading. Vertical lines delimit the sweet spot and the range of application run-lengths.</i>	44
14	<i>Effect of remote cache miss latency on prefetching. Vertical lines delimit the sweet spot and the range of application run-lengths.</i>	45
15	<i>Effect of context-switch overhead on multithreading. Vertical lines delimit the sweet spot and the range of application run-lengths.</i>	46
16	<i>Effect of prefetch instruction overhead on prefetching. Vertical lines delimit the sweet spot and the range of application run-lengths.</i>	47

17 *Effect of number of processors on multithreading run-lengths.* 48

18 *Modeling the performance of all-software multithreading.* 49

19 *Modeling the performance of all-software prefetching.* 50

20 *Time lines illustrating the overlap of communication latency $T(p', p)$, with computation pt , when multithreading with $p' = 2$ contexts. In the second case, $T(p', p)$ is too long to be completely hidden by prefetching, but can be hidden by multithreading. In the third case, both prefetching and multithreading fail to completely hide $T(p', p)$.* 51

21 *Modeling the performance of the combination of multithreading and prefetching on Alewife.* 52

List of Tables

I	<i>Definitions of the parameters used in the multithreading model.</i>	53
II	<i>Definitions of the parameters used in the prefetching model.</i>	54
III	<i>Main application input parameters.</i>	55
IV	<i>Definitions of the parameters used in the model for the combination of multithreading and prefetching.</i>	56

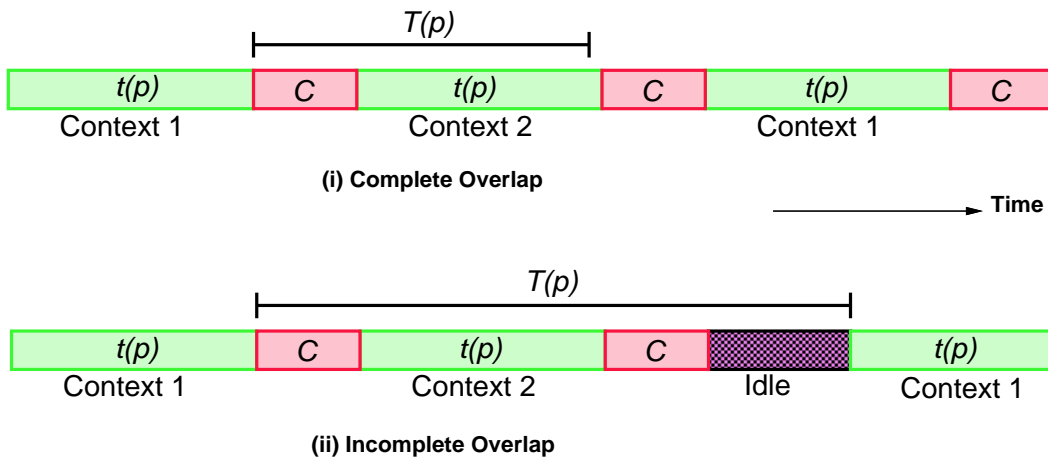


Figure 1: Time lines illustrating the overlap of communication latency $T(p)$, with computation $t(p)$, when multithreading with $p = 2$ contexts. In the second case, $T(p)$ is too long to be completely hidden and results in some idle processor cycles.

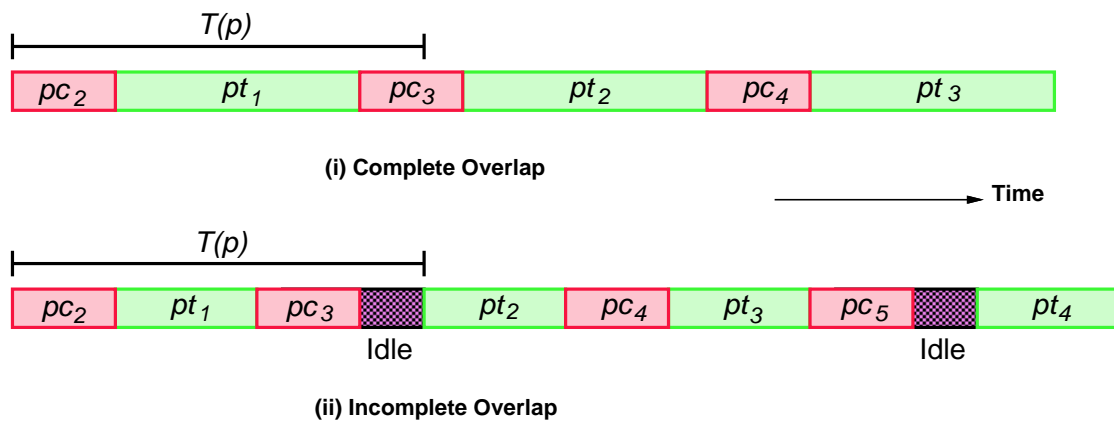


Figure 2: Time lines illustrating the overlap of communication latency $T(p)$, with computation pt , when prefetching p blocks at a time. pc_i prefetches cache blocks for computation block pt_i . In the second case, $T(p)$ is too long to be completely hidden and results in some idle processor cycles in every other computation block.

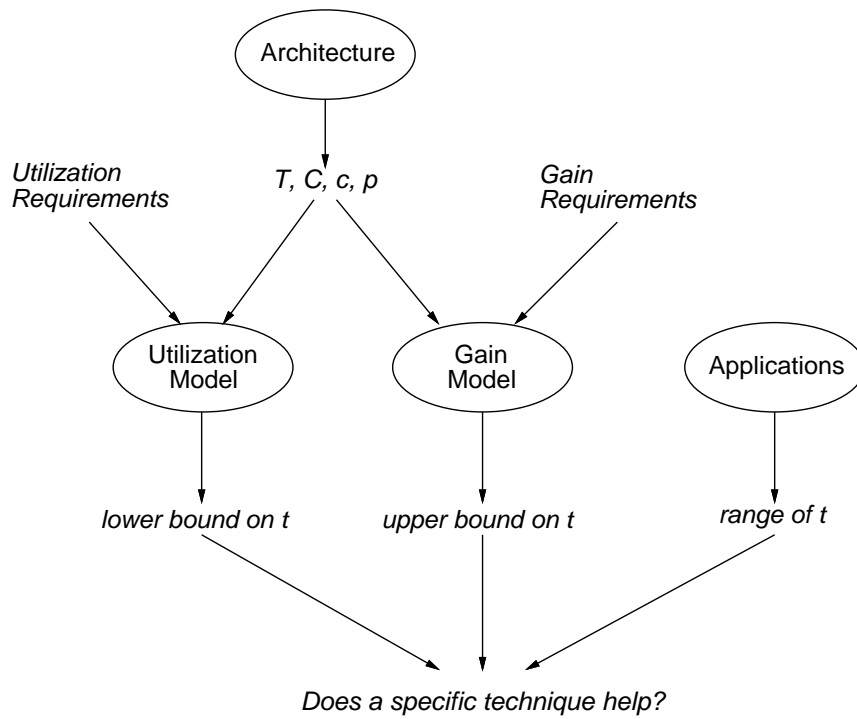


Figure 3: Applying the models and application parameters to estimate the potential benefits of latency-tolerating techniques.

32 Nodes
64 KB Cache/Node
2 MB Shared-Memory/Node
2-D Mesh Wormhole-Routed Interconnect
20 MHz Clock Speed

4 contexts/processor
 $T(p) = 42-63$ cycles
 $C = 18$ cycles




Figure 4: *Architecture of the MIT Alewife Machine.*




Figure 5: Modeling the performance of multithreading on Alewife.

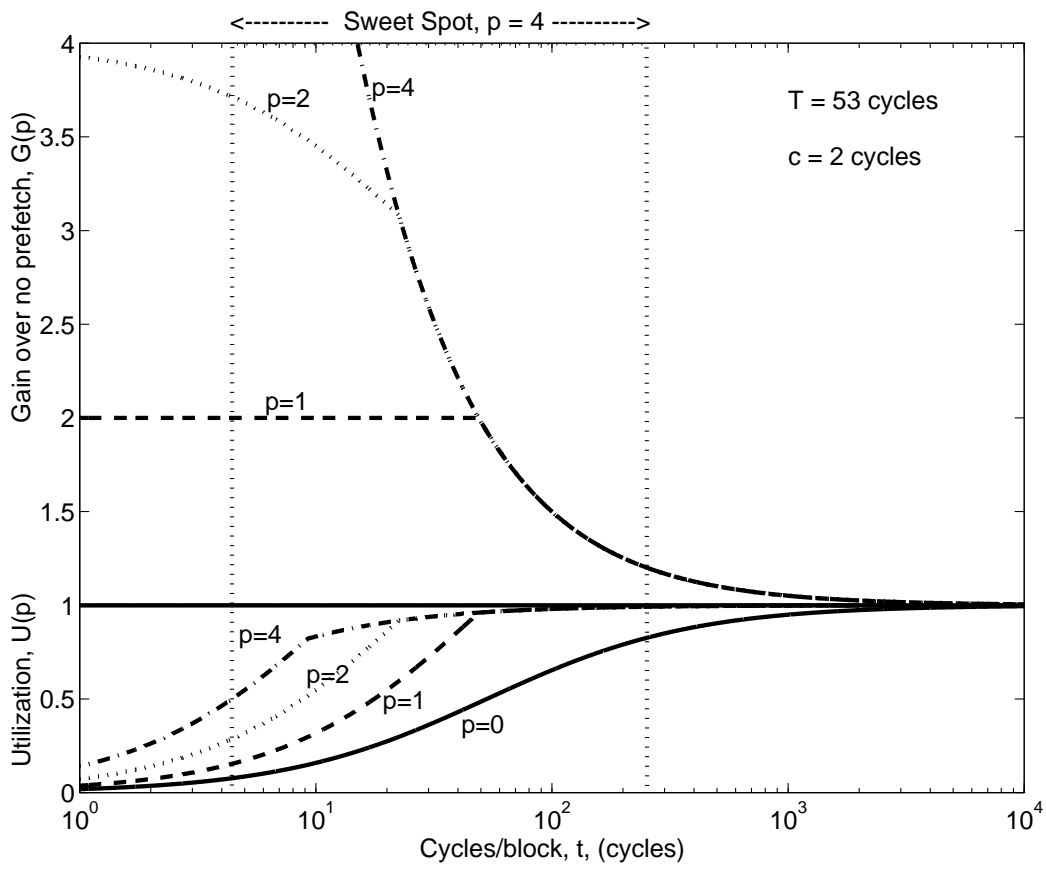


Figure 6: Modeling the performance of prefetching on Alewife.




Figure 7: Potential gain based on multithreading run-lengths on Alewife.

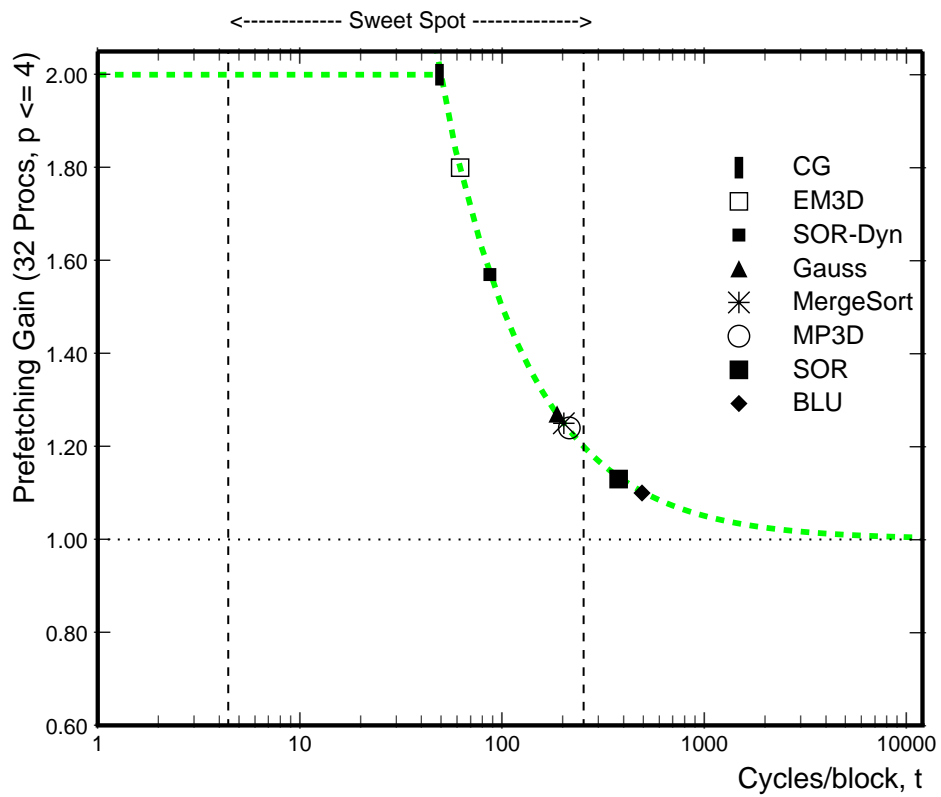


Figure 8: *Potential gain based on prefetching run-lengths on Alewife.*

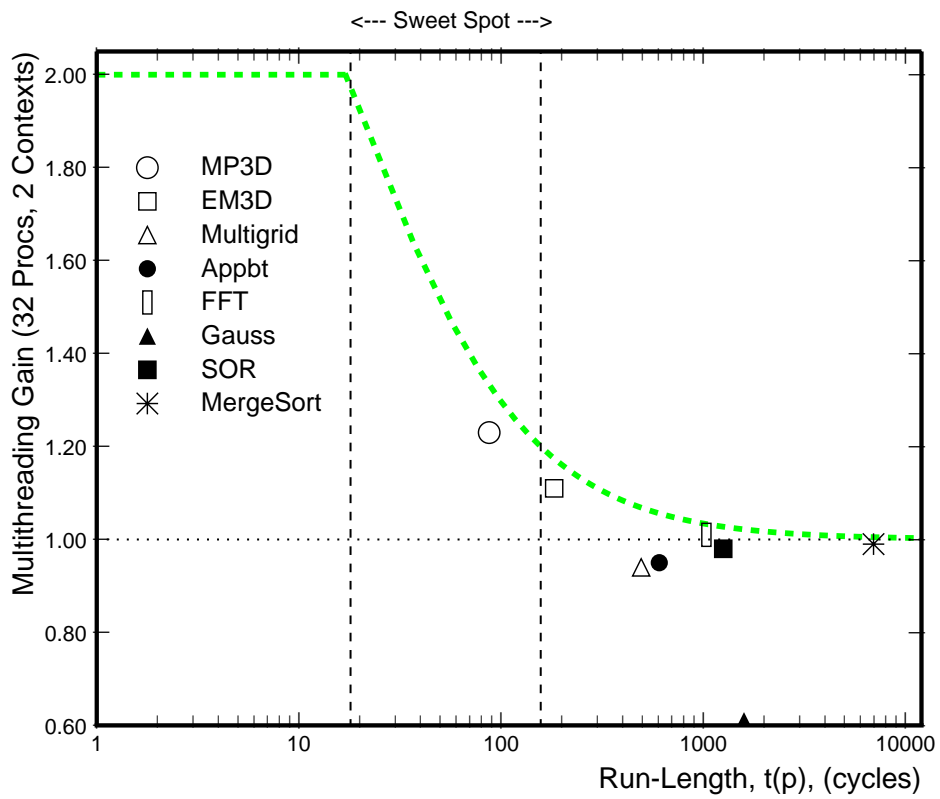


Figure 9: Measured gain on Alewife due to multithreading.




Figure 10: Measured gain on Alewife due to prefetching.




Figure 11: Modeling the performance of multithreading under higher latency. Vertical lines delimit the sweet spot and the range of application run-lengths.

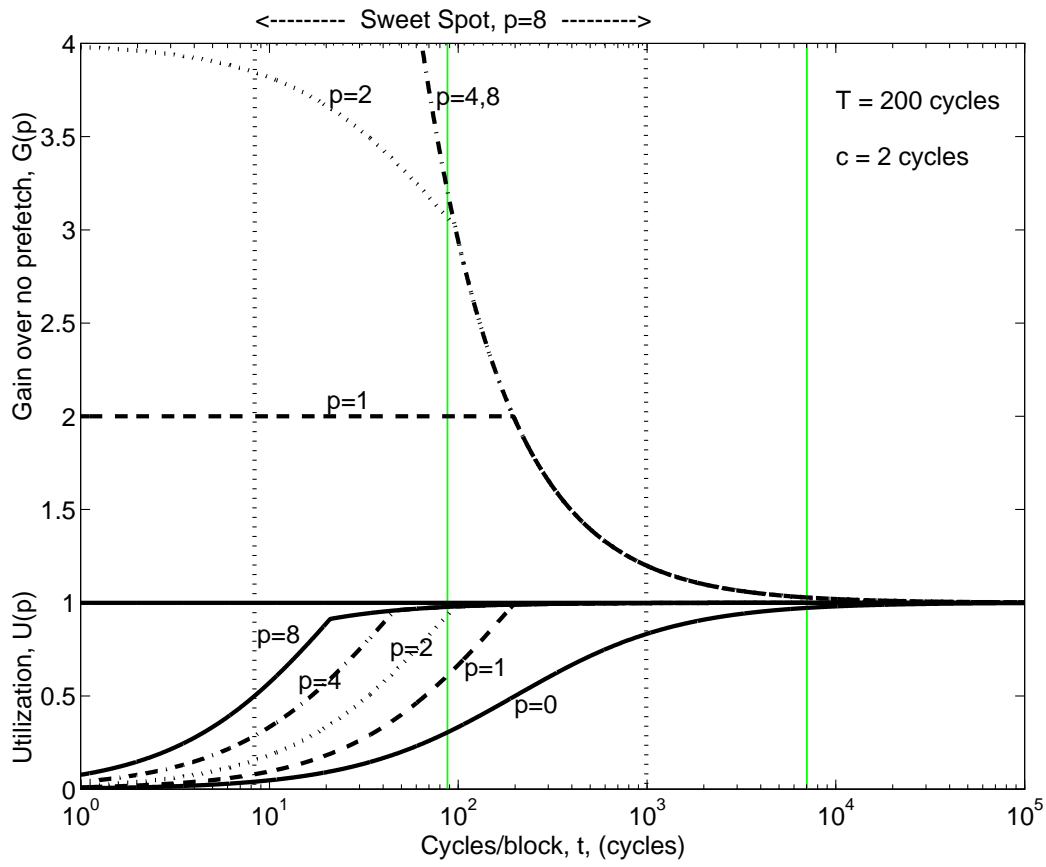


Figure 12: Modeling the performance of prefetching under higher latency. Vertical lines delimit the sweet spot and the range of application run-lengths.




Figure 13: Effect of remote cache miss latency on multithreading. Vertical lines delimit the sweet spot and the range of application run-lengths.




Figure 14: *Effect of remote cache miss latency on prefetching. Vertical lines delimit the sweet spot and the range of application run-lengths.*




Figure 15: *Effect of context-switch overhead on multithreading. Vertical lines delimit the sweet spot and the range of application run-lengths.*




Figure 16: *Effect of prefetch instruction overhead on prefetching. Vertical lines delimit the sweet spot and the range of application run-lengths.*

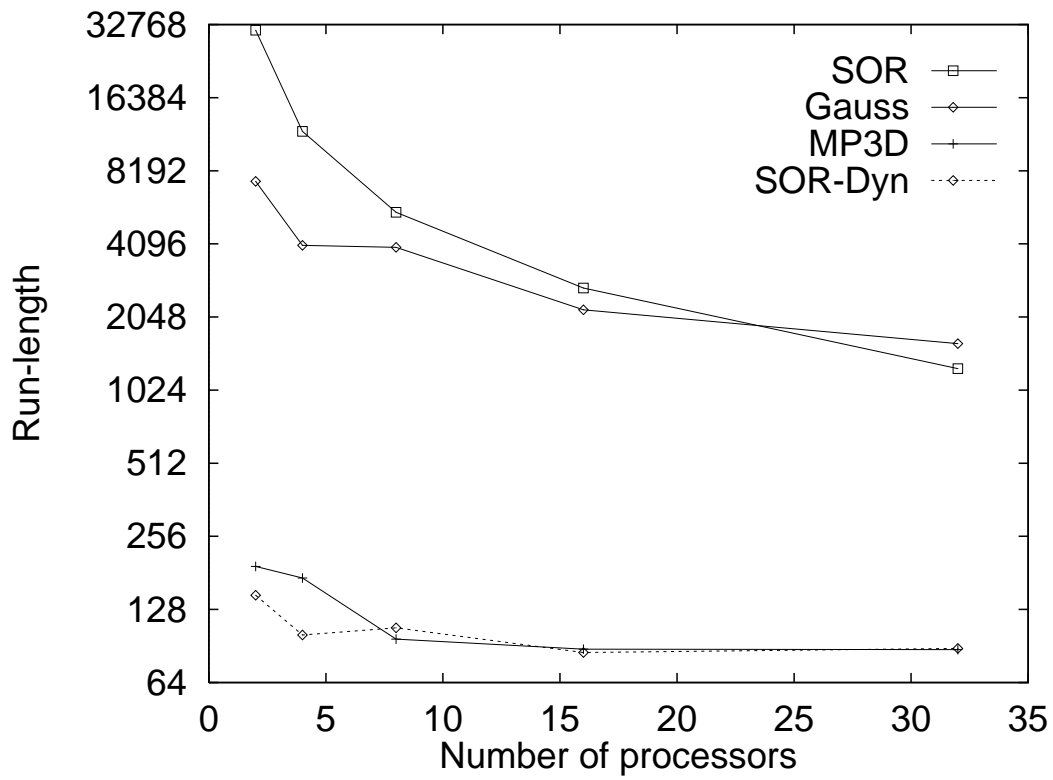


Figure 17: *Effect of number of processors on multithreading run-lengths.*




Figure 18: Modeling the performance of all-software multithreading.




Figure 19: Modeling the performance of all-software prefetching.

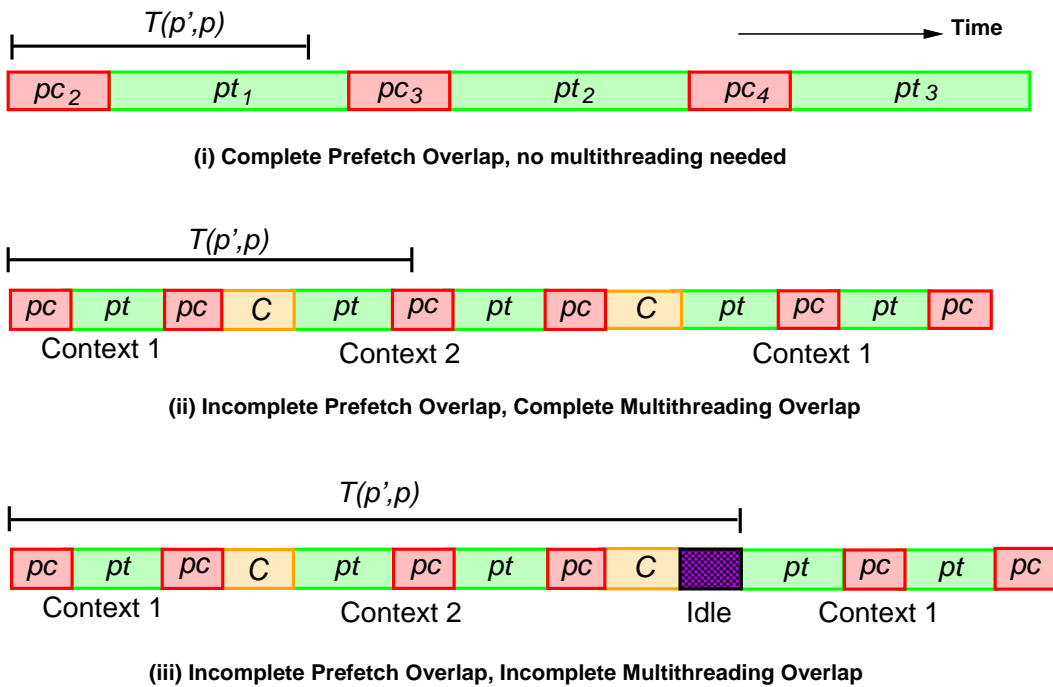


Figure 20: Time lines illustrating the overlap of communication latency $T(p', p)$, with computation pt , when multithreading with $p' = 2$ contexts. In the second case, $T(p', p)$ is too long to be completely hidden by prefetching, but can be hidden by multithreading. In the third case, both prefetching and multithreading fail to completely hide $T(p', p)$.




Figure 21: Modeling the performance of the combination of multithreading and prefetching on Alewife.

p	Number of contexts on a processor
C	Context-switch overhead (time to switch between processor-resident threads)
$T(p)$	Remote cache miss latency (time to service a cache miss involving remote nodes)
$t(p)$	Average run-length (time between remote cache misses)
$U(p)$	Processor utilization (percentage of time doing useful computation)
$G(p)$	Performance gain due to multithreading ($U(p)/U(1)$)

Table I: *Definitions of the parameters used in the multithreading model.*

p	Number of cache blocks prefetched at a time
c	Prefetch overhead (time taken by a prefetch instruction)
$T(p)$	Remote cache miss latency (time to service a cache miss involving remote nodes)
t	Average computation per prefetched cache block
$U(p)$	Processor utilization with p prefetches at a time
$G(p)$	Performance gain due to prefetching ($U(p)/U(0)$)

Table II: *Definitions of the parameters used in the prefetching model.*

Program	Input Size
Barnes-Hut	16K bodies, 4 iterations
Cholesky	$3,948 \times 3,948$ floats, 56,934 non-zeros (TK15)
LocusRoute	3,817 wires, 20 routing channels (Primary2)
MP3D	18K particles, 6 iterations
Water	125 molecules, 10 iterations
Appbt	$20 \times 20 \times 20$ floats
Multigrid	$56 \times 56 \times 56$ floats
CG	$1,400 \times 1,400$ doubles, 78,148 non-zeros
EM3D	20,000 nodes, 20% remote neighbors
FFT	80,000 floats
Gauss	512×512 floats
SOR	512×512 floats, 50 iterations
SOR-Dyn	512×512 floats, 50 iterations
BLU	512×512 floats, 16×16 blocks
MergeSort	64,000 integers

Table III: *Main application input parameters.*

p	Number of cache blocks prefetched at a time
c	Prefetch overhead (time to execute a prefetch instruction)
t	Average computation per prefetched cache block
p'	Number of contexts on a processor
C	Context-switch overhead (time to switch between processor-resident threads)
$t(p')$	Average run-length (time between remote cache misses)
$T(p', p)$	Remote cache miss latency (time to service a cache miss involving remote nodes)
$U(p', p)$	Processor utilization (percentage of time doing useful computation)
$G(p', p)$	Performance gain due to multithreading + prefetching

Table IV: *Definitions of the parameters used in the model for the combination of multithreading and prefetching.*