Improving Hash Join Performance through Prefetching

Shimin Chen Anastassia Ailamaki
Carnegie Mellon University
{*chensm,natassa,tcm*}*@cs.cmu.edu*

Phillip B. Gibbons† Todd C. Mowry
†*Intel Research Pittsburgh*
phillip.b.gibbons@intel.com

Abstract

Hash join algorithms suffer from extensive CPU cache stalls. This paper shows that the standard hash join algorithm for disk-oriented databases (i.e. GRACE) spends over 73% of its user time stalled on CPU cache misses, and explores the use of prefetching *to improve its cache performance. Applying prefetching to hash joins is complicated by the data dependencies, multiple code paths, and inherent randomness of hashing. We present two techniques,* group prefetching *and* software-pipelined prefetching*, that overcome these complications. These schemes achieve 2.0–2.9X speedups for the join phase and 1.4–2.6X speedups for the partition phase over GRACE and simple prefetching approaches. Compared with previous cache-aware approaches (i.e. cache partitioning), the schemes are at least 50% faster on large relations and do not require exclusive use of the CPU cache to be effective.*

1. Introduction

Hash join [11, 15, 17, 24, 28, 30] has been studied extensively over the past two decades, and it is commonly used in today's commercial database systems to implement equijoins efficiently. In its simplest form, the algorithm first builds a hash table on the smaller (*build*) relation, and then probes this hash table using tuples of the larger (*probe*) relation to find matches. However, the random access patterns inherent in the hashing operation have little spatial or temporal locality. When the main memory available to a hash join is too small to hold the build relation and the hash table, the simplistic algorithm suffers excessive random disk accesses. To avoid this problem, the *GRACE* hash join algorithm [15] begins by partitioning the two joining relations such that each build partition and its hash table can fit within memory; pairs of build and probe partitions are then joined separately as in the simple algorithm. This *I/O partitioning* technique limits the random accesses to objects that fit within main memory and results in nice sequential I/Os for source relations and intermediate partitions. As a result, the I/O costs no longer dominate. For example, our experiments on a quad-processor Pentium III show that a hash join of two several GB relations is already CPU-bound with only 4 disks, and it becomes increasingly CPU bound with each additional disk (details in Section 7).

1.1 Hash Joins Suffer from CPU Cache Stalls

So where *do* hash joins spend most of their time? Previous studies have demonstrated that hash joins can suffer from excessive CPU cache stalls [5, 20, 29]. The lack of spatial or temporal locality means the GRACE hash join algorithm cannot take advantage of the multiple levels of CPU cache in modern processors, and hence it repeatedly suffers the full latency to main memory during building and probing. Figure 1 provides a breakdown of the simulated user-level performance on a state-of-the-art machine (details in Section 7). The "partition" experiment divides a 1GB relation into 800 partitions, while the "join" experiment joins a 50MB build partition with a 100MB probe partition. Each bar is broken down into four categories: busy time, data cache stalls, TLB miss stalls, and other stalls. As we see in Figure 1, both the *partition* and *join* phases spend a significant fraction of their time—82% and 73%, respectively—stalled on data cache misses!

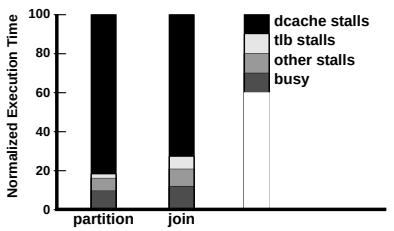

Figure 1. Execution time breakdown for hash join.

Given the success of I/O partitioning in avoiding random *disk* accesses, the obvious question is whether a similar technique can be used to avoid random *memory* accesses. *Cache partitioning*, in which the joining relations are partitioned such that each build partition and its hash table can fit within the (largest) CPU cache, has been shown to be effective in improving performance in memory-resident and main-memory databases [5, 20, 29]. However, cache partitioning suffers from two important practical limitations. First, for traditional disk-oriented databases, generating cache-sized partitions while scanning from disk requires a large number of concurrently active partitions. Experiences with the IBM DB2 have shown that storage managers can handle only hundreds of active partitions per join [17]. Given a 1MB CPU cache and (optimistically) 1000 partitions, the *maximum relation size that can be handled is only 1 GB*. Beyond that hard limit, any cache partitioning must be done using additional passes through the data — as will be shown in Section 7, this results in at least a 50% slowdown compared to the techniques we propose. Second, cache partitioning assumes *exclusive use* of the cache, but this assumption is unlikely to be valid in an environment with multiple ongoing

activities. Once the cache is too busy with other requests to effectively retain its partition, *the performance may degrade significantly* (up to 67% in the experiments in Section 7). Hence, we would like to explore an alternative technique that does not suffer from these limitations.

1.2 Our Approach: Cache Prefetching

Rather than trying to *avoid* CPU cache misses by building tiny (cache-sized) hash tables, we instead propose to *hide* the cache miss latency associated with accessing normal (memory-sized) hash tables, by overlapping these misses with computation. Modern processors allow multiple cache misses to be in flight simultaneously in the memory hierarchy (e.g., the Compaq ES40 [9] supports 32 in-flight loads, 32 in-flight stores, and 8 outstanding off-chip cache misses per processor), and the trend has been toward supporting more and more simultaneous misses. To enable software to fully exploit this parallelism, modern processors also provide explicit *prefetch* instructions for moving data into the cache ahead of its use. Software-based prefetching has been successfully applied in the past to array-based programs [23], pointer-based programs [18], and database B^+-Trees [7, 8], but it has not been applied to hash joins.

Challenges in Applying Prefetching to Hash Join. A naive approach to prefetching for hash join might simply try to hide the latency within the processing of a single tuple. For example, to improve hash table probing performance, one might try to prefetch hash bucket headers, hash buckets, build tuples, etc. Unfortunately, such an approach would have little benefit because later memory references often depend upon previous ones (via pointer dereferences). Existing techniques for overcoming this *pointer-chasing problem* [18] will not work because the randomness of hashing makes it impossible to predict the memory locations to be prefetched.

The good news is that although there are many dependencies *within* the processing of a single tuple, dependencies are less common *across* subsequent tuples due to the random nature of hashing. Hence our approach is to exploit *inter-tuple parallelism* to overlap the cache misses of one tuple with the computation and cache misses associated with other tuples.

A natural question is whether either the hardware or the compiler could accomplish this inter-tuple cache prefetching automatically; if so, we would not need to modify the hash join software. Unfortunately, the answer is no. Hardware-based prefetching techniques [2] rely upon recognizing regular and predictable (e.g., strided) patterns in the data address stream, but the inter-tuple hash table probes do not exhibit such behavior. In many modern processors, the hardware also attempts to overlap cache misses by speculating ahead in the instruction stream; while this approach is useful for hiding the latency of primary data cache misses that hit in the secondary cache, the amount of lookahead buffering (in the *reorder buffers*) is far too small to fully hide the latency of cache misses to main memory [10] (e.g., 128 vs. 600 entries for the Compaq ES40 [9]), and is even small compared with the amount of processing required for a single tuple. While our prefetching approaches (described below) are inspired by compiler-based scheduling techniques, existing compiler techniques for scheduling prefetches [18, 23] cannot handle the ambigous data dependencies present in the hash join code (as discussed in detail in Sections 4.3 and 5.2).

Overcoming these Challenges. To effectively hide the cache miss latencies in hash join, we propose and evaluate two new prefetching techniques: *group prefetching* and *software-pipelined prefetching*. For group prefetching, we apply modified forms of compiler transformations called *strip mining* and *loop distribution* (illustrated later in Section 4) to restructure the code such that hash probe accesses resulting from groups of G consecutive probe tuples can be pipelined.[1] The potential drawback of group prefetching is that cache miss stalls can still occur during the transition between groups. Hence our second prefetching scheme leverages a compiler scheduling technique called *software pipelining* [16] to avoid these intermittent stalls.

A key challenge that required us to extend existing compiler-based techniques in both cases is that although we expect dependencies across tuples to be unlikely, they are still possible, and we must take them into account to preserve correctness. If we did this conservatively (as the compiler would), it would severely limit our potential performance gain. Hence we optimistically schedule the code assuming that there are no inter-tuple dependencies, but we perform some extra bookkeeping at runtime to check whether dependencies actually occur: if so, we temporarily stall the consumer of the dependence until it can be safely resolved. Additional challenges arose from the multiple levels of indirection and multiple code paths in hash table probing.

A surprising result in our study is that contrary to the conventional wisdom in the compiler optimization community that software pipelining outperforms strip mining, group prefetching appears to be more attractive than software-pipelined prefetching for hash joins. A key reason for this difference is that the code in the hash join loop is far more complex than the typical loop body of a numeric application (where software pipelining is more commonly used [16]).

1.3 Contributions of This Paper

This paper makes the following contributions. First, to our knowledge, this is the first study to explore how prefetching can be used to accelerate both the join and partition phases of hash join by exploiting inter-tuple parallelism. Second, we propose two prefetching techniques, *group prefetching* and *software-pipelined prefetching*, and show how they can be applied to significantly improve hash join performance. Overall, the techniques achieve 2.0–2.9X speedups for the join phase and 1.4–2.6X speedups for the partition phase over GRACE and simple prefetching approaches. Moreover, they are at least 50% faster than cache partitioning on large relations and do not require exclusive use of the cache to be effective. Finally, we make extensive comparisons between group prefetching and software-pipelined prefetching, demonstrating that group prefetching

[1] In our experimental set-up in Section 7, $G = 19$ is optimal.

is 1%–15% faster than software-pipelined prefetching.

The paper is organized as follows. Section 2 discusses the related work in greater detail. Section 3 analyzes the dependencies in the join phase, the more complicated of the two phases, while Sections 4 and 5 use group prefetching and software-pipelined prefetching to improve the join phase performance. Section 6 discusses prefetching for the partition phase. Experimental results appear in Section 7 and conclusions in Section 8.

2. Related Work

Since the GRACE hash join algorithm was first introduced [15], many refinements of this algorithm have been proposed for the sake of avoiding I/O by keeping as many intermediate partitions in memory as possible [11, 17, 24, 28, 30]. All of these hash join algorithms, however, share two common building blocks: (1) *partitioning* and (2) *joining* with in-memory hash tables. To cleanly separate these two phases, we use GRACE as our baseline algorithm throughout this paper. We point out, however, that our techniques should be directly applicable to the other hash join algorithms.

CPU cache performance has been identified as a major performance bottleneck for database systems [1, 3, 13] and many recent studies have focused on improving the cache performance of core database algorithms [4, 5, 7, 8, 25, 26, 29]. Several papers have developed techniques to improve the cache performance of hash joins [5, 20, 29]. Shatdal *et al.* showed that cache partitioning achieved 6-10% improvement for joining memory-resident relations with 100B tuples [29]. Boncz, Manegold and Kersten proposed using multiple passes in cache partitioning to avoid cache and TLB thrashing [5, 20]. They showed large performance improvements on real machines for joining vertically-partitioned relations in the Monet main memory database, under exclusive use of the CPU caches. They considered neither disk-oriented databases, more traditional physical layouts, multiple activities trashing the cache, nor the use of prefetching. They also proposed a variety of code optimizations (e.g., using shift-based hash functions) to reduce CPU time; these optimizations may be beneficial for our techniques as well.

As mentioned earlier, software prefetching has been used successfully in other scenarios [7, 8, 18, 23]. While software pipelining has been used to schedule prefetches in array-based programs [23], we have extended that approach to deal with more complex data structures, multiple code paths, and the read-write conflicts present in hash join.

Previous work demonstrated that TLB misses may degrade performance [5, 20], particularly when TLB misses are handled by software. Because the vast majority of modern processors (including those from Intel) handle TLB misses in hardware, we model hardware-based TLB miss handling in our simulations. In addition, our simulator supports TLB prefetching [27] by treating TLB misses caused by prefetches as normal TLB misses. Hence, using our prefetching techniques, the TLB misses are overlapped with computation, minimizing TLB stall time.

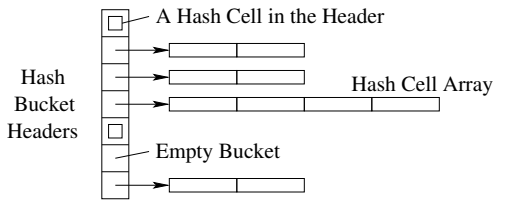

Figure 2. An in-memory hash table structure

3. Dependencies in the Join Phase

In this section, we analyze the dependencies in a hash table visit in the join phase to show why a naive prefetching algorithm would fail. We study a concrete implementation of the in-memory hash table, as shown in Figure 2. It consists of an array of hash buckets, each composed of a header and (possibly) an array of hash cells pointed to by the header. A hash cell represents a build tuple hashed to the bucket. It contains the tuple pointer and a fixed-length (e.g., 4 byte) hash code computed from the join key, which serves as a filter for the actual key comparisons.[2] A single hash cell is put into the bucket header. When more tuples are hashed to the bucket, a hash cell array is allocated, the size of which can be dynamically increased.

A naive prefetching algorithm would try to hide cache miss latencies *within* a single hash table visit by prefetching for potential cache misses, including hash bucket headers, hash cell arrays, and/or build tuples. However, this approach would fail because there are a lot of dependencies in a hash table visit. For example, the memory address of the bucket header is determined by the hashing computation. The address of the hash cell array is stored in the bucket header. The memory reference for a build tuple is dependent on the corresponding hash cell (in a probe). These dependencies essentially form a critical path; a previous computation or memory reference generates the memory address of the next reference, and must finish before the next one can start. Addresses would be generated too late for prefetching to hide miss latencies. Moreover, the randomness of hashing makes it almost impossible to predict memory addresses for hash table visits. These arguments are true for all hash-based structures.[3] Therefore, applying prefetching to the join phase algorithm is not a straightforward task.

4. Group Prefetching

Although dependencies *within* a hash table visit prevent effective prefetching, the join phase algorithm processes a large number of tuples and dependencies are less common *across* subsequent tuples due to the randomness of hashing. Therefore, our approach is to exploit *inter-tuple parallelism* to overlap cache miss latencies of one tuple with computations and miss latencies of other tuples. To ensure correctness, we must systematically intermix multiple hash

[2] Hash codes are usually good distinguishers of the join keys.

[3] The structure in Figure 2 improves upon chained bucket hashing, which uses a linked list of hash cells in a bucket. It avoids the pointer chasing problem of linked lists [19, 7].

```
foreach tuple in probe partition
{
    compute hash bucket number;
    visit the hash bucket header;
    visit the hash cell array;
    visit the matching build tuple to
      compare keys and produce output tuple;
}
```

(a) A simplified probing algorithm

```
foreach group of tuples in probe partition
{
   foreach tuple in the group {
     compute hash bucket number;
     prefetch the target bucket header;
   }
   foreach tuple in the group {
     visit the hash bucket header;
     prefetch the hash cell array;
   }
   foreach tuple in the group {
     visit the hash cell array;
     prefetch the matching build tuple;
   }
   foreach tuple in the group {
     visit the matching build tuple to
       compare keys and produce output tuple;
   }
}
```

(b) Group prefetching

Figure 3. The idea of group prefetching

```
for i=0 to N-1 do
{
    code 0;
    visit ($m_i^1$); code 1;
    visit ($m_i^2$); code 2;
    ... ...
    visit ($m_i^k$); code k;
}
```

(a) Processing an element per iteration

```
for j=0 to N-1 step G do
{
   for i=j to j+G-1 do {
     code 0;
     prefetch ($m_i^1$);
   }
   for i=j to j+G-1 do {
     visit ($m_i^1$); code 1;
     prefetch ($m_i^2$);
   }
   for i=j to j+G-1 do {
     visit ($m_i^2$); code 2;
     prefetch ($m_i^3$);
   }
   ... ...
   for i=j to j+G-1 do {
     visit ($m_i^k$); code k;
   }
}
```

(b) Group prefetching

Figure 4. General group prefetching algorithm

table visits, reorder their memory references, and schedule prefetch instructions sufficiently early. In this section, we propose group prefetching to achieve these objectives.

4.1 Group Prefetching for a Simplified Probing Algorithm

We use a simplified probing algorithm to describe the idea of group prefetching. As shown in Figure 3(a), the algorithm assumes that all hash buckets have hash cell arrays and every probe tuple matches exactly one build tuple. It performs a probe per loop iteration.

As shown in Figure 3(b), the group prefetching algorithm combines multiple iterations of the original loop into a single loop body, and rearranges the probe operations into stages[4]. Each stage performs one computation or memory reference on the critical path for all the tuples in the group and then issues prefetch instructions for the memory references of the next stage. For example, the first stage computes the hash bucket number for every tuple and issues prefetch instructions for the hash bucket headers, which will be visited in the second stage. In this way, the cache miss to read the hash bucket header of a probe will be overlapped with hashing computations and cache misses for other probes. Prefetching is used similarly in the other stages except the last stage. Note that the dependent memory operations of the same probe are still performed one after another as before. However, the memory operations of different probes are now overlapped.

[4]Technically, what we do are modified forms of compiler transformations called *strip-mining* and *loop distribution* [14].

4.2 Understanding Group Prefetching

To better understand group prefetching, we generalize the previous algorithms of Figure 3 in Figure 4. Suppose we need to process N independent elements. For each element i, we need to make k dependent memory references, $m_i^1, m_i^2, \cdots, m_i^k$. As shown in Figure 4(a), a straightforward algorithm processes an element per loop iteration. The loop body is natually divided into $k+1$ stages by the k memory references. *Code 0* (if it exists) computes the first memory address m_i^1. *Code 1* uses the contents in m_i^1 to compute the second memory address m_i^2. Generally *code* l uses the contents in m_i^l to compute the memory address m_i^{l+1}, where $l = 1, \cdots, k-1$. Finally, *code* k performs some processing using the contents in m_i^k. If every memory reference m_i^l incurs a cache miss, the algorithm will suffer kN expensive, fully exposed cache misses.

Because the elements are independent of each other, we can use group prefetching to overlap cache miss latencies across multiple elements, as shown in Figure 4(b). The group prefetching algorithm combines the processing of G elements into a single loop body. It processes *code* l for all the elements in the group before moving on to *code* $l+1$. As soon as an address is computed, it issues a prefetch instruction for the memory location so that the reference will be overlapped across the processing of other elements.

Now we determine the condition for fully hiding all cache miss latencies. Suppose the execution time of *code* l is C_l, the full latency of fetching a cache line from main memory is T, and the additional latency of fetching the next cache line in parallel is T_{next}, which is the inverse of the memory bandwidth. (Table 1 shows the terminology used throughout

Table 1. Terminology used throughout this paper.

Name	Definition
k	# of dependent memory references for an element
G	group size in group prefetching
D	prefetch distance in software-pipelined prefetching
T	full latency of a cache miss
T_{next}	latency of an additional pipelined cache miss
C_l	execution time for code l, $l = 0, 1, \ldots, k$

the paper.) Assume every m_i^l incurs a cache miss and there are no cache conflicts [5]. Then, a sufficient condition for fully hiding all cache miss latencies is as follows [6]:

$$\begin{cases} (G-1) \cdot C_0 \geq T \text{ and} \\ (G-1) \cdot \max\{C_l, T_{next}\} \geq T, \; l = 1, 2, \cdots, k \end{cases}$$

For an intuitive explanation, we focus on the first element in a group, element j. The prefetch for m_j^1 is overlapped with the processing of the remaining $G-1$ elements at code stage 0. The first inequality ensures that this memory reference will complete before the visit operation for m_j^1 in code stage 1. Similarly, the prefetch for m_j^{l+1} is overlapped with the processing of the remaining $G-1$ elements at code stage l. The second inequality ensures that its latency is fully hidden. Here, T_{next} corresponds to the visit operations in the processing of the $G-1$ elements. For elements $j+1, \ldots, j+G-1$ in a group, the lth memory references will be overlapped with operations at code stages l and $l+1$. We can prove the above condition is sufficient for hiding these reference latencies by simple combinations of the inequalities. [6]

We can always choose a G large enough to satisfy the second inequality because T_{next} is always greater than 0. However, if *code 0* were empty, m_j^1 could not be fully hidden. Fortunately, in the simplified probing algorithm, *code 0* computes the hash bucket number and is not empty. Therefore, we can choose a G to hide all the cache miss penalties.

In the above, cache conflict misses are ignored for simplicity of analysis. However, we will show in Section 7 that conflict misses are a problem when G is too large. Therefore, among all possible G's that satisfy the inequalities, we should choose the smallest in order to minimize the number of concurrent prefetches and the conflict miss penalty.

4.3 Dealing with Complexities

Previous research showed how to prefetch for two dependent memory references for array-based codes [22]. Our group prefetching algorithm solves the problem of prefetching for an arbitrary fixed number k of dependent memory references.

We have implemented group prefetching for both hash table building and probing. In contrast to the simplified probing algorithm, the actual probing algorithm contains multiple code paths: there could be zero or multiple matches, hash buckets could be empty, and there may not be a hash cell array in a bucket. To cope with this complexity, we keep state information for the G tuples of a group. We divide each possible code path into code pieces on the boundaries of dependent memory references. Then we combine the code pieces at the same position of different code paths into a single stage using conditional tests on the tuple states. Figure 5 shows the idea of this process. Note that the common starting point of all code paths is in *code 0*. The first code piece including a branch sets the state of an element. Then subsequent code stages test the state and execute the code pieces for the corresponding code paths. The total number of group prefetching stages ($k+1$) is the largest number of code pieces along any original code path.

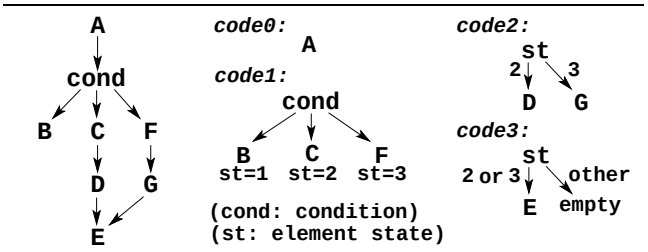

(a) An example (b) Code stages in group prefetching

Figure 5. Dealing with multiple code paths.

When multiple independent cache lines are visited at a stage (e.g., to visit multiple build tuples), our algorithm issues multiple independent prefetches in the previous stage.

The group prefetching algorithm must also cope with read-write conflicts. Though quite unlikely, it is possible that two build tuples in a group may be hashed into the same bucket. However, in our algorithm, hash table visits are interleaved and no longer atomic. Therefore, a race condition could arise; the second tuple might see an inconsistent hash bucket being changed by the first one. Note that this complexity occurs because of the read-write nature of hash table building. To cope with this problem, we set a busy flag in a hash bucket header before inserting a tuple. If a tuple is to be inserted into a busy bucket, we delay its processing until the end of the group prefetching loop body. At this natural group boundary, the previous access to the busy hash bucket must have finished. Interestingly, the previous access has also warmed up the cache for the bucket header and hash cell array, so we insert the delayed tuple without prefetching. The algorithm can deal with any number of delayed tuples (to tolerate skews in the key distribution).

5. Software-pipelined Prefetching

In this section, we describe our technique of exploiting software pipelining to schedule prefetches for hash joins. We then compare our two prefetching schemes.

Figure 6 shows the difference between group prefetching and software-pipelined prefetching intuitively. Group prefetching hides cache miss latencies within a group of elements and there is no overlapping memory operation between groups. In contrast, software-pipelined prefetching

[5] We use these assumptions only to simplify the derivation of the conditions. Note that our experimental evaluations include all the possible effects of locality and conflicts in hash joins.

[6] Please see the extended version of the paper [6] for the critical path analysis that derives the condition.

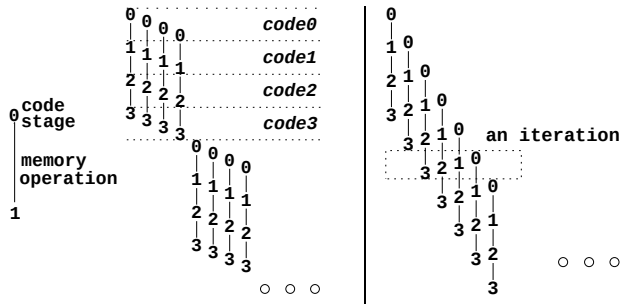

(a) Group prefetching ($G = 4$)

(b) Software-pipelined prefetching ($D = 1$)

Figure 6. Intuitive pictures of the prefetching schemes

combines different code stages of different elements into an iteration and hides latencies across iterations. It keeps running without gaps and therefore may potentially achieve better performance.

5.1 Understanding Software-pipelined Prefetching

Figure 7 shows the software-pipelined prefetching for the simplified probing algorithm. The subsequent stages for a particular tuple are processed D iterations away. (D is called the *prefetch distance* [22].) Figure 6(b) depicts the intuitive picture when $D = 1$. Suppose the left-most line in the dotted rectangle corresponds to tuple j. Then, an iteration combines the processing of stage 0 for tuple $j + 3D$, stage 1 for tuple $j + 2D$, stage 2 for tuple $j + D$, and stage 3 for tuple j.

Figure 8 shows the generalized algorithm for software-pipelined prefetching. In the steady state, the pipeline has $k + 1$ stages. The loop body processes a different element for every stage. The subsequent stages for a particular element are processed D iterations away. Intuitively, if we make the intervals between code stages for the same element sufficiently large, we will be able to hide cache miss latencies. Under the same assumption as in Section 4.2, a sufficient condition for hiding all cache miss latencies in the steady state is as follows: [7]

$$D \cdot (max\{C_0 + C_k, T_{next}\} + \sum_{l=1}^{k-1} \max\{C_l, T_{next}\}) \geq T$$

We can always choose a D large enough to satisfy this condition. Section 7 will show that conflict misses are a problem when D is too large. Thus, we should choose the smallest D in order to minimize the number of concurrent prefetches.

5.2 Dealing with Complexities

We have implemented software-pipelined prefetching by modifying our group prefetching algorithm. The code stages are kept almost unchanged. To apply the general model in Figure 8, we use a circular array for state information. Since *code 0* and *code k* of the same element is processed kD iterations away, we ensure the array size s is at least $kD + 1$.

[7] Please see the extended version of the paper [6] for the critical path analysis that derives this condition.

```
prologue;
for j=0 to N-3D-1 do
{
    tuple j+3D:
      compute hash bucket number;
      prefetch the target bucket header;

    tuple j+2D:
      visit the hash bucket header;
      prefetch the hash cell array;

    tuple j+D:
      visit the hash cell array;
      prefetch the matching build tuple;

    tuple j:
      visit the matching build tuple to
        compare keys and produce output tuple;

}
epilogue;
```

Figure 7. Software-pipelined prefetching for simplified probing

```
prologue;
for j=0 to N-kD-1 do
{
    i=j+kD;
    code 0 for element i;
    prefetch (m_i^1);

    i=j+(k-1)D;
    visit (m_i^1); code 1 for element i;
    prefetch (m_i^2);

    i=j+(k-2)D;
    visit (m_i^2); code 2 for element i;
    prefetch (m_i^3);

    ... ...

    i=j;
    visit (m_i^k); code k for element i;
}
epilogue;
```

Figure 8. General software-pipelined prefetching

The index j in the general model is implemented as the array index $j \bmod s$. To reduce overhead, we choose s to be a power of 2 so that $j \bmod s$ can be computed using a bit mask operation.

The read-write conflict problem in hash table building is solved in a more sophisticated way. Because there is no place (like the end of a group in group prefetching) to conveniently process all the conflicts, we deal with the conflicts in the pipeline stages themselves. We build a waiting queue for each busy hash bucket. The hash bucket header contains the array index of the tuple t updating the bucket; only tuple t can update the hash bucket header and the hash cell array. The state information of a tuple contains a pointer to the next tuple waiting for the same bucket. When a tuple is to be inserted into a busy bucket, it is appended to the waiting queue and a "waiting" flag is set in its state information. Tuples with the waiting flag set are ignored in subsequent stages, until the last stage. In the meantime, tuple t proceeds through its code stages. At the end of the last stage for tuple t, we check its waiting queue. If the queue is not empty, we record the array index of the first waiting tuple in the bucket header, and perform the ignored code stages for

it (without prefetching, because t has already prefetched the needed lines into the cache). When this tuple subsequently gets to the last stage, it will handle the next tuple in the waiting queue if it exists, and so on.

5.3 Group vs. Software-pipelined Prefetching

Both prefetching schemes try to increase the interval between a prefetch and the corresponding visit, in order to hide cache miss latency. According to the sufficient conditions, software-pipelined prefetching can always hide all miss latencies, while group prefetching achieves this only when *code 0* is not empty (as is the case of the join phase). When *code 0* is empty, the first cache miss cannot be hidden. However, with a large group of elements, the amortized performance impact can be small.

In practice, group prefetching is easier to implement. The natural group boundary provides a place to do any necessary "clean-up" processing (e.g., for read-write conflicts). Moreover, the join phase can pause at group boundaries and send outputs to the parent operator to support pipelined query processing. Although a software pipeline may also be paused, the restart costs will diminish its performance advantage. Furthermore, software-pipelined prefetching has larger bookkeeping overhead because of its use of modular index operations and its larger maintained state (such as the waiting queue for read-write conflicts).

6. Prefetching for the Partition Phase

Having studied how to prefetch for the join phase of the hash join algorithm, in this section, we discuss prefetching for the partition phase. In the partition phase, an input relation is divided into multiple output partitions by hashing on the join keys. Typically an output buffer per partition and an input buffer are allocated in main memory. Disk pages from the input relation are streamed through the input buffer. Every input tuple is examined. Its partition number is computed from the join key. The relevant columns of the input tuple are then extracted (projected) and copied to the target output buffer. When an output buffer is full, it is written out.

Clearly, we should employ different prefetching techniques depending on the number of partitions generated. If the number of partitions is small enough so that all the buffers and relevant data structures fit in *cache*, we only need to prefetch for the input page to bring the input data into cache faster after every disk page read. This constitutes our simple prefetching scheme for the partition phase.

When the number of partitions is large, however, there could be cache thrashing during the partition phase; every output buffer visit may incur a cache miss. Similar to the join phase, the processing of a tuple needs to make several dependent memory references, whereas the processings of subsequent tuples are mostly independent due to the randomness of hashing. Therefore, we employ group prefetching and software-pipelined prefetching under this situation.

Note that there are read-write conficts in visiting the output buffers. Imagine that two tuples are hashed to the same partition. When processing the second tuple, the algorithm may find that the output buffer has no space and needs to be written out. However, it is possible that the data from the first tuple has not been copied into the output buffer yet because of the reorganization of processing. To solve this problem, in group prefetching, we wait until the end of the loop body to write out the buffer and process the second tuple. In software-pipelined prefetching, we use waiting queues similar to those for hash table building in the join phase.

7. Experimental Results

In this section, we show that hash join is CPU bound through real-machine experiments. We then evaluate the CPU cache performance of our prefetching techniques by simulation.

7.1 Experimental Setup

Implementation Details. We have implemented our own hash join engine. For real machine experiments, we implemented a buffer manager that stripes pages across multiple disks and performs I/O prefetching with background worker threads. For CPU performance simulation studies, we store relations and intermediate partitions as disk files for simplicity. We employ the slotted page structure and support fixed length and variable length attributes in tuples. Schemas and statistics are kept in separate description files for simplicity, the latter of which are used by the hash join algorithms to compute hash table sizes and numbers of partitions.

Our baseline algorithm is the GRACE hash join algorithm [28]. The in-memory hash table structure follows Figure 2 in Section 3. A simple XOR and shift based hash function is used to convert join keys of any length to 4-byte hash codes. Typically the same hash codes are used in both the partition and the join phase. Partition numbers in the partition phase are the hash codes modulo the total number of partitions. Hash bucket numbers in the join phase are the hash codes modulo the hash table size. Our algorithms ensure that the hash table size is a relative prime to the number of partitions. Because the same hash codes are used in both phases, we avoid the memory access and computational overheads of reading the join keys and hashing them a second time, by storing hash codes in the page slot area in the intermediate partitions and reusing them in the join phase. Note that changing the page structure of intermediate partitions is relatively easy because the partitions are only used in hash joins.

We implemented three prefetching schemes for both the partition phase and the join phase algorithm: simple prefetching, group prefetching, and software-pipelined prefetching. As suggested by the name, simple prefetching uses straightforward prefetching techniques, such as prefetching an entire input page after a disk read. We use simple prefetching as an enhanced baseline in order to show the additional benefit achieved using our more sophiscated prefetching schemes. Prefetch instructions are inserted into C++ source codes using gcc inline ASM macros.

Table 2. Simulation parameters

Processor pipeline parameters	
Clock Rate	1 GHz
Issue Width	4 insts/cycle
Functional Units	2 Integer, 1 integer divide, 2 Memory, 1 Branch, 2 FP
Reorder Buffer Size	128 insts
Integer Multiply/Divide	15/56 cycles
All Other Integer	1 cycle
Branch Prediction Scheme	gshare [21]

Memory parameters	
Line Size	64 bytes
Primary Instruction Cache	64 KB, 2-way set-assoc.
Primary Data Cache	64 KB, 4-way set-assoc.
Miss Handlers	32 for data, 2 for inst.
DTLB	64 entries, fully-assoc.
DTLB Miss Handlers	1
Page Size	8 KB
Unified Secondary Cache	1 MB, 4-way set-assoc.
Primary-to-Secondary Miss Latency	15 cycles (plus any delays caused by contention)
DTLB Miss Latency	20 cycles
Primary-to-Memory Miss Latency	150 cycles (plus any delays caused by contention)
Main Memory Bandwidth	1 access per 10 cycles

Cache Partitioning. Cache partitioning generates cache-sized build partitions so that every build partition and its hash table can fit in cache, greatly reducing the cache misses in the join phase. It has been shown to be effective in main-memory and memory-resident database environments [5, 29]. We have implemented two cache partitioning algorithms for disk-oriented databases. In the first, we increase the number of partitions and generate cache-sized partitions directly in the I/O partition phase. In the second, we partition twice: the I/O partition phase generates memory-sized partitions, which are subsequently partitioned again in memory as a preprocessing step for the join phase. We call the first scheme "direct cache" and the second "two-step cache".

Experimental Design. In all our experiments (except for Figure 14(a)), we assume the available memory size for the join phase is 50MB [8] and the partition phase generates partitions that tightly fit in 50MB. That is, in the baseline and our prefetching schemes, a build partition and its hash table fit tightly in the available memory. In the cache partitioning schemes, the partition sizes are also computed to satisfy the algorithm constraints and best utilize available memory.

Build relations and probe relations have the same schemas: a tuple consists of a 4-byte join key and a fixed-length payload. We believe that selection and projection are orthogonal issues to our study and we do not perform these operations in our experiments. An output tuple contains all the fields of the matching build and probe tuples. The join keys are randomly generated. A build tuple may match zero or more probe tuples and a probe tuple may match zero or one build tuple. In our experiments, we vary the tuple size, the number of probe tuples matching a build tuple, and the percentage of tuples that have matches, in order to show the benefits of our solutions in various situations.

[8] This is the memory allocated for joining a pair of build and probe partitions. It is limited by the simulation environment. However, we set the memory to cache size ratio to be 50:1, which is reasonable for joins on a modern computer system. Therefore, we expect the experiments to reflect real-world hash join cache behaviors.

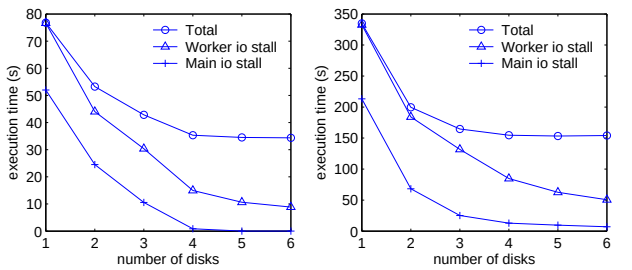

(a) Partition Performance (b) Join Performance

Figure 9. Hash join is CPU-bound with reasonable I/O bandwidth

Measurement Methodology. We first measure GRACE hash join performance on a real machine with multiple disks to show that hash join is CPU-bound with reasonable I/O bandwidth. Therefore, it is important to study hash join cache performance.

We then evaluate the CPU cache performance (of user mode executions) of all the schemes through simulation in order to get good prefetching support. We generate fully-functional executables with gcc and run the programs with detailed cycle-by-cycle simulations. The simulator models a dynamically-scheduled, superscalar processor running at a clock rate of 1 GHz. The memory hierarchy is based on the Compaq ES40 [9]. Because the Alpha processor (in the ES40) supports only software-simulated integer divide, we use instead the integer divide latency from the Intel Pentium4 [12]. The simulator does not drop prefetches when miss handlers are all busy. Moreover, the simulator supports TLB prefetching [27] by treating TLB misses caused by prefetches as normal TLB misses. Therefore, our prefetching schemes can overlap TLB miss latencies with computations and cache misses. Important simulator parameters are shown in Table 2.

7.2 Is Hash Join I/O-Bound or CPU-Bound?

Figure 9 shows the performance of GRACE hash join on a machine running Linux 2.4.18 with four 550MHz PentiumIII processors and 512MB RAM. Our experiments use 6 Seagate Cheetah X15 36LP SCSI disks, each with a maximum transfer rate of 68 MBytes/sec. We imitate raw disk partitions by allocating a large file on each disk and managing the mapping from page IDs to file offsets ourselves. To get good I/O performance, we stripe a relation across all the disks in 256KB units. Our buffer manager has a dedicated worker thread for each of the disks, which performs I/O operations on behalf of the main hash join thread. The buffer manager implements I/O prefetching and background writing so that I/O operations can be overlapped with computations as much as possible. We measure total elapsed times with gettimeofday() system calls and I/O stall times with pro-

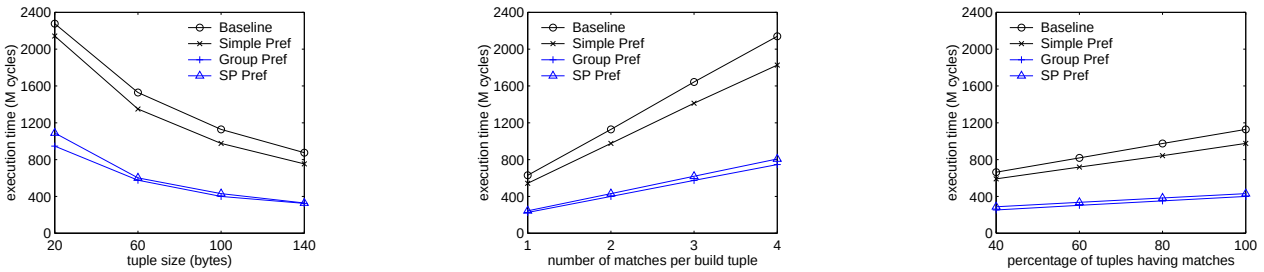

(a) Varying tuple size (b) Varying number of matches per build tuple (c) Varying join selectivity

Figure 10. Join phase performance

cessor cycle counter and the PAPI package. All the reported points are the average of 10 measurements with standard deviations less than 10% of the averages or less than 1 second.

The experiments join a 1.5GB build relation with a 3GB probe relation by producing 31 intermediate partitions for both relations. Tuples are 100 bytes long. Figure 9(a) shows the partition phase performance for the build relation, and Figure 9(b) shows the join phase performance of joining all the build and probe partitions.

The figures vary the number of disks used and report the total elapsed time for the operations, the maximum I/O stall time of all the background worker threads, and the stall time of the main thread waiting for workers. The worker I/O stall time shows the time to finish all the I/Os in background, which decreases dramatically as the number of disks increases. With four or more disks, hash join is clearly CPU-bound: the total elapsed time becomes flat, and the main thread spends less than 10% of the total time waiting for worker threads. Since there are typically 10 disks per processor on a balanced DB server, we expect that hash join is CPU-bound in a large number of real-world systems.

Moreover, the large gap between the top and the middle curves highlights the opportunity for dramatically reducing the total time by improving the CPU performance, e.g., a 3-fold potential improvement when there are 6 disks.

7.3 Join Phase Performance

Figure 10 shows the join phase performance of the baseline and the prefetching schemes. The experiments model the processing of a pair of partitions in the join phase. In all experiments, the build partition fits tightly in the 50MB memory. The default settings are 100B tuples and that every build tuple matches two probe tuples. As shown in the figure, group and software-pipelined prefetching achieve 2.4-2.9X and 2.1-2.7X speedups over the baseline, respectively. Because the central part of the join phase algorithm is hash table visits, simple prefetching only obtains marginal benefit, a 1.1-1.2X speedup over the baseline. By exploring the inter-tuple parallelism, group and software-pipelined prefetching achieve additional 2.3-2.5X and 2.0-2.3X speedups over the simple prefetching scheme, respectively.

In Figure 10(a), as we increase the tuple size, the number of tuples processed decreases, leading to the decreasing trend of the curves. In Figure 10(b) and (c), the total number of matches increases as we increase the number of matches per build tuple or the percentage of tuples having matches. This explains the upward trends. Moreover, the probe partition size also increases in Figure 10(b), contributing to the much steeper curves than those in Figure 10(c).

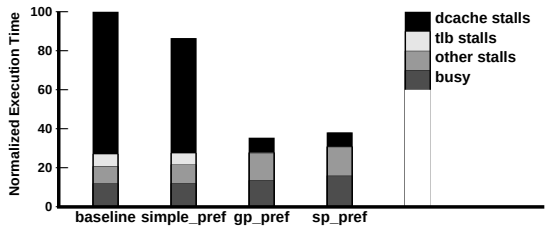

Figure 11. Execution time breakdown for join phase performance (Figure 10(a), 100B tuples)

Figure 11 shows the execution time breakdowns for the 100B points in Figure 10(a). The baseline case is shown as the "join" bar in Figure 1. Group prefetching and software-pipelined prefetching indeed successfully hide most of the data cache miss latencies. The simulator outputs confirm that the remaining cache misses are mostly L1 cache misses (but L2 hits) due to cache conflicts. The (transformation, bookkeeping, and prefetching) overheads of the techniques lead to larger portions of busy times. Software-pipelined prefetching is 7.7% more costly than group prefetching in this experiment. Interestingly, other stalls also increase. A possible reason is that some secondary causes of stalls show up when the data cache stalls are reduced.

Figure 12 shows the relationship between the cache performance and the parameters of our prefetching algorithms. We perform the same experiment as in Figure 10(a) when tuples are 20 bytes. We show the tuning results for only the probing loop here but the curves for the building loop have similar shapes. The optimal values for probing are $G = 19$ and $D = 1$. These values are used in all experiments unless otherwise noted.

The top curves in Figure 12 show the performance when the memory latency T is set to 1000 cycles in the simulator. The optimal points shift right; larger group size and prefetch distance are needed to hide the increased latencies, as expected by our models. Interestingly, software-pipelined prefetching becomes better than group prefetching. More importantly, software-pipelined prefetching achieves similar

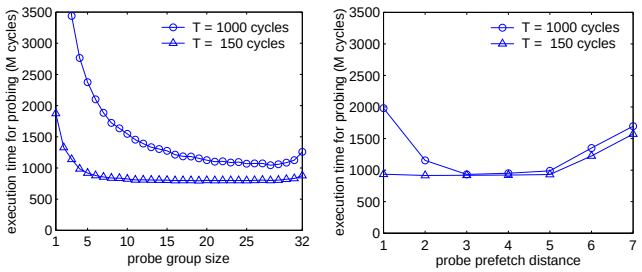

(a) group (b) software-pipelined

Figure 12. Tuning parameters of group and software-pipelined prefetching for join phase

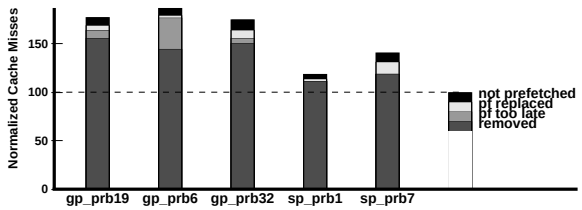

Figure 13. Breakdowns of cache misses to understand the tuning curves of the join phase

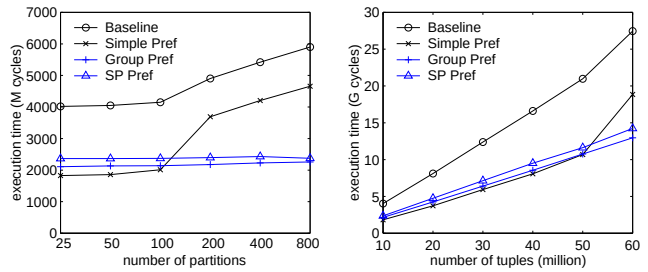

(a) varying number of partitions (b) varying relation size

Figure 14. Partition phase performance

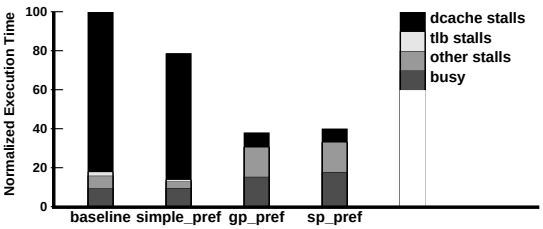

Figure 15. Execution time breakdown for partition phase performance (Figure 14(a), 800 partitions)

performance when we change T from 150 to 1000 cycles. This means that the prefetching algorithm will still keep up when the processor/memory speed gap increases even more (6 times in our experiments) as expected to happen in the future by the technology trend.

The curves all have concave shapes; performance becomes worse when the parameters are too small or too large. According to our models, the group size and the prefetch distance must be large enough to hide cache miss latencies. This explains the poor performance with small parameters. To verify this and to understand the cases with large parameters, we analyze the breakdowns of cache misses in Figure 13. The bars correspond to optimal points ($G = 19, D = 1$), too-small points ($G = 6$), and too-large points ($G = 32, D = 7$) from the $T = 150$ curves in Figure 12.

The bars are normalized to the number of cache misses in the baseline GRACE case (100 in the figure). The total heights of the bars correspond to the number of cache misses when we apply the code transformations but without inserting any prefetches. When we insert the prefetches, a lot of cache misses disappear, which are captured by the bottom part of the bars ("removed"). The other categories are i) "pf too late", i.e. prefetching is too late to hide all the latency; ii) "pf replaced", i.e. prefetching is too early and the prefetched cache line has already been replaced from the cache by other memory references or prefetches; and iii) "not prefetched".

From Figure 13, the too-small case shows a large "pf too late" portion. This confirms the above analysis. In the too-large cases, a lot of prefetches have been replaced, meaning that the poor performance is caused by cache conflicts. The larger the parameters, the more prefetches and other memory references are executed between a prefetch and its visit instruction, and therefore the greater chance that a prefetch is replaced from cache.

7.4 Partition Phase Performance

Figure 14(a) shows the partition phase performance varying the number of partitions from 25 to 800. The source relation has 10 million 100B tuples. (Unlike all the other experiments, the generated partitions may be much smaller than 50MB.) The figure is divided into two different regions. When the partition number is 25, 50, and 100, simple prefetching achieves the best performance. However, when the number of partitions becomes larger, the performance of simple prefetching deteriorates dramatically, while group prefetching and software-pipelined prefetching win. Since a 1MB L2 cache can hold 128 pages of 8KB each, the output buffers (and other miscellaneous data structures) fit in the L2 cache in the left region, and hence more sophisticated prefetching techniques are unnecessary. However, when the output buffers can not fit in cache as in the right region, simple prefetching suffers from excessive cache misses. Group and software-pipelined prefetching exploit inter-tuple parallelism to successfully hide most of the cache miss latencies.

Figure 14(b) varies the number of 100B tuples in the input relation while keeping the partition size the same (to fit tightly in 50MB memory). Hence the number of partitions also increases; 26, 51, 76, 102, 127, and 152 partitions are generated, respectively. Essentially the graph shows the same tradeoff as Figure 14(a) but in a more natural setting.

To get the best of both situations, we choose the prefetching algorithm based on the cache size and the number of partitions. Overall, this combined prefetching achieves 1.9-2.6X speedups over the baseline.

Figure 15 shows the execution time breakdown for Figure 14(a) where 800 partitions are generated. Group prefetching and software-pipelined prefetching successfully

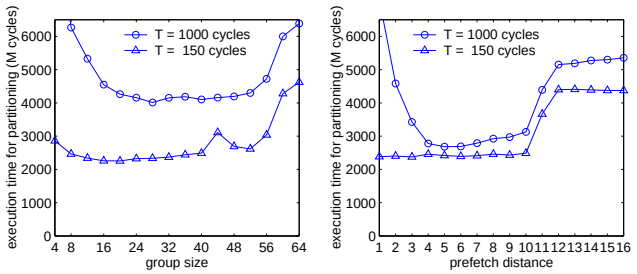

(a) group (b) software-pipelined

Figure 16. Tuning parameters of group and software-pipelined prefetching for the partition phase

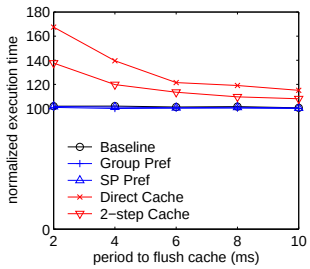

Figure 17. Impact of cache flushing on the different techniques.

hide most of the data cache miss latencies. Figure 16 shows the relationships between parameters and the cache performance of group prefetching and software-pipelined prefetching. The optimal values for partition phase are $G = 20$ and $D = 3$. The figure and a cache miss analysis as in Figure 13 show similar curve shapes and trends as in the join phase.

7.5 Comparison with Cache Partitioning

Problems with Large Relations. The number of I/O partitions is upper bounded by the available memory of the partition phase and by the requirements of the storage manager. Experiences with the IBM DB2 have shown that storage managers can handle only hundreds of active partitions per hash join [17]. Given a 1 MB CPU cache and (optimistically) 1000 partitions, the maximum relation size for "direct cache" is only 1 GB. "Two-step cache" solves this problem by introducing an additional partition pass. However, this additional copying cost results in 50-150% slowdown compared to our prefetching schemes, as discussed below.

Robustness. Cache partitioning assumes exclusive use of the cache, which is unlikely to be valid in a dynamic environment with multiple concurrent activities. Although a smaller "effective" cache size can be used, cache conflicts may still be a big problem and cause poor performance. Figure 17 shows the performance degradation of all the schemes when the cache is periodically flushed, which is the worst case interference. We vary the period to flush the cache from 2 ms to 10 ms. "100" corresponds to the join phase execution time when there is no cache flush. "Direct cache" and "two-step cache" suffer from 15-67% and 8-38% performance degradation, respectively. The reason that "two-step cache" suffers from less degradation is that cache flushes may occur during in-memory partition operations and be less harmful. Although the figure shows the worst-case cache interference, it certainly reflects the robustness problem of cache partitioning. In contrast, our prefetching schemes do not assume hash tables and build partitions remain in the cache. As shown in the figure, they are very robust against even cache flushes.

Experiments when "direct cache" also applies. Figure 18 compares our prefetching schemes with cache partitioning. Note that the I/O partition phases of all schemes use the combined prefetching scheme discussed above; the major difference is in the different numbers of I/O partitions generated. The second partition step in "two-step cache" is shown as part of the join phase performance. Moreover, we employ prefetching in the join phase to enhance the cache partitioning schemes wherever possible.

Figure 18(a)-(c) show experiments joining a 200MB build relation with a 400MB probe relation. Every build tuple matches two probe tuples. We increase the tuple size, which results in decreasing numbers of tuples in the relations and the downward trends of the curves. "Direct cache" achieves the best performance in the join phase by avoiding most cache misses. However, it suffers from larger overheads in the partition phase because it generates many more partitions. "Two-step cache" suffers from the overhead of the additional partition step and is 50-150% worse than the prefetching schemes. Overall, our prefetching schemes are the best (slightly better than "direct cache" even under exclusive use of the cache). In Figure 18(d), we keep the tuple size to be 100B and vary the percentage of tuples that have matches. Again, we see similar trends as in Figure 18(c). In Figure 18, our prefetching techniques achieve 1.4-2.5X speedups for the partition phase, 2.1-2.9X speedups for the join phase, and 1.9-2.7X speedups overall compared to the baseline algorithm.[9]

8. Conclusions

While prefetching is a promising technique for improving CPU cache performance, applying it to the hash join algorithm is not straightforward due to the dependencies within the processing of a single tuple and the randomness of hashing. In this paper, we have explored the potential for exploiting *inter-tuple* parallelism to schedule prefetches effectively. Our prefetching techniques—*group prefetching* and *software-pipelined prefetching*—systematically reorder the memory references of hash joins and schedule prefetches so that cache miss latencies in the processing of a tuple can be overlapped with computation and miss latencies of other tuples. We developed generalized models to better understand these techniques and successfully overcame the complexities involved with prefetching the hash join algorithm.

Our results demonstrated that hash join cache performance can be improved dramatically through prefetching. More interestingly, the techniques will still be effective even

[9]The experiments shown in the prefetching curves generate the same number of partitions as in the baseline curve.

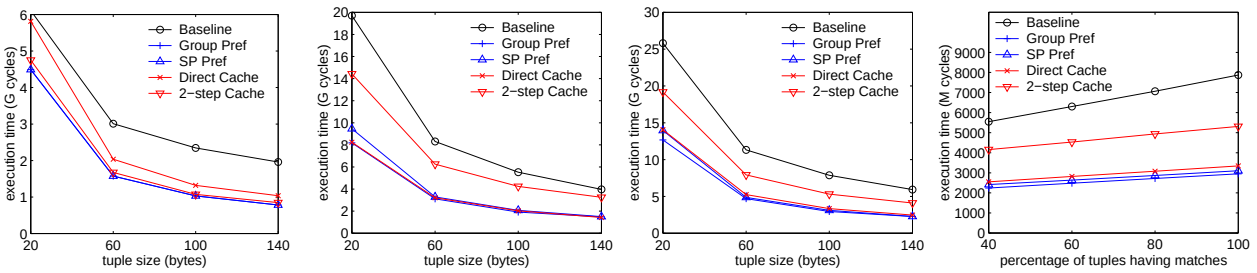

(a) I/O partition phase of (c) using combined prefetching

(b) Join phase of (c)

(c) Overall performance (varying tuple size)

(d) Overall performance (varying join selectivity)

Figure 18. Comparisons with cache partitioning when it applies

when the speed gap between processors and memory increases significantly in the future (e.g., by a factor of six). Moreover, we believe that our techniques can improve other hash-based algorithms such as hash-based group-by and aggregation algorithms, and other algorithms that have inter-element parallelism.

Acknowledgement. This research is supported by a grant from the NSF. The third author thanks P. Bohannon, S. Ganguly, H. F. Korth, and P. P. S. Narayan for helpful discussions.

References

[1] A. Ailamaki, D. J. DeWitt, M. D. Hill, and D. A. Wood. DBMSs on a Modern Processor: Where Does Time Go? In *the 25th VLDB*, pages 266–277, Sept. 1999.

[2] J.-L. Baer and T.-F. Chen. An effective on-chip preloading scheme to reduce data access penalty. In *Proceedings of Supercomputing '91*, 1991.

[3] L. A. Barroso, K. Gharachorloo, and E. D. Bugnion. Memory System Characterization of Commercial Workloads. In *the 25th ISCA*, pages 3–14, June 1998.

[4] P. Bohannon, P. McIlroy, and R. Rastogi. Improving Main-Memory Index Performance with Partial Key Information. In *the SIGMOD Conference*, May 2001.

[5] P. A. Boncz, S. Manegold, and M. Kersten. Database Architecture Optimized for the New Bottleneck: Memory Access. In *the 25th VLDB*, pages 54–65, Sept. 1999.

[6] S. Chen, A. Ailamaki, P. B. Gibbons, and T. C. Mowry. Improving Hash Join Performance through Prefetching. Technical Report CMU-CS-03-157, School of Computer Science, Carnegie Mellon University, Oct. 2003.

[7] S. Chen, P. B. Gibbons, and T. C. Mowry. Improving Index Performance through Prefetching. In *the SIGMOD Conference*, pages 235–246, May 2001.

[8] S. Chen, P. B. Gibbons, T. C. Mowry, and G. Valentin. Fractal Prefetching B^+-Trees: Optimizing Both Cache and Disk Performance. In *the SIGMOD Conference*, June 2002.

[9] Z. Cvetanovic and R. E. Kessler. Performance Analysis of the Alpha 21264-Based Compaq ES40 System. In *the 27th ISCA*, pages 192–202, June 2000.

[10] C. Gniady, B. Falsafi, and T. N. Vijaykumar. Is SC + ILP = RC? In *the 26th ISCA*, May 1999.

[11] G. Graefe. Query Evaluation Techniques for Large Databases. *ACM Computing Surveys*, 25(2):73–170, 1993.

[12] Intel Corp. *Intel Pentium4 and Intel Xeon Processor Optimization*. Order Number: 248966-007.

[13] K. Keeton, D. A. Patterson, Y. Q. He, R. C. Raphael, and W. E. Baker. Performance Characterization of a Quad Pentium Pro SMP using OLTP Workloads. In *the 25th ISCA*, pages 15–26, June 1998.

[14] K. Kennedy and K. McKinley. Loop Distribution With Arbitrary Control Flow. In *Supercomputing'90*, pages 407–416.

[15] M. Kitsuregawa, H. Tanaka, and T. Moto-Oka. Application of Hash to Data Base Machine and its Architecture. *New Generation Computing*, 1(1):63–74, 1983.

[16] M. S. Lam. *A Systolic Array Optimizing Compiler*. PhD thesis, Carnegie Mellon University, May 1987.

[17] B. Lindsay. Hash Joins in DB2 UDB: the Inside Story. *Carnegie Mellon DB Seminar*, Mar 2002.

[18] C.-K. Luk and T. C. Mowry. Compiler-Based Prefetching for Recursive Data Structures. In *the 7th ASPLOS*, pages 222–233, Oct. 1996.

[19] C.-K. Luk and T. C. Mowry. Automatic Compiler-Inserted Prefetching for Pointer-Based Applications. *IEEE Transactions on Computers*, 48(2):134–141, Feb. 1999.

[20] S. Manegold, P. A. Boncz, and M. L. Kersten. What Happens During a Join? Dissecting CPU and Memory Optimization Effects. In *the 26th VLDB*, Sept. 2000.

[21] S. McFarling. Combining Branch Predictors. Technical Report WRL Technical Note TN-36, Digital Equipment Corporation, June 1993.

[22] T. C. Mowry. *Tolerating Latency Through Software-Controlled Data Prefetching*. PhD thesis, Stanford University, Mar. 1994.

[23] T. C. Mowry, M. S. Lam, and A. Gupta. Design and Evaluation of a Compiler Algorithm for Prefetching. In *the 5th ASPLOS*, pages 62–73, Oct. 1992.

[24] M. Nakayama, M. Kitsuregawa, and M. Takagi. Hash-Partitioned Join Method Using Dynamic Destaging Strategy. In *the 14th VLDB*, pages 468–478, Aug. 1988.

[25] C. Nyberg, T. Barclay, Z. Cvetanovic, J. Gray, and D. Lomet. AlphaSort: A RISC Machine Sort. In *the SIGMOD Conference*, pages 233–242, May 1994.

[26] J. Rao and K. A. Ross. Making B^+-Trees Cache Conscious in Main Memory. In *the SIGMOD Conference*, pages 475–486, May 2000.

[27] A. Saulsbury, F. Dahlgren, and P. Stenstrom. Recency-based TLB preloading. In *the 27th ISCA*, pages 117–127, May 2000.

[28] L. D. Shapiro. Join Processing in Database Systems with Large Main Memories. *TODS*, 11(3):239–264, 1986.

[29] A. Shatdal, C. Kant, and J. F. Naughton. Cache Conscious Algorithms for Relational Query Processing. In *the 20th VLDB*, pages 510–521, Sept. 1994.

[30] H. Zeller and J. Gray. An Adaptive Hash Join Algorithm for Multiuser Environments. In *the 16th VLDB*, pages 186–197, Aug. 1990.