

Shared Memory Parallelization of Data Mining Algorithms: Techniques, Programming Interface, and Performance

Ruoming Jin, Ge Yang, and Gagan Agrawal, *Member, IEEE Computer Society*

Abstract—With recent technological advances, shared memory parallel machines have become more scalable, and offer large main memories and high bus bandwidths. They are emerging as good platforms for data warehousing and data mining. In this paper, we focus on shared memory parallelization of data mining algorithms. We have developed a series of techniques for parallelization of data mining algorithms, including full replication, full locking, fixed locking, optimized full locking, and cache-sensitive locking. Unlike previous work on shared memory parallelization of specific data mining algorithms, all of our techniques apply to a large number of popular data mining algorithms. In addition, we propose a reduction-object-based interface for specifying a data mining algorithm. We show how our runtime system can apply any of the techniques we have developed starting from a common specification of the algorithm. We have carried out a detailed evaluation of the parallelization techniques and the programming interface. We have experimented with apriori and fp-tree-based association mining, k-means clustering, k-nearest neighbor classifier, and decision tree construction. The main results from our experiments are as follows. 1) Among full replication, optimized full locking, and cache-sensitive locking, there is no clear winner. Each of these three techniques can outperform others depending upon machine and dataset parameters. These three techniques perform significantly better than the other two techniques. 2) Good parallel efficiency is achieved for each of the four algorithms we experimented with, using our techniques and runtime system. 3) The overhead of the interface is within 10 percent in almost all cases. 4) In the case of decision tree construction, combining different techniques turned out to be crucial for achieving high performance.

Index Terms—Shared memory parallelization, programming interfaces, association mining, clustering, decision tree construction.

1 INTRODUCTION

WITH the availability of large data sets in application areas like bioinformatics, medical informatics, scientific data analysis, financial analysis, telecommunications, retailing, and marketing, it is becoming increasingly important to execute data mining tasks in parallel. At the same time, technological advances have made shared memory parallel machines commonly available to organizations and individuals. SMP machines with two to four processors are frequently used as desktops. Clusters of SMPs are very popular for high-end computing, and offer shared memory parallelism within each node. Shared memory machines are also becoming more scalable. Large shared memory machines with high bus bandwidth and very large main memory are being sold by several vendors. Vendors of these machines are targeting data warehousing and data mining as major markets.

Using these machines for speeding up and scaling up data mining implementations, however, involves a number of challenges. First, the appropriate parallelization strategy could depend upon the data mining task and algorithm that is being parallelized. Second, with an increasingly complicated memory hierarchy, achieving high performance on

SMP machines often requires subtle optimizations. Finally, maintaining, debugging, and performance tuning a parallel application is an extremely time consuming task. As parallel architectures evolve, or architectural parameters change, it is not easy to modify existing codes to achieve high performance on new systems. As new performance optimizations are developed, it is useful to be able to apply them to different parallel applications. Currently, this cannot be done for parallel data mining implementations without a high programming effort.

We believe that the above problems can be alleviated by developing parallelization techniques and runtime support that applies across a variety of data mining algorithms. In our research, we are developing a middleware for rapid development of data mining implementations on large SMPs and clusters of SMPs. Our system is called FREERIDE (FRamework for Rapid Implementation of Datamining Engines). It is based on the observation that parallel versions of several well-known data mining techniques, including a priori association mining [2], k-means clustering [25], k-nearest neighbor classifier [22], and artificial neural networks [22] share a relatively similar structure. The middleware performs distributed memory parallelization across the cluster and shared memory parallelization within each node. It enables high I/O performance by minimizing disk seek time and using asynchronous I/O operations. Thus, it can be used for developing efficient parallel data mining applications that operate on disk-resident data sets.

• The authors are with the Department of Computer and Information Sciences, Ohio State University Columbus, OH 43210.
E-mail: {jinr, yangg, agrawal}@cis.ohio-state.edu.

Manuscript received 15 July 2002; revised 17 Apr. 2003; accepted 27 Aug. 2003.

For information on obtaining reprints of this article, please send e-mail to: tkde@computer.org, and reference IEEECS Log Number 116955.

This paper focuses on shared memory parallelization. We have developed a series of techniques for runtime parallelization of data mining algorithms, including full replication, full locking, fixed locking, optimized full locking, and cache-sensitive locking. Unlike previous work on shared memory parallelization of specific data mining algorithms, all of our techniques apply across a large number of common data mining algorithms. The techniques we have developed involve a number of trade offs between memory requirements, opportunity for parallelization, and locking overheads. Thus, the relative performance of these techniques depends upon machine parameters, as well as the characteristics of the algorithm and the data set.

In addition, we present and evaluate the middleware interface and the underlying runtime support for shared memory parallelization. We describe how a programmer can perform minor modifications to a sequential code and specify a data mining algorithm using the reduction object interface we have developed. We allow complex reduction objects and user-defined reduction functions, which are not available in OpenMP. We show how the runtime system can apply any of the techniques we have developed starting from a common specification that uses the reduction object interface.

We initially evaluated our techniques and programming interface using our implementations of a priori and fp-tree-based association mining, k-means clustering, and k-nearest neighbor classifier. The main results from these experiments are as follows:

1. Among full replication, optimized full locking, and cache-sensitive locking, there is no clear winner. Each of these three techniques can outperform others depending upon machine and data set parameters. These three techniques perform significantly better than the other two techniques.
2. Our techniques scale well on a large SMP machine.
3. The overhead of the interface is within 10 percent in almost all cases.

We have also carried out a detailed case study of applying our techniques and runtime system for decision tree construction. We have particularly focused on parallelization of the *RainForest* approach originally proposed by Gehrke et al. [14]. Typically, the techniques used for parallelizing decision tree construction have been quite different than the techniques used for association mining and clustering. Here, we have demonstrated that the full replication and either optimized full locking or cache-sensitive locking can be combined to achieve an efficient parallel implementation for decision tree construction.

Overall, our work has shown that a common collection of techniques can be used to efficiently parallelize algorithms for a variety of mining tasks. Moreover, a high-level interface can be supported to allow the programmers to rapidly create parallel implementations.

The rest of this paper is organized as follows: Section 2 reviews parallel versions of several common data mining techniques. Parallelization techniques are presented in Section 3. The middleware interface and implementation of different techniques starting from the common specification is described in Section 4. Experimental results from association mining, clustering, and k-nearest neighbor

search are presented in Section 5. A detailed case study, decision tree construction, is presented in Section 6. We compare our work with related research efforts in Section 7. Our broader vision, which shows the application of the techniques described in this paper, is presented in Section 8. We conclude in Section 9.

2 PARALLEL DATA MINING ALGORITHMS

In this section, we describe how several commonly used data mining techniques can be parallelized on a shared memory machine in a very similar way. Our discussion focuses on five important data mining techniques: associating mining [2], k-means clustering [25], k-nearest neighbors [22], artificial neural networks [22], and Bayesian networks [9].

2.1 Association Mining

Association rule mining is the process of analyzing a set of transactions to extract *association rules* and is a very commonly used and well-studied data mining problem [3], [57]. Given a set of transactions¹ (each of them being a set of items), an association rule is an expression $X \rightarrow Y$, where X and Y are the sets of items. Such a rule implies that transactions in databases that contain the items in X also tend to contain the items in Y .

Formally, the goal is to compute the sets L_k . For a given value of k , the set L_k comprises the frequent itemsets of length k . A well accepted algorithm for association mining is the a priori mining algorithm [3]. The main observation in the a priori technique is that if an itemset occurs with frequency f , all the subsets of this itemset also occur with at least frequency f . In the first iteration of this algorithm, transactions are analyzed to determine the frequent 1-itemsets. During any subsequent iteration k , the frequent itemsets L_{k-1} found in the $(k-1)$ th iteration are used to generate the candidate itemsets C_k . Then, each transaction in the data set is processed to compute the frequency of each member of the set C_k . K-itemsets from C_k that have a certain prespecified minimal frequency (called the *support level*) are added to the set L_k .

A simple shared memory parallelization scheme for this algorithm is as follows: One processor generates the complete C_k using the frequent itemset L_{k-1} created at the end of the iteration $k-1$. The transactions are scanned, and each transaction (or a set of transactions) is assigned to one processor. This processor evaluates the transaction(s) and updates the counts of candidates itemsets that are found in this transaction. Thus, by assigning different sets of transactions to processors, parallelism can be achieved. The only challenge in maintaining correctness is avoiding the possible race conditions when multiple processors may want to update the count of the same candidate.

The same basic parallelization strategy can be used for parallelization of a number of other association mining algorithms and variations of a priori, including SEAR [38], DHP [42], Partition [49], and DIC [8]. These algorithms differ from the a priori algorithm in the data structure used for representing candidate itemsets, candidate space prun-

1. We use the terms *transactions*, *data items*, and *data instances* interchangeably.

ing, or in reducing passes over the set of transactions, none of which require a significant change in the parallelization strategy. More importantly, the same strategy can be used for the *tree-growth* phase of the relatively new FP-tree-based frequent itemset mining algorithm [21].

2.2 K-Means Clustering

The second data mining algorithm we describe is the k-means clustering technique [25], which is also very commonly used. This method considers transactions or data instances as representing points in a high-dimensional space. Proximity within this space is used as the criterion for classifying the points into clusters.

Four steps in the sequential version of this algorithm are as follows:

1. Start with k given centers for clusters.
2. Scan the data instances, for each data instance (point), find the center closest to it, and assign this point to the corresponding cluster.
3. Determine the k centroids from the points assigned to the corresponding center.
4. Repeat this process until the assignment of points to cluster does not change.

It is important to note that the convergence of the algorithm is dependent upon the initial choice of k centers.

This method can also be parallelized in a fashion very similar to the method we described for a priori association mining. The data instances are read, and each data instance (or a set of instances) are assigned to one processor. This processor performs the computations associated with the data instance, and then updates the center of the cluster this data instance is closest to. Again, the only challenge in maintaining correctness is avoiding the race conditions when multiple processors may want to update center of the same cluster. A number of other clustering algorithms, like k-harmonic means and expectation-maximization (EM), can also be parallelized using the same basic scheme [12].

2.3 K-Nearest Neighbors

K-nearest neighbor classifier is based on learning by analogy [22]. The training samples are described by an n-dimensional numeric space. Given an unknown sample, the k-nearest neighbor classifier searches the pattern space for k training samples that are closest, using the Euclidean distance, to the unknown sample.

Again, this technique can be parallelized as follows: Each training sample is processed by one processor. After processing the sample, the processor determines if the list of k current nearest neighbors should be updated to include this sample. Again, the correctness issue is the race conditions if multiple processors try to update the list of nearest neighbors at the same time.

2.4 Artificial Neural Networks

An artificial neural network is a set of connected input/output units where each connection has a weight associated with it. During the learning phase, the network learns by adjusting the weights so as to be able to predict the correct class labels of the input samples. A very commonly used algorithm for training a neural network is *backpropagation*

```

{ * Outer Sequential Loop * }
While() {
  { * Reduction Loop * }
  Foreach(element e) {
    (i, val) = process(e);
    Reduc(i) = Reduc(i) op val;
  }
}

```

Fig. 1. Structure of common data mining algorithms.

[22]. For each training sample, the weights are modified so as to minimize the difference between the network's prediction and the actual class label. These modifications are made in the backward direction.

The straight forward method for parallelizing this technique on a shared memory machine is as follows: Each transaction is assigned to one processor. This processor performs the computations associated with this transaction and updates the weights for each connection in the network. Again, the only correctness consideration is the possibility of race conditions when the weights are updated.

2.5 Bayesian Network

Bayesian network is an approach to unsupervised classification [9]. Each transaction or data instance X_i is represented as an ordered vector of attribute values $\{X_{i1}, \dots, X_{ik}\}$. Given a set of data instances, the goal is to search for the best class descriptions that predict the data. Class membership is expressed probabilistically, i.e., a data instance probabilistically belongs to a number of possible classes. The classes provide probabilities for all attribute values of each instance. Class membership probabilities are then determined by combining all these probabilities.

The two most time consuming steps in computing the classification are `update_wts` and `update_parameters`. `update_wts` computes the weight of each class, which is the sum of the probabilities of each data instance being in that class. `update_parameters` uses the weights computed to update the parameters for classification used during the next phase.

A parallelization strategy that can be used for both of these steps is as follows: The data instances are partitioned across processors. In the `update_wts` phase, each processor updates the weight of each class after processing each data instance. The sequential version of `update_parameters` is composed of three nested loops. The outermost loop iterates over all the classes, the next loop iterates over all attributes, and the innermost loop iterates over the data instances. The innermost loop uses the values of all data instances to compute the class parameters. Since the data instances have been partitioned across processors, parallelization is done at the innermost loop. Processors update class parameters after processing each data instance. For both the phases, the correctness challenge is the race condition when weights of the class or class parameters are updated.

3 PARALLELIZATION TECHNIQUES

This section focuses on parallelization techniques we have developed for data mining algorithms.

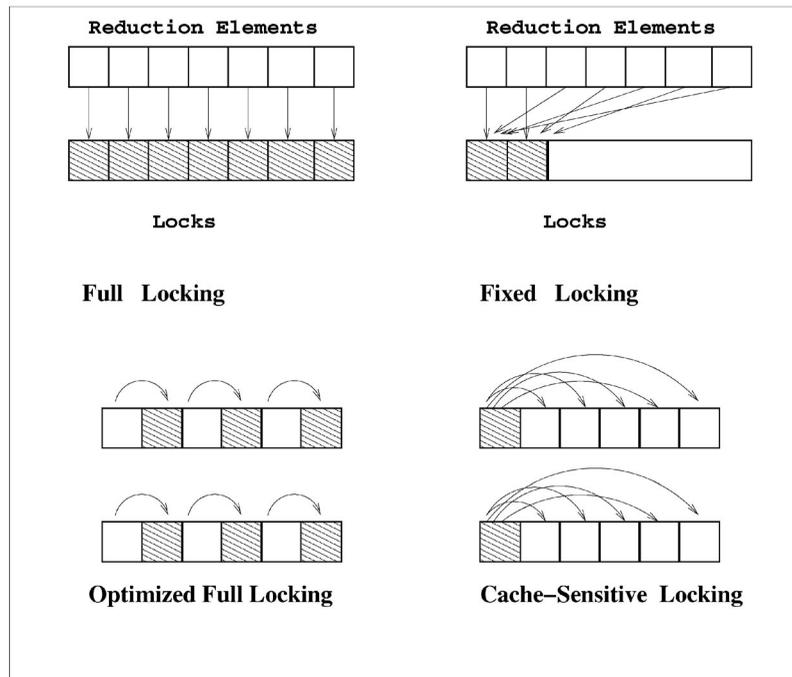


Fig. 2. Memory layout for various locking schemes.

3.1 Overview of the Problem

In the previous section, we argued how several data mining algorithms can be parallelized in a very similar fashion. The common structure behind these algorithms is summarized in Fig. 1. The function op is an associative and commutative function. Thus, the iterations of the foreach loop can be performed in any order. The data-structure $Reduc$ is referred to as the reduction object.

The main correctness challenge in parallelizing a loop like this on a shared memory machine arises because of possible race conditions when multiple processors update the same element of the reduction object. The element of the reduction object that is updated in a loop iteration (i) is determined only as a result of the processing. For example, in the a priori association mining algorithm, the data item read needs to be matched against all candidates to determine the set of candidates whose counts will be incremented. In the k-means clustering algorithm, first the cluster to which a data item belongs is determined. Then, the center of this cluster is updated using a reduction operation.

The major factors that make these loops challenging to execute efficiently and correctly are as follows:

- It is not possible to statically partition the reduction object so that different processors update disjoint portions of the collection. Thus, race conditions must be avoided at runtime.
- The execution time of the function $process$ can be a significant part of the execution time of an iteration of the loop. Thus, runtime preprocessing or scheduling techniques cannot be applied.
- In many algorithms, the size of the reduction object can be quite large. This means that the reduction object cannot be replicated or privatized without significant memory overheads.

- The updates to the reduction object are *fine-grained*. The reduction object comprises a large number of elements that take only a few bytes, and the foreach loop comprises a large number of iterations, each of which may take only a small number of cycles. Thus, if a locking scheme is used, the overhead of locking and synchronization can be significant.

3.2 Techniques

We have implemented five approaches for avoiding race conditions as multiple threads may want to update the same elements in the reduction object. These techniques are, *full replication*, *full locks*, *optimized full locks*, *fixed locks*, and *cache-sensitive locks*.

3.2.1 Full Replication

One simple way of avoiding race conditions is to replicate the reduction object and create one copy for every thread. The copy for each thread needs to be initialized in the beginning. Each thread simply updates its own copy, thus avoiding any race conditions. After the local reduction has been performed using all the data items on a particular node, the updates made in all the copies are *merged*.

The other four techniques are based upon locking. The memory layout for these four techniques is shown in Fig. 2.

3.2.2 Full Locking

One obvious solution to avoiding race conditions is to associate one lock with every element in the reduction object. After processing a data item, a thread needs to acquire the lock associated with the element in the reduction object it needs to update. For example, in the a priori association mining algorithm, there will be a lock associated with the count for each candidate, which will need to be acquired before updating that count. If two

TABLE 1
Trade Off among the Techniques

	Full Replication	Optimized Full Locks	Cache Sensitive Locks	Full Locks	Fixed Locks
Memory Requirement	very high	high	low	high	low
Parallelism	very high	high	medium	high	low
Locking Overhead	none	medium	high	medium	high
Cache Misses	low	medium	low	high	medium
False Sharing	none	yes	none	yes	yes
Merging Costs	yes	none	none	none	none

threads need to update the count of the same candidate, one of them will need to wait for the other one to release the lock. In a priori association mining, the number of candidates considered during any iteration is typically quite large, so the probability of one thread needing to wait for another one is very small.

In our experiment with a priori, with 2,000 distinct items and a support level of 0.1 percent, up to 3 million candidates were generated. In full locking, this means supporting 3 million locks. Supporting such a large number of locks results in overheads of three types. The first overhead is the high memory requirement associated with a large number of locks. The second overhead comes from cache misses. Consider an update operation. If the total number of elements is large and there is no locality in accessing these elements, then the update operation is likely to result in two cache misses, one for the element and second for the lock. This cost can slow down the update operation significantly.

The third overhead is of *false sharing* [23]. In a cache-coherent shared memory multiprocessor, false sharing happens when two processors want to access different elements from the same cache block. In full locking scheme, false sharing can result in cache misses for both reduction elements and locks.

We have designed three schemes to overcome one or more of these three overheads associated with full locking. These three techniques are, *optimized full locks*, *fixed locks*, and *cache-sensitive locks*.

3.2.3 Optimized Full Locks

The next scheme we describe is optimized full locks. Optimized full locks scheme overcomes the the large number of cache misses associated with full locking scheme by allocating a reduction element and the corresponding lock in consecutive memory locations, as shown in Fig. 2. By appropriate alignment and padding, it can be ensured that the element and the lock are in the same cache block. Each update operation now results in at most one cold or capacity cache miss. The possibility of false sharing is also reduced. This is because there are fewer elements (or locks) in each cache block. This scheme does not reduce the total memory requirements.

3.2.4 Fixed Locking

To alleviate the memory overheads associated with the large number of locks required in the full locking and optimized full locking schemes, we designed the fixed

locking scheme. As the name suggests, a fixed number of locks are used. The number of locks chosen is a parameter to this scheme. If the number of locks is l , then the element i in the reduction object is assigned to the lock $i \bmod l$. So, in the a priori association mining algorithm, if a thread needs to update, the support count for the candidate i , it needs to acquire the lock $i \bmod l$. In Fig. 2, two locks are used. Alternate reduction elements use each of these two locks.

Clearly, this scheme avoids the memory overheads associated with supporting a large number of locks in the system. The obvious tradeoff is that as the number of locks is reduced, the probability of one thread having to wait for another one increases. Also, each update operation can still result in two cache misses. Fixed locking can also result in even more false sharing than the full locking scheme. This is because now there is a higher possibility of two processors wanting to acquire locks in the same cache block.

3.2.5 Cache-Sensitive Locking

The final technique we describe is *cache-sensitive locking*. This technique combines the ideas from optimized full locking and fixed locking. Consider a 64 byte cache block and a 4 byte reduction element. We use a single lock for all reduction elements in the same cache block. Moreover, this lock is allocated in the same cache block as the elements. So, each cache block will have 1 lock and 15 reduction elements. This scheme results in lower memory requirements than the full locking and optimized full locking schemes. Similar to the fixed locking scheme, this scheme could have the potential of limiting parallelism. However, in cache-coherent multiprocessors, if different CPUs want to concurrently update distinct elements of the same cache block, they incur several cache misses. This is because of the effect of false sharing. This observation is exploited by the cache-sensitive locking scheme to have a single lock with all elements in a cache block.

Cache-sensitive locking reduces each of three types of overhead associated with full locking. Each update operation results in at most one cache miss, as long as there is no contention between the threads. The problem of false sharing is also reduced because there is only one lock per cache block.

One complication in implementation of cache-sensitive locking scheme is that modern processors have two or more levels of cache and the cache block size is different at different levels. Our implementation and experiments have been done on machines with two levels of cache, denoted by L1 and L2. Our observation is that if the reduction object exceeds the size of L2 cache, L2 cache misses are a more dominant overhead. Therefore, we have used the size of L2

cache in implementing the cache-sensitive locking scheme.

3.3 Comparing the Techniques

We now compare the five techniques we have presented along six criteria. The comparison is summarized in Table 1. The six criteria we use are:

- Memory requirements, denoting the extra memory required by a scheme because of replication or locks.
- Parallelism, denoting if parallelism is limited because of contention for locks.
- Locking overhead, which includes the cost of acquiring and releasing a lock and any extra computational costs in computing the address of the lock associated with the element.
- Cache misses, which only includes cold and capacity cache misses, and excludes coherence cache misses and false sharing.
- False sharing, which occurs when two processors want to access reduction elements or locks on the same cache block.
- Merge costs denotes the cost of merging updates made on replicated copies.

Full replication has the highest memory requirement. If the size of reduction object is S and there are T threads, then the total memory requirement is $S \times T$. Full locking and optimized full locking are next in the level of memory requirement. If each lock takes the same number of bytes as each reduction element, the total memory requirement for both these schemes is $2 \times S$. Fixed locking and cache-sensitive locking have the lowest memory requirements. If one lock is used for every r reduction element, total memory requirement in both the schemes is $(1 + 1/r) \times S$.

The full replication scheme does not limit parallelism in any way because each thread updates its own copy of reduction object. Full locks and optimized full locks are next in the level of parallelism. The probability of two threads trying to acquire the same lock is very small in both these schemes. Cache-sensitive and fixed locks can both limit parallelism because of sharing of locks.

The next criteria we address is the computation overhead because of locking. Full replication does not have such overhead. All locking schemes involve the cost of acquiring and releasing a lock. In addition, fixed locking and cache-sensitive locking require extra computation to determine the locks that needs to be used for an element.

Cache misses can be a significant cost when the total memory required by reduction object and locks is large. In full replication, there is no increase in working set size or the size of memory accessed by each processor because of locks. Only the reduction element needs to be accessed during the update operation. Replicating the reduction object does not increase the number of cold or capacity misses, because each processor has its own cache and accesses its own copy of the reduction object. As we discussed earlier, full locking scheme is the worst with respect to cold and capacity misses. The locks double the working set size at each processor. Further, each update operation results in accesses to two cache blocks.

We have denoted the cache misses for both fixed locking and optimized full locking as *medium*. In fixed locking, the

```

void Kmeans::initialize() {
    for (int i = 0; i < k; i++) {
        clusterID[i]=reducobject->alloc(ndim + 2);
    }
    /* Initialize Centers */
}

void Kmeans::reduction(void *point) {
    for (int i=0; i < k; i++) {
        dis=distance(point, i);
        if (dis < min) {
            min=dis;
            min_index=i;
        }
    }
    objectID=clusterID[min_index];
    for (int j=0; j < ndim; j++)
        reducobject->Add(objectID, j, point[j]);
    reducobject->Add(objectID, ndim, 1);
    reducobject->Add(objectID, ndim + 1, dis);
}

```

Fig. 3. Initialization and local reduction functions for k-means.

total working set size increases only marginally because of locks. However, each update operation can still result in two cache misses. In optimized full locking, the working set size is the same as in full locking. However, each update operation needs to access only a single cache block.

Cache-sensitive locking is the best one among all locking schemes with respect to cache misses. Like fixed locking, the working set size is increased only marginally because of locks. Also, each update operation needs to access only a single cache block.

In our implementation, read-only data is segregated from the reduction object. Therefore, there is no false sharing in full replication. In full locking, optimized full locking, and fixed locking, two threads can access the same cache block comprising either reduction elements or locks, and incur the cost of false sharing. Cache-sensitive locking does not have any false sharing. This is because there is only one lock in a cache block and two threads cannot simultaneously acquire the same lock. Any cache misses incurred by a thread waiting to acquire a lock are considered a part of the waiting cost.

Our final criteria is the cost of merging. This cost is only incurred by the full replication scheme.

4 PROGRAMMING INTERFACE

In this section, we explain the interface we offer to the programmers for specifying a parallel data mining algorithm. We also describe how each of the five techniques described in the previous section can be implemented starting from a common specification.

4.1 Middleware Interface

As we stated earlier, this work is part of our work on developing a middleware for rapid development of data mining implementations on large SMPs and clusters of SMPs [27], [26]. Our middleware exploits the similarity in both shared memory and distributed memory paralleliza-

```

template<class T>
inline void ReducObject<T>::Reduc(int ObjectID,
    int Offset, void (*func)(void *, void *),
    int *param) {
    T *group_address =reducgroup[ObjectID];
    SWITCH (TECHNIQUE) {
        CASE FULL_REPLICATION:
            func(group_address[Offset],param);
            break;
        CASE FULL_LOCKS:
            offset = abs_offset(ObjectID,Offset);
            S_LOCK(&locks[offset]);
            func(group_address[Offset],param);
            S_UNLOCK(&locks[offset]);
            break;
        CASE FIXED_LOCKS:
            offset = abs_offset(ObjectID,Offset);
            S_LOCK(&locks[offset%num_locks]);
            func(group_address[Offset],param);
            S_UNLOCK(&locks[offset%num_locks]);
            break;
        CASE OPTIMIZE_FULL_LOCKS:
            S_LOCK(&group_address[Offset*2]);
            func(group_address[Offset*2+1],value);
            S_UNLOCK(&group_address[Offset*2]);
            break;
        CASE CACHE_SENSITIVE_LOCKS:
            cache_index = divide15(Offset);
            S_LOCK(&group_address[cache_index*16]);
            func(group_address[Offset+cache_index+1],value);
            S_UNLOCK(&group_address[cache_index*16]);
            break;
    }
}

```

Fig. 4. Implementation of different parallelization techniques starting from a common specification.

tion strategy to offer a high-level interface to the programmers. For shared memory parallelization, the programmer is responsible for creating and initializing a reduction object. Further, the programmer needs to write a local reduction function that specifies the processing associated with each transaction. The initialization and local reduction functions for k-means are shown in Fig. 3.

As we discussed in Section 3.1, a common aspect of data mining algorithms is the *reduction object*. Declaration and allocation of a reduction object is a significant aspect of our middleware interface. There are two important reasons why reduction elements need to be separated from other data-structures. First, by separating them from read-only data-structures, false sharing can be reduced. Second, the middleware needs to know about the reduction object and its elements to optimize memory layout, allocate locks, and potentially replicate the object.

Consider, as an example, a priori association mining algorithm. Candidate itemsets are stored in a prefix or hash tree. During the reduction operation, the interior nodes of the tree are only read. Associated with each leaf node is the support count of the candidate itemset. All such counts need to be allocated as part of the reduction object. To facilitate updates to the counts while traversing the tree, pointers from leaf node to appropriate elements within the reduction object need to be inserted. Separate allocation of candidate counts allows the middleware to allocate appropriate number of locks depending upon the parallelization scheme used and optimize the memory layout of counts and locks. If full replication is used, the counts are replicated, without replicating the candidate tree. Another important benefit is avoiding or reducing the problem of false sharing. Separate allocation of counts ensures that the nodes of the tree and the counts of candidates are in separate cache blocks. Thus, a thread cannot incur false sharing misses while traversing the nodes of the tree, which is now a read-only data-structure. A disadvantage of separate allocation is that extra pointers need to be stored as part of the tree. Further, there is extra pointer chasing as part of the computation.

Two granularity levels are supported for reduction objects, the *group* level and the *element* level. One group is allocated at a time and comprises a number of elements. The goal is to provide programming convenience, as well as high performance. In a priori, all k itemsets that share the same parent $k - 1$ itemsets are typically declared to be in the same group. In k-means, a group represents a center, which has $ndim + 2$ elements, where $ndim$ is the number of dimensions in the coordinate space.

After the reduction object is created and initialized, the runtime system may *clone* it and create several copies of it. However, this is transparent to the programmer, who views a single copy of it.

The reduction function shown in Fig. 3 illustrates how updates to elements within a reduction object are performed. The programmer writes sequential code for processing, except the updates to elements within a reduction object are performed through member functions of the reduction object. A particular element in the reduction object is referenced by a group identifier and an offset within the group. In this example, *add* function is invoked for all elements. Besides supporting the commonly used reduction functions, like addition, multiplication, maximum, and minimum, we also allow user defined functions. A function pointer can be passed a parameter to a generic reduction function. The reduction functions are implemented as part of our runtime support. Several parallelization strategies are supported, but their implementation is kept transparent from application programmers.

After the reduction operation has been applied on all transactions, a *merge* phase may required, depending upon the parallelization strategy used. If several copies of the reduction object have been created, the merge phase is responsible for creating a single correct copy. We allow the application programmer to choose between one of the standard merge functions (like add corresponding elements from all copies), or to supply their own function.

4.2 Implementation from the Common Interface

Outline of the implementation of these five techniques is shown in Fig. 4. Implementation of a general reduction function *reduc*, which takes a function pointer *func* and a pointer to a parameter *param*, is shown in this figure.

The reduction element is identified by an *ObjectID* and an *Offset*. The operation *reducgroup[ObjectID]* returns the starting address of the group to which the element belongs. This value is stored in variable *group_address*. The function *abs_offset* returns the offset of an element from the start of allocation of first reduction group.

Implementation of full replication is straightforward. The function *func* with the parameter *param* is applied to the reduction element.

Implementation of full locks is also simple. If *offset* is the offset of an element from start of the allocation of reduction objects, *locks[offset]* denotes the lock that can be used for this element. We use simple *spin locks* to reduce the locking overhead. Since we do not consider the possibility of executing more than one thread per processor, we do not need to block a thread that is waiting to acquire a lock. In other words, a thread can simply keep spinning until it acquires the lock. This allows us to use much simpler locks than the ones used in posix threads. The number of bytes taken by each spin lock is either eight or the number of bytes required for storing each reduction element, whichever is smaller. These locks reduce the memory requirements and the cost of acquiring and releasing a lock.

The only difference in the implementation of fixed locking from full locking is that a *mod* operation is performed to determine which lock is used.

However, the implementations of *optimized full locking* and *cache-sensitive locking* are more involved. In both cases, each reduction group is allocated combining the memory requirements of the reduction elements and locks. In optimized full locking scheme, given an element with a particular *Offset*, the corresponding lock is at $group_address + Offset * 2$ and the element is at $group_address + Offset * 2 + 1$. In cache-sensitive locking, each reduction object is allocated at the start of a cache block. For simplification of our presentation, we assume that a cache block is 64 bytes and each element or lock takes 4 bytes. Given an element with a particular *Offset*, $Offset/15$ determines the cache block number within the group occupied by this element. The lock corresponding to this element is at $group_address + Offset/15 * 16$. The element itself is at

$$group_address + Offset + Offset/15 + 1.$$

Implementation of cache-sensitive locking involves a division operation that cannot be implemented using shifts. This can add significant overhead to the cache-sensitive locking scheme. To reduce this overhead, we use special properties of 15 (or 7) to implement a *divide15* (or *divide7*) function.

5 EXPERIMENTAL RESULTS

We have so far implemented five data mining algorithms using our interface for shared memory parallelization. These algorithms are a priori association mining, fp-tree

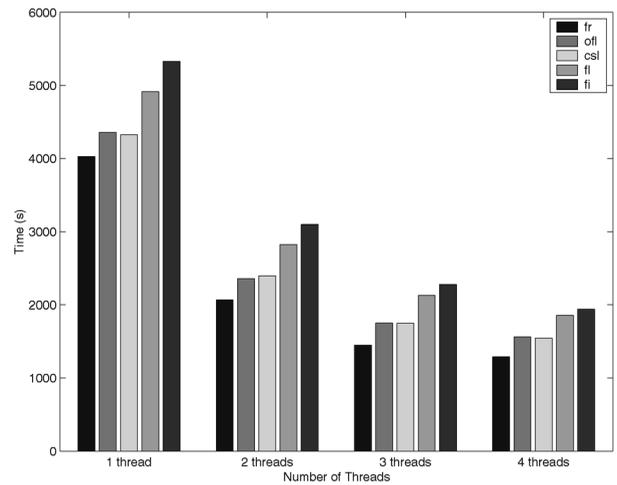


Fig. 5. Comparing all five techniques for a priori.

based association mining, k-means clustering, k-nearest neighbor classifier, and RainForest-based decision tree construction. This section reports experimental results from the first four applications. We have conducted a series of experiments to evaluate the following:

- Parallel performance achieved using the techniques we have developed for parallelization of data mining algorithms.
- The overhead introduced by the interface we have developed, i.e., the relative difference in the performance between the versions that use our interface and the versions that apply the same parallelization technique manually.

Throughout this section, the program versions in which a parallelization technique was implemented by hand are referred to as *manual* versions, and versions where parallelization was done using the middleware interface are referred to as *interface* versions.

5.1 Experimental Platform

We used three different SMP machines for our experiments.

The first machine is a Sun Microsystems Ultra Enterprise 450, with 4 250MHz Ultra-II processors and 1 GB of 4-way interleaved main memory. This configuration represents a common SMP machine available as a desktop or as part of a cluster of SMP workstations.

The second machine used for our experiments is a 24 processor SunFire 6800. Each processor in this machine is a 64bit, 900MHz Sun UltraSparc III. Each processor has a 96 KB L1 cache and a 64 MB L2 cache. The total main memory available is 24 GB. The Sun Fireplane interconnect provides a bandwidth of 9.6 GB per second. This configuration represents a state-of-the-art server machine that may not be affordable to all individuals or organizations interested in data mining.

Finally, to demonstrate that our techniques can be ported on a diverse set of architectures, we ported our software on Pentium-based 4 processor SMP machines. Our new results on k-means clustering on high-dimensional data sets use such a machine.

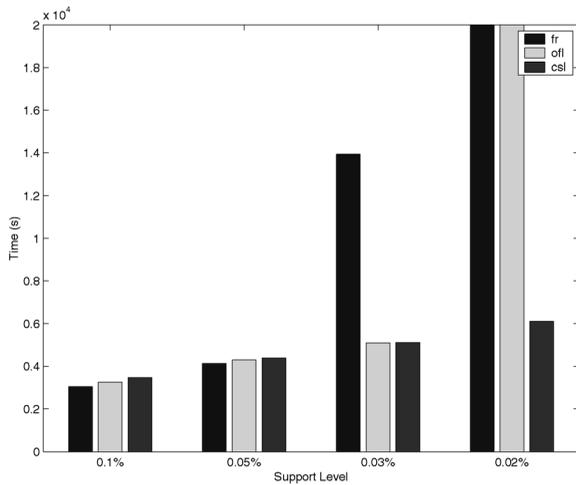


Fig. 6. Relative performance of full replication, optimized full locking, and cache-sensitive locking: four threads, different support levels.

5.2 Results from A Priori

Our first experiment focused on evaluating all five techniques. Since we were interested in seeing the best performance that can be obtained from these techniques, we used only manual versions of each of these techniques. We used a 1 GB data set, generated using a tool available from IBM [3]. The total number of distinct items was 1,000 and the average number of items in a transaction was 15. A confidence of 90 percent and support of 0.5 percent is used. The results are presented in Fig. 5.

The versions corresponding to the full replication, optimized full locking, cache-sensitive locking, full locking, and fixed locking are denoted by *fr*, *ofl*, *csl*, *f1*, and *fi*, respectively. Execution times using one, two, three, and four threads are presented on the four processor SMP machine. With 1 thread, *fr* does not have any significant overheads as compared to the sequential version. Therefore, this version is used for reporting all speedups. With one thread, *ofl* and *csl* are slower by nearly 7 percent, *f1* is slower by 22 percent, and *fi* is slower by 30 percent. For this data set, even after replicating four times, the reduction object did not exceed the main memory. Therefore, *fr* has the best speedups. The speedup with four threads is 3.12 for *fr*, 2.58 with *ofl*, 2.60 with *csl*, 2.16 with *f1*, and 2.07 with *fi*.

From this experiment and our discussion in Section 3.3, *f1* and *fi* do not appear to be competitive schemes. Though the performance of *ofl* and *csl* is considerably lower than *fr*, they are promising for the cases when sufficient memory for supporting full replication may not be available. Therefore, in the rest of this section, we only focus on full replication, optimized full locking, and cache-sensitive locking.

Our second experiment demonstrates that each of these three techniques can be the winner, depending upon the problem and the data set. We use a data set with 2,000 distinct items, where the average number of items per transaction is 20. The total size of the data set is 500 MB and a confidence level of 90 percent is used. We consider four support levels, 0.1, 0.05, 0.03, and 0.02 percent. Again, since we were only interested in relative performance of the three techniques, we experimented with manual version only.

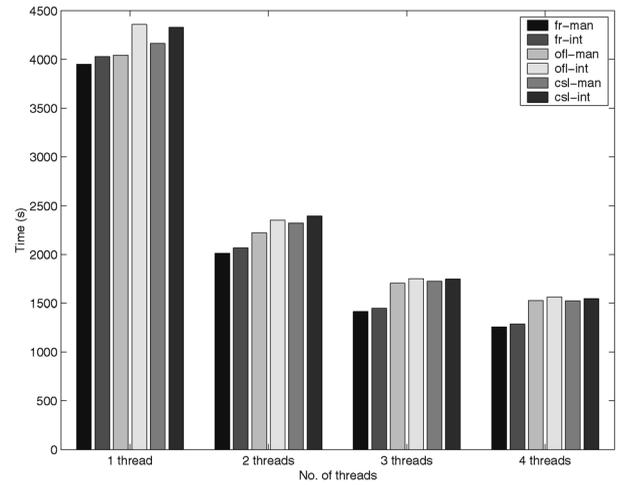


Fig. 7. Scalability and middleware overhead for a priori: 4 processor SMP machine.

The results are shown in Fig. 6. We present results on the 4 processor SMP using four threads. In a priori association mining, the total number of candidate itemsets increases as the support level is decreased. Therefore, the total memory requirement for the reduction objects also increases. When support level is 0.1 or 0.05 percent, sufficient memory is available for reduction object, even after replicating four times. Therefore, *fr* gives the best performance. At the support level of 0.1 percent, *ofl* is slower by 7 percent and *csl* is slower by 14 percent. At the support level of 0.05 percent, they are slower by 4 and 6 percent, respectively. When the support level is 0.03 percent, the performance of *fr* degrades dramatically. This is because replicated reduction object does not fit in main memory and memory thrashing occurs. Since the memory requirements of locking schemes are lower, they do not see the same effect. *ofl* is the best scheme in this case, though *csl* is slower by less than 1 percent. When the support level is 0.02 percent, the available main memory is not even sufficient for *ofl*. Therefore, *csl* has the best performance. The execution time for *csl* was 6,117 seconds, whereas the execution time for *ofl* and *fr* was more than 80,000 seconds.

The next two experiments evaluated scalability and middleware overhead on 4 processor and large SMP, respectively. We use the same data set as used in the first experiment. We created manual and interface versions of each of the three techniques, full replication, optimized full locking, and cache sensitive locking. Thus, we had six versions, denoted by *fr-man*, *fr-int*, *ofl-man*, *ofl-int*, *csl-man*, and *csl-int*.

Results on 4 processor SMP are shown in Fig. 7. The results of manual versions are the same as ones presented in Fig. 5. The overhead of middleware's general interface is within 5 percent in all but two cases, and within 10 percent in all cases. The overhead of middleware comes primarily because of extra function calls and pointer chasing.

Results on the large SMP machine are shown in Fig. 8. We were able to use only up to 16 processors of this machine at any time. We have presented experimental data on 1, 2, 4, 8, 12, and 16 threads. Because the total memory available is very large, sufficient memory is always available for *fr*. Therefore,

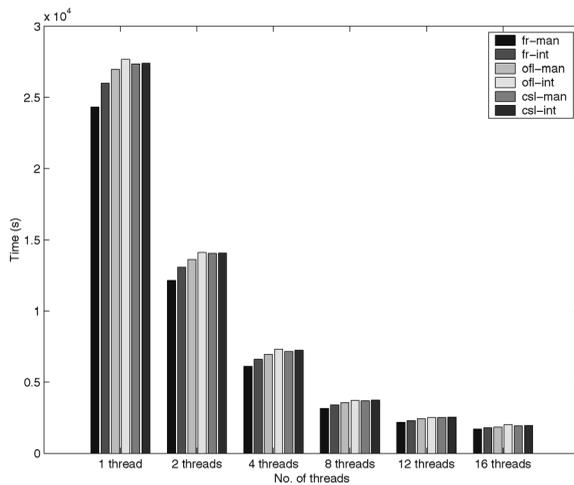


Fig. 8. Scalability and middleware overhead for a priori: large SMP machine.

fr always gives the best performance. However, the locking versions are slower by at most 15 percent.

One interesting observation is that all versions have a uniformly high relative speedup from 1 to 16 threads. The relative speedups for six versions are 14.30, 14.69, 14.68, 13.85, 14.29, and 14.10, respectively. This shows that different versions incur different overheads with 1 thread, but they all scale well. The overhead of middleware is within 10 percent in all cases. In some cases, it is as low as 2 percent.

5.3 Results from K-Means

For evaluating our implementation of k-means, we used a 200 MB data set comprising nearly 16 million three-dimensional points. The value of k we used was 1,000. This value was chosen because we wanted to see the impact from a larger reduction object. Our previous work has shown good speedups even with smaller values of k [27]. In parallelizing k-means, the ratio of communication and computation or the load-balance is not impacted by the specific values of points. Since our focus was on measuring

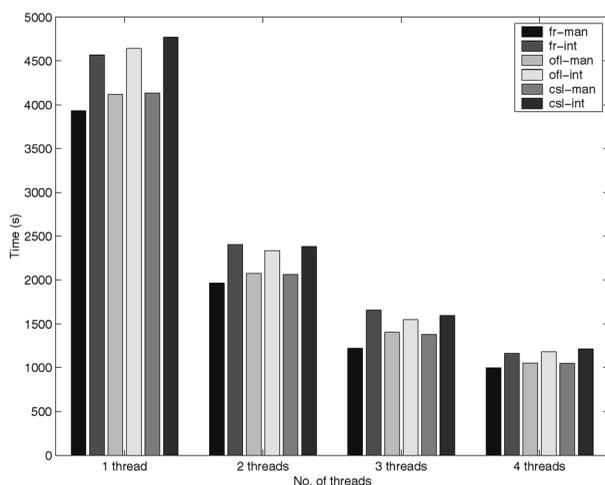


Fig. 9. Scalability and middleware overhead for k-means: 4 processor SMP machine.

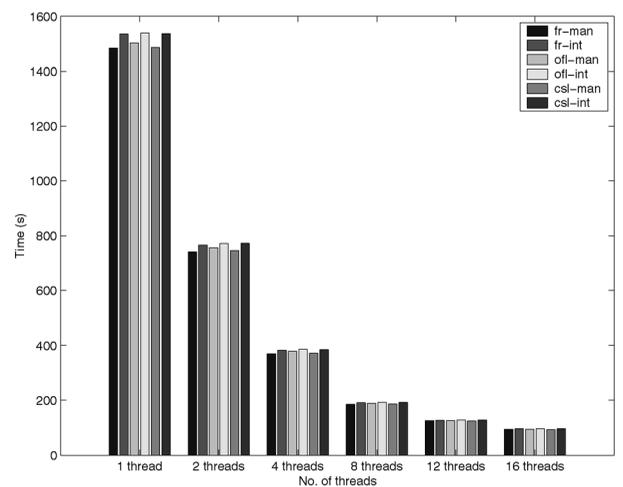


Fig. 10. Scalability and middleware overhead for k-means: large SMP machine.

parallel speedups, and not the rate of convergence, the data-points were chosen randomly.

We focused on three techniques that produced competitive performance for a priori, i.e. full replication, optimized full locking, and cache-sensitive locking. We conducted experiments to evaluate scalability, relative performance, and middleware overheads on the 4 processor and large SMP machines.

The results on 4 processor machine are presented in Fig. 9. As the memory requirements of reduction object are relatively small, full replication gives the best performance. However, the locking versions are within 5 percent. The relative speedups of six versions are 3.94, 3.92, 3.92, 3.93, 3.94, and 3.92, respectively. Thus, after the initial overhead on one thread versions, all versions scale almost linearly. The middleware overhead is up to 20 percent with k-means, which is higher than that from a priori. This is because the main computation phase of k-means involves accessing coordinates of centers, which are part of the reduction object. Therefore, extra point chasing is involved when middleware is used. The manual versions can simply declare an array comprising all centers, and avoid the extra cost.

The results on large SMP machine are presented in Fig. 10. Six versions are run with 1, 2, 4, 8, 12, and 16 threads. Though full replication gives the best performance, locking versions are within 2 percent. The relative speedups in going from 1 thread to 16 threads for the six versions are 15.78, 15.98, 15.99, 16.03, 15.98, and 16.01, respectively. In other words, all versions scale linearly up to 16 threads. The overhead of the interface is significantly lower as compared to the 4 processor machine. We believe this is because the newer UltraSparc III processor performs aggressive out-of-order issues and can hide some latencies.

As mentioned earlier, to demonstrate the portability of our system, we have ported it to Pentium-based SMPs. Also, we wanted to demonstrate that the performance of parallel k-means does scale to relatively higher dimensional data sets. Therefore, we experimented with two 10-dimensional data sets on a 4 processor SMP with 700 MHz Pentium IV processors and 256 MB of main memory. The two data sets are 250 MB and 500 MB, respectively. The results are shown in

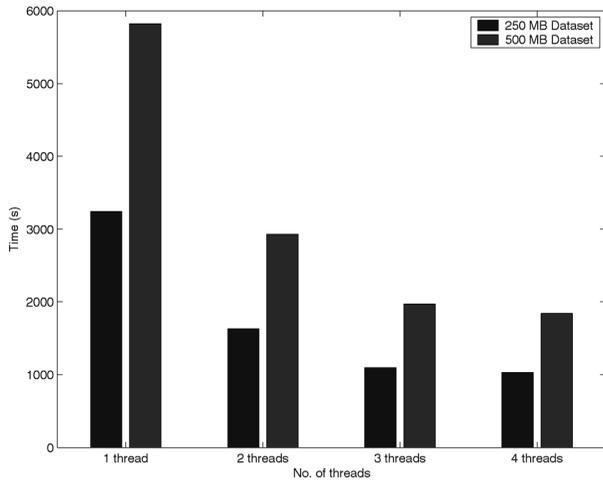


Fig. 11. K-Means on a 10-dimensional data sets and Pentium-based SMP.

Fig. 11 and correspond to the `of1-int` version. The results are very similar to those observed on the small SUN SMP machine. For the 250 MB data set, the speedups on two, three, and four threads are 1.99, 2.95, and 3.15, respectively. For the 500 MB data set, the speedups on two, three, and four threads are 1.98, 2.95, and 3.16, respectively.

5.4 Results from FP-Tree-Based Association Mining

To demonstrate that our techniques and the system can be used for newer algorithms for association mining, we have parallelized the FP-tree-based algorithm for association rule mining [21]. This algorithm requires only two passes over the entire data set. In the first step, an ordered set of frequent items is determined from the set of transactions. In the second step, a tree is grown dynamically. The candidate frequent itemsets and their counts are implicitly stored in the tree.

Because of the nature of this algorithm, we tried only two versions. Both the versions used full replication for the first step, i.e., each thread independently updated counts of the items, which were later merged. In the first version, called full replication, the tree is completely replicated in the second step, i.e., each thread independently grows and updates its own copy of the tree. Merging these trees requires some

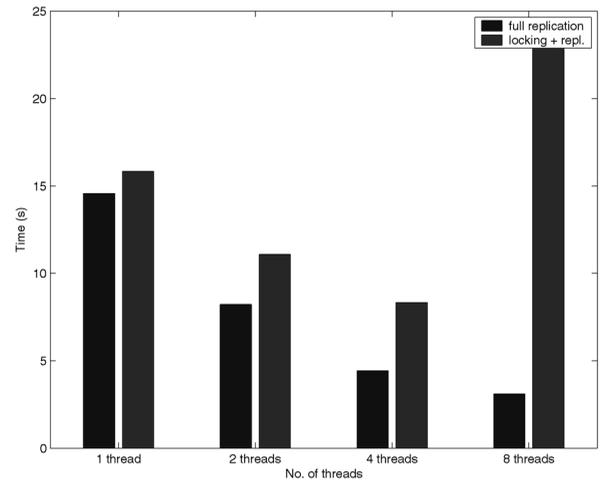


Fig. 13. Parallelization of FP-tree, 200 MB data set.

nontrivial steps, which we do not elaborate on here. The second version combines replication and locking. Only one copy of the tree is maintained, but the counts at each of the nodes are replicated. When the tree needs to be expanded, a thread must acquire a lock. If the tree is expanded frequently, this can potentially serialize some of the computation.

Our experiments evaluated the scalability of these two versions on two data sets, which are 200 MB and 800 MB, respectively. Again, these data sets were generated using the tool available from IBM [3]. We used up to eight threads on the large SUN SMP. The results are presented in Figs. 13 and 14. For both the data sets, use of full replication gives good scalability, whereas the performance of the second version is quite poor. With the 200 MB data set and full replication, the speedups on two, four, and eight threads are 1.77, 3.29, and 4.72, respectively. The second version gives speedups of 1.43, and 1.90 on two and four threads, but the parallel code on eight threads is slower than the sequential code. This is because of the lock contention, which becomes quite severe with a larger number of threads. The trends are very similar for the larger data set.

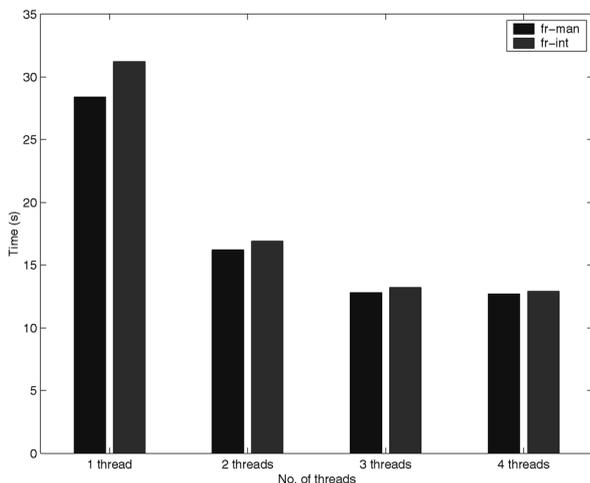


Fig. 12. Parallelization of k-nearest neighbors.

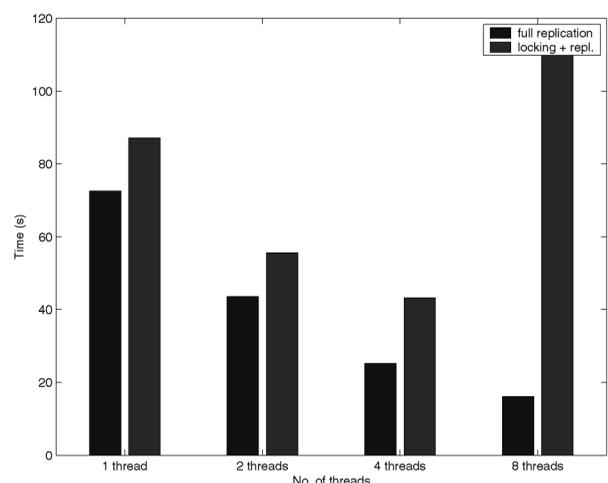


Fig. 14. Parallelization of FP-tree, 800 MB data set.

5.5 Results from K-Nearest Neighbors

The last data mining algorithm we consider is k-nearest neighbors. We have experimented with a 800 MB data set. Again, the set of points were chosen randomly, as it does not impact the parallel speedups. The reduction object in this algorithm's parallel implementation is the list of k-nearest neighbors. This is considered a single element. Because of this granularity, only a full replication scheme was implemented for this algorithm. With small values of k , this algorithm becomes memory or I/O bound. Since our goal was to show that our library can allow us to get speedups when the parallelism is inherent in the algorithm, the value of k used in our experiments was 2,000. Our previous work has shown that good distributed memory parallelism can be achieved for small values of k [27].

Fig. 12 presents experimental results from *fr-man* and *fr-int* versions. The speedups of manual version are 1.75, 2.22, and 2.24 with two, three, and four threads, respectively. The speedups are limited because only a small amount of computation is associated with each transaction. The overhead of the use of the interface is within 10 percent in all cases. Because of the limited computation in this code, we did not experiment further with the large SMP machine.

6 A DETAILED CASE STUDY: DECISION TREE CONSTRUCTION

We now describe our experience in parallelizing decision tree construction using the parallelization techniques and middleware described earlier in the paper.

Decision tree construction is a very well-studied problem in data mining, machine learning, and statistics communities [14], [13], [39], [41], [45], [46], [47]. The input to a decision tree construction algorithm is a database of *training records*. Each record has several *attributes*. An attribute whose underlying domain is totally ordered is called an *ordered*, *numerical*, or *continuous* attribute. Other attributes are called *categorical* attributes. One particular attribute is called *class label*, and typically can hold only two values, true and false. All other attributes are referred to as the *predictor* attributes. A decision tree construction algorithm processes the set of training records, and builds a model which is used for predicting the class label of new records. A number of algorithms for decision tree construction have been proposed. In recent years, particular attention has been given to developing algorithms that can process data sets that do not fit in main memory [14], [29], [50].

A lot of effort has been put into developing parallel algorithms for decision tree construction [4], [15], [29], [37], [50], [53], [55]. Most of these efforts have targeted distributed memory parallel machines. To the best of our knowledge, there is only one effort on shared memory parallelization of decision tree construction on disk-resident data sets, which is by Zaki et al. [55].

Usually, the parallelization approach taken for decision tree construction is quite different than the approach taken for other common data mining algorithms, like association mining and clustering. Parallelization of decision tree construction typically involves sorting of numerical attributes and/or frequently writing back of input data. There-

fore, an important question is, "Can the parallelization techniques and runtime support that are suitable for association mining and clustering also be effective for decision tree construction?" Here, we demonstrate that this is indeed the case. We particularly focus on parallelizing the *RainForest* framework for scalable decision tree construction [14]. We believe that our effort is the first on parallelizing RainForest-based decision tree construction.

6.1 Decision Tree Construction Using RainForest Framework

RainForest is a general framework for scaling decision tree construction [14]. The key observation that motivates this approach is as follows. Though a large number of decision tree construction approaches have been used in the past, they are common in an important way. The decision tree is constructed in a top-down, recursive fashion. Initially, all training records are associated with the root of the tree. A criteria for splitting the root is chosen, and two or more children of this node are created. The training records are partitioned (physically or logically) between these children. This procedure is recursively applied until either all training records associated with a node have the same class label, or the number of training records associated with a node is below a certain threshold. The different approaches for decision tree construction differ in the way criteria for splitting a node is selected, and the data-structures required for supporting the partitioning of the training sets.

The RainForest approach scales the decision tree construction process to larger (disk-resident) data sets, while also effectively exploiting the available main memory. This is done by isolating an AVC (Attribute-Value, Classlabel) set for a given attribute and a node being processed. The size of the AVC-set for a given node and attribute is proportional to the number of distinct values of the attribute and the number of distinct class labels. For example, in a SPRINT-like approach, AVC-set for a categorical attribute will simply be the count of occurrence of each distinct value the attribute can take. Therefore, the AVC-set can be constructed by taking one pass through the training records associated with the node.

Given a node of the decision tree, AVC-group is the combination of AVC-set for all attributes. The key observation is that, though AVC-group does not contain sufficient information to reconstruct the training data set, it contains all information that is required for selecting the criteria for splitting the node. Since the number of attributes and the distinct values they can take is usually not very large, one can expect the AVC-group for a node to easily fit in main memory. With this observation, processing for selecting the splitting criteria for the root node can be easily performed even if the data set is disk-resident. By reading the training data set once, AVC-group of the root is constructed. Then, the criteria for splitting the node is selected.

A number of algorithms have been proposed within the RainForest framework to split decision tree nodes at lower levels. In the algorithm RF-read, the data set is never partitioned. The algorithm progresses level by level. In the first step, AVC-group for the root node is built and a splitting criteria is selected. At any of the lower levels, all nodes at that level are processed in a single pass if the AVC-

```

RF-Read(dataset  $D$ )
global Tree  $root$ , Queue  $Q$ ,  $AQ$ ;
local Node  $node$ ;
 $Q \leftarrow NULL$ ;  $AQ \leftarrow NULL$ ;
 $add(root, AQ)$ ;
while not ( $empty(Q)$  and  $empty(AQ)$ )
  { Loop1: build AVC-group for nodes in  $AQ$  }
  foreach (chunk  $C \in D$ )
    foreach (tuple  $t \in C$ )
       $node \leftarrow classify(root, t)$ ;
      if  $node \in AQ$ 
        foreach (attribute  $a \in t$ )
           $reduction(node.avc\_group, a, t.class)$ ;
  { Loop2: split the nodes in  $AQ$  }
  foreach ( $node \in AQ$ )
    if not  $satisfy\_stop\_condition(node)$ 
       $find\_best\_split(node)$ ;
      foreach (Node  $child \in create\_children(node)$ )
         $add(child, Q)$ ;
  { Loop3: build new  $AQ$  }
   $AQ \leftarrow NULL$ ;
   $done \leftarrow false$ ;
  while not  $empty(Q)$  and not  $done$ 
     $get(node, Q)$ ;
    if  $enough\_memory(Q)$ 
       $remove(node, Q)$ ;
       $add(node, AQ)$ ;
    else  $done \leftarrow true$ ;

```

Fig. 15. Algorithm for parallelizing RF-read using our middleware.

group for all the nodes fit in main memory. If not, multiple passes over the input data set are made to split nodes at the same level of the tree. Because the training data set is not partitioned, this can mean reading each record multiple times for one level of the tree.

Another algorithm, RF-write, partitions and rewrites the data set after each pass. The algorithm RF-hybrid combines the previous two algorithms. Overall, RF-read and RF-hybrid algorithms are able to exploit the available main memory to speedup computations, but without requiring the data set to be main memory resident.

In our work, we will mainly focus on parallelizing the RF-read algorithm. There are several reasons for this. First, the only difference between RF-read, RF-write, and RF-hybrid is the frequency of writing back. RF-read could be looked as the main computing subroutine to be called for every physical partition, and includes the dominant processing for them, i.e., building AVC-groups and finding the criteria for splitting. Second, the main memory size of SMP machines has been growing at a rapid speed in recent years. This makes RF-read more practical even when the decision tree constructed is quite deep. Finally, because RF-read does not require the data to be written back, it fits into the structure of algorithms for association mining and clustering, and suits our middleware.

The next two sections describe our existing middleware and the parallelization techniques supported and our approach for parallelizing RF-read.

6.2 Parallel RainForest Algorithm and Implementation

In this section, we will present the algorithm and implementation details for parallelizing RF-read using our middleware. The algorithm is presented in Fig. 15.

The algorithm takes several passes over the input data set D . The data set is organized as a set of chunks. During every pass, there are a number of nodes that are *active* or belong to the set AQ . These are the nodes for which AVC-group is built and splitting criteria is selected.

This processing is performed over three consecutive loops. In the first loop, the chunks in the data set are read. For each training record or *tuple* in each chunk that is read, we determine the *node* at the current level to which it belongs. Then, we check if the node belongs to the set AQ . If so, we increment the elements in the AVC-group of the node.

The second loop finds the best splitting criteria for each of the active nodes and creates the children. Before that, however, it must check if a *stop condition* holds for this node and, therefore, it need not be partitioned. For the nodes that are partitioned, no physical rewriting of data needs to be done. Instead, just the tree should be updated, so that future invocations to *classify* point to the appropriate children. The nodes that have been split are removed from the set AQ and the newly created children are added to the set Q .

At the end of the second loop, the set AQ is empty and the set Q contains the nodes that still need to be processed.

The third loop determines the set of the nodes that will be processed in the next phase. We iterate over the nodes in the set Q , remove a node from Q and move it to AQ . This is done until either no more memory is available for AVC-groups, or Q is empty.

The last loop contains only a very small part of the overall computing. Therefore, we focus on parallelizing the first and the second loop. Parallelization of the second loop is straightforward and discussed first.

A simple multithreaded implementation is used for the second loop. There is one thread per processor. This thread gets a node from the set AQ and processes the corresponding AVC-group to find the best splitting criteria. The computing done for each node is completely independent. The only synchronization required is for getting a node from AQ to process. This is implemented by simple locking.

Parallelization of the first loop is facilitated by the producer consumer framework we support. The producer thread reads the chunks and puts them in a queue. Consumer threads grab chunks from the queue and perform the computing associated with these. The main problem is ensuring correctness as consumer threads process training records from different chunks. Note that the first loop fits nicely with the structure of the canonical loop we had shown in Fig. 1. The set of AVC-groups for all nodes that are currently active is the reduction object. As different consumer threads try to update the same element in an AVC-set, race conditions can arise. The elements of the reduction object that are updated after processing a tuple cannot be determined without processing the tuple.

Therefore, the parallelization techniques we have developed are applicable to parallelizing the first loop. Both memory overheads and locking costs are important considerations in selecting the parallelization strategy. At lower levels of the tree, the total size of the reduction object can be very large. Therefore, memory overhead of the parallelization technique used is an important consideration. Also, the updates to the elements of the reduction object are fine-grained. After getting a lock associated with an element or a set of elements, the only computing performed is incrementing one value. Therefore, locking overheads can also be significant.

Next, we discuss the application of the techniques we have developed to parallelization of the first loop. Recall that the memory requirements of the three techniques are very different. If R is the size of reduction object, N is the size of consumer threads, and L is the number of elements per cache line, the memory requirement of full replication, optimized full locking and cache sensitive locking are $N \times R$, $2 \times R$, and $\frac{N}{N-1} \times R$, respectively. This has an important implication for our parallel algorithm. Choosing a technique with larger memory requirements means that the set AQ will be smaller. In other words, a larger number of passes over the data set will be required.

An important property of the reduction object in RF-read is that updates to each AVC-set are independent. Therefore, we can apply different parallelization techniques to nodes at different levels, and for different attributes. Based upon this observation, we developed a number of approaches for applying one or more of the parallelization techniques we have. These approaches are *pure*, *horizontal*, *vertical*, and *mixed*.

In the pure approach, the same parallelization approach is used for all AVC-sets, i.e., for nodes at different levels and for both categorical and numerical attributes.

The vertical approach is motivated by the fact that the sum of sizes of AVC-groups for all nodes at a level is quite small at upper levels of the tree. Therefore, full replication can be used for these levels without incurring the overhead of additional passes. Moreover, because the total number of elements in the reduction object is quite small at these levels, locking schemes can result in high overhead of waiting for locks and coherence cache misses. Therefore, in the vertical approach, replication is used for the first few levels (typically, between 3 to 5) in the tree, and either optimized full locking or cache-sensitive locking is used at lower levels.

In determining the memory overheads, the cost of waiting for locks and coherence cache misses, one important consideration is the number of distinct values of an attribute. If the number of the distinct values of an attribute is small, the corresponding AVC-set is small. Therefore, the memory overhead in replicating such AVC-sets may not be a significant consideration. At the same time, because the number of elements is small, the cost of waiting for locks and coherence cache misses can be significant. Note that, typically, categorical attributes have a small number of distinct values and numerical attributes can have a large number of distinct values in a training set.

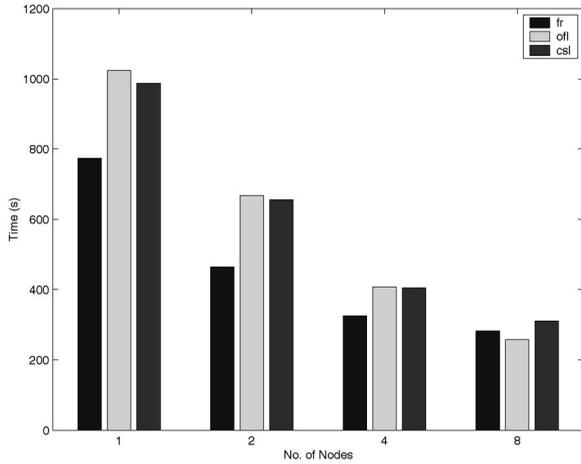
Therefore, in the horizontal approach, full replication is used for attributes with a small number of distinct values, and one of the locking schemes is used for attributed with a large number of distinct values. For any attribute, the same technique is used at all levels of the tree.

Finally, the mixed strategy combines the two approaches. Here, full replication is used for all attributes at the first few levels, and for attributes with small number of distinct values at the lower levels. One of the locking schemes is used for the attributes with a large number of distinct values at lower levels of the tree.

6.3 Experimental Results

In this section, we evaluate our implementation of decision tree construction. Since our primary consideration was evaluating scalability, we only evaluated the performance on SunFire 6800.

We used two data sets for our experiments, generated using a tool described by Agrawal et al. [1]. Both the data sets were nearly 1.3 GB, with 32 million records in the training set. Each record has nine attributes, of which three are categorical and the other 6 are numerical. Every record belongs to one of two classes. Agrawal et al. use a series of classification functions of increasing complexity to classify records into different groups. Tuples in the training set are assigned the group label (classes) by first generating the tuple and then applying the classification function to the tuple to determine the group to which the tuple belongs. The two data sets we used correspond to the use of functions 1 and 7, respectively. The use of function 1 generates a relatively small decision tree, whereas the tree generated by function 7 is large. In our experiments, the *stop point* for the node size is 10,000, i.e., if the subtree includes fewer than 10,000 tuples, we do not expand it any further. The use of function 1 results in a tree with 10 levels,

Fig. 16. Performance of *pure* versions, data set 1.

whereas the use of function 7 generates a tree with 16 levels. The data sets corresponding to the use of functions 1 and 7 are referred to as data set 1 and data set 2, respectively.

In Section 6.2, we had described *pure*, *vertical*, *horizontal*, and *mixed* approaches for using one or more of the parallelization techniques we support in the middleware. Based upon these, a total of nine different versions of our parallel implementation were created. Obviously, there are three pure versions, corresponding to the use of full replication (*fr*), optimized full locking (*ofl*), and cache sensitive locking (*cs1*). Optimized full locking can be combined with full replication using vertical, horizontal, and mixed approach, resulting in three versions. Similarly, cache sensitive locking can be combined with full replication using vertical, horizontal, and mixed approach, resulting in three additional versions, for a total of nine versions.

Figs. 16 and 17 show the performance of pure versions on two data sets. With one thread, *fr* gives the best performance for both the data sets. This is because there is no locking overhead. The relative speedups on eight threads for *fr* are only 2.73 and 1.86 for the first and the second data set, respectively. In fact, with the second data set, there is a slowdown observed in going from four to eight threads. This is because the the use of full replication for AVC-sets at all levels

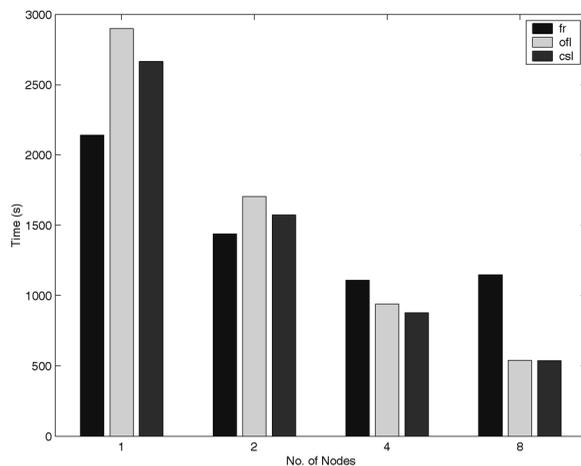
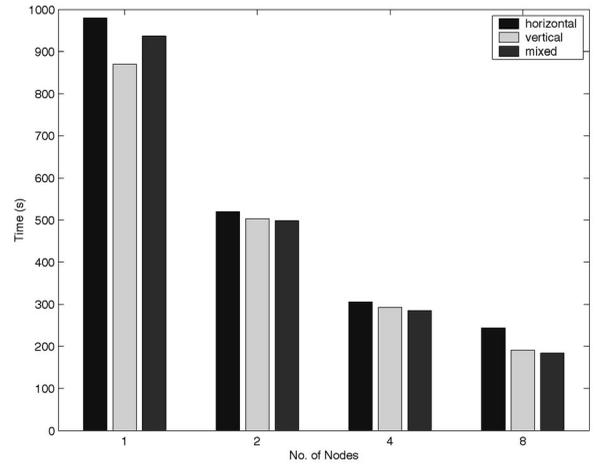
Fig. 17. Performance of *pure* versions, data set 2.

Fig. 18. Combining full replication and full locking, data set 1.

results in very high memory requirements. The second data set produces a deeper tree, so the sum of the sizes of AVC-set for all nodes at a level can be even higher.

Locking schemes result in a 20 to 30 percent overhead on one thread, but the relative speedups are better. Using eight threads, the relative speedups for *ofl* and *cs1* are 5.37 and 4.95, respectively, for the second data set. However, for data set 1, the relative speedups for both the versions are less than four. There are two types of overhead in the use of locking schemes. The first is the additional cost of locking while incrementing every value in the AVC-set. Second, at the upper levels of the tree, the total number of elements associated with AVC-sets of all nodes at the level is quite small. This results in waiting for locks and coherence cache misses when different threads want to update these elements. With data set 1, the tree generated is not very deep, and therefore, a larger fraction of the time is spent doing the processing for the upper levels.

Figs. 18 and 19 present experimental results from combining *fr* and *ofl*, from data sets 1 and 2, respectively. As stated earlier, the two schemes can be combined in three different ways, horizontal, vertical, and mixed. The performance of these three versions is quite similar. With data set 1, *horizontal* is consistently the slowest, and *mixed* gives the best performance on two, four, and eight threads. With

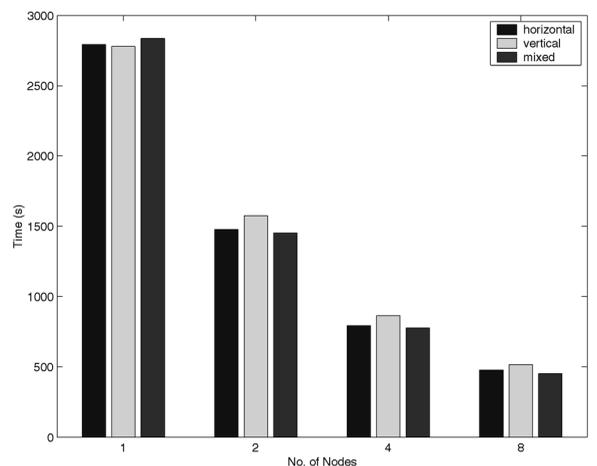


Fig. 19. Combining full replication and full locking, data set 2.

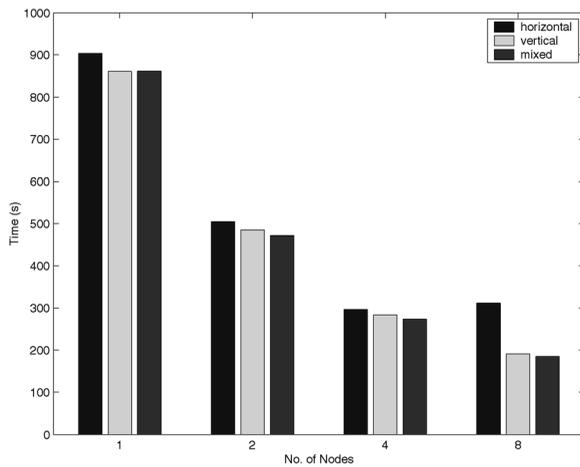


Fig. 20. Combining full replication and cache-sensitive data set 1.

data set 2, *vertical* is the slowest on two, four, and eight nodes, whereas *mixed* is the best on two, four, and eight nodes.

In the *horizontal* approach, the use of locking for continuous attributes at upper levels of the tree can slow down the computation because of waiting for locks and coherence cache misses. In contrast, in the *vertical* approach, the use of locking for categorical attributes at the lower levels results in waiting time for locks and coherence cache misses. As a result, the *mixed* approach results in the best performance on two, four, or eight threads, for either data set.

Figs. 20 and 21 present experimental results from combining *fr* and *cs1*, from data sets 1 and 2, respectively. Again, the *mixed* version is the best among the three versions, for two, four, and eight threads, for either data set.

It is interesting to compare the relative performance between combining optimized full locking with full replication and combining cache sensitive locking with full replication. As compared to optimized full locking, cache sensitive locking has lower memory requirements. However, the overhead of waiting for locks could also be higher in cache sensitive locking.

For the first data set, cache sensitive locking results in worse performance with *horizontal* approach and almost identical performance with *vertical* and *mixed* approaches. This is because the first data set results in a tree with only 10 levels, and memory overhead is not a significant issue. With the second data set, cache sensitive locking results in significantly better performance with all three approaches.

Combining cache sensitive locking and full replication using the *mixed* strategy results in the best performance. With the data set 2, the relative speedup of this version on eight threads is 5.9. Comparing this version against the best 1 thread version (which is *fr*), the speedup is 5.2.

6.4 Discussion

This case study has led to a number of interesting observations. First, we have shown that a RainForest-based decision tree construction algorithm can be parallelized in a way which is very similar to the way association mining and clustering algorithms have been parallelized. Therefore, a general middleware framework for decision tree construction can simplify the parallelization of algorithms for a variety of mining tasks. Second, unlike the algorithms for other mining tasks, a combination of locking and replication-based

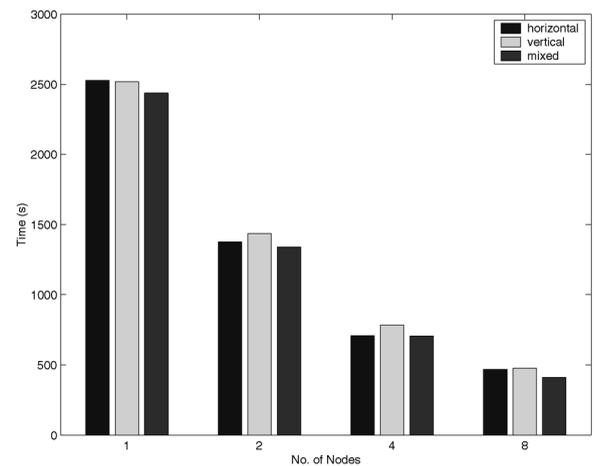


Fig. 21. Combining full replication and cache-sensitive locking data set 2.

techniques results in the best speedups for decision tree construction. Thus, it is important that the framework used supports a variety of parallelization techniques.

The best relative speedup we obtained was 5.9 using eight threads. This compares well with the speedups that have been obtained by the researchers developing stand-alone shared memory or distributed memory decision tree implementations. Thus, our work also shows that a general framework for parallel data mining implementations can achieve high performance while significantly simplifying the programmer's task.

7 RELATED WORK

We now compare our work with related research efforts.

A significant amount of work has been done on parallelization of individual data mining techniques. Most of the work has been on distributed memory machines, including association mining [2], [19], [20], [57], k-means clustering [10], and decision tree classifiers [4], [15], [29], [50], [53]. Recent efforts have also focused on shared memory parallelization of data mining algorithms, including association mining [56], [43], [44] and decision tree construction [55]. Our work is significantly different, because we offer an interface and runtime support to parallelize a number of data mining algorithms. Our shared memory parallelization techniques are also different, because we focus on a common framework for parallelization of a number of algorithms.

Since we have used *a priori* as an example in our implementation, we do a detailed comparison of our approach with the most recent work on parallelizing *a priori* on a shared memory machine [44]. One limitation of our approach is that we do not parallelize the candidate generation part of the algorithm in our framework. We have at least two advantages, however. First, we dynamically assign transactions to threads, whereas their parallel algorithm works on a static partition of the data set. Second, our work on memory layout of locks and reduction elements also goes beyond their techniques. There are also many similarities in the two approaches. Both approaches segregate read-only data to reduce false sharing and consider replicating the reduction elements.

Decision tree construction for disk-resident data sets and on parallel machines has also been studied by many

researchers. We describe these efforts and compare them with our approach.

One of the first decision tree construction methods for disk-resident data sets was SLIQ [37]. SLIQ requires sorting of ordered attributes and separation of the input data set into attribute lists. In addition, it requires a data-structure called *class list*, whose size is proportional to the number of records in the data set. Class list is accessed frequently and randomly, therefore, it must be kept in main memory. In parallelizing SLIQ, either the class list must be replicated, or a very high communication overhead is imposed.

A somewhat related approach is SPRINT [50]. SPRINT does not require class lists, but instead requires a hash table which is proportional to the number of records associated with a node in the decision tree. Like SLIQ, SPRINT requires attributes lists and sorting of ordered attributes. Moreover, SPRINT requires partitioning of attribute lists whenever a node in the decision tree is split. Thus, it can have a significant overhead of rewriting a disk-resident data set.

The algorithm by Srivastava et al. [53] has many similarities with the RF-write algorithm. An important difference, however, is that it reduces the number of potential splitting points of ordered attributes by clustering their values. A somewhat similar idea is used in the CLOUDS method [4].

The only previous work on shared memory parallelization of decision tree construction on disk-resident data sets is by Zaki et al. [55]. They have carried out a shared memory parallelization of SPRINT algorithm. Our work is distinct in parallelizing a very different method for decision tree construction. In parallelizing SPRINT, each attribute list is assigned to a separate processor. In comparison, we parallelize updates to reduction objects with fine sharing between the processors. Narlikar has used a fine-grained threaded library for parallelizing a decision tree algorithm [40], but the work is limited to memory-resident data sets.

Becuzzi et al. [5] have used a structured parallel programming environment PQE2000/SkIE for developing parallel implementation of data mining algorithms. However, they only focus on distributed memory parallelization, and I/O is handled explicitly by the programmers. The similarity among parallel versions of several data mining techniques has also been observed by Skillicorn [52]. Our work is different in offering a middleware to exploit the similarity and ease parallel application development. Goil and Choudhary have developed PARSIMONY, which is an infrastructure for analysis of multi-dimensional data sets, including OLAP and data mining [16]. PARSIMONY does not focus on shared memory parallelization.

Note that locking has been extensively used in databases. There are some similarities between the mining algorithms and aggregation queries used in relational databases. Shatdal has studied SMP parallelization of aggregation queries [51]. However, they have not focused on memory hierarchy impact of the use of locks (or latches), which, in our experience, has become a significant issue for modern SMP machines. Some of the ideas in our parallelization techniques have been independently developed in the computer architecture field. Kagi et al. have used the idea of *collocation*, in which a lock and the reduction object are placed in the same cache block [31], [30]. Our focus has been on cases where a large number of small reduction elements exist and false sharing can be a significant problem. In addition, we have presented an interface for data mining algorithms, and evaluated different techniques specifically in the context of data mining algorithms.

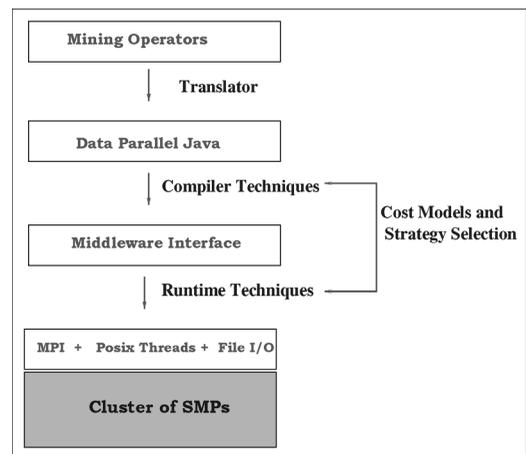


Fig. 22. Long term vision in providing high-level interfaces for scalable data mining.

OpenMP is the general accepted standard for shared memory programming. OpenMP currently only supports scalar reduction variables and a small number of simple reduction operations, which makes it unsuitable for data mining algorithms we focus on. Inspector/executor approach for shared memory parallelization is based upon using an inspector that can examine certain values at runtime to determine an assignment of iterations to processors [6], [18], [34], [35], [17], [48], [54]. For data mining algorithms, the inspector will have to perform almost all the computation associated with local reductions, and will not be practical. Cilk, a language for shared memory programming developed at MIT [7], also does not target applications with reductions.

8 AN OVERALL VISION

In this section, we describe our overall vision toward offering high-level support for creating scalable data mining implementations, and how the techniques and the system described in this paper can be used in longer term.

As we stated earlier, the work reported in this paper is part of a middleware system called FREERIDE (FRamework for Rapid Implementation of Datamining Engines). The target environment of this system includes both clusters of SMPs and large scale SMP machines. The middleware performs both distributed memory and shared memory parallelization. Moreover, it enables high I/O performance by minimizing disk seek time and using asynchronous I/O operations. Thus, it can be used for developing efficient parallel data mining applications that operate on disk-resident data sets. The middleware support for distributed memory parallelization and scaling to disk-resident data sets was reported in an earlier publication [27].

Moreover, we envision the current FREERIDE system as only an intermediate step in providing high-level support for specifying or developing scalable mining implementations. Our longer term vision is shown in Fig. 22. Here, we briefly describe our current work or ideas towards the other components shown in this figure.

Data Parallel Language Interface: Our recent work [32], [33] has shown that data parallel languages are well-suited for specifying mining algorithms. Specifically, we have used a dialect of Java for expressing data mining algorithms at a higher level than our middleware interface. Our compiler work has leveraged our earlier work on compiling scientific data-intensive applications [11]. Our initial work has shown

that existing compilation techniques can be used for translating from a data-parallel language to the FREERIDE interface, with at most 10 percent performance penalty.

Cost Models and Strategy Selection: A number of different parallelization strategies are usually possible for any data mining algorithm. Our middleware and compilation system can be used for making a selection between different parallelization strategies. We have initiated work on modeling performance factors like locking overhead, waiting time for locks, cold and capacity cache misses, coherence cache misses, and false sharing. Our initial work in this area was presented in the ACM SIGMETRICS Conference in June 2002 [28].

Declarative Interface: Declarative interfaces have always been popular with the users of database systems. Ideally, the end users will like to specify data mining tasks they want to perform at a very high-level (e.g., in a declarative or nonprocedural fashion), and yet achieve scalable execution. Commercial products like IBM DB2 [24] and Oracle [36] are very successful in parallelizing relational queries. Like the relational database operations, commonly used data mining algorithms are well-suited for parallel execution, but no tools and techniques are currently available for supporting their parallel execution. A significant effort, however, has been put on developing a set of mining operators [22]. We believe that by having a library of Java implementations of common data mining algorithms, we can easily create data parallel Java versions from a suitable declarative interface.

To summarize, our initial work on providing high-level language support and performance modeling has been very encouraging. Our work in both these areas builds on the fact that a common set of techniques and a common interface can be used for parallelizing a number of data mining algorithms, which were presented in this paper.

9 CONCLUSION

In this paper, we have focused on shared memory parallelization of data mining algorithms. By exploiting the similarity in the parallel algorithms for several data mining tasks, we have been able to develop a programming interface and a set of techniques. Our techniques offer a number of trade offs between memory requirements, parallelism, and locking overhead. In using our programming interface, the programmers only need to make simple modifications to a sequential code and use our reduction object interface for updates to elements of a reduction object.

We have reported experimental results from implementations of a priori and FP-tree-based association mining, k-means clustering, k-nearest neighbors, and decision tree construction. These experiments establish the following:

1. Among full replication, optimized full locking, and cache-sensitive locking, there is no clear winner. Each of these three techniques can outperform others depending upon machine and data set parameters. These three techniques perform significantly better than the other two techniques.
2. Good parallel efficiency is achieved for each of the four algorithms we experimented with, using our techniques and runtime system.
3. The overhead of the interface is within 10 percent in almost all cases.
4. In the case of decision tree construction, combining different techniques turned out to be crucial for achieving high performance.

ACKNOWLEDGMENTS

This research was supported by the US National Science Foundation CAREER award ACI-9733520, US National Science Foundation grant CCR-9808522, and US National Science Foundation grant ACR-9982087. The equipment for this research was purchased under US National Science Foundation grant EIA-9703088.

REFERENCES

- [1] R. Agrawal, T. Imielinski, and A. Swami, "Database Mining: A Performance Perspective," *IEEE Trans. Knowledge and Data Eng.*, vol. 5, no. 6, pp. 914-925, Dec. 1993.
- [2] R. Agrawal and J. Shafer, "Parallel Mining of Association Rules," *IEEE Trans. Knowledge and Data Eng.*, vol. 8, no. 6, pp. 962-969, June 1996.
- [3] R. Agrawal and R. Srikant, "Fast Algorithms for Mining Association Rules," *Proc. 1994 Int'l Conf. Very Large Databases (VLDB '94)*, pp. 487-499, Sept. 1994.
- [4] K. Alsabti, S. Ranka, and V. Singh, "Clouds: Classification for Large or Out-of-Core Datasets," <http://www.cise.ufl.edu/ranka/dm.html>, 1998.
- [5] P. Becuzzi, M. Coppola, and M. Vanneschi, "Mining of Association Rules in Very Large Databases: A Structured Parallel Approach," *Proc. EuroPar-99*, vol. 1685, pp. 1441-1450, Aug. 1999.
- [6] W. Blume, R. Doallo, R. Eigenman, J. Grout, J. Hoelflinger, T. Lawrence, J. Lee, D. Padua, Y. Paek, B. Pottenger, L. Rauchwerger, and P. Tu, "Parallel Programming with Polaris," *Computer*, vol. 29, no. 12, pp. 78-82, Dec. 1996.
- [7] R.D. Blumofe, C.F. Joerg, et al., "Cilk: An Efficient Multithreaded Runtime System," *Proc. Fifth ACM Conf. Principles and Practices of Parallel Programming (PPoPP)*, 1995.
- [8] S. Brin, R. Motwani, J. Ullman, and S. Tsur, "Dynamic Itemset Counting and Implication Rules for Market Basket Data," *Proc. ACM SIGMOD Conf. Management of Data*, May 1997.
- [9] P. Cheeseman and J. Stutz, "Bayesian Classification (Autoclass): Theory and Practice," *Advanced in Knowledge Discovery and Data Mining*, pp. 61-83, 1996.
- [10] I.S. Dhillon and D.S. Modha, "A Data-Clustering Algorithm on Distributed Memory Multiprocessors," *Proc. Workshop Large-Scale Parallel KDD Systems, in conjunction with the Fifth ACM SIGKDD Int'l Conf. Knowledge Discovery and Data Mining (KDD '99)*, pp. 47-56, Aug. 1999.
- [11] R. Ferreira, G. Agrawal, and J. Saltz, "Compiling Object-Oriented Data Intensive Computations," *Proc. 2000 Int'l Conf. Supercomputing*, May 2000.
- [12] G. Forman and B. Zhang, "Distributed Data Clustering Can be Efficient and Exact," *Proc. SIGKDD Explorations*, vol. 2, no. 2, Dec. 2000.
- [13] J. Gehrke, V. Ganti, R. Ramakrishnan, and W. Loh, "Boat—Optimistic Decision Tree Construction," *Proc. ACM SIGMOD Conf. Management of Data*, June 1999.
- [14] J. Gehrke, R. Ramakrishnan, and V. Ganti, "Rainforest—A Framework for Fast Decision Tree Construction of Large Datasets," *Proc. Conf. Very Large Databases (VLDB)*, 1998.
- [15] S. Goil and A. Choudhary, "Efficient Parallel Classification Using Dimensional Aggregates," *Proc. Workshop Large-Scale Parallel KDD Systems, with ACM SIGKDD-99*, Aug. 1999.
- [16] S. Goil and A. Choudhary, "PARSIMONY: An Infrastructure for Parallel Multidimensional Analysis and Data Mining," *J. Parallel and Distributed Computing*, vol. 61, no. 3, pp. 285-321, Mar. 2001.
- [17] E. Gutierrez, O. Plata, and E.L. Zapata, "A Compiler Method for the Parallel Execution of Irregular Reductions in Scalable Shared Memory Multiprocessors," *Int'l Conf. Supercomputing (ICS '00)*, pp. 78-87, May 2000.
- [18] M. Hall, S. Amarsinghe, B. Murphy, S. Liao, and M. Lam, "Maximizing Multiprocessor Performance with the SUIF Compiler," *Computer*, vol. 12, Dec. 1996.
- [19] E.-H. Han, G. Karypis, and V. Kumar, "Scalable Parallel Datamining for Association Rules," *Proc. ACM SIGMOD 1997*, May 1997.
- [20] E.-H. Han, G. Karypis, and V. Kumar, "Scalable Parallel Datamining for Association Rules," *IEEE Trans. Data and Knowledge Eng.*, vol. 12, no. 3, May/June 2000.

- [21] J. Han, J. Pei, and Y. Yin, "Mining Frequent Patterns without Candidate Generation," *Proc. ACM SIGMOD Conf. Management of Data*, 2000.
- [22] J. Han and M. Kamber, *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers, 2000.
- [23] J.L. Hennessy and D.A. Patterson, *Computer Architecture: A Quantitative Approach*, 2nd ed. Morgan Kaufmann, Inc., 1996.
- [24] IBM. Db2, "Universal Database Goes Parallel with Enterprise and Enterprise-Extended Editions," <http://www-4.ibm.com/software/data/db2/udb/98eeebrochure>, 1999.
- [25] A.K. Jain and R.C. Dubes, *Algorithms for Clustering Data*. Prentice Hall, 1988.
- [26] R. Jin and G. Agrawal, "An Efficient Implementation of Apriori Association Mining on Cluster of SMPs," *Proc. Workshop High Performance Data Mining, (IPDPS 2001)*, Apr. 2001.
- [27] R. Jin and G. Agrawal, "A Middleware for Developing Parallel Data Mining Implementations," *Proc. First SIAM Conf. Data Mining*, Apr. 2001.
- [28] R. Jin and G. Agrawal, "Performance Prediction for Random Write Reductions: A Case Study in Modeling Shared Memory Programs," *Proc. ACM SIGMETRICS*, June 2002.
- [29] M.V. Joshi, G. Karypis, and V. Kumar, "Scalparc: A New Scalable and Efficient Parallel Classification Algorithm for Mining Large Datasets," *Proc. Int'l Parallel Processing Symp.*, 1998.
- [30] A. Kagi, "Mechanisms for Efficient Shared-Memory, Lock-Based Synchronization," PhD thesis, Univ. of Wisconsin, Madison, 1999.
- [31] A. Kagi, D. Burger, and J.R. Goodman, "Efficient Synchronization: Let Them Eat QOLB," *Proc. 24th Ann. Int'l Symp. Computer Architecture*, pp. 170-180, June 1997.
- [32] X. Li, R. Jin, and G. Agrawal, "A Compilation Framework for Distributed Memory Parallelization of Data Mining Algorithms," submitted for publication, 2002.
- [33] X. Li, R. Jin, and G. Agrawal, "Compiler and Runtime Support for Shared Memory Parallelization of Data Mining Algorithms," *Proc. Conf. Language and Compilers for Parallel Computing*, Aug. 2002.
- [34] Y. Lin and D. Padua, "On the Automatic Parallelization of Sparse and Irregular Fortran Programs," *Proc. Workshop Languages, Compilers, and Runtime Systems for Scalable Computers (LCR-98)*, May 1998.
- [35] H. Lu, A.L. Cox, S. Dwarkadas, R. Rajamony, and W. Zwaenepoel, "Compiler and Software Distributed Shared Memory Support for Irregular Applications," *Proc. Sixth ACM SIGPLAN Symp. Principles and Practice of Parallel Programming (PPOPP)*, pp. 48-56, June 1997.
- [36] T. Mahapatra and S. Mishra, *Oracle Parallel Processing*. O'Reilly Publishers, 2000.
- [37] M. Mehta, R. Agrawal, and J. Rissanen, "SLIQ: A Fast Scalable Classifier for Data Mining," *Proc. Fifth Int'l Conf. Extending Database Technology*, 1996.
- [38] A. Mueller, "Fast Sequential and Parallel Algorithms for Association Rule Mining: A Comparison," Technical Report CS-TR-3515, Univ. of Maryland, College Park, Aug. 1995.
- [39] S.K. Murthy, "Automatic Construction of Decision Trees from Data: A Multi-Disciplinary Survey," *Data Mining and Knowledge Discovery*, vol. 2, no. 4, pp. 345-389, 1998.
- [40] G.J. Narlikar, "A Parallel, Multithreaded Decision Tree Builder," Technical Report CMU-CS-98-184, School of Computer Science, Carnegie Mellon Univ., 1998.
- [41] C.R. Palmer and C. Faloutsos, "Density Biases Sampling: An Improved Method for Data Mining and Clustering," *Proc. 2000 ACM SIGMOD Int'l Conf. Management of Data*, June 2000.
- [42] J.S. Park, M. Chen, and P.S. Yu, "An Effective Hash Based Algorithm for Mining Association Rules," *Proc. ACM SIGMOD Int'l Conf. Management of Data*, May 1995.
- [43] S. Parthasarathy, M. Zaki, and W. Li, "Memory Placement Techniques for Parallel Association Mining," *Proc. Fourth Int'l Conf. Knowledge Discovery and Data Mining (KDD)*, Aug. 1998.
- [44] S. Parthasarathy, M. Zaki, M. Ogihara, and W. Li, "Parallel Data Mining for Association Rules on Shared-Memory Systems," *Knowledge and Information Systems*, to appear, 2000.
- [45] F. Provost and V. Kolluri, "A Survey of Methods for Scaling up Inductive Algorithms," *Knowledge Discovery and Data Mining*, vol. 3, 1999.
- [46] J.R. Quinlan, *C4.5: Programs for Machine Learning*, San Mateo, Calif.: Morgan Kaufmann, 1993.
- [47] S. Ruggieri, "Efficient C4.5," Technical Report TR-00-01, Dept. of Information, Univ. of Pisa, Feb. 1999.
- [48] J.H. Saltz, R. Mirchandaney, and K. Crowley, "Run-Time Parallelization and Scheduling of Loops," *IEEE Trans. Computers*, vol. 40, no. 5, pp. 603-612, May 1991.
- [49] A. Savasere, E. Omiecinski, and S. Navathe, "An Efficient Algorithm for Mining Association Rules in Large Databases," *Proc. 21st Conf. Very Large Databases (VLDB)*, 1995.
- [50] J. Shafer, R. Agrawal, and M. Mehta, "SPRINT: A Scalable Parallel Classifier for Data Mining," *Proc. 22nd Int'l Conf. Very Large Databases (VLDB)*, pp. 544-555, Sept. 1996.
- [51] A. Shatdal, "Architectural Considerations for Parallel Query Evaluation Algorithms," Technical Report CS-TR-1996-1321, Univ. of Wisconsin, 1999.
- [52] D.B. Skillicorn, "Strategies for Parallel Data Mining," *IEEE Concurrency*, Oct./Dec. 1999.
- [53] A. Srivastava, E. Han, V. Kumar, and V. Singh, "Parallel Formulations of Decision-Tree Classification Algorithms," *Proc. 1998 Int'l Conf. Parallel Processing*, 1998.
- [54] H. Yu and L. Rauchwerger, "Adaptive Reduction Parallelization Techniques," *Proc. 2000 Int'l Conf. Supercomputing*, pp. 66-75, May 2000.
- [55] M.J. Zaki, C.-T. Ho, and R. Agrawal, "Parallel Classification for Data Mining on Shared-Memory Multiprocessors," *Proc. IEEE Int'l Conf. Data Eng.*, pp. 198-205, May 1999.
- [56] M.J. Zaki, M. Ogihara, S. Parthasarathy, and W. Li, "Parallel Data Mining for Association Rules on Shared Memory Multiprocessors," *Proc. Conf. Supercomputing '96*, Nov. 1996.
- [57] M.J. Zaki, "Parallel and Distributed Association Mining: A Survey," *IEEE Concurrency*, vol. 7, no. 4, pp. 14-25, 1999.



Ruoming Jin received the MS degree in computer science from the University of Delaware in 2001, and the ME and BE degrees in computer engineering from Beijing University of Aeronautics and Astronautics, China in 1999 and 1996, respectively. He is a PhD candidate in the Computer Information Science Department at the Ohio State University. His research focuses on data mining, streaming data processing, and high performance computing.



Ge Yang received the masters degree in computer sciences from Ohio State University in 2003, where his research focused on parallel OLAP algorithms and scalable data mining. He also received the masters degree in applied mathematics at Zhongshan University in 2000. He is currently a masters student in the financial engineering program at the Kent State University.



Gagan Agrawal received the BTech degree from the Indian Institute of Technology, Kanpur, in 1991, and the MS and PhD degrees from the University of Maryland, College Park, in 1994 and 1996, respectively. He is an associate professor of computer and information science at the Ohio State University. His research interests include parallel and distributed computing, compilers, data mining, OLAP, grid computing, and processing of streaming data. He has

published more than 75 refereed papers in these areas. He is a member of ACM and IEEE Computer Society. He received a US National Science Foundation CAREER award in 1998.

► For more information on this or any computing topic, please visit our Digital Library at www.computer.org/publications/dlib.