

Scalable Parallel Data Mining for Association Rules *

Eui-Hong (Sam) Han George Karypis, Vipin Kumar

Department of Computer Science
University of Minnesota
4-192 EECS Bldg., 200 Union St. SE
Minneapolis, MN 55455, USA
{han,karypis,kumar}@cs.umn.edu

Last updated on July 15, 1997 at 10:31am

Abstract

One of the important problems in data mining is discovering association rules from databases of transactions where each transaction consists of a set of items. The most time consuming operation in this discovery process is the computation of the frequency of the occurrences of interesting subset of items (called candidates) in the database of transactions. To prune the exponentially large space of candidates, most existing algorithms consider only those candidates that have a user defined minimum support. Even with the pruning, the task of finding all association rules requires a lot of computation power and memory. Parallel computers offer a potential solution to the computation requirement of this task, provided efficient and scalable parallel algorithms can be designed. In this paper, we present two new parallel algorithms for mining association rules. The *Intelligent Data Distribution* algorithm efficiently uses aggregate memory of the parallel computer by employing intelligent candidate partitioning scheme and uses efficient communication mechanism to move data among the processors. The *Hybrid Distribution* algorithm further improves upon the *Intelligent Data Distribution* algorithm by dynamically partitioning the candidate set to maintain good load balance. The experimental results on a Cray T3D parallel computer show that the *Hybrid Distribution* algorithm scales linearly, exploits the aggregate memory better, and can generate more association rules with a single scan of database per pass.

Keywords: Data mining, parallel processing, association rules, load balance, scalability.

*This work was supported by NSF grant ASC-9634719, Army Research Office contract DA/DAAH04-95-1-0538, Cray Research Inc. Fellowship, and IBM partnership award, the content of which does not necessarily reflect the policy of the government, and no official endorsement should be inferred. Access to computing facilities was provided by AHPARC, Minnesota Supercomputer Institute, Cray Research Inc., and NSF grant CDA-9414015. See <http://www.cs.umn.edu/~han> for other related papers.

1 Introduction

One of the important problems in data mining [SAD⁺93] is discovering association rules from databases of transactions, where each transaction contains a set of items. The most time consuming operation in this discovery process is the computation of the frequencies of the occurrence of subsets of items, also called candidates, in the database of transactions. Since usually such transaction-based databases contain a large number of distinct items, the total number of candidates is prohibitively large. Hence, current association rule discovery techniques [AS94, HS95, SON95, SA95] try to prune the search space by requiring a minimum level of support for candidates under consideration. Support is a measure of the number of occurrences of the candidates in database transactions. *Apriori* [AS94] is a recent state-of-the-art algorithm that aggressively prunes the set of potential candidates of size k by using the following observation: a candidate of size k can meet the minimum level of support only if all of its subsets also meet the minimum level of support. In the k^{th} iteration, this algorithm computes the occurrences of potential candidates of size k in each of the transactions. To do this task efficiently, the algorithm maintains all potential candidates of size k in a hash tree. This algorithm does not require the transactions to stay in main memory, but requires the hash trees to stay in main memory.

Even with the highly effective pruning method of *Apriori*, the task of finding all association rules can require a lot of computation power that is available only in parallel computers. Furthermore, the size of the main memory in the serial computer puts an upper limit on the number of the candidates that can be considered in any iteration without requiring multiple scans of the data. This effectively puts a lower bound on the minimum level of support imposed on candidates under consideration. Parallel computers also offer increased memory to solve this problem.

Two parallel algorithms, *Count Distribution* and *Data Distribution* were proposed in [AS96]. The *Count Distribution* algorithm scales linearly and has excellent speedup and sizeup behavior with respect to the number of transactions [AS96]. However, this algorithm works only when the entire hash tree in each pass of the algorithm fits into the main memory of single processor of the parallel computers. Hence, the *Count Distribution* algorithm, like its sequential counterpart *Apriori*, is unscalable with respect to the increasing size of candidate set. The *Data Distribution* algorithm addresses the memory problem of the *Count Distribution* algorithm by partitioning the candidate set and assigning a partition to each processor in the system. However, this algorithm suffers from three types of inefficiency. First, the algorithm results in high communication overhead due to data movement. Second, the schedule for interactions among processors is such that it can cause processors to idle. Third, each transaction has to visit multiple hash trees causing redundant computation.

In this paper, we present two parallel algorithms for mining association rules. We first present *Intelligent Data Distribution* algorithm that improves upon the *Data Distribution* algorithm by minimizing communication and idling overhead, and by eliminating redundant computation. The *Hybrid Distribution* algorithm further improves upon the *Intelligent Data Distribution* algorithm by dynamically grouping processors and partitioning the candidate set accordingly to maintain good load balance. The experimental results on a Cray T3D parallel computer show that the *Hybrid Distribution* algorithm scales linearly, exploits the aggregate memory efficiently, and can generate more association rules with a single scan of database per pass.

The rest of this paper is organized as follows. Section 2 provides an overview of the serial algorithm for mining association rules. Section 3 describes existing and proposed parallel algorithms. Section 4 presents the performance analysis of the algorithms. Experimental results are shown in Section 5. Section 6 contains conclusions. A preliminary version of this paper appeared in [HKK97].

TID	Items
1	Bread, Coke, Milk
2	Beer, Bread
3	Beer, Coke, Diaper, Milk
4	Beer, Bread, Diaper, Milk
5	Coke, Diaper, Milk

Table 1: Transactions from supermarket.

2 Basic Concepts

Let T be the set of transactions where each transaction is a subset of the item-set I . Let C be a subset of I , then we define the *support count* of C with respect to T to be:

$$\sigma(C) = |\{t | t \in T, C \subseteq t\}|.$$

Thus $\sigma(C)$ is the number of transactions that contain C . For example, consider a set of transactions from supermarket as shown in Table 1. The items set I for these transactions is {Bread, Beer, Coke, Diaper, Milk}. The support count of {Diaper, Milk} is $\sigma(\text{Diaper, Milk}) = 3$, whereas $\sigma(\text{Diaper, Milk, Beer}) = 2$.

An *association rule* is an expression of the form $X \xrightarrow{s, \alpha} Y$, where $X \subseteq I$ and $Y \subseteq I$. The *support* s of the rule $X \xrightarrow{s, \alpha} Y$ is defined as $\sigma(X \cup Y)/|T|$, and the *confidence* α is defined as $\sigma(X \cup Y)/\sigma(X)$. For example, consider a rule {Diaper, Milk} \implies {Beer}, i.e. presence of diaper and milk in a transaction tends to indicate the presence of beer in the transaction. The support of this rule is $\sigma(\text{Diaper, Milk, Beer})/5 = 40\%$. The confidence of this rule is $\sigma(\text{Diaper, Milk, Beer})/\sigma(\text{Diaper, Milk}) = 66\%$. A rule that has a very high confidence (i.e., close to 1.0) is often very important, because it provides an accurate prediction on the association of the items in the rule. The support of a rule is also important, since it indicates how frequent the rule is in the transactions. Rules that have very small support are often uninteresting, since they do not describe significantly large populations. This is one of the reasons why most algorithms [AS94, HS95, SON95] disregard any rules that do not satisfy the minimum support condition specified by the user. This filtering due to the minimum required support is also critical in reducing the number of derived association rules to a manageable size. Note that the total number of possible rules is proportional to the number of subsets of the item-set I , which is $2^{|I|}$. Hence the filtering is absolutely necessary in most practical settings.

The task of discovering an association rule is to find all rules $X \xrightarrow{s, \alpha} Y$, such that s is greater than a given minimum support threshold and α is greater than a given minimum confidence threshold. The association rule discovery is composed of two steps. The first step is to discover all the frequent item-sets (candidate sets that have more support than the minimum support threshold specified). The second step is to generate association rules from these frequent item-sets. The computation of finding the frequent item-sets is much more expensive than finding the rules from these frequent item-sets. Hence in this paper, we only focus on the first step. The parallel implementation of the second step is straightforward and is discussed in [AS96].

A number of algorithms have been developed for discovering frequent item-sets [AIS93, AS94, HS95]. Our parallel algorithms are based on the *Apriori* algorithm [AS94] that has smaller computational complexity compared to other algorithms. In the rest of this section, we briefly describe the *Apriori* algorithm. The reader should refer to [AS94] for further details.

The high level structure of the *Apriori* algorithm is given in Figure 1. The *Apriori* algorithm consists of a number of

```

1.  $F_1 = \{ \text{frequent 1-item-sets} \}$  ;
2. for (  $k = 2$ ;  $F_{k-1} \neq \phi$ ;  $k++$  ) {
3.    $C_k = \text{apriori\_gen}(F_{k-1})$ 
4.   for all transactions  $t \in T$  {
5.      $\text{subset}(C_k, t)$ 
6.   }
7.    $F_k = \{ c \in C_k \mid c.\text{count} \geq \text{minsup} \}$ 
8. }
9.  $\text{Answer} = \bigcup F_k$ 

```

Figure 1: Apriori Algorithm

passes. Initially F_1 contains all the items (i.e., item set of size one) that satisfy the minimum support requirement. During pass k , the algorithm finds the set of frequent item-sets F_k of size k that satisfy the minimum support requirement. The algorithm terminates when F_k is empty. In each pass, the algorithm first generates C_k , the candidate item-sets of size k . Function $\text{apriori_gen}(F_{k-1})$ constructs C_k by extending frequent item-sets of size $k - 1$. This ensures that all the subsets of size $k - 1$ of a new candidate item-set are in F_{k-1} . Once the candidate item-sets are found, their frequencies are computed by counting how many transactions contain these candidate item-sets. Finally, F_k is generated by pruning C_k to eliminate item-sets with frequencies smaller than the minimum support. The union of the frequent item-sets, $\bigcup F_k$, is the frequent item-sets from which we generate association rules.

Computing the counts of the candidate item-sets is the most computationally expensive step of the algorithm. One naive way to compute these counts is to perform string-matching of each transaction against each candidate item-set. A faster way of performing this operation is to use a candidate hash tree in which the candidate item-sets are hashed [AS94]. Here we explain this via an example to facilitate the discussions of parallel algorithms and their analysis.

Figure 2 shows one example of the candidate hash tree with candidates of size 3. The internal nodes of the hash tree have hash tables that contain links to child nodes. The leaf nodes contain the candidate item-sets. A hash tree of candidate item-sets is constructed as follows. Initially, the hash tree contains only a root node, which is a leaf node containing no candidate item-set. When each candidate item-set is generated, the items in the set are stored in sorted order. Each candidate item-set is inserted into the hash tree by hashing each successive item at the internal nodes and then following the links in the hash table. Once a leaf is reached, the candidate item-set is inserted at the leaf if the total number of candidate item-sets are less than the maximum allowed. If the total number of candidate item-sets at the leaf exceeds the maximum allowed and the depth of the leaf is less than k , the leaf node is converted into an internal node and child nodes are created for the new internal node. The candidate item-sets are distributed to the child nodes according to the hash values of the items. For example, the candidate item set $\{1\ 2\ 4\}$ is inserted by hashing item 1 at the root to reach the left child node of the root, hashing item 2 at that node to reach the middle child node, hashing item 3 to reach the left child node which is a leaf node.

The *subset* function traverses the hash tree from the root with every item in a transaction as a possible starting item of a candidate. In the next level of the tree, all the items of the transaction following the starting item are hashed. This is done recursively until a leaf is reached. At this time, all the candidates at the leaf are checked against the transaction and their counts are updated accordingly. Figure 2 shows the subset operation at the first level of the tree with transaction $\{1\ 2\ 3\ 5\ 6\}$. The item 1 is hashed to the left child node of the root and the following transaction $\{2\ 3\ 5\ 6\}$ is applied

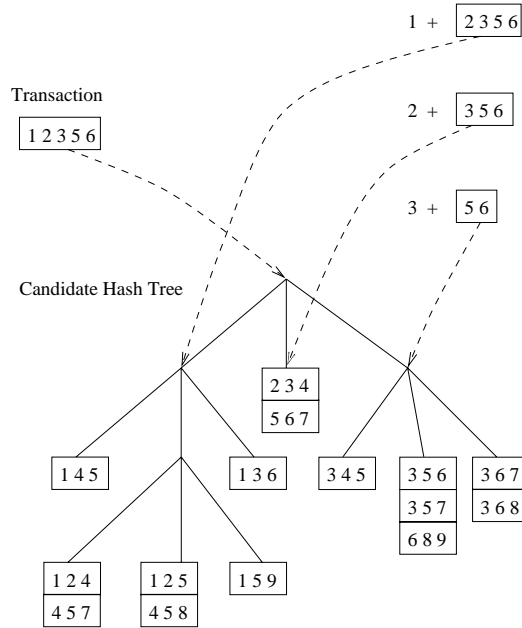



Figure 2: Subset operation on the root of a candidate hash tree.

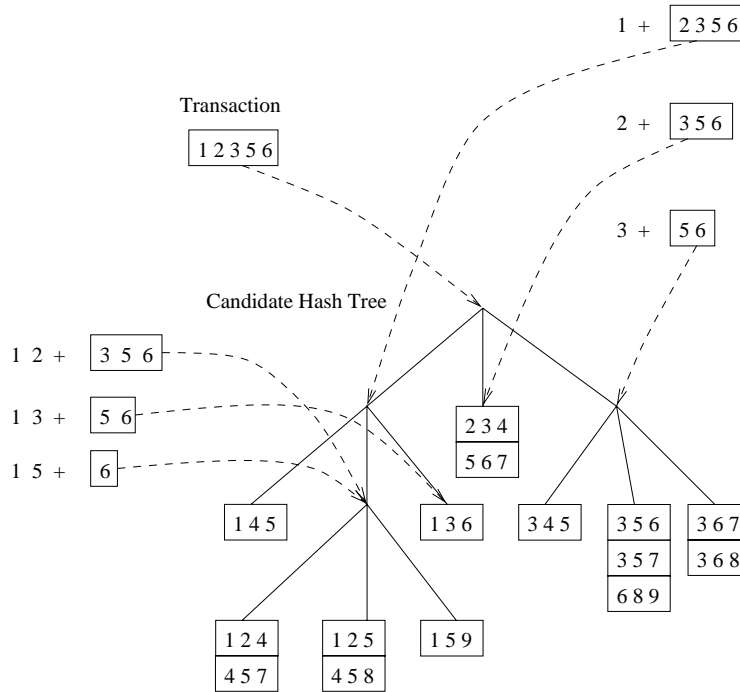
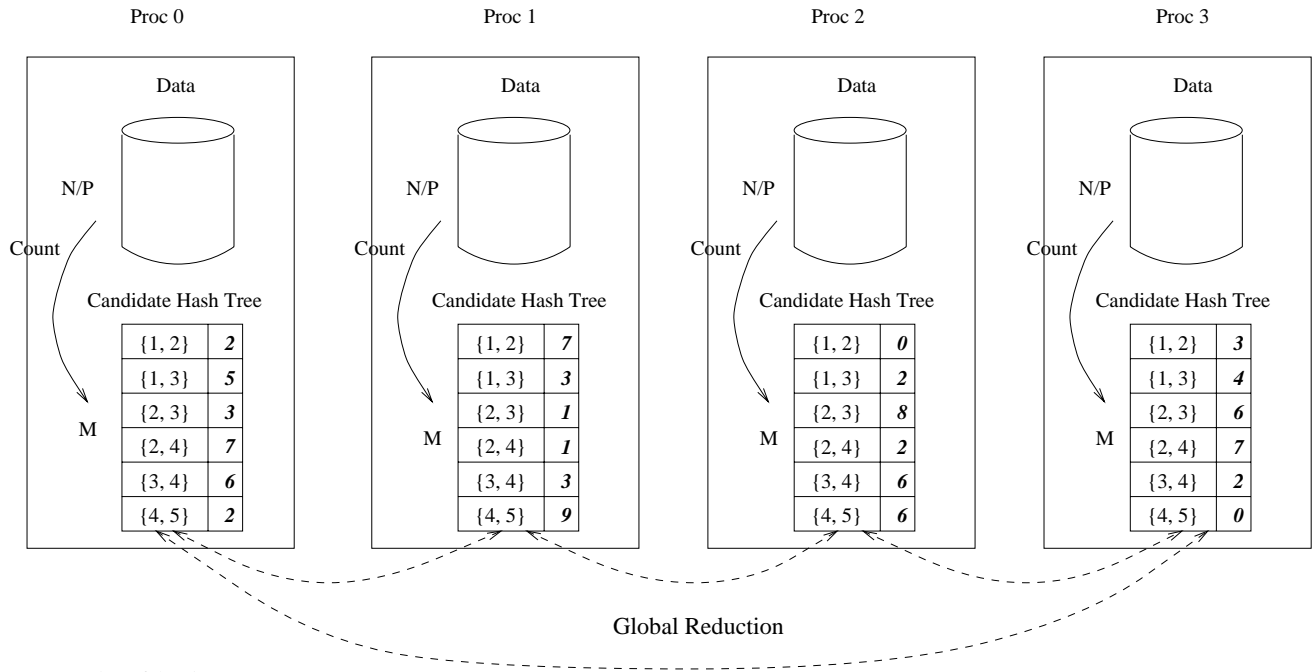


Figure 3: Subset operation on the left most subtree of the root of a candidate hash tree.



N: number of data items
M: size of candidate set
P: number of processors

Figure 4: Count Distribution (CD) Algorithm

recursively to the left child node. The item 2 is hashed to the middle child node of the root and the whole transaction is checked against two candidate item-sets in the middle child node. Then item 3 is hashed to the right child node of the root and the following transaction {5 6} is applied recursively to the right child node. Figure 3 shows the subset operation on the left child node of the root. Here the items 2 and 5 are hashed to the middle child node and the following transactions {3 5 6} and {6} respectively are applied recursively to the middle child node. The item 3 is hashed to the right child node and the remaining transaction {5 6} is applied recursively to the right child node.

3 Parallel Algorithms

In this section, we will focus on the parallelization of the task that finds all frequent item-sets. We first discuss two parallel algorithms proposed in [AS96] to help motivate our parallel formulations. In all our discussions, we assume that the transactions are evenly distributed among the processors.

3.1 Count Distribution Algorithm

In the *Count Distribution (CD)* algorithm proposed in [AS96], each processor computes how many times all the candidates appear in the locally stored transactions. This is done by building the entire hash tree that corresponds to all the candidates and then performing a single pass over the locally stored transactions to collect the counts. The global counts of the candidates are computed by summing these individual counts using a global reduction operation [KGGK94]. This algorithm is illustrated in Figure 4. Note that since each processor needs to build a hash tree for all the candidates, these hash trees are identical at each processor. Thus, excluding the global reduction, each processor in the *CD* algorithm executes the serial *Apriori* algorithm on the locally stored transactions.

This algorithm has been shown to scale linearly with the number of transactions [AS96]. This is because each processor can compute the counts independently of the other processors and needs to communicate with the other processors only once at the end of the computation step. However, this algorithm works well only when the hash trees can fit into the main memory of each processor. If the number of candidates is large, then the hash tree does not fit into the main memory. In this case, this algorithm has to partition the hash tree and compute the counts by scanning the database multiple times, once for each partition of the hash tree. Note that the number of candidates increases if either the number of distinct items in the database increases or if the minimum support level of the association rules decreases. Thus the *CD* algorithm is effective for small number of distinct items and a high minimum support level.

3.2 Data Distribution Algorithm

The *Data Distribution (DD)* algorithm [AS96] addresses the memory problem of the *CD* algorithm by partitioning the candidate item-sets among the processors. This partitioning is done in a round robin fashion. Each processor is responsible for computing the counts of its locally stored subset of the candidate item-sets for all the transactions in the database. In order to do that, each processor needs to scan the portions of the transactions assigned to the other processors as well as its locally stored portion of the transactions. In the *DD* algorithm, this is done by having each processor receive the portions of the transactions stored in the other processors as follows. Each processor allocates P buffers (each one page long and one for each processor). At processor P_i , the i^{th} buffer is used to store transactions from the locally stored database and the remaining buffers are used to store transactions from the other processors, such that buffer j stores transactions from processor P_j . Now each processor P_i checks the P buffers to see which one contains data. Let k be this buffer (ties are broken in favor of buffers of other processors and ties among buffers of other processors are broken arbitrarily). The processor processes the transactions in this buffer and updates the counts of its own candidate subset. If this buffer corresponds to the buffer that stores local transactions (i.e., $k = i$), then it is sent to all the other processors (via asynchronous sends), and a new page is read from the local database. If this buffer corresponds to a buffer that stores transactions from another processor (i.e., $k \neq i$), then it is cleared and an asynchronous receive request is issued to processor P_k . This continues until every processor has processed all the transactions. Having computed the counts of its candidate item-sets, each processor finds the frequent item-sets from its candidate item-set and these frequent item-sets are sent to every other processor using an all-to-all broadcast operation [KGGK94]. Figure 5 shows the high level operations of the algorithm. Note that each processor has a different set of candidates in the candidate hash tree.

This algorithm exploits the total available memory better than *CD*, as it partitions the candidate set among processors. As the number of processors increases, the number of candidates that the algorithm can handle also increases. However, as reported in [AS96], the performance of this algorithm is significantly worse than the *CD* algorithm. The run time of this algorithm is 10 to 20 times more than that of the *CD* algorithm on 16 processors [AS96]. The problem lies with the communication pattern of the algorithm and the redundant work that is performed in processing all the transactions.

The communication pattern of this algorithm causes three problems. First, during each pass of the algorithm each processor sends to all the other processors the portion of the database that resides locally. In particular, each processor reads the locally stored portion of the database one page at a time and sends it to all the other processors by issuing $P - 1$ send operations. Similarly, each processor issues a receive operation from each other processor in order to receive these pages. If the interconnection network of the underlying parallel computer is fully connected (i.e., there is a direct link between all pairs of processors) and each processor can receive data on all incoming links simultane-

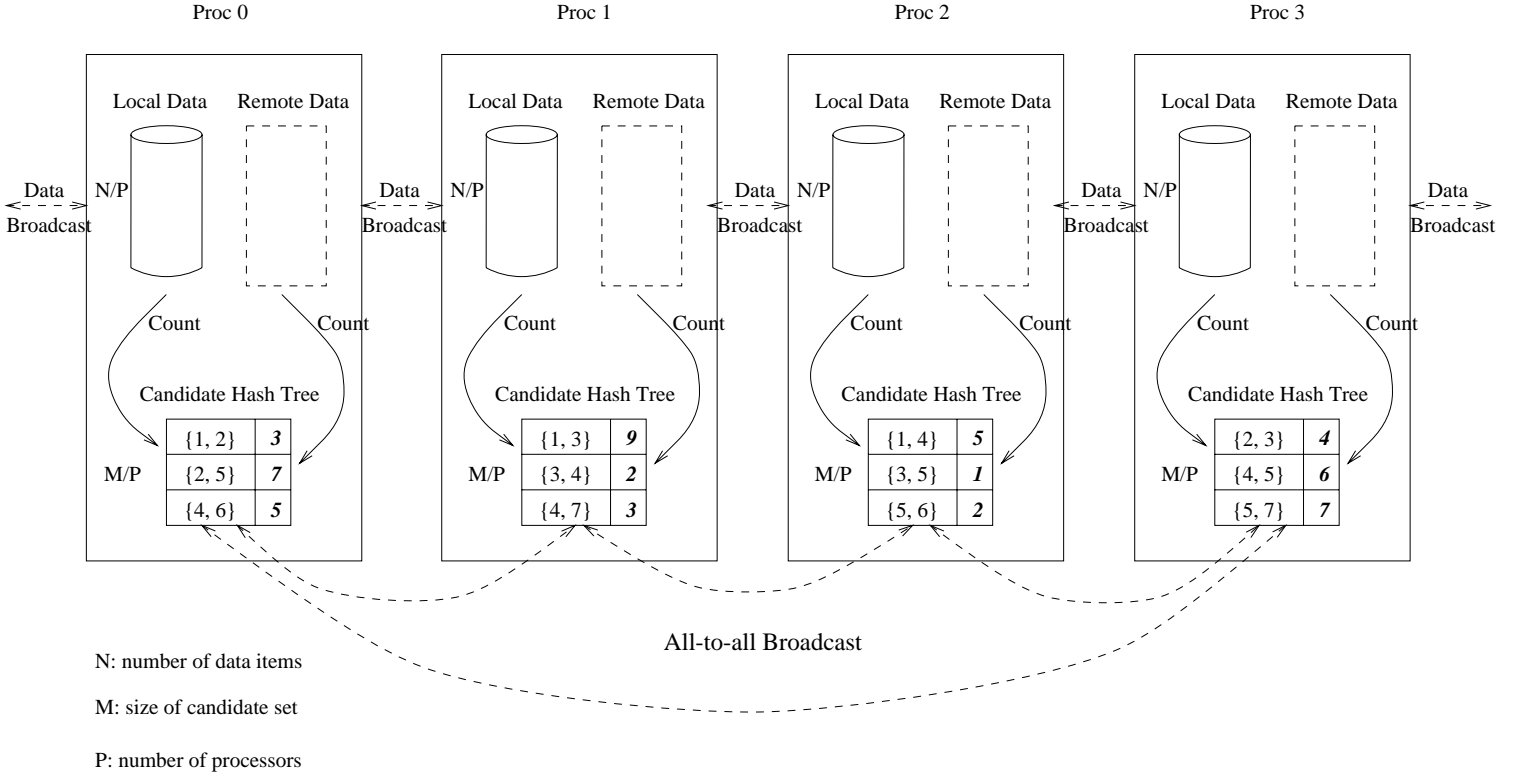


Figure 5: Data Distribution (DD) Algorithm

ously, then this communication pattern will lead to a very good performance. In particular, if $O(N/P)$ is the size of the database assigned locally to each processor, the amount of time spent in the communication will be $O(N)$. However, on all realistic parallel computers, the processors are connected via a sparser networks (such as 2D, 3D or hypercube) and a processor can receive data from (or send data to) only one other processor at a time. On such machines, this communication pattern will take significantly more than $O(N)$ time because of contention.

Second, with finite number of communication buffers in each processor, the proposed all-to-all communication scheme causes processors to idle. For instance, consider the case when one processor finishes its operation on local data and sends the buffer to all other processors. Now if the communication buffer of any receiving processors is full, the send operation is blocked.

Third, if we look at the size of the candidate sets as a function of the number of passes of the algorithm, we see that in the first few passes, the size of the candidate sets increases and after that it decreases. In particular, during the last several passes of the algorithm, there are only a small number of items in the candidate sets. However, each processor in the *DD* algorithm still sends the locally stored portions of the database to all the other processors. Thus, even though the computation decreases, the amount of communication remains the same.

The redundant work is introduced due to the fact that every processor has to process every single transaction in the database. In *CD* (see Figure 4), only N/P transactions go through each hash tree of M candidates, whereas in *DD* (see Figure 5), all N transactions have to go through each hash tree of M/P candidates. Although, the number of candidates stored at each processor has been reduced by a factor of P , the amount of computation performed for each transaction has not been proportionally reduced. If the amount of work required for each transaction to be checked against the hash tree of M/P candidates is $1/P$ of that of the hash tree of M candidates, then there is no extra work. As discussed in Section 4, in general, the amount of work per transaction will go down by a factor much smaller than P .


```

while (!done) {
  FillBuffer(fd, SBuf);
  for (k = 0; k < P-1; ++k) {
    /* send/receive data in non-blocking pipeline */
    MPI_Irecv(RBuf, left);
    MPI_Isend(SBuf, right);

    /* process transactions in SBuf and update hash tree */
    Subset(HTree, SBuf);

    MPI.Waitall();

    /* swap two buffers */
    tmp = SBuf;
    SBuf = RBuf;
    RBuf = tmp;
  }
  /* process transactions in SBuf and update hash tree */
  Subset(HTree, SBuf);
}

```

Figure 6: Pseudo Code for Data Movements

3.3 Intelligent Data Distribution Algorithm

We developed the *Intelligent Data Distribution (IDD)* algorithm that solves the problems of the *DD* algorithm discussed in Section 3.2. In *IDD*, the locally stored portions of the database are sent to all the other processors by using a ring-based all-to-all broadcast described in [KGGK94]. This operation does not suffer from the contention problems of the *DD* algorithm and it takes $O(N)$ time on any parallel architecture that can be embedded in a ring. Figure 6 shows the pseudo code for this data movement operation. In our algorithm, the processors form a logical ring and each processor determines its right and left neighboring processors. Each processor has one send buffer (SBuf) and one receive buffer (RBuf). Initially, the SBuf is filled with one block of local data. Then each processor initiates an asynchronous send operation to the right neighboring processor with SBuf and an asynchronous receive operation to the left neighboring processor with RBuf. While these asynchronous operations are proceeding, each processor processes the transactions in SBuf and collects the counts of the candidates assigned to the processor. After this operation, each processor waits until these asynchronous operations complete. Then the roles of SBuf and RBuf are switched and the above operations continue for $P - 1$ times. Compared to *DD*, where all the processors send data to all other processors, we perform only a point-to-point communication between neighbors, thus eliminating any communication contention. Furthermore, if the time to process a buffer does not vary much, then there is little time lost in idling.

In order to eliminate the redundant work due to the partitioning of the candidate item-sets, we must find a fast way to check whether a given transaction can potentially contain any of the candidates stored at each processor. This cannot be done by partitioning C_k in a round-robin fashion. However, if we partition C_k among processors in such a way that each processor gets item-sets that begin only with a subset of all possible items, then we can check the items of a transaction against this subset to determine if the hash tree contains candidates starting with these items. We traverse the hash tree with only the items in the transaction that belong to this subset. Thus, we solve the redundant work problem of *DD* by

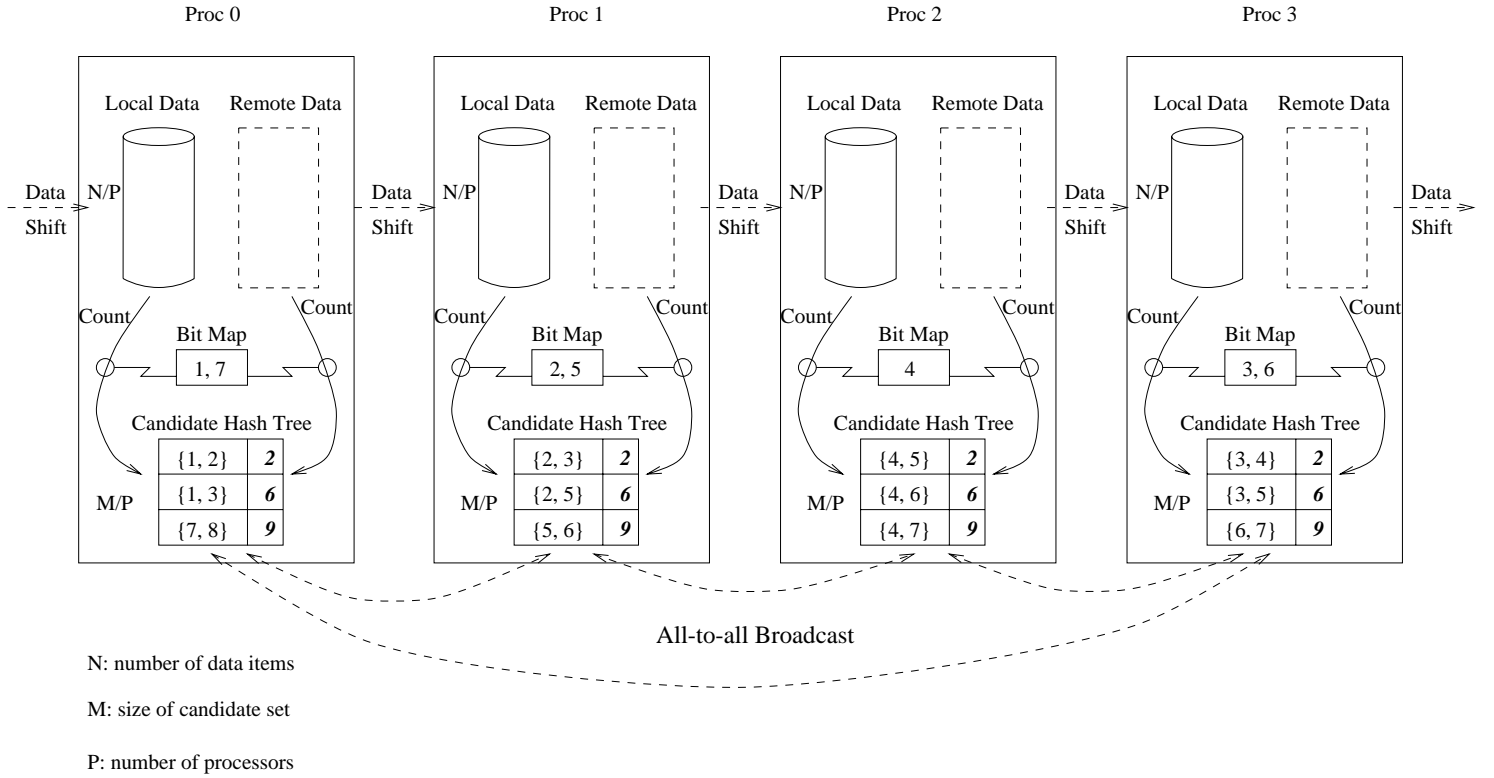


Figure 7: Intelligent Data Distribution (IDD) Algorithm

the intelligent partitioning of C_k .

Figure 7 shows the high level picture of the algorithm. In this example, Processor 0 has all the candidates starting with items 1 and 7, Processor 1 has all the candidates starting with 2 and 5, and so on. Each processor keeps the first items of the candidates it has in a bit-map. In the *Apriori* algorithm, at the root level of hash tree, every item in a transaction is hashed and checked against the hash tree. However, in our algorithm, at the root level, each processor filters every item of the transaction by checking against the bit-map to see if the processor contains candidates starting with that item of the transaction. If the processor does not contain the candidates starting with that item, the processing steps involved with that item as the first item in the candidate can be skipped. This reduces the amount of transaction data that has to go through the hash tree; thus, reducing the computation. For example, let $\{1\ 2\ 3\ 4\ 5\ 6\ 7\ 8\}$ be a transaction that processor 0 is processing in the *subset* function discussed in Section 2. At the top level of the hash tree, processor 0 will only proceed with items 1 and 7 (i.e., $1 + 2\ 3\ 4\ 5\ 6\ 7\ 8$ and $7 + 8$). When the page containing this transaction is shifted to processor 1, this processor will only process items starting with 2 and 5 (i.e., $2 + 3\ 4\ 5\ 6\ 7\ 8$ and $5 + 6\ 7\ 8$). Figure 8 shows how this scheme works when a processor contains only those candidate item-sets that start with 1, 3 and 5. Thus for each transaction in the database, our approach partitions the amount of work to be performed among processors, thus eliminating any redundant work. Note that both the judicious partitioning of the hash tree (indirectly caused by the partitioning of candidate item-set) and the filtering step are required to eliminate this redundant work.

The intelligent partitioning of the candidate set used in *IDD* requires our algorithm to have a good load balancing. One of the criteria of a good partitioning involved here is to have an equal number of candidates in all the processors. This gives about the same size hash tree in all the processors and thus provides good load balancing among processors. Note that in the *DD* algorithm, this was accomplished by distributing candidates in a round robin fashion. A naive

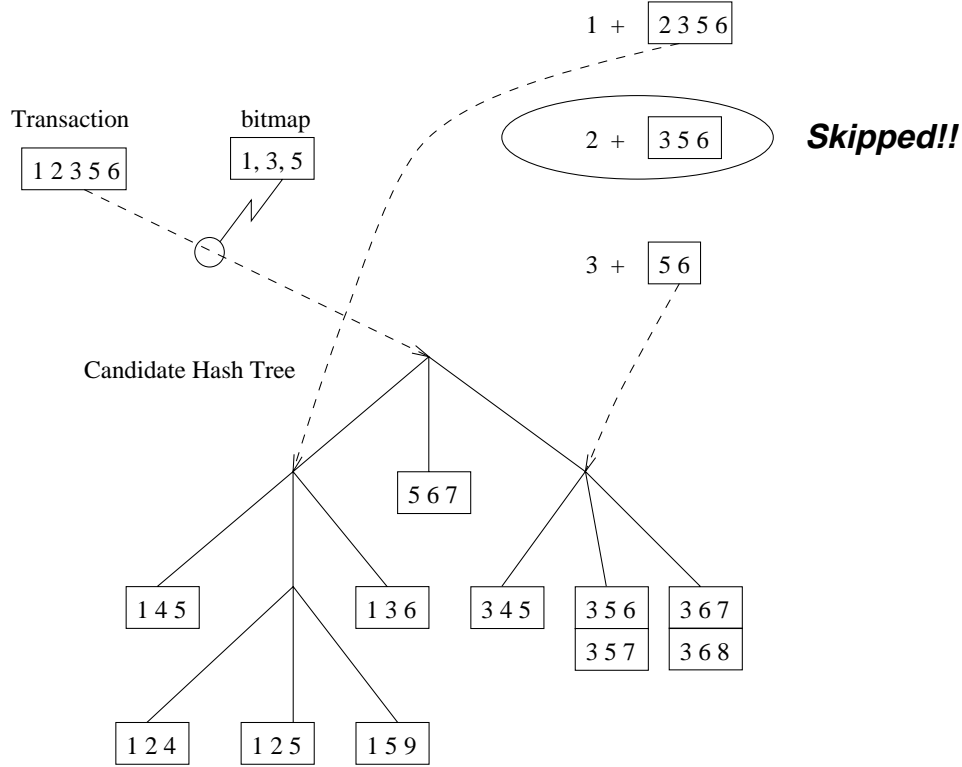


Figure 8: Subset operation on the root of a candidate hash tree in IDD.

method for assigning candidates to processors can lead to a significant load imbalance. For instance, consider a database with 100 distinct items numbered from 1 to 100 and that the database transactions have more data items numbered with 1 to 50. If we partition the candidates between two processors and assign all the candidates starting with items 1 to 50 to processor P_0 and candidates starting with items 51 to 100 to processor P_1 , then there would be more work for processor P_0 .

To achieve a load-balanced distribution of the candidate item-sets, we use a partitioning algorithm that is based on bin-packing [PS82]. For each item, we first compute the number of candidate item-sets starting with this particular item. Note that at this time we do not actually store the candidate item-sets, but just store the number of candidate item-sets starting with each item. We then use a bin-packing algorithm to partition these items in P buckets such that the numbers of the candidate item-sets starting with these items in each bucket are equal. To remove any data skew, our bin-packing algorithm randomly selects the item to be assigned next in a bin. Once the location of each candidate item-set is determined, then each processor locally regenerates and stores candidate item-sets that are assigned to this processor. Figure 7 shows the partitioned candidate hash tree and its corresponding bitmaps in each processor. We were able to achieve less than 5% of load imbalance with the bin packing method described here.

3.4 Hybrid Algorithm

The *IDD* algorithm exploits the total system memory by partitioning the candidate set among all processors. The average number of candidates assigned to each processor is M/P , where M is the number of total candidates. As more processors are used, the number of candidates assigned to each processor decreases. This has two implications. First, with fewer number of candidates per processor, it is much more difficult to balance the work. Second, the smaller number of candidates gives a smaller hash tree and less computation work per transaction. Eventually the amount of

computation may become less than the communication involved. This would be more evident in the later passes of the algorithm as the hash tree size further decreases dramatically. This reduces overall efficiency of the parallel algorithm. This will be an even more serious problem in a system that cannot perform asynchronous communication.

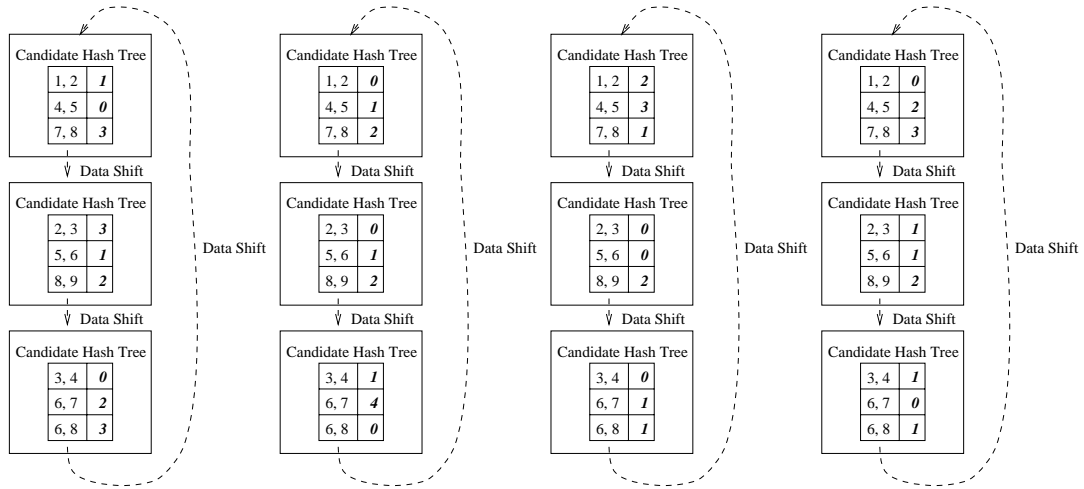
The *Hybrid Distribution (HD)* algorithm addresses the above problem by combining the *CD* and the *IDD* algorithms in the following way. Consider a P -processor system in which the processors are split into G equal size groups, each containing P/G processors. In the *HD* algorithm, we execute the *CD* algorithm as if there were only P/G processors. That is, we partition the transactions of the database into P/G parts each of size $N/(P/G)$, and assign the task of computing the counts of the candidate set C_k for each subset of the transactions to each one of these groups of processors. Within each group, these counts are computed using the *IDD* algorithm. That is, the transactions and the candidate set C_k are partitioned among the processors of each group, so that each processor gets roughly $|C_k|/G$ candidate item-sets and N/P transactions. Now, each group of processors computes the counts using the *IDD* algorithm, and the overall counts are computed by performing a reduction operation among the P/G groups of processors.

The *HD* algorithm can be better visualized if we think of the processors as being arranged in a two dimensional grid of G rows and P/G columns. The transactions are partitioned equally among the P processors. The candidate set C_k is partitioned among the processors of each column of this grid. This partitioning of C_k is identical for each column of processors; i.e., the processors along each row of the grid get the same subset of C_k . Figure 9 illustrates the *HD* algorithm for a 3×4 grid of processors. In this example, the *HD* algorithm executes the *CD* algorithm as if there were only 4 processors, where the 4 processors correspond to the 4 processor columns. That is, the database transactions are partitioned in 4 parts, and each one of these 4 hypothetical processors computes the local counts of all the candidate item-sets. Then the global counts can be computed by performing the global reduction operation discussed in Section 3.1. However, since each one of these hypothetical processors is made up of 3 processors, the computation of local counts of the candidate item-sets in a hypothetical processor requires the computation of the counts of the candidate item-sets on the database transactions sitting on the 3 processors. This operation is performed by executing the *IDD* algorithm within each of 4 hypothetical processors. This is shown in the step 1 of Figure 9. Note that processors in the same row have exactly the same candidates, and candidate sets along the each column partition the total candidate set. At the end of this operation, each processor has complete count of its local candidates for all the transactions located in the processors of the same column (i.e., of a hypothetical processor). Now a reduction operation is performed along the rows such that all processors in each row have the sum of the counts for the candidates in the same row. At this point, the count associated with each candidate item-set corresponds to the entire database of transactions. Now each processor finds frequent item-sets by dropping all those candidate item-sets whose frequency is less than the threshold for minimum support. These candidate item-sets are shown as shaded in Figure 9(b). In the next step, each processor performs all-to-all broadcast operation along the columns of the processor mesh. At this point, all the processors have the frequent sets and are ready to proceed to the next pass.

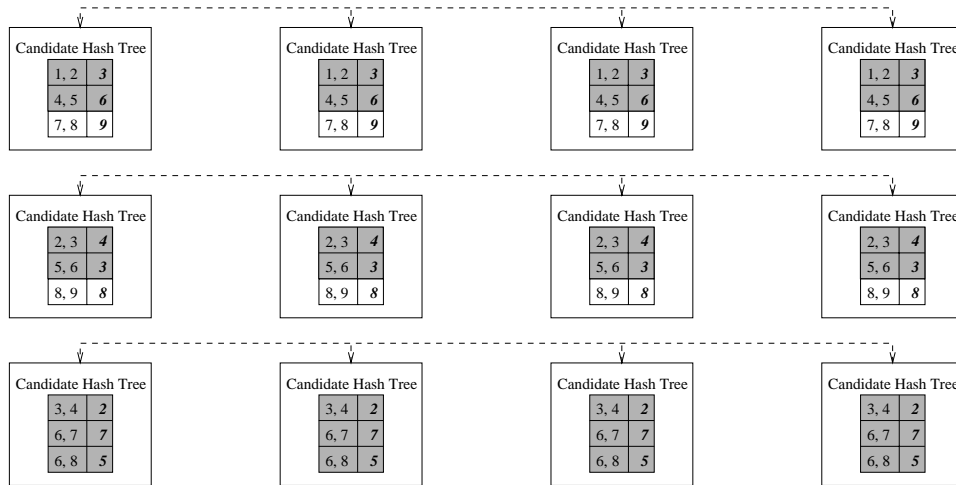
The *HD* algorithm determines the configuration of the processor grid dynamically. In particular, the *HD* algorithm partitions the candidate set into a big enough section and assign a group of processors to each partition. Let m be the number of candidate item-sets such that each processor can hold the hash tree constructed from m candidate item-sets in memory. If the total number of candidates M is less than m , then the *HD* algorithm makes G equal to 1, which means that the *CD* algorithm is run on all the processors. Otherwise G is set to $\lceil M/m \rceil$.

The *HD* algorithm inherits all the good features of the *IDD* algorithm. It also provides good load balance and enough computation work by maintaining minimum number of candidates per processor. At the same time, the amount of data movement in this algorithm has been cut down to $1/G$ of the *IDD*.

Step 1: Partitioning of Candidate Sets and Data Movement Along the Columns



Step 2: Reduction Operation Along the Rows



Step 3: All-to-all Broadcast Operation Along the Columns

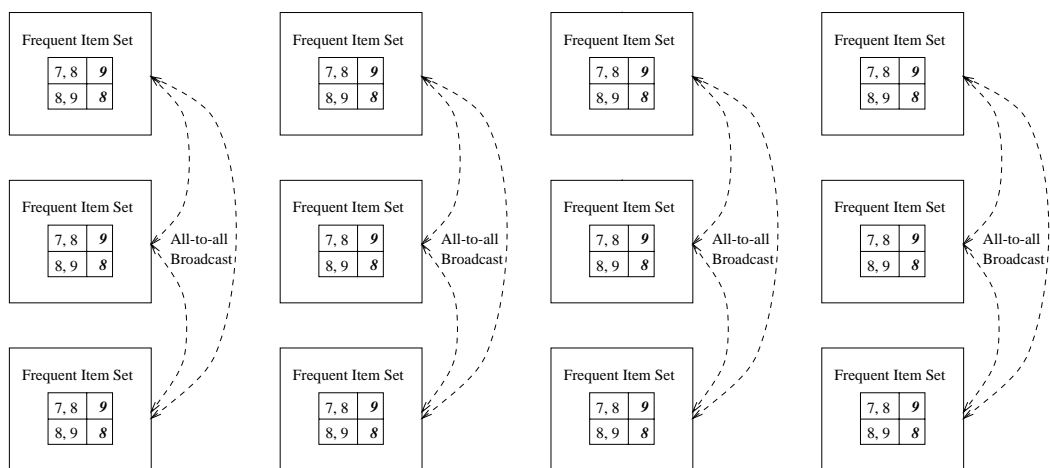


Figure 9: Hybrid Distribution (HD) Algorithm in 3×4 Processor Mesh ($G = 3, P = 12$)

symbol	definition
N	Total number of transactions
P	Number of processors
M	Total number of candidates
G	Number of partitions of candidates in the <i>HD</i> algorithm
k	Pass number in Apriori algorithm
I	Average number of items in a transaction
C	Average number of potential candidates in a transaction
S	Average number of candidates at the leaf node
L	Average number of leaves in the hash tree for the serial <i>Apriori</i> algorithm
$t_{travers}$	Cost of hash tree traversal per potential candidate
t_{check}	Cost of checking at the leaf with S candidates
$V_{i,j}$	Expected number of leaves visited with i potential candidates and j leaves

Table 2: Symbols used in the analysis.

4 Performance Analysis

In this section, we analyze the amount of work done by each algorithm. Throughout the analysis, we ignore the cost of communication. Table 2 describes the symbols used in this section.

As discussed in Section 2, the bulk of the computation is performed by the *subset* function. Consider a transaction that has I items. During the k^{th} pass of the algorithm, this transaction has $C = \binom{I}{k}$ potential candidates that need to be checked against the candidate hash tree. Note that for a given transaction, if one potential candidate visits a leaf node, then all the candidates of this transaction are checked against the leaf node. As a result, if this node is revisited due to a different candidate from the same transaction, no checking is performed. Clearly the total cost of checking at the leaf nodes is directly proportional to the number of distinct leaf nodes visited with the transaction. We assume that the average number of candidate item-sets at the leaf nodes is S . Hence the average number of leaf nodes in a hash tree is $L = M/S$. In the implementation of the algorithm, the desired value of S can be obtained by adjusting the branching factor of the hash tree. In general, the cost of traversal for each potential candidate will depend on the depth of the leaf node in the hash tree reached by the traversal. To simplify the analysis, we assume that the cost of each traversal is the same. Hence, the total traversal cost is directly proportional to C . For each potential candidate, we define $t_{travers}$ to be the cost associated with the traversal of the hash tree and t_{check} to be the cost associated with checking the candidate item-sets of the reached leaf node.

Note that the number of distinct leaves visited by a transaction is in general smaller than the number of potential candidates C . This is because different potential candidates may lead to the same leaf node. In general, if C is relatively large with respect to the number of leaf nodes in the hash tree, then the number of distinct leaf nodes visited will be smaller than C . We can compute the expected number of distinct leaf nodes visited as follows. To simplify the analysis, we assume that each traversal of the hash tree due to a different potential candidate is equally likely to lead to any leaf node of the hash tree.

Let $V_{i,j}$ be the expected number of distinct leaf nodes visited when the transaction has i potential candidates, and the hash tree has j leaf nodes.

$$V_{1,j} = 1$$

$$V_{i,j} = V_{i-1,j} \times (\text{Prob. of reaching a previously visited node}) + (V_{i-1,j} + 1) \times (\text{Prob. of reaching a new node})$$

$$\begin{aligned}
&= V_{i-1,j} \frac{V_{i-1,j}}{j} + (V_{i-1,j} + 1) \frac{j - V_{i-1,j}}{j} \\
&= 1 + \frac{j-1}{j} V_{i-1,j} \\
&= \frac{1 - \left(\frac{j-1}{j}\right)^i}{1 - \left(\frac{j-1}{j}\right)} \\
&= \frac{j^i - (j-1)^i}{j^{i-1}}
\end{aligned} \tag{1}$$

Note that for large j , $V_{i,j} \simeq i$. This can be shown by taking limit on Equation 1:

$$\begin{aligned}
\lim_{j \rightarrow \infty} V_{i,j} &= \lim_{j \rightarrow \infty} \frac{j^i - (j-1)^i}{j^{i-1}} \\
&= \frac{[i(i-1) \cdots 3 \cdot 2] j - [i(i-1) \cdots 3 \cdot 2](j-1)}{(i-1)(i-2) \cdots 2 \cdot 1} \\
&= ij - i(j-1) \\
&= i
\end{aligned} \tag{2}$$

This shows that if the hash tree size is much larger than the number of potential candidates in a transaction, then each potential candidate is likely to visit a distinct leaf node in the hash tree.

Serial Apriori algorithm Recall that in the serial *Apriori* algorithm, the average number of leaf nodes in the hash tree is $L = M/S$. Hence the number of distinct leaf visited per transaction is $V_{C,L}$, and the computation time per transaction for visiting the hash tree is:

$$T_{trans} = C \times t_{travers} + V_{C,L} \times t_{check}$$

So the run time of the serial algorithm for processing N transactions is:

$$\begin{aligned}
T_{comp}^{serial} &= N \times T_{trans} \\
&= N \times C \times t_{travers} + N \times V_{C,L} \times t_{check}
\end{aligned} \tag{3}$$

The CD algorithm In the *CD* algorithm the entire set of candidates is replicated at each processor. Hence the average number of leaf nodes in the local hash tree at each processor is $L = M/S$, which is the same as in the serial *Apriori* algorithm. Thus the *CD* algorithm performs the same computation per transactions as the serial algorithm, but each processor handles only N/P number of transactions. Hence the run time of the *CD* algorithm is:

$$\begin{aligned}
T_{comp}^{CD} &= \frac{N}{P} \times T_{trans} \\
&= \frac{N}{P} \times C \times t_{travers} + \frac{N}{P} \times V_{C,L} \times t_{check}
\end{aligned} \tag{4}$$

Comparing Equation 4 to Equation 3, we see that *CD* performs no redundant computation. In particular, both the time for traversal and for checking scales down by a factor of P .

The *DD* algorithm In the *DD* algorithm, the number of candidates per processor is M/P , as the candidate set is partitioned. Hence the average number of leaf nodes in the local hash tree of each processor is L/P . Therefore, the number of distinct leaf nodes visited per transaction is $V_{C, \frac{L}{P}}$, and the computation time per transaction is:

$$T_{trans}^{DD} = C \times t_{travers} + V_{C, \frac{L}{P}} \times t_{check}$$

The number of transactions processed by each processor is N , as the transactions are shifted around the processors. Hence, the computation per processor of the *DD* algorithm is:

$$\begin{aligned} T_{comp}^{DD} &= N \times T_{trans}^{DD} \\ &= N \times C \times t_{travers} + N \times V_{C, \frac{L}{P}} \times t_{check} \end{aligned} \quad (5)$$

Comparing Equation 5 with the serial complexity (Equation 3), we see that the *DD* algorithm does not reduce the computation associated with the hash tree traversal. For both the serial *Apriori* and the *DD* algorithm, this cost is $N \times C \times t_{travers}$. However, the *DD* algorithm is able to reduce the cost associated with the checking at the leaf nodes. In particular, it reduces the serial cost of $N \times V_{C, L} \times t_{check}$ down to $N \times V_{C, \frac{L}{P}} \times t_{check}$. However, because $V_{C, \frac{L}{P}} > V_{C, L}/P$, the reduction achieved in this part is less than a factor of P . We can easily see this if we consider the case when L is very large. In this case, $V_{C, \frac{L}{P}} \simeq C$ and $V_{C, L}/P \simeq C/P$ by Equation 2. Thus, the number of leaf nodes checked over all the processors by the *DD* algorithm is higher than that of the serial algorithm. This is why the *DD* algorithm performs redundant computation.

The *IDD* algorithm In the *IDD* algorithm, just like the *DD* algorithm, the average number of leaf nodes in the local hash tree of each processor is L/P . However, the average number of potential candidates that need to be checked for each transaction at each processor is much less than *DD*, because of the intelligent partitioning of candidates set and the use of bitmap to prune at the root of the hash tree. More precisely, the number of potential candidates that need to be checked for a transaction is roughly C/P assuming that we have a good balanced partition. So the computation per transaction is:

$$T_{trans}^{IDD} = \frac{C}{P} \times t_{travers} + V_{\frac{C}{P}, \frac{L}{P}} \times t_{check}$$

Thus the computation per processor is:

$$\begin{aligned} T_{comp}^{IDD} &= N \times T_{trans}^{IDD} \\ &= N \times \frac{C}{P} \times t_{travers} + N \times V_{\frac{C}{P}, \frac{L}{P}} \times t_{check} \end{aligned} \quad (6)$$

Comparing Equation 6 to Equation 3, we see that the *IDD* algorithm is successful in reducing the cost associated with the hash tree traversal linearly. It also reduces the checking cost from $N \times V_{C, L} \times t_{check}$ down to $N \times V_{\frac{C}{P}, \frac{L}{P}} \times t_{check}$. Note that for sufficiently large L , $V_{C, L} \simeq C$ and $V_{\frac{C}{P}, \frac{L}{P}} \simeq C/P$. This shows that *IDD* is also able to linearly reduce the cost of checking at the leaf nodes, and thus unlike *DD*, it performs no redundant work. We instrumented our algorithms to verify that the average number of distinct leaf node visited by *IDD* is indeed much less than *DD*. Figure 10 shows that $V_{\frac{C}{P}, \frac{L}{P}}$ of *IDD* goes down by factor of P , but $V_{C, \frac{L}{P}}$ of *DD* does not go down by factor of P . However, P must be relatively small for *IDD* to have a good load balance. If P becomes large with fixed M , the problem of load imbalance discussed in Section 3 makes some processors work on more than $1/P$ of items in a transaction at the root of the hash




Figure 10: Comparison of *DD* and *IDD* in terms of the average number of distinct leaf node visited per transaction with 50K transactions per processor and 0.2% minimum support.

tree.

The *HD* algorithm In the *HD* algorithm, the number of potential candidates per transactions is C/G and the number of candidates per processor is M/G . So the computation time per transaction is:

$$T_{trans}^{HD} = \frac{C}{G} \times t_{travers} + V_{\frac{C}{G}, \frac{L}{G}} \times t_{check}$$

The total number of transactions each processor has to process is GN/P . Thus the computation per processor is:

$$\begin{aligned} T_{comp}^{HD} &= \frac{G \times N}{P} \times T_{trans}^{HD} \\ &= \frac{G \times N}{P} \times \frac{C}{G} \times t_{travers} + \frac{G \times N}{P} \times V_{\frac{C}{G}, \frac{L}{G}} \times t_{check} \end{aligned} \quad (7)$$

Compared to the serial algorithm, Equation 7 shows that the *HD* algorithm reduces the computation linearly with respect to the hash tree traversal cost. The traversal cost is reduced from $N \times C \times t_{travers}$ down to $N \times C \times \frac{t_{travers}}{P}$. The cost of checking at the leaf nodes is reduced from $N \times V_{C,L} \times t_{check}$ down to $(G \times N \times V_{\frac{C}{G}, \frac{L}{G}} \times t_{check})/P$. Note that for sufficiently large L , $N \times V_{C,L} \simeq NC$ and $G \times N \times V_{\frac{C}{G}, \frac{L}{G}}/P \simeq N \times C/P$. Thus, the *HD* algorithm has a linear speedup with respect to the cost of checking at the leaf nodes.

5 Experimental Results

We implemented our parallel algorithms on a 128-processor Cray T3D parallel computer. Each processor on the T3D is a 150Mhz Dec Alpha (EV4), and has 64Mbytes of memory. The processors are interconnected via a three dimensional torus network that has a peak unidirectional bandwidth of 150Mbytes per second, and a small latency. For communication we used the message passing interface (MPI). Our experiments have shown that for 16Kbyte messages we obtain a bandwidth of 74Mbytes/seconds and an effective startup time of 150 microseconds.




Figure 11: Scaleup result with 100K transactions and 0.25% minimum support.

We generated a synthetic dataset using a tool provided by [Pro96] and described in [AS94]. The parameters for the data set chosen are average transaction length of 15 and average size of frequent item sets of 6. Data sets with 1000 transactions (63KB) were generated for different processors. Due to the absence of a true parallel I/O system on the T3D system, we kept a set of transactions in a main memory buffer and read the transactions from the buffer instead of the actual disks. For the experiments involving larger data sets, we read the same data set multiple times. We also performed similar experiments on an IBM SP2 in which the entire database resided on disks. Our experiments (not reported here) show that the I/O requirements do not change the relative performance of the various schemes. We do present the results of one experiment on 16-processor SP2 for comparing *CD* to *IDD* and *HD* when *CD* scans database multiple times due to the partitioned hash tree.

To compare the scalability of the four schemes (*CD*, *DD*, *IDD* and *HD*), we performed scaleup tests with 100K transactions per processor and minimum support of 0.25%. We could not use minimum support smaller than 0.25% because the *CD* algorithm ran out of main memory for storing the hash tree. For this experiment, in the *HD* algorithm we have set the threshold on the number of candidates for switching to the *CD* algorithm very low to show the validity of our approach. With 0.25% support, the *HD* algorithm switched to *CD* algorithm in pass 7 of total 12 passes and 90.7% of the overall response time of the serial code was spent in the first 6 passes. These scaleup results are shown in Figure 11.

As noted in [AS96], the *DD* algorithm scales very poorly. However, the performance achieved by *IDD* is much better than that of the *DD* algorithm. In particular, on 32 processors, *IDD* is faster than *DD* by a factor of 4.4. It can be seen that the performance gap between *IDD* and *DD* widens as the number of processors increases. *IDD* performs better than *DD* because of the better communication mechanism for data movements and the intelligent partitioning of the candidate set. To show the effects of these two improvements, we replaced the communication mechanism of the *DD* algorithm with that of the *IDD*. The scaleup result of this improvement is shown as “*DD+comm*” in Figure 11. Hence the response time reduction from *DD* to *DD+comm* is due to the the better communication mechanism for data movements, and the reduction from *DD+comm* to *IDD* is due to the intelligent partitioning of the candidate set.

Note that the response time of *IDD* increases as we increase the number of processors. This is due to the load balanc-




Figure 12: Sizeup result with 16 processors and 0.25% minimum support.

ing problem discussed in Section 3, where the number of candidates per processor decreases as the number of processors increases. Looking at the performance of the *HD* algorithm, we see that response time remains almost constant as we increase the number of processors while keeping the number of transactions per processor and the minimum support fixed. Comparing against *CD*, we see that *HD* actually performs better as the number of processors increases. Its performance on 128 processors is 9.5% better than *CD*. This performance advantage of *HD* over *CD* is due to the smaller cost of reduction in *HD*, because the number of processors involved in global reduction operation of counts is much less in *HD* than in *CD*.

We measured how our algorithms perform as we increase the number of transactions per processor from 50K(3.2MB) to 800K(50.4MB). For these experiments, we fixed the number of processors at 16 and the minimum support at 0.25%. These results are shown in Figure 12. From this figure, we can see that *CD* and *HD* perform almost identically. For both algorithms, the response time increases linearly with the number of transactions. *IDD* also scales linearly, but because of its load imbalance problem, its performance is somewhat worse.

Our experiments so far have shown that the performance of *HD* and *CD* are quite comparable if the entire hash tree can reside on each processor. However, the real advantage of *HD* (and *IDD*) over *CD* is that they do not require the whole hash tree to reside on each processor, and thus better exploit the available memory. This allows us to use a smaller minimum support in the *Apriori* algorithm.

To verify this, we performed experiments in which we fixed the number of transactions per processor to 50K and successively decreased the minimum support level. These experiments for 16 and 64 processors are shown in Figures 13 and 14 respectively. A couple of interesting observations can be made from these results. First, both *IDD* and *HD* successfully ran using lower support levels that *CD* could not run with. In particular, *IDD* and *HD* ran down to a support level of 0.06% on 16 processors and 0.04% on 64 processors. In contrast, *CD* could only run down to a support level of 0.25% and ran out of memory for the lower supports. The support level for *IDD* and *HD* becomes smaller with the increasing number of processors, because the *IDD* and *HD* algorithms can exploit the aggregate memory of the larger number of processors.

The second thing to notice is that *HD* performs better than *IDD* both on 16 and 64 processors, and the relative perfor-




Figure 13: Response time on 16 processors with 50K transactions as the minimum support varies. At each support level, the total number of candidate item-sets is shown in parenthesis

mance of *HD* compared to *IDD* get better as the number of processors increases. As discussed earlier, this performance difference is due to the load imbalance. As the number of processors increases, this load imbalance gets worse. However, on 16 processors, *IDD* is 37% worse than *HD* for support level 0.25%, but only 18% worse for support of 0.06%. This is because as the support level decreases, the number of candidates (shown in parenthesis in Figures 13 and 14) increases which improves the load balance.

Figures 13 and 14 also show the performance of a simple hybrid algorithm obtained by combining *CD* and *IDD*. In this scheme, in each pass of the *Apriori* algorithm, we perform *CD* if the hash table can fit in the memory of each processors or *IDD* if it can not. As we can see from these results, this simple hybrid algorithm performs worse than *HD*. In particular, the relative performance of this scheme compared to *HD* gets worse as the number of processors increases. For example, for a support level of 0.06%, it is 6% worse on 16 processors and 17% worse on 64 processors. Thus the *HD* algorithm achieves better performance by gradually adjusting the subsets of processors that perform *IDD* and *CD*. This is because of the following two reasons. First, the candidate set is split among fewer number of processors which minimizes load imbalance and second, the reduction operation to obtain the counts in *CD* is performed among fewer processors, which decreases the communication overhead.

Note that an alternative method to handle large candidate set in *CD* is to partition them such that each partition fits in the main memory. Now the entire set of local transactions have to be read at each processor as many time as the number of partitions. This method increases the I/O cost. On the system in which I/O is scalable and fast (e.g., IBM SP2), this cost may be acceptable. We implemented the *CD* algorithm to partition the hash tree and read database multiple times in case the hash tree does not fit into main memory. Figure 15 shows the performance comparison of *CD*, *IDD* and *HD* on 16-processor IBM SP2 machine as the number of candidates increases by lowering minimum support. Unlike the earlier experiments on Cray T3D machine, the whole transactions were read in from the file. Figure 15 shows that as the number of candidates increases both *IDD* and *HD* outperform *CD*. This is due to the increased I/O time required for multiple scan of the database and increased communication time required for global reduction operation of multiple partitions of the candidate frequencies. Note that even on IBM SP2, the penalty due to higher I/O cost is about 8% for 1 million candidates, 11% for 3 million candidates and 25% for 11 million candidates. The I/O penalty increases as the




Figure 14: Response time on 64 processors with 50K transactions as the minimum support varies. At each support level, the total number of candidate item-sets is shown in parenthesis

Number of processors	1	2	4	8	16	32	64
Successful down to	0.25	0.2	0.15	0.1	0.06	0.04	0.03
Ran out of memory at	0.2	0.15	0.1	0.06	0.04	0.03	0.02

Table 3: Minimum support (%) reachable with different number of processors in our algorithms.

number of candidates increases. On systems with slower I/O, the penalty can be much larger.

In another experiment, we varied the number of processors from 2 to 64 and measured how low we can go with minimum support for the *IDD* and *HD* algorithms. Table 3 shows the result for these algorithms. The result shows that as we have more processors, these algorithms can handle lower minimum support. Table 4 shows how the *HD* algorithm chose the processor configuration based on the number of candidates at each pass with 64 processors and 0.04% minimum support.

6 Conclusion

In this paper, we proposed two new parallel algorithms for mining association rules. The *IDD* algorithm utilizes total main memory available more effectively than the *CD* algorithm. This algorithm improves over the *DD* algorithm which has high communication overhead and redundant work. As shown in Section 4, for each transaction the *DD*

Pass	2	3	4	5	6	7	8	9	10
Configuration	8 × 8	64 × 1	4 × 16	2 × 32	2 × 32	2 × 32	2 × 32	2 × 32	1 × 64
No of Cand.	351K	4348K	115K	76K	56K	34K	16K	6K	2K

Table 4: Processor configuration and number of candidates of the *HD* algorithm with 64 processors and 0.04% minimum support for each pass. Note that 64 × 1 configuration is the same as the *DD* algorithm and 1 × 64 is the same as the *CD* algorithm. The total number of pass was 13 and all passes after 9 had 1 × 64 configuration.




Figure 15: Response time on 16 processor IBM SP2 with 100K transactions as the minimum support varies from 0.1% to 0.025%.

algorithm performs substantially more work overall than the serial *Apriori* algorithm. The communication and idling overheads were reduced using a better data movement communication mechanism, and redundant work was reduced by partitioning the candidate set intelligently and using bit maps to prune away unnecessary computation. Another useful feature of *IDD* is that it is well suited for the system environment with single source of data base. For instance, when all the data is coming from a database server or a single file system, one processor can read data from the single source and pass the data along the communication pipe line defined in the algorithm. However, as the number of available processors increases, the efficiency of this algorithm decreases due to load imbalance, unless the amount of work is increased by having more number of candidates.

The *HD* combines the advantages of *CD* and *IDD*. This algorithm partitions candidate sets just like the *IDD* to exploit the aggregate main memory, but dynamically determines the number of partitions such that the partitioned candidate set fits into the main memory of each processor and each processor has enough number of candidates for computation. It also exploits the advantage of *CD* by just exchanging counts information and moving around the minimum number of transactions among the smaller subset of processors.

The experimental results on a 128-processor Cray T3D parallel machine show that the *HD* algorithm scales just as well as the *CD* algorithm with respect to the number of transactions. However, it exploits the aggregate main memory more efficiently and thus is able to discover more association rules with much smaller minimum support with a single scan of database per pass. This parallel formulation is particularly useful in domains where the size of transaction is very large (e.g., similarity tables in molecular biology [HKKM97], generalized association rules [HF95, SA95], or sequential patterns [MTV95, SA96]).

References

- [AIS93] R. Agrawal, T. Imielinski, and A. Swami. Mining association rules between sets of items in large databases. In *Proc. of 1993 ACM-SIGMOD Int. Conf. on Management of Data*, Washington, D.C., 1993.
- [AS94] R. Agrawal and R. Srikant. Fast algorithms for mining association rules. In *Proc. of the 20th VLDB Conference*, pages 487–499, Santiago, Chile, 1994.

- [AS96] R. Agrawal and J.C. Shafer. Parallel mining of association rules. *IEEE Transactions on Knowledge and Data Eng.*, 8(6):962–969, December 1996.
- [HF95] J. Han and Y. Fu. Discovery of multiple-level association rules from large databases. In *Proc. of the 21st VLDB Conference*, Zurich, Switzerland, 1995.
- [HKK97] E.H. Han, G. Karypis, and V. Kumar. Scalable parallel data mining for association rules. In *Proc. of 1997 ACM-SIGMOD Int. Conf. on Management of Data*, Tucson, Arizona, 1997.
- [HKKM97] E.H. Han, G. Karypis, V. Kumar, and B. Mobasher. Clustering based on association rule hypergraphs. Technical Report TR-97-019, Department of Computer Science, University of Minnesota, Minneapolis, 1997.
- [HS95] M. A. W. Houtsma and A. N. Swami. Set-oriented mining for association rules in relational databases. In *Proc. of the 11th Int'l Conf. on Data Eng.*, pages 25–33, Taipei, Taiwan, 1995.
- [KGGK94] Vipin Kumar, Ananth Grama, Anshul Gupta, and George Karypis. *Introduction to Parallel Computing: Algorithm Design and Analysis*. Benjamin Cummings/ Addison Wesley, Redwood City, 1994.
- [MTV95] H. Mannila, H. Toivonen, and A. I. Verkamo. Discovering frequent episodes in sequences. In *Proc. of the First Int'l Conference on Knowledge Discovery and Data Mining*, pages 210–215, Montreal, Quebec, 1995.
- [Pro96] IBM Quest Data Mining Project. Quest synthetic data generation code. <http://www.almaden.ibm.com/cs/quest/syndata.html>, 1996.
- [PS82] C. H. Papadimitriou and K. Steiglitz. *Combinatorial Optimization: Algorithms and Complexity*. Prentice-Hall, Englewood Cliffs, NJ, 1982.
- [SA95] R. Srikant and R. Agrawal. Mining generalized association rules. In *Proc. of the 21st VLDB Conference*, pages 407–419, Zurich, Switzerland, 1995.
- [SA96] R. Srikant and R. Agrawal. Mining sequential patterns: Generalizations and performance improvements. In *Proc. of the Fifth Int'l Conference on Extending Database Technology*, Avignon, France, 1996.
- [SAD⁺93] M. Stonebraker, R. Agrawal, U. Dayal, E. J. Neuhold, and A. Reuter. DBMS research at a crossroads: The vienna update. In *Proc. of the 19th VLDB Conference*, pages 688–692, Dublin, Ireland, 1993.
- [SON95] A. Savasere, E. Omiecinski, and S. Navathe. An efficient algorithm for mining association rules in large databases. In *Proc. of the 21st VLDB Conference*, pages 432–443, Zurich, Switzerland, 1995.

Contents

1	Introduction	1
2	Basic Concepts	2
3	Parallel Algorithms	5
3.1	Count Distribution Algorithm	5
3.2	Data Distribution Algorithm	6
3.3	Intelligent Data Distribution Algorithm	8
3.4	Hybrid Algorithm	10
4	Performance Analysis	13
5	Experimental Results	16
6	Conclusion	20

List of Figures

1	Apriori Algorithm	3
2	Subset operation on the root of a candidate hash tree.	4
3	Subset operation on the left most subtree of the root of a candidate hash tree.	4
4	Count Distribution (CD) Algorithm	5
5	Data Distribution (DD) Algorithm	7
6	Pseudo Code for Data Movements	8
7	Intelligent Data Distribution (IDD) Algorithm	9
8	Subset operation on the root of a candidate hash tree in IDD.	10
9	Hybrid Distribution (HD) Algorithm in 3×4 Processor Mesh ($G = 3, P = 12$)	12
10	Comparison of <i>DD</i> and <i>IDD</i> in terms of the average number of distinct leaf node visited per transaction with 50K transactions per processor and 0.2% minimum support.	16
11	Scaleup result with 100K transactions and 0.25% minimum support.	17
12	Sizeup result with 16 processors and 0.25% minimum support.	18
13	Response time on 16 processors with 50K transactions as the minimum support varies. At each support level, the total number of candidate item-sets is shown in parenthesis	19
14	Response time on 64 processors with 50K transactions as the minimum support varies. At each support level, the total number of candidate item-sets is shown in parenthesis	20
15	Response time on 16 processor IBM SP2 with 100K transactions as the minimum support varies from 0.1% to 0.025%.	21

List of Tables

1	Transactions from supermarket.	2
2	Symbols used in the analysis.	13

3	Minimum support (%) reachable with different number of processors in our algorithms.	20
4	Processor configuration and number of candidates of the <i>HD</i> algorithm with 64 processors and 0.04% minimum support for each pass. Note that 64×1 configuration is the same as the <i>DD</i> algorithm and 1×64 is the same as the <i>CD</i> algorithm. The total number of pass was 13 and all passes after 9 had 1×64 configuration.	20