

# Adaptive and Resource-Aware Mining of Frequent Sets

S. Orlando<sup>1</sup>, P. Palmerini<sup>1,2</sup>, R. Perego<sup>2</sup>, F. Silvestri<sup>2,3</sup>

<sup>1</sup>Dipartimento di Informatica, Università Ca' Foscari, Venezia, Italy

<sup>2</sup>Istituto ISTI, Consiglio Nazionale delle Ricerche (CNR), Pisa, Italy

<sup>3</sup>Dipartimento di Informatica, Università di Pisa, Italy

## Abstract

The performance of an algorithm that mines frequent sets from transactional databases may severely depend on the specific features of the data being analyzed. Moreover, some architectural characteristics of the computational platform used – e.g. the available main memory – can dramatically change its runtime behavior. In this paper we present DCI (Direct Count & Intersect), an efficient algorithm for discovering frequent sets from large databases. Due to the multiple heuristics strategies adopted, DCI can adapt its behavior not only to the features of the specific computing platform, but also to the features of the dataset being mined, so that it results very effective in mining both short and long patterns from sparse and dense datasets. Finally we also discuss the parallelization strategies adopted in the design of ParDCI, a distributed and multi-threaded implementation of DCI.

## 1 Introduction

Association Rule Mining (ARM), one of the most popular topic in the KDD field, regards the extractions of association rules from a database of transactions  $\mathcal{D}$ . Each rule has the form  $X \Rightarrow Y$ , where  $X$  and  $Y$  are sets of items (itemsets), such that  $(X \cap Y) = \emptyset$ . A rule  $X \Rightarrow Y$  holds in  $\mathcal{D}$  with a minimum confidence  $c$  and a minimum support  $s$ , if at least the  $c\%$  of all the transactions containing  $X$  also contains  $Y$ , and  $X \cup Y$  is present in at least the  $s\%$  of all the transactions of the database. In this paper we are interested in the most computationally expensive phase of ARM, i.e. the Frequent Set Counting (FSC) one. During this phase, the set of all the frequent itemsets is built. An itemset of  $k$  items ( $k$ -itemset) is frequent if its support is greater than the fixed threshold  $s$ , i.e. the itemset occurs in at least  $minsup$  transactions ( $minsup = s/100 \cdot n$ , where  $n$  is the number of transactions in  $\mathcal{D}$ ).

The computational complexity of the FSC problem de-

rives from the exponential size of its search space  $\mathcal{P}(M)$ , i.e. the power set of  $M$ , where  $M$  is the set of items contained in the various transactions of  $\mathcal{D}$ . A way to prune  $\mathcal{P}(M)$  is to restrict the search to itemsets whose subsets are all frequent. The *Apriori* algorithm [5] exactly exploits this pruning technique, and visits breadth-first  $\mathcal{P}(M)$  for counting itemset supports. At each iteration  $k$ , *Apriori* generates  $C_k$ , a set of candidate  $k$ -itemsets, and counts the occurrences of these candidates in the transactions. The candidates in  $C_k$  for which the the minimum support constraint holds are then inserted into  $F_k$ , i.e. the set of frequent  $k$ -itemsets, and the next iteration is started. Other algorithms adopt instead a depth-first visit of  $\mathcal{P}(M)$  [8, 2]. In this case the goal is to discover long frequent itemsets first, thus saving the work needed for discovering frequent itemsets included in long ones. Unfortunately, while it is simple to derive all the frequent itemsets from the maximal ones, the same does not hold for their supports, which require a further counting step. In the last years several variations to the original *Apriori* algorithm, as well as many parallel implementations, have been proposed. We can recognize two main methods for determining the supports of the various itemsets present in  $\mathcal{P}(M)$ : a *counting*-based approach [3, 5, 10, 13, 8, 1, 7], and an *intersection*-based one [15, 9, 16]. The former one, also adopted in *Apriori*, exploits a *horizontal* dataset and *counts* how many times each candidate  $k$ -itemset occurs in every transaction. The latter method, on the other hand, exploits a *vertical* dataset, where a *tidlist*, i.e. a list of transaction identifiers (tids), is associated with items (or itemsets), and itemset supports are determined through tidlist intersections. The support of a  $k$ -itemset  $c$  can thus be computed either by a *k-way* intersection, i.e. by intersecting the  $k$  tidlists associated with the  $k$  items included in  $c$ , or by a *2-way* intersection, i.e. by intersecting the tidlists associated with a pair of frequent  $(k - 1)$ -itemsets whose union is equal to  $c$ . Recently another category of methods, i.e. the *pattern growth ones*, have been proposed [1, 11, 14]. **FP-growth** [11] is the best representant of this kind of algorithms. It is not based on

candidate generation as *Apriori*, but builds in memory a compact representation of the dataset, where repeated patterns are represented once along with the associated repetition counters. FP-growth does not perform well on sparse datasets [17], so the same authors recently proposed a new pattern-growth algorithm, H-mine [14], based on an innovative hyper-structure that allows the in-core dataset to be recursively projected by selecting those transactions that include a given pattern prefix.

In this paper we discuss DCI (Direct Count & Intersect), a new algorithm to solve the FSC problem. We also introduce ParDCI, a parallel version of DCI, which explicitly targets clusters of SMPs. Several considerations concerning the features of real datasets to be mined and the characteristics of modern hw/sw system have motivated the design of DCI. On the one hand, transactional databases may have different peculiarities in terms of the correlations among items, so that they may result either dense or sparse. Hence, a desirable characteristic of a new algorithm should be the ability to adapt its behavior to these features. DCI, which supports this kind of adaptiveness, thus constitutes an innovation in the arena of previously proposed FSC algorithms, which often outperformed others only for specific datasets. On the other hand, modern hw/sw systems need high locality for effectively exploiting memory hierarchies and achieving high performances. Large dynamic data structures with pointers may lack in locality due to unstructured memory references. Other sources of performance limitations may be unpredictable branches. DCI tries to take advantages of modern systems by using simple array data structures, accessed by tight loops which exhibit high spatial and temporal locality. In particular, DCI exploits such techniques for intersecting tidlists, which are actually represented as bit-vectors that can be intersected very efficiently with primitive bitwise *and* instructions. Another issue regards I/O operations, which must be carefully optimized in order to allow DM algorithms to efficiently manage large databases. Even if the disk-stored datasets to be mined may be very large, DM algorithms usually access them sequentially with high spatial locality, so that suitable out-of-core techniques to access them can be adopted, also taking advantage of prefetching and caching features of modern OSs [6]. DCI adopts these out-of-core techniques to access large databases, prunes them as execution progresses, and starts using in-core strategies as soon as possible.

Once motivated the design requirements of DCI, we can now detail how it works. As *Apriori*, at each iteration DCI generates the set of candidates  $C_k$ , determines their supports, and produces the set  $F_k$  of the frequent  $k$ -itemsets. However, DCI adopts a hybrid approach to determine the support of the candidates. During its first iterations, DCI exploits a novel counting-based technique, accompanied by an effective pruning of the dataset, stored to disk in horizon-

tal form. During the following iterations, DCI adopts a very efficient intersection-based technique. DCI starts using this technique as soon as the pruned dataset fits into the main memory.

DCI deals with dataset peculiarities by dynamically choosing between distinct *heuristic strategies*. For example, when a dataset is dense, identical sections appearing in several bit-vectors are aggregated and clustered, in order to reduce the number of intersections actually performed. Conversely, when a dataset is sparse, the runs of zero bits in the bit-vectors to be intersected are promptly identified and skipped. We will show how the sequential implementation of DCI significantly outperforms previously proposed algorithms. In particular, under a number of different tests and independently of the dataset peculiarities, DCI results to be faster than *Apriori* and FP-growth. By comparing our experimental results with the published ones obtained on the same sparse dataset, we deduced that DCI is also faster than H-mine [14]. DCI performs very well on both synthetic and real-world datasets characterized by different density features, i.e. datasets from which, due to the different correlations among items, either short or long frequent patterns can be mined.

The rest of the paper is organized as follows. Section 2 describes the DCI algorithm and discusses the various adaptive heuristics adopted, while Section 3 sketches the solutions adopted to design ParDCI, the parallel version of DCI. In Section 4 we report our experimental results. Finally in Section 5 we present some concluding remarks.

## 2 The DCI algorithm

During its initial counting-based phase, DCI exploits an out-of-core, *horizontal* database with variable length records. DCI, by exploiting effective pruning techniques inspired by DHP [13], trims the transaction database as execution progresses. In particular, a pruned dataset  $\mathcal{D}_{k+1}$  is written to disk at each iteration  $k$ , and employed at the next iteration. Let  $m_k$  and  $n_k$  be the number of items and transactions that are included in the pruned dataset  $\mathcal{D}_k$ , where  $m_k \geq m_{k+1}$  and  $n_k \geq n_{k+1}$ . Pruning the dataset may entail a reduction in I/O activity as the algorithm progresses, but the main benefits come from the reduced computation required for subset counting at each iteration  $k$ , due to the reduced number and size of transactions. As soon as the pruned dataset becomes small enough to fit into the main memory, DCI adaptively changes its behavior, builds a *vertical* layout database in-core, and starts adopting an intersection-based approach to determine frequent sets. Note, however, that DCI continues to have a level-wise behavior. At each iteration, DCI generates the candidate set  $C_k$  by finding all the pairs of  $(k-1)$ -itemsets that are included in  $F_{k-1}$  and share a common  $(k-2)$ -

prefix. Since  $F_{k-1}$  is lexicographically ordered, the various pairs occur in close positions, and candidate generation is performed with high spatial and temporal locality. Only during the DCI counting-phase,  $C_k$  is further pruned by checking whether also all the other subsets of a candidate are included in  $F_{k-1}$ . Conversely, during the intersection-based phase, since our intersection method is able to quickly determine the support of a candidate itemset, we found much more profitable to avoid this further check. While during its counting-based phase DCI has to maintain  $C_k$  in main memory to search candidates and increment associated counters, this is no longer needed during the intersection-based phase. As soon as a candidate  $k$ -itemset is generated, DCI determines its support on-the-fly by intersecting the corresponding tidlists. This is an important improvement over other *Apriori*-like algorithms, which suffer from the possible huge memory requirements due to the explosion of the size of  $C_k$  [11].

DCI makes use of a large body of out-of-core techniques, so that it is able to adapt its behavior also to machines with limited main memory. Datasets are read/written in blocks, to take advantage of I/O prefetching and system pipelining [6]. The outputs of the algorithm, e.g. the various frequent sets  $F_k$ , are written to files that are mmap-ped into memory during the next iteration for candidate generation.

## 2.1 Counting-based phase

Since the counting-based approach is used only for few iterations, in the following we only sketch its main features. Further details about DCI counting technique can be found in [12], where we proposed an effective algorithm for mining short patterns. In the first iteration, similarly to all FSC algorithms, DCI exploits a vector of counters, which are directly addressed through item identifiers. For  $k \geq 2$ , instead of using complex data structures like hash-trees or prefix-trees, DCI uses a novel *Direct Count technique* that can be thought as a generalization of the technique used for  $k = 1$ . The technique uses a *prefix table*,  $\text{PREFIX}_k[\ ]$ , of size  $\binom{m_k}{2}$ . In particular, each entry of  $\text{PREFIX}_k[\ ]$  is associated with a distinct *ordered prefix* of two items. For  $k = 2$ ,  $\text{PREFIX}_k[\ ]$  can directly contain the counters associated with the various candidate 2-itemsets, while, for  $k > 2$ , each entry of  $\text{PREFIX}_k[\ ]$  contains the pointer to the contiguous section of ordered candidates in  $C_k$  sharing the same prefix. To permit the various entries of  $\text{PREFIX}_k[\ ]$  to be directly accessed, we devised an order preserving, minimal perfect hash function. This prefix table is thus used to count the support of candidates in  $C_k$  as follows. For each transaction  $t = \{t_1, \dots, t_{|t|}\}$ , we select all the possible 2-prefixes of all  $k$ -subsets included in  $t$ . We then exploit  $\text{PREFIX}_k[\ ]$  to find the sections of  $C_k$  which must be visited in order to check set-inclusion of candidates in transaction  $t$ .

## 2.2 Intersection-based phase

Since the counting-based approach becomes less efficient as  $k$  increases [15], DCI starts its intersection-based phase as soon as possible. Unfortunately, the intersection-based method needs to maintain in memory the vertical representation of the pruned dataset. So, at iteration  $k$ ,  $k \geq 2$ , DCI checks whether the pruned dataset  $\mathcal{D}_k$  may fit into the main memory. When the dataset becomes small enough, its vertical in-core representation is built on the fly, while the transactions are read and counted against  $C_k$ . The intersection-based method thus starts at the next iteration.

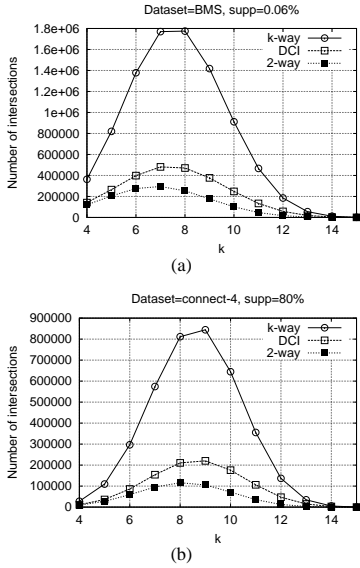
The vertical layout of the dataset is based on fixed length records (tidlists), stored as *bit-vectors*. The whole vertical dataset can thus be seen as a bidimensional bit-array  $\mathcal{VD}[\ ][\ ]$ , whose rows correspond to the bit-vectors associated with non pruned items. Therefore, the amount of memory required to store  $\mathcal{VD}[\ ][\ ]$  is  $m_k \times n_k$  bits.

At each iteration of its intersection-based phase, DCI computes  $F_k$  as follows. For each candidate  $k$ -itemset  $c$ , we *and*-intersect the  $k$  bit-vectors associated with the items included in  $c$  ( $k$ -way intersection), and count the 1's present in the resulting bit-vector. If this number is greater or equal to *minsup*, we insert  $c$  into  $F_k$ . Consider that a bit-vector intersection can be carried out efficiently and with high spatial locality by using primitive bitwise *and* instructions with word operands. As previously stated, this method does not require  $C_k$  to be kept in memory: we can compute the support of each candidate  $c$  on-the-fly, as soon as it is generated.

The strategy above is, in principle, highly inefficient, because it always needs a  $k$ -way intersection to determine the support of each candidate  $c$ . Nevertheless, a caching policy could be exploited in order to save work and speed up our  $k$ -way intersection method. To this end, DCI uses a small “cache” buffer to store all the  $k - 2$  intermediate intersections that have been computed for the last candidate evaluated. Since candidate itemsets are generated in lexicographic order, with high probability two consecutive candidates, e.g.  $c$  and  $c'$ , share a common prefix. Suppose that  $c$  and  $c'$  share a prefix of length  $h \geq 2$ . When we process  $c'$ , we can avoid performing the first  $h - 1$  intersections since their result can be found in the cache.

To evaluate the effectiveness of our caching policy, we counted the actual number of intersections carried out by DCI on two different datasets: BMS, a real-world sparse dataset, and connect-4, a dense dataset (the characteristics of these two datasets are reported in Table 1). We compared this number with the best and the worst case. The best case corresponds to the adoption of a 2-way intersection approach, which is only possible if we can fully cache the tidlists associated with all the frequent  $(k - 1)$ -itemsets in  $F_{k-1}$ . The worst case regards the adoption of a pure  $k$ -way intersection method, i.e. a method that does not exploit

caching at all. Figure 1.(a) plots the results of this analysis on the sparse dataset for support threshold  $s = 0.06\%$ , while Figure 1.(b) regards the dense dataset mined with support threshold  $s = 80\%$ . In both the cases the caching policy of DCI turns out to be very effective, since the actual number of intersections performed results to be very close to the best case. Moreover, DCI requires orders of magnitude less memory than a pure 2-way intersection approach, thus better exploiting memory hierarchies.

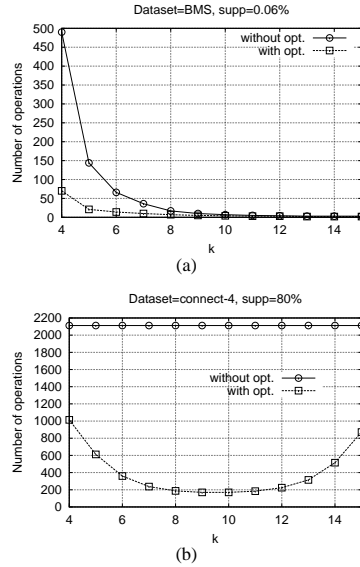


**Figure 1. Evaluation of DCI intersection caching policy.**

We have to consider that while caching reduces the number of tidlist intersections, we also need to reduce intersection cost. To this end, further heuristics, differentiated w.r.t. *sparse* or *dense* datasets, are adopted by DCI. In order to apply the right optimization, the vertical dataset is tested for checking its density as soon as it is built. In particular, we compare the bit-vectors associated with the *most frequent items*, i.e., the vectors which likely need to be intersected several times since the associated items occur in many candidates. If large sections of these bit-vectors turn out to be identical, we deduce that the items are highly correlated and that the dataset is dense. In this case we adopt a specific heuristics which exploits similarities between these vectors. Otherwise the technique for sparse datasets is adopted. In the following we illustrate the two heuristics in more detail.

**Sparse datasets.** Sparse or moderately dense datasets originate bit-vectors containing long runs of 0's. To speedup computation, while we compute the intersection of the bit-vectors relative to the first two items  $c_1$  and  $c_2$  of a generic candidate itemset  $c = \{c_1, c_2, \dots, c_k\} \in C_k$ , we also identify and maintain information about the runs of 0's appearing in the resulting bit-vector stored in cache. The

further intersections that are needed to determine the support of  $c$  (as well as intersections needed to process other  $k$ -itemsets sharing the same 2-item prefix) will skip these runs of 0's, so that only vector segments which may contain 1's are actually intersected. Since information about the runs of 0's are computed once, and the same information is reused many times, this optimization results to be very effective. Moreover, sparse and moderately dense datasets offer the possibility of further pruning vertical datasets as computation progresses. The benefits of pruning regard the reduction in the length of the bit-vectors and thus in the cost of intersections. Note that a transaction, i.e. a column of  $\mathcal{VD}$ , can be removed from the vertical dataset when it does not contain any of the itemsets included in  $F_k$ . This check can simply be done by *or*-ing the intersection bit-vectors computed for all the frequent  $k$ -itemsets. However, we observed that dataset pruning is expensive, since vectors must be compacted at the level of single bits. Hence DCI prunes the dataset only if turns out to be profitable, i.e. if we can obtain a large reduction in the vector length, and the number of vectors to be compacted is small with respect to the cardinality of  $C_k$ .



**Figure 2. Evaluation of DCI optimization heuristics for sparse and dense datasets.**

**Dense datasets.** If the dataset turns out to be dense, we expect to deal with a dataset characterized by strong correlations among the most frequent items. This not only means that the bit-vectors associated with the *most frequent items* contain long runs of 1's, but also that they turn out to be very similar. The heuristic technique adopted by DCI for dense dataset thus works as follows: A) we reorder the columns of the vertical dataset, in order to move identical segments of the bit-vectors associated with the most frequent items to the first consecutive positions; B) since each

**Table 1. Datasets used in the experiments.**

Dataset	Description
T25I10D10K	Synthetic dataset with 1K items and 10K transactions [11]. The average size of transactions is 25, and the average size of the maximal potentially frequent itemsets is 10.
T25I20D100K	Synthetic dataset with 10K items and 100K transactions [11]. The average size of transactions is 25, and the average size of the maximal potentially frequent itemsets is 20.
400k.t10.p8.m10k	10K items and 400K transactions. The average size of transactions is 10, and the average size of the maximal potentially frequent itemsets is 8. Synthetic dataset created with the IBM dataset generator [5].
400k.t30.p16.m1k	1K items and 400K transactions. The average size of transactions is 30, and the average size of the maximal potentially frequent itemsets is 16. Synthetic dataset created with the IBM dataset generator [5].
t20.p8.m1k	With this notation we identify a series of synthetic datasets characterized by 1K items. The average transaction size is 20, and the average size of maximal potentially frequent itemsets is 8. The number of transactions is varied for scaling measurements.
t50.p32.m1k	A series of three synthetic datasets with the same number of items (1K), average transaction size of 50, and average size of maximal potentially frequent itemsets equal to 32. We used three datasets of this series with 1000k, 2000k and 3000k transactions.
connect-4	Publicly available dense dataset with 130 items and about 60K transactions. The maximal transaction size is 45.
BMS	Publicly available sparse dataset also known as <i>Gazelle</i> . 497 items and 59K transactions containing click-stream data from an e-commerce web site <i>gazelle.com</i> .

candidate is likely to include several of these most frequent items, we avoid repeatedly intersecting the identical segments of the corresponding vectors. This technique may save a lot of work because (1) the intersection of identical vector segments is done once, (2) the identical segments are usually very large, and (3), long candidate itemsets presumably contains several of these most frequent items.

The plots reported in Figure 2 show the effectiveness of the heuristic optimizations discussed above in reducing the average number of bitwise *and* operations needed to intersect a pair of bit-vectors. In particular, Figure 2.(a) regards the *sparse* BMS dataset mined with support threshold  $s = 0.06\%$ , while Figure 2.(b) regards the *dense* dataset connect-4 mined with support threshold  $s = 80\%$ . In both cases, we plotted the per-iteration cost of each bit-vector intersection in terms of bitwise *and* operations when either our heuristic optimizations are adopted or not. The two plots show that our optimizations for both sparse and dense datasets have the effect of reducing the intersection cost up to an order of magnitude. Note that when no optimizations are employed, the curves exactly plot the bit-vector length (in words). Finally, from the plot reported in Figure 2.(a), we can also note the effect of the pruning technique used on sparse datasets. Pruning has the effect of reducing the length of the bit-vectors as execution progresses. On the other hand, when datasets are dense, the vertical dataset is not pruned, so that the length of bit-vectors remains the same for all the DCI iterations.

### 3 ParDCI

In the following we describe the different parallelization techniques exploited for the counting- and intersection-based phases of ParDCI, the parallel version of DCI. Since

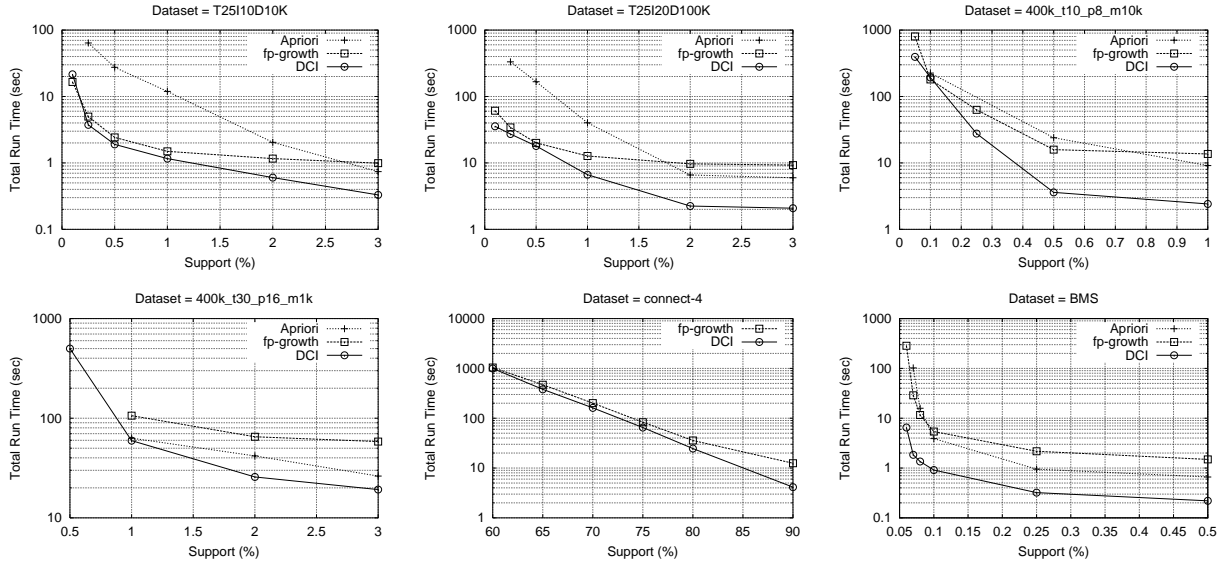
our target architecture is a cluster of SMP nodes, in both phases we distinguish between *intra-node* and *inter-node* levels of parallelism. At the inter-node level we used the message-passing paradigm through the MPI communication library, while at the intra-node level we exploited multi-threading through the *Posix Thread* library. A *Count Distribution* approach is adopted to parallelize the counting-based phase, while the intersection-based phase exploits a very effective *Candidate Distribution* approach [4].

**The counting-based phase.** At the inter-node level, the dataset is statically split in a number of partitions equal to the number of SMP nodes available. The size of partitions depend on the relative powers of nodes. At each iteration  $k$ , an identical copy of  $C_k$  is independently generated by each node. Then each node  $p$  reads blocks of transactions from its own dataset partition  $\mathcal{D}_{p,k}$ , performs subset counting, and writes pruned transactions to  $\mathcal{D}_{p,k+1}$ . At the end of the iteration, an all-reduce operation is performed to update the counters associated to all candidates of  $C_k$ , and all the nodes produce an identical set  $F_k$ .

At the intra-node level each node uses a pool of threads, each holding a private set of counters associated with candidates. They have the task of checking in parallel candidate itemsets against chunks of transactions read from  $\mathcal{D}_{p,k}$ .

At the end of each iteration, a global reduction of counters take place, and a copy of  $F_k$  is produced on each node.

**The intersection-based phase.** During the intersection-based phase, a Candidate Distribution approach is adopted at both the inter- and intra-node levels. This parallelization schema makes the parallel nodes completely independent: inter-node communications are no longer needed for all the following iterations of ParDCI.



**Figure 3. Total execution times for DCI, *Apriori*, and FP-growth on various datasets as a function of the support threshold.**

Let us first consider the inter-node level, and suppose that the intersection-based phase is started at iteration  $\bar{k} + 1$ . Therefore, at iteration  $\bar{k}$  the various nodes build on-the-fly the bit-vectors representing their own in-core portions of the vertical dataset. Before starting the intersection-based phase, the partial vertical datasets are broadcast to obtain a complete replication of the whole vertical dataset on each node. The frequent set  $F_{\bar{k}}$  (i.e., the set computed in the last counting-based iteration) is then partitioned on the basis of itemset prefixes. A disjoint partition  $F_{p,\bar{k}}$  of  $F_{\bar{k}}$  is thus assigned to each node  $p$ , where  $\bigcup_p F_{p,\bar{k}} = F_{\bar{k}}$ . It is worth remarking that this partitioning entails a Candidate Distribution schema for all the following iterations, according to which each node  $p$  will be able to generate a unique  $C_k^p$  ( $k > \bar{k}$ ) independently of all the other nodes, where  $C_k^p \cap C_k^{p'} = \emptyset$  if  $p \neq p'$ , and  $\bigcup_p C_k^p = C_k$ .

At the intra-node level, a similar Candidate Distribution approach is employed, but at a finer granularity by using dynamic scheduling to ensure load balancing.

## 4 Experimental Results

The DCI algorithm is currently available in two versions, a MS-Windows one, and a Linux one. ParDCI, which exploits the MPICH MPI and the *pthread* libraries, is currently available only for the Linux platform. We used the MS-Windows version of DCI to compare its performance with other FSC algorithms. For test comparisons we used the

FP-growth algorithm<sup>1</sup>, and the Christian Borgelt's implementation of *Apriori*<sup>2</sup>. For the sequential tests we used a Windows-NT workstation equipped with a Pentium II 350 MHz processor, 256 MB of RAM memory and a SCSI-2 disk. For testing ParDCI performance, we employed a small cluster of three Pentium II 233MHz 2-way SMPs, for a total of six processors. Each SMP is equipped with 256 MBytes of main memory and a SCSI disk. For the tests, we used both synthetic and real datasets by varying the minimum support threshold  $s$ . The characteristics of the datasets used are reported in Table 1.

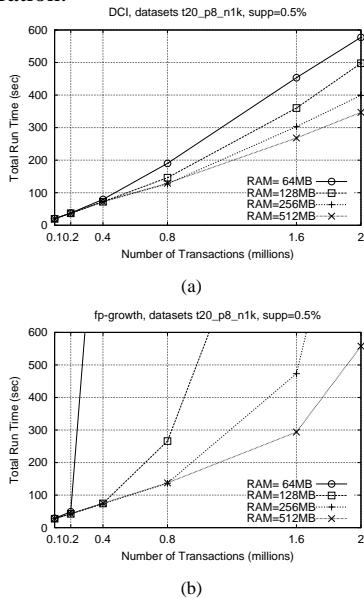
**DCI performances and comparisons.** Figure 3 reports the total execution times obtained running *Apriori*, FP-growth, and our sequential DCI algorithm on some datasets described in Table 1 as a function of the support threshold  $s$ . In all the tests conducted, DCI outperforms FP-growth with speedups up to 8. Of course, DCI also remarkably outperforms *Apriori*, in some cases for more than one order of magnitude. For connect-4, the dense dataset, the curve of *Apriori* is not shown, due to the relatively too long execution times. Note that, accordingly to [17], on the real-world sparse dataset BMS (also known as *Gazelle*), *Apriori* turned out to be faster than FP-growth. To overcome such bad performance results on sparse datasets, the same authors of FP-growth recently proposed a new pattern-growth algorithm, H-mine [14]. By comparing our experimental results with

<sup>1</sup>We acknowledge Prof. Jiawei Han for kindly providing us the latest, fully optimized, binary version of FP-growth.

<sup>2</sup><http://fuzzy.cs.uni-magdeburg.de/~borgelt>

the published execution times on the BMS dataset, we deduced that DCI is also faster than H-mine. For  $s = 0.06\%$ , we obtained an execution time of about 7 sec., while H-mine completes in about 40 sec. on a faster machine.

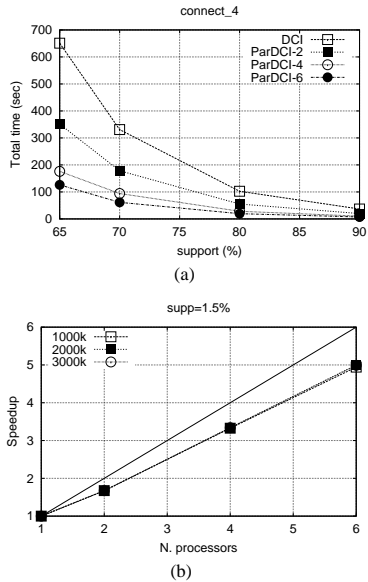
The encouraging results obtained with DCI are due to both the efficiency of the counting method exploited during early iterations, and the effectiveness of the intersection-based approach used when the pruned vertical dataset fits into the main memory. For only a dataset, namely T25I10D10K, FP-growth turns out to be slightly faster than DCI for  $s = 0.1\%$ . The cause of this behavior is the size of  $C_3$ , which in this specific case results much larger than the final size of  $F_3$ . Hence, DCI has to carry out a lot of useless work to determine the support of many candidate itemsets, which will eventually result to be not frequent. In this case FP-growth is faster than DCI since it does not require candidate generation.



**Figure 4. Total execution times of (a) DCI, and (b) FP-growth, on datasets in the series t20.p8.m1k ( $s = 0.5\%$ ) on a PC equipped with different RAM sizes as a function of the number of transactions (ranging from 100K to 2M).**

We also tested the scale-up behavior of DCI when both the size of the dataset and the size of RAM installed in the PC vary. The datasets employed for these tests belong to the series t20.p8.m1k (see Table 1) mined with support threshold  $s = 0.5\%$ , while the available RAM was changed from 64MB to 512MB by physically plugging additional memory into the PC main board. Figure 4.(a) and 4.(b) plot several curves representing the execution times of DCI and FP-growth, respectively, as a function of the number of transactions contained in the dataset processed. Each curve plotted refers to a series of tests conducted with the same

PC equipped with a different amount of memory. As it can be seen from Figure 4.(a), DCI scales linearly also on machines with a few memory. Due to its adaptiveness and the use of efficient out-of-core techniques, it is able to modify its behavior in function of the features of the dataset mined and the computational resources available. For example, in the tests conducted with the largest dataset containing two millions of transactions, the in-core intersection-based phase was started at the sixth iteration when only 64MB of RAM were available, and at the third iteration when the available memory was 512MB. On the other hand the results reported in Figure 4.(b) show that FP-growth requires much more memory than DCI, and is not able to adapt itself to memory availability. For example, in the tests conducted with 64MB of RAM, FP-growth requires less than 30 seconds to mine the dataset with 200k transactions, but when we double the size of the dataset to 400k transactions, FP-growth execution time becomes 1303 seconds, more than 40 times higher, due to a heavy page swapping activity.



**Figure 5. (a): Dense dataset connect-4: completion times of DCI and ParDCI varying the minimum support threshold. (b): Speedup for sparse datasets 1000K, 2000K and 3000K with  $s = 1.5\%$**

**Performance evaluation of ParDCI.** We evaluated ParDCI on both dense and sparse datasets. First we compared the performance of DCI and ParDCI on the dense dataset connect-4, for which we obtained very good speedups. Figure 5.(a) plots total execution times as functions of the support thresholds  $s$ . ParDCI-2 corresponds to the pure multithread version running on a single 2-way SMP, while ParDCI-4 and ParDCI-6 also exploit inter-node parallelism, and run, respectively, on two and three 2-way SMPs. For what regard sparse datasets, we used

the synthetic dataset series identified as *t50\_p32\_m1k* in Table 1. We varied the total number of transactions from 1000k to 3000k. In the following we will identify the various synthetic datasets on the basis of their number of transactions, i.e. 1000k, 2000k, and 3000k. Figure 5.(b) plots the speedups obtained on the three synthetic datasets for a given support threshold ( $s = 1.5\%$ ), as a function of the number of processors used. Consider that, since our cluster is composed of three 2-way SMPs, we mapped tasks on processors always using the minimum number of SPMP nodes (e.g., when we used 4 processors, we actually employed 2 SMP nodes). This implies that experiments performed on either 1 or 2 processors actually have identical memory and disk resources available, whereas the execution on 4 processors benefit from a double amount of such resources. According to the tests above, ParDCI showed a speedup that, in some cases, is close to the optimal one. Considering the results obtained with one or two processors, one can note that the slope of the speedup curve is relatively worse than its theoretical limit, due to resource sharing and thread implementation overheads at the inter-node level. Nevertheless, when additional SMPs are employed, the slope of the curve improves. The strategies adopted for partitioning dataset and candidates on our homogeneous cluster of SMPs sufficed for balancing the workload. In our tests we observed a very limited imbalance. The differences in the execution times of the first and last node to end execution were always below the 0.5%.

## 5 Conclusions

DCI uses different approaches for extracting frequent patterns: counting-based during the first iterations and intersection-based for the following ones. Adaptiveness and resource awareness are the main innovative features of the algorithm. On the basis of the characteristics of the dataset mined, DCI chooses at run-time which optimization to adopt for reducing the cost of mining. Dataset pruning and effective out-of-core techniques are exploited during the counting-based phase, while the intersection-based phase works in core, and is started only when the pruned dataset can fit into the main memory. As a result, our algorithm can manage efficiently, also on machines with limited physical memory, very large datasets from which, due to the different correlations among items, either short or long frequent patterns can be mined.

The experimental evaluations demonstrated that DCI significantly outperforms *Apriori* and *FP-growth* on both synthetic and real-world datasets. In many cases the performance improvements are impressive. Moreover, ParDCI, the parallel version of DCI, exhibits excellent scaleups and speedups on our homogeneous cluster of SMPs. The variety of datasets used and the large amount of tests con-

ducted permit us to state that the performances of DCI are not influenced by dataset characteristics, and that our optimizations are very effective and general. To share our efforts with the data mining community, we made the DCI binary code available for research purposes at <http://www.miles.cnuce.cnr.it/~palmeri/datam/DCI>.

## References

- [1] R. C. Agarwal, C. C. Aggarwal, and V.V.V. Prasad. A Tree Projection Algorithm for Generation of Frequent Itemsets. *JPDC*, 2000. Special Issue on High Performance Data Mining.
- [2] R. C. Agarwal, C. C. Aggarwal, and V.V.V. Prasad. Depth first generation of long patterns. In *Proc. of the 6th ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, pages 108–118, 2000.
- [3] R. Agrawal, H. Mannila, R. Srikant, H. Toivonen, and A. Inkeri Verkamo. Fast Discovery of Association Rules in Large Databases. In *Advances in Knowledge Discovery and Data Mining*, pages 307–328. AAAI Press, 1996.
- [4] R. Agrawal and J. C. Shafer. Parallel mining of association rules. *IEEE TKDE*, 8:962–969, 1996.
- [5] R. Agrawal and R. Srikant. Fast Algorithms for Mining Association Rules in Large Databases. In *Proc. of the 20th VLDB Conf.*, pages 487–499, 1994.
- [6] R. Baraglia, D. Laforenza, S. Orlando, P. Palmerini, and R. Perego. Implementation issues in the design of I/O intensive data mining applications on clusters of workstations. In *Proc. of the 3rd HPDM Workshop, (IPDPS-2000), Cancun, Mexico*, pages 350–357. LNCS 1800 Springer-Verlag, 2000.
- [7] Y. Bastide, R. Taouil, N. Pasquier, G. Stumme, and L. Lakhal. Mining frequent patterns with counting inference. *ACM SIGKDD Explorations Newsletter*, 2(2):66–75, December 2000.
- [8] R. J. Bayardo Jr. Efficiently Mining Long Patterns from Databases. In *Proc. of the ACM SIGMOD Int. Conf. on Management of Data*, pages 85–93, Seattle, Washington, USA, 1998.
- [9] Brian Dunkel and Nandit Soparkar. Data organization and access for efficient data mining. In *Proc. of the 15th Int. Conf. on Data Engineering*, pages 522–529, Sydney, Australia, 1999. IEEE Computer Society.
- [10] E. H. Han, G. Karypis, and Kumar V. Scalable Parallel Data Mining for Association Rules. *IEEE TKDE*, 12(3):337–352, May/June 2000.
- [11] J. Han, J. Pei, and Y. Yin. Mining Frequent Patterns without Candidate Generation. In *Proc. of the ACM SIGMOD Int. Conf. on Management of Data*, pages 1–12, Dallas, Texas, USA, 2000.
- [12] S. Orlando, P. Palmerini, and R. Perego. Enhancing the Apriori Algorithm for Frequent Set Counting. In *Proc. of the 3<sup>rd</sup> Int. Conf. on Data Warehousing and Knowledge Discovery, DaWaK 2001, LNCS 2114*, pages 71–82, Munich, Germany, 2001.
- [13] J. S. Park, M.-S. Chen, and P. S. Yu. An Effective Hash Based Algorithm for Mining Association Rules. In *Proc. of the 1995 ACM SIGMOD Int. Conf. on Management of Data*, pages 175–186, 1995.
- [14] J. Pei, J. Han, H. Lu, S. Nishio, and D. Tang, S. amd Yang. H-Mine: Hyper-Structure Mining of Frequent Patterns in Large Databases. In *Proc. of the 2001 IEEE ICDM Conf.*, San Jose, CA, USA, 2001.
- [15] A. Savasere, E. Omiecinski, and S. B. Navathe. An Efficient Algorithm for Mining Association Rules in Large Databases. In *Proc. of the 21th VLDB Conf.*, pages 432–444, Zurich, Switzerland, 1995.
- [16] M. J. Zaki. Scalable algorithms for association mining. *IEEE TKDE*, 12:372–390, May/June 2000.
- [17] Z. Zheng, R. Kohavi, and L. Mason. Real World Performance of Association Rule Algorithms. In *Proc. of KDD-2001*, 2001.