Intelligent Data Analysis 11 (2007) 49–73
IOS Press

Approximate mining of frequent patterns on streams

Claudio Silvestri and Salvatore Orlando
Dipartimento di Informatica, Università Ca' Foscari, Via Torino 155, Venezia, Italy
E-mail: {silvestri, orlando}@dsi.unive.it

Abstract. Many critical applications, like intrusion detection or stock market analysis, require a nearly immediate result based on a continuous and infinite stream of data. In most cases finding an exact solution is not compatible with limited availability of resources and real time constraints, but an approximation of the exact result is enough for most purposes.

This paper introduces a new algorithm for approximate mining of frequent itemsets from streams of transactions using a limited amount of memory. The proposed algorithm is based on the computation of frequent itemsets in recent data and an effective method for inferring the global support of previously infrequent itemsets. Both upper and lower bounds on the support of each pattern found are returned along with the interpolated support. An extensive experimental evaluation shows that $\mathtt{AP}_{\mathtt{stream}}$, the proposed algorithm, yields a good approximation of the exact global result considering both the set of patterns found and their supports.

1. Introduction

Association Rule Mining (ARM), one of the most popular topic in the KDD field [2,14], regards the extraction association rules from a database of transactions $\mathcal{D}$. In this paper we are interested in the most computationally expensive phase of ARM, i.e., the Frequent Itemset Mining (*FIM*) one, during which the set $\mathcal{F}$ of all the itemsets that occur in at least a user specified number of transactions is discovered. Those itemsets are named *Frequent Itemsets*.

The computational complexity of the FIM problem derives from the exponential size of its search space $\mathcal{P}(\mathcal{I})$, i.e. the power set of $\mathcal{I}$, where $\mathcal{I}$ is the set of items contained in the various transactions of $\mathcal{D}$. A way to prune $\mathcal{P}(\mathcal{I})$ is to restrict the search to itemsets whose subsets are all frequent. The `Apriori` algorithm [2], and other derived algorithms for non dynamic datasets, exactly exploits this pruning technique, based on the Apriori anti-monotonic principle.

In a stream setting, new transactions are continuously added to the dataset. The infinite nature of stream data sources is a serious obstacle to the use of most of the traditional methods, since available computing resources are limited, whereas the amount of previously happened events is usually overwhelming. Thus, one of the first effects is the need to process data as they arrive, due to the impossibility of storing them. The results extracted evolve continuously along with data. In our case, since we adopt a landmark window model, these results refer to the whole data stream arrived so far, from a given past time (when we started collecting data) to the current time.

Obviously, an algorithm suitable for stream data should be able to compute the 'next step' solution on-line, starting from the previously known one and the current data, if necessary with some additional information stored along with the past solution. In our case, this information is the count of a significant

1088-467X/07/$17.00 © 2007 – IOS Press and the authors. All rights reserved

part of frequent single items, and a transaction hash table used for improving deterministic bounds on supports returned by the algorithm.

Unfortunately, even the apparently simple discovery of frequent items in a stream is challenging [5]. Some items, initially frequent, may eventually become infrequent. On the other hand, other items may appear initially in a sporadic way and then become frequent. Thus the only way to exactly compute the support of these items is to maintain a counter since the first appearance of each of them. This could be acceptable when the number of distinct items is reasonably bounded. If the stream contains a large and potentially unbounded number of spurious items, as in case of data with probabilities of occurrence that follow a Zipf's law, like internet traffic data, this approach may lead to a huge waste of memory.

In this paper we discuss a *stream* algorithm for approximate mining of frequent itemsets, $\texttt{AP}_{\text{stream}}$ (Approximate Partition for Stream), which exploits DCI [27], a state-of-the-art algorithm for FIM, as the miner engine for recent data. The $\texttt{AP}_{\text{stream}}$ algorithm uses techniques similar to those that we have already exploited in $\texttt{AP}_{\text{Interp}}$ [34], an algorithm for approximate distributed mining of frequent itemsets. Both $\texttt{AP}_{\text{stream}}$ and $\texttt{AP}_{\text{Interp}}$ use a computation method inspired by the Partition algorithm [31].

Partition relies on a horizontally partitioned dataset, and consists in independently computing *local* results from each partition, merging the local sets of frequent itemsets, and then recounting each potentially frequent pattern over the whole dataset to discover the *global* results. In order to extend this approach to a stream setting, blocks of data received from the stream are used as an infinite set of partitions.

Others stream association mining algorithms, such as LOSSY COUNT [24] for frequent itemsets, use a similar approach with some variation. Obviously, all of them avoid recounting potentially frequent itemsets over the whole dataset, which is not feasible with streaming data. $\texttt{AP}_{\text{stream}}$ applies the same heuristic used by the previously introduced $\texttt{AP}_{\text{Interp}}$ algorithm. The infinite flow of data block in the stream is processed pairwise, using past processed data and recent data as two partitions. Upon new data arrival, as many transactions as possible are buffered and processed in-core. The amount of buffered transactions obviously depends on their lengths, but also on the size of main memory available. The past approximate solution is then merged with the frequent pattern set obtained from recent data.

Since a second pass on the whole stream is impossible, we use an approximate support inference heuristic during the merge phase in order to improve the support accuracy. Along with each interpolated support value, this method yields a pair of deterministic upper and lower bounds. The proposed inference heuristic can be easily replaced with a different one, more complex or better fitting a particular application context. In particular, our method is based on a simple, yet effective, interpolation schema based on the knowledge of the supports of the sub-patterns of a given infrequent pattern. Despite its simplicity, it entails good approximation results in experimental evaluation. So we expect that, for specific application contexts, a more focused inference method also based on domain knowledge would yield even better results.

In data streams, the underlying data distribution may change. Hence the models built on old data might become inaccurate. This problem, known as concept drift, complicates the task of interpolating the count of past occurrences of a given pattern. The method we propose is in some way concept drift resilient, in particular when the drift concerns only the single item probability distributions and not the joint distributions. In Section 7, we propose an extension able to deal with this issue in more challenging cases.

This paper is organized as follows. Section 2 formally introduces the FIM problem on streams. Then Section 3 describes the $\texttt{AP}_{\text{stream}}$ algorithm, and the Partition algorithm that inspired $\texttt{AP}_{\text{Interp}}$ and $\texttt{AP}_{\text{stream}}$. Before presenting and discussing our experimental results in Section 5, we introduce, in

Section 5.1, some similarity measures that we use in order to evaluate the quality of the approximate results. Section 6 surveys the main related works in the field. Finally, in Section 7 we discuss some interesting extensions of the proposed method, and, in Section 8, we draw some conclusions.

2. The problem

In this section we formally define the FIM problem in both non-evolving databases and stream ones.

Definition 1. (TRANSACTION DATASET) Let $\mathcal{I} = \{i_1, \ldots, i_m\}$ be a set of items. A non-evolving transaction dataset $\mathcal{D}$ is a collection of *input sets* or *transactions*:

$$\mathcal{D} = \{\bar{t} | \, \bar{t} = (tid, t)\},$$

where *tid* is a transaction identifier, and $t = \{i_1, \ldots, i_k\} \subseteq \mathcal{I}$ is a set of distinct items. The size of $\mathcal{D}$ is the number n of transactions contained in $\mathcal{D}$, i.e., $n = |\mathcal{D}|$.

The support of an itemset is a measure of its interestingness as a pattern, and is based on its frequency.

Definition 2. (SUPPORT OF AN ITEMSET) Let $p \subseteq \mathcal{I}$ be an itemset. The *support* $\sigma(p)$ of itemset p in dataset $\mathcal{D}$ is defined as

$$\sigma(p) = |\{(tid, t) \in \mathcal{D} \mid p \subseteq t\}|$$

i.e., the number of transactions in $\mathcal{D}$ that contain pattern p. The relative support *sup*(*p*) of pattern p is instead expressed as a fraction of the total number of transactions:

$$sup(p) = \frac{\sigma(p)}{|\mathcal{D}|}$$

Even if a transaction represents a set of items, with no particular order, it is convenient to assume that there exists some kind of total order R among them. Such order makes unequivocal the way in which an itemset is written, e.g., if we adopt an alphanumeric order we cannot write $\{B, A\}$ since the correct way is $\{A, B\}$.

Definition 3. (FREQUENT ITEMSET MINING) Let $minsup$ be a user chosen threshold. An itemset p is frequent in $\mathcal{D}$ if its support $\sigma(p)$ is not less than $\sigma_{\min} = \mathrm{minsup} \cdot |\mathcal{D}|$, i.e., if $sup(p) \geqslant \mathrm{minsup}$. A k-itemset is a pattern composed of k items, $\mathcal{F}_k$ is the set of all frequent k-itemsets, and $\mathcal{F}$ is the set of all frequent itemsets.

The *Frequent Itemset Mining* (*FIM*) problem consists in discovering $\mathcal{F}$ in $\mathcal{D}$.

In a stream setting, since new transactions are continuously added to the dataset, we need a notation for indicating that a particular dataset or result refers to a particular part of the stream. To this end, we write the interval as a subscript after the entity.

Definition 4. (TRANSACTION STREAM DATASET) Let $\mathcal{I} = \{i_1, \ldots, i_m\}$ be a set of items. A transaction data stream $\mathcal{D}$ is an infinite sequence of *input sets* or *transactions*:

$$\mathcal{D} = \{\bar{t} | \, \bar{t} = (bid, tid, t)\}$$

where $t = \{i_1, \ldots, i_k\} \subseteq \mathcal{I}$ is a set of distinct items, while *tid* and *bid* are monotonically increasing identifiers, which are respectively associated with single transactions and blocks. The block identifier *bid* is chosen at reception time. In particular, all the transactions labeled with the same $bid = i$ arrived before all the transactions labeled with $bid = i + 1$. The transactions in the i^{th} *block*, denoted as $\mathcal{D}_i$, are processed at the same time. The notation $\mathcal{D}_{[i,j)}$, $i < j$, identifies the part of the stream containing only the transactions whose *bid*s are included in the interval $[i, j)$, i.e., $i \leqslant bid < j$.

Thus $\mathcal{D}_{[1,j]}$ denotes the part of the stream from the starting block until block j. If the j^{th} block is the current one, and the notation is not ambiguous, we will just write $\mathcal{D}$ instead of $\mathcal{D}_{[1,j]}$.

Due do the continuous evolution of stream datasets, a solution to the FIM problem must be tied to a part of the stream, indicated as a block interval $\mathcal{D}_{[i,j]}$. Depending on the part of stream involved, the problem presents different challenges, and is named differently. In particular, the *landmark model* [22, 24] considers the entire stream, the *sliding window model* [7] refers to its most recent part, and, finally, the *tilted-time window model* [17], is obtained by composing several distinct sliding windows, in order to maintain multiple time-granularities. In this paper we discuss an algorithm for the solution of the FIM in the landmark model. We formally introduce this problem in the following.

Definition 5. (FREQUENT ITEMSET MINING IN DATA STREAMS) Let *minsup* be a user chosen threshold. An itemset *p* is frequent in $\mathcal{D}_{[1,i]}$ if its support $\sigma_{[1,i]}(p)$ is not less than $\sigma_{\min_{[1,i]}} = \text{minsup} \cdot |\mathcal{D}_{[1,i]}|$. A *k-itemset* is a pattern composed of k items, $\mathcal{F}_{k_{[1,i]}}$ is the set of all frequent k-itemsets, and $\mathcal{F}_{[1,i]}$ is the set of all frequent itemsets.

The problem of *Frequent Itemset Mining* (*FIM*) *in Data Streams* consists in discovering $\mathcal{F}_{[1,i]}$ in $\mathcal{D}_{[1,i]}$, for increasing values of i.

3. The `Partition` algorithm and its extensions

The $\texttt{AP}_{\text{stream}}$ (Approximate Partition for Stream) algorithm uses a technique similar to the one that we have introduced with our algorithm $\texttt{AP}_{\text{Interp}}$ [34] for approximate mining of frequent itemsets in a distributed setting. Both algorithms are inspired by `Partition` [31], a sequential algorithm which divides the dataset into several partitions processed independently, and then merges the *local* solutions to produce *the* global result. In this paper we will also use the terms local and global, as referred to stream input data or associated results. Local indicates something just concerning a contiguous part of the stream, hereinafter called a block of transactions, whereas global indicates something pertaining to the whole stream seen so far.

In this section we will describe the `Partition` algorithm and its naïve distributed and streaming versions, which we have used as a starting point for designing our approximate algorithms.

3.1. The original `Partition` algorithm

The basic idea exploited by `Partition` is the following: if the dataset is divided into several partitions, then each *globally* frequent itemset must be *locally* frequent in at least one partition. This guarantees that the union of all local solutions is a superset of the global solution. `Partition` sequentially reads the dataset, one partition at a time. For each partition it extracts the locally frequent itemset, and adds them to a set of potential globally frequent itemsets. After this phase, the result set contains every globally frequent itemset, mixed with several infrequent ones (*false positives*). Thus the dataset is read again, counting the exact occurrences of each candidate pattern, i.e., the ones that turned out to be frequent in only a proper subset of all the dataset partitions. At the end of the second scan all the infrequent patterns are removed, so that the result set only contains the FIM problem solution.

3.2. *The* Distributed *algorithm*

Obviously, Partition can be straightforwardly implemented in a distributed setting with a master/slave paradigm [26]. Each slave becomes responsible of a local partition, while the master performs the sum-reduction of local counters (first phase) and orchestrates the slaves for computing the missing local supports for potential globally frequent patterns (second phase) to remove patterns having global support less than minsup (false positive patterns collected during the first phase).

While the Distributed algorithm gives the exact values for supports, it has pros and cons with respect to other distributed algorithms. The *pros* are related to the number of communications/synchronizations: other methods like count-distribution [19,38] require several communications/synchronizations, while the Distributed algorithm only requires two communications from the slaves to the master, a single message from the master to the slaves and synchronization after the first scan. The *cons* concern the size of the messages exchanged, and the possible additional computation performed by the slaves when the first phase of the algorithm produces many false positives. Consider that, when low absolute minimum supports are used, it is likely to produce a lot of false positives due to data skew present in the various dataset partitions [30]. This has a large impact also on the cost of the second phase of the algorithm: most of the slaves will participate in counting the local supports of these false positives, thus wasting a lot of time.

A naïve technique to work around this problem is to stop Distributed after the first-pass. We call the algorithm that adopts this simple technique Distributed One-pass Partition. So, in Distributed One-pass Partition, each slave independently computes locally frequent patterns and sends them to the master which sum-reduces the support for each pattern, and writes in the result set only the patterns having the sum of the known supports not less than $minsup \cdot |\mathcal{D}|$. Distributed One-pass Partition has obvious performance advantages over Distributed. On the other hand, it yields a result which may be approximate, since it is possible that some globally frequent pattern occur in a partition where it resulted to be locally infrequent, so that its local support count is unknown. In several cases this may cause the erroneous omission of globally frequent patterns. However, Distributed One-pass Partition ensures that at least the number of occurrences reported for each returned pattern exists.

3.3. *The* Stream Partition *algorithm*

The infinite sequence of blocks of data that arrive from the data stream can be considered as an infinite set of partitions. This allows us to adopt the Distributed approach also in a stream setting. Since the stream is infinite, however, it is impossible to collect and merge the results obtained from the various blocks. Thus the partial results must be merged repeatedly, and each time the result set needs to be updated. A block of data is processed as soon as "enough" transactions are available, and the local result set of the current block is merged with the previous approximate result set, which refers to the past part of the stream. Unfortunately, due to memory constraints, in the stream case only recent raw data – i.e., the last block of transactions – can be maintained available for processing. Thus, in this case we can perform a second scan of them to check the support count of frequent patterns that resulted to be frequent in the past, but that are locally infrequent in the current block.

Only the partial results extracted so far from previous blocks of the stream, plus some other additional information, can be available for determining the global result set, i.e. the frequent itemsets and their supports. Therefore, in the stream case it is impossible to perform a second scan on the past data to check the support count of a pattern that is locally frequent in the current block $\mathcal{D}_j$, but that resulted infrequent

in the past stream $\mathcal{D}_{[1,j)}$. A naïve technique to work-around this problem is to keep in the global result set only those patterns having the sum of the known supports not less than $minsup \cdot |\mathcal{D}_{[1,j]}|$. We call the algorithm that adopts this simple technique `Stream Partition`. The known support counts are only the ones corresponding to those blocks in which the patterns resulted to be locally frequent. The first time an itemset x is reported, its support count corresponds to the support computed in the current block. In case it appeared previously, this means introducing an error. If $\mathcal{D}_i$ is the first block where x is frequent, then this error can be at most $\sum_{b \in [1,i)} (\sigma_{\min_b} - 1)$.

4. The `AP` algorithm family

The two naïve algorithms discussed above for distributed and stream settings, both inspired by `Partition`, have serious shortcomings. In particular, the weakness of `Stream Partition` is common to several other stream FIM algorithms. When a previously ignored pattern becomes interesting, its exact support is largely underestimated. In order to overcome this issue, we propose a general framework that corrects the known supports of itemsets that result frequent in the current block of transactions, by using an interpolation schema based on other knowledge gathered from past data. The kind of interpolation used can be substituted seamlessly, in order to better fit the particular application context. In this article and in our previous works [34] we have used a really simple, yet effective, interpolation based on the reduction factor with respect to the supports of the subsets of the considered pattern.

The $\texttt{AP}_{\text{stream}}$ algorithm is derived from the distributed algorithm $\texttt{AP}_{\text{Interp}}$ [34], using a method similar to the one used to build `Stream Partition` from `Distributed`.

In the following subsection we will quickly describe $\texttt{AP}_{\text{Interp}}$, than we will introduce the $\texttt{AP}_{\text{stream}}$ algorithm.

4.1. The $\texttt{AP}_{\text{Interp}}$ *algorithm*

One of the most evident issues in `Distributed` is the generation of several false positives, which in turn cause an increment of both resource utilization and execution time, especially when data skew between data partitions is high. The $\texttt{AP}_{\text{Interp}}$ algorithm addresses this issue by means of global pruning based on good approximate knowledge of the global $\mathcal{F}_2$: each locally frequent k-pattern which contains a globally non-frequent 2-pattern will be locally removed from the set of frequents patterns before sending it to the master, and generating the next $k+1$-candidate patterns.

On the other hand, `Distributed One-pass Partition` uses a very conservative estimate for the support of patterns, since it always chooses the lower bounds (known support counts) to approximate the results. This causes underestimated support values, but also several false negatives, often for those patterns whose global supports are close to the threshold. The data skew, indeed, might cause a globally frequent k-pattern x to result infrequent on a given partition $\mathcal{D}_i$ only. In other words, since $\sigma_i(x) < \text{minsup} \cdot |\mathcal{D}_i|$, x will not be returned as a frequent pattern by the i^{th} slave. As a consequence, the master of `Distributed One-pass Partition` cannot count on the knowledge of $\sigma_i(x)$, and thus cannot exactly compute the global support of x. Unfortunately, in `Distributed One-pass Partition`, the master might also deduce that x is not globally frequent, because $\sum_{j, j \neq i} \sigma_j(x) < \text{minsup} \cdot |\mathcal{D}|$. In order to limit this issue, in $\texttt{AP}_{\text{Interp}}$ the master infers an approximate value for this unknown $\sigma_i(x)$ by exploiting an *interpolation method*. The master bases its interpolation reasoning on the knowledge of:

– the exact support of *single items* in each partition;

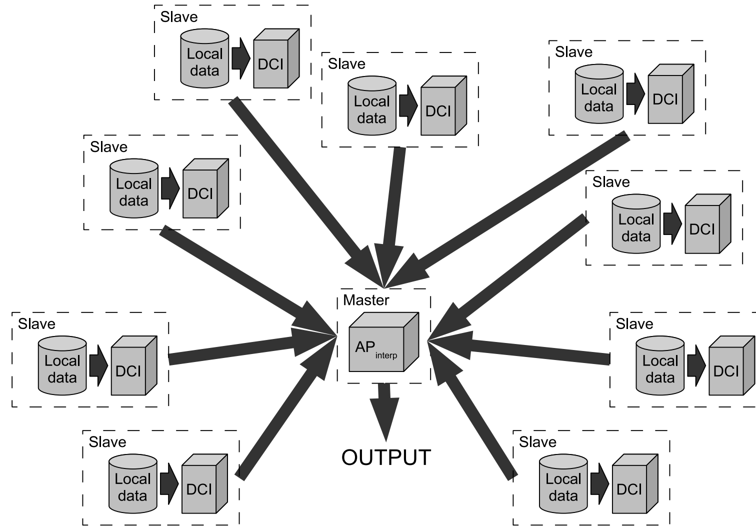

Fig. 1. $\mathtt{AP}_{\mathrm{Interp}}$ overview.

– the *reduction factor* $r(x)$ with respect to the known supports of the items and subsets contained in the considered pattern x.

Note that the support of some subset of x in a partition could be unknown too. This mean that it has been interpolated and discarded because globally infrequent during the $k-1$ iteration, otherwise an approximation of its support would be known. In this case x can be discarded as well.

The master can thus deduce the *unknown* support $\sigma_i(x)$ on the basis of $r(x)$, in turn derived from the supports of x in those partitions $\mathcal{D}_i$ where x resulted to be frequent. Figure 1 shows an overview of the data flows in the distributed $\mathtt{AP}_{\mathrm{Interp}}$ algorithm.

When the number of distributed dataset partitions is really high, the computation cost for collecting and merging the local solutions could become considerable, since the complexity of the merge operation is linear in the amount of input data. To limit this issue, the nodes can be organized in a hierarchy, where each node fetches and merges the results of its direct descendant, and returns the result of the merge to the parent node.

4.2. *The* $\mathtt{AP}_{\mathrm{stream}}$ *algorithm*

The streaming algorithm we propose in this paper, $\mathtt{AP}_{\mathrm{stream}}$, tries to overcome some of the problems encountered by `Stream Partition` and other similar algorithms for association mining on streams, when the data skew between different incoming blocks is high.

This skew might cause a globally frequent itemset x to result infrequent on a given data block $\mathcal{D}_i$. In other words, since $\sigma_i(x) < minsup \cdot |\mathcal{D}_i|$, x will not be found as a frequent itemset in the i^{th} block. As a consequence, we will not be able to count on the knowledge of $\sigma_i(x)$, and thus exactly compute the support of x. Unfortunately, `Stream Partition` might also deduce that x is not globally frequent, because $\sum_{j,j \neq i} \sigma_j(x) < minsup \cdot |\mathcal{D}|$.

$\mathtt{AP}_{\mathrm{stream}}$ addresses this issue in different ways, as summarized in Table 1. In particular, the table shows all the possible cases regarding the knowledge of $\sigma(x)$ on the current block D_i and the previous part of the stream $D_{[1,i)}$.

Table 1
$\texttt{AP}_{\text{stream}}$: Computing the support of x in the whole data stream $\mathcal{D}_{[1,i]}$

$\sigma_{[1,i)}(x)$	$\sigma_i(x)$	Action
Known	Known	$\sigma_{[1,i]}(x) = \sigma_{[1,i)}(x) + \sigma_i(x)$.
Known	Unknown	Recount support $\sigma_i(x)$ on recent, still available, data. Then $\sigma_{[1,i]}(x) = \sigma_{[1,i)}(x) + \sigma_i(x)$.
Unknown	Known	Interpolate past support $\sigma_{[1,i)}^{interp}(x)$. Then $\sigma_{[1,i]}(x) = \sigma_{[1,i)}^{interp}(x) + \sigma_i(x)$.

The first case is the simplest to handle: the new support $\sigma_{[1,i]}(x)$ will be the sum of $\sigma_{[1,i)}(x)$ and $\sigma_i(x)$. The second one is similar, except that we need to look at recent data for computing $\sigma_i(x)$. The key difference with `Stream Partition` is the handling of the last case. $\texttt{AP}_{\text{stream}}$, instead of supposing that x never appeared in the past, tries to interpolate $\sigma_{[1,i)}(x)$. The interpolation is based on the knowledge of:

– the exact support of each *item* in $\mathcal{D}_{[1,i)}$ (or, optionally, just the approximate support of a fixed number of the most frequent items);
– the *reduction factors* $r(x)$ of the support count of subsets of x in the current block with respect to its interpolated support over the past part of the stream.

The algorithm will thus infer the *unknown* support $\sigma_{[1,i)}(x)$ of itemset x on the part of the stream preceding the i^{th} block as follows:

$$\sigma_{[1,i)}^{\text{interp}}(x) = \sigma_i(x) \cdot r(x)$$

where

$$r(x) = \min_{item \in x} \left(\min \left(\frac{\sigma_{[1,i)}(item)}{\sigma_i(item)}, \frac{\sigma_{[1,i)}(x \setminus item)}{\sigma_i(x \setminus item)} \right) \right) \tag{1}$$

The rationale of Eq. (1) is that, given two itemsets x and x', $x' \subset x$, if the exact value of $\sigma_{[1,i)}(x)$ is unknown, its interpolated value $\sigma_{[1,i)}^{\text{interp}}(x)$ is approximated by using the following proportion:

$$\sigma_i(x) : \sigma_i(x') \;=\; \sigma_{[1,i)}^{interp}(x) : \sigma_{[1,i)}(x')$$

so that

$$\sigma_{[1,i)}^{\text{interp}}(x) = \sigma_i(x) \;\cdot \frac{\sigma_{[1,i)}(x')}{\sigma_i(x')}$$

Note that also $\sigma_{[1,i)}(x')$ might be an approximate value previously interpolated.

Given a k-itemset x, the reduction factor $r(x)$ defined by Eq. (1) is thus computed by considering all x', $x' \subset x$, such that x' is either one of the single items belonging to x, or a $k-1$-itemset set-included in x. Finally, the value chosen for $r(x)$ is the minimum one.

Note that, since the merge of the results is performed level-wise starting first from shorter itemsets, when we try to approximate $\sigma_{[1,i)}^{\text{interp}}(x)$, the exact or approximate value of $\sigma_{[1,i)}(x \setminus item)$ must surely be known or already interpolated, for all $item \in x$. This is because all the $k-1$-itemsets included in x must be globally frequent. Otherwise, x could not be a valid candidate.

Figure 2 shows an overview of the data flows in the $\texttt{AP}_{\text{stream}}$ algorithm.

Table 2
Sample supports and reduction ratios

x	$\sigma_i(x)$	$\sigma_{[1,i)}(x)$	$\frac{\sigma_{[1,i)}(x)}{\sigma_i(x)}$
{A,B,C}	6	?	
{A,B }	8	50	6.2
{A,C }	12	30	**2.5**
{B,C }	10	100	10
{A }	17	160	9.4
{B }	14	140	10
{C }	18	160	8.9
{}	40	400	–

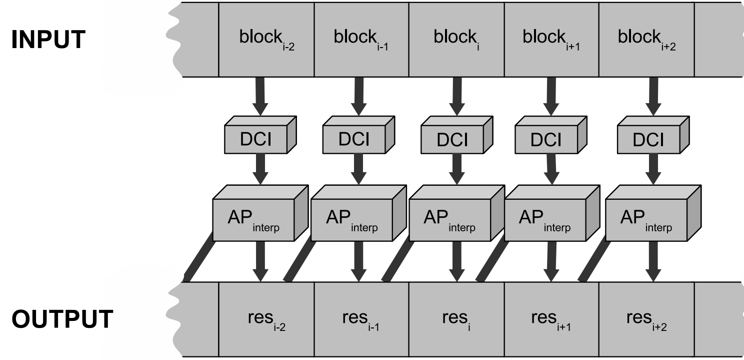

Fig. 2. $\texttt{AP}_{\text{stream}}$ overview.

Example of interpolation. Suppose that we have received 440 transactions so far, and that 40 of these are in the current block $\mathcal{D}_i$. The itemset $x = \{A, B, C\}$, briefly indicated as ABC, is locally frequent, whereas it was infrequent in previous data. Table 2 reports the support of every subset involved in the computation. The first columns contains the itemsets, the second and third columns contain the known supports of the patterns in the current block $\mathcal{D}_i$ and in the past part of the stream $\mathcal{D}_{[1,i)}$. Finally, the last column shows the reduction factor implied by each pattern.

According to Eq. (1), the algorithm chooses the reduction factor $r(x)$ for $x = \{A, B, C\}$ by considering all the itemsets x', $x' \subset x$, of size one and two. In this case the chosen minimum ratio $\frac{\sigma_{[1,i)}(x')}{\sigma_i(x')}$ is 2.5, corresponding to the subset $x' = \{A, C\}$. Since in $\mathcal{D}_i$ the support of $x = \{A, B, C\}$ is $\sigma_i(x) = 6$, the interpolated support will be $\sigma^{\text{interp}}_{[1,i)}(x) = 6 \cdot 2.5 = 15$.

It is worth remarking that this method works if the support of larger itemsets decreases similarly in most parts of the stream, so that a reduction factor (different for each itemset) can be used to interpolate unknown values. Finally note that, as regards the interpolated value above, we expect that the following inequality should hold: $\sigma^{\text{interp}}_{[1,i)}(x) < \text{minsup} \cdot |\mathcal{D}_{[1,i)}|$. So, if we obtain it is not satisfied, this interpolated result should not be accepted. If it was true, the exact value $\sigma_{[1,i)}(x)$ should have already been found. Hence, in those few cases where the above inequality does not hold, the interpolated value will be: $\sigma^{\text{interp}}_{[1,i)}(x) = (\text{minsup} \cdot |\mathcal{D}_{[1,i)}|) - 1$.

Implementation. We can finally introduce the pseudo-code of $\texttt{AP}_{\text{stream}}$. As in `Stream Partition` the transactions are received and buffered. `DCI`, the algorithm used for the local computations, exactly knows the amount of transactions that can be processed in-core.

```
processBlock(buffer, globFreq)
    locFreq[1] = <frequent items>;
    k = 2;
    while size(locFreq[k − 1]) >= k do
        locFreq[k] = computeFrequent(k, locFreq, globFreq);
        commitInsert(k, locFreq, globFreq);
    end while
 end;

commitInsert(k, locFreq, globFreq)
    for all pat in globFreq[k] and not in locFreq[k] do
        <count support of pat in recent data>
        if <pat is frequent> then
            <pre-insert pat in globFreq[k]>
        end if
    end for
    <update globFreq>
end;

computeFrequent(k, locFreq, globFreq)
    < compute local frequent itemsets >
    for all pat locally frequent do
        <compute global interpolated support and bounds>
        if <pat is frequent> then
            <insert pat in locFreq[k]>
            <pre-insert pat in globFreq[k]>
        end if
    end for
    return F_k;
end;
```

Fig. 3. $\texttt{AP}_{\text{stream}}$ pseudo-code.

Thus we can use this knowledge in order to maximize the size of each block of transactions processed at a time. Since frequent itemsets are processed sequentially and can be offloaded to disk, we can ignore the memory occupied by the mined results.

Figure 3 contains the pseudo-code of $\texttt{AP}_{\text{stream}}$. For the sake of simplicity we will neglect the quite obvious main loop with code related to buffering, and concentrate our attention on the processing of each data block. The interpolation formula has been omitted too for the same reason.

Each block is processed, visiting the search space level-wise, for discovering frequent itemsets. In this way itemsets are sorted according to their length and the interpolated support for frequent subpatterns is always available when required. The processing of itemsets of length k is performed in two steps. First frequent itemsets are computed in the current block, and then the actual insertion into the past set of frequent itemsets is carried out. When a pattern is found to be frequent in the current block, its support on past data is immediately checked: if it was already known then the local support is summed to previous support and previous bounds. Otherwise a support and a pair of bounds are inferred for past data, and summed to the support in the current block. In both cases, if the resulting support passes the support test, the pattern is queued for insertion. After every locally frequent itemset of length k has been processed, the support of every previously known itemset which, on the other hand, resulted to be locally infrequent

must be computed on recent data. Itemsets passing the support test are queued for insertion too. Then the pre-inserted itemsets in the queue are sorted and the actual insertion takes place.

4.3. Tighter bounds

As a consequence of using an interpolation method to guess an approximate support value in the past part of the stream, it is very important to establish some bounds on the support found for each pattern. In the previous subsection we have already indicated a pair of really loose bounds: each support cannot be negative, and if a pattern was found infrequent in the past data $\mathcal{D}_{[1,i)}$, then its interpolated support should be less than $\mathrm{minsup} \cdot |\mathcal{D}_{[1,i)}|$. This criteria is completely true for a non-evolving distributed dataset (*distributed frequent pattern mining*). In the stream case, however, the results are approximate and may be affected by false negatives. When a pattern is erroneously discarded as infrequent, its future upper bounds might be underestimated. Anyhow, this issue concerns just a limited number of patterns and, also in these cases, the bounds represent a useful approximation of the exact ones.

4.3.1. Bounds based on pattern subset

The first bounds that interpolated supports should obey, derive from the *Apriori property*: no set can have a support greater than those of any of its subset. Since recent results are merged level-wise with previously known ones, the interpolation can exploit already interpolated subset support. When a subpattern is missing during interpolation, it means that it has been examined during a previous level and discarded. In this case all of its superset may be discarded as well. The computed bound is thus affected by the approximation of past results: an itemset with an erroneous support will affect the bounds for each of its superset. To avoid this issue it is possible to compute the upper bound for an itemset x using the upper bounds of its sub-patterns instead of their support. In this way the upper bounds will be weaker, but there will be less false negatives due to erroneous bounds enforcement.

4.3.2. Bounds based on transaction hash

In order to address the issue of error propagation in support bounds we need to devise some other kind of bounds, which are computed exclusively from received data, and thus are independent of any previous results. Such bounds can be obtained using inverted transaction hashes. The technique discussed below was first introduced in the algorithm IHP [21], an association mining algorithm, where it is used for finding an upper bound for the support of candidates in order to prune infrequent ones. As we will show, this method can also be used for lower bounds.

The key idea is to use a number H of arrays of item counters where each array is associated with a disjoint set of input transactions. When a transaction $\bar{t} = (bid, tid, t)$ is processed, we only modify the counters in the h^{th} array, where h is the result of a hash function applied to tid. Since $tids$ are consecutive integer numbers, a trivial hash function, like $hf(tid) = tid \mod H$, will guarantee an equal distribution of transactions among all hash bins. Thus, when the transaction $\bar{t} = (bid, tid, t)$ is processed, we update the array associated with the current tid

$$(\forall item \in t)\; Count_h[item] + +$$

where $h = tid \mod H$.

Let $H = 1$, i.e., a single array of counters is used. Let A and B be two items, and $Count_0[A]$ and $Count_0[B]$ the associated counters, i.e. $Count_0[A]$ and $Count_0[B]$ are the number of occurrences of items A and B in the whole dataset. According to the Apriori principle

$$\sigma(\{A, B\}) \leqslant min(Count_0[A], Count_0[B])$$

Furthermore we are able to indicate a lower bound for the same support. Let n be the total number of transactions n. We know from the inclusion/exclusion principle that

$$\sigma(\{A,B\}) \geqslant max(0, Count_0[A] + Count_0[B] - n)$$

In fact, if $n - Count_0[A]$ transactions does not contain the item A, then at least $Count_0[B] - (n - Count_0[A])$ of the $Count_0[B]$ transactions containing B will also contain A. Suppose that $n = 30$, $Count_0[A] = 18$, $Count_0[B] = 18$. If we represent with an `X` each transaction supporting a pattern, and with a dot any other transaction, we obtain the following diagrams:

```
      Best case(ub(AB)= 18)               Worst case(lb(AB)=6)
A:    XXXXXXXXXX XXXXXXXX.. ..........    XXXXXXXXXX XXXXXXXX.. ..........
B:    XXXXXXXXXX XXXXXXXX.. ..........    .......... ..XXXXXXXX XXXXXXXXXX
AB:   XXXXXXXXXX XXXXXXXX.. ..........    .......... ..XXXXXX.. ..........
supp                 18                                  6
```

Then no more than 18 transactions will contain both A and B. At the same time at least $18 + 18 - 30 = 6$ transactions will satisfy that constraint. Since each counter represents a set of transaction, this operation is equivalent to the computation of the minimal and maximal intersections of the tid-lists associated with the single items.

Usually, however, $H > 1$. In this case, for each transaction tid, we will increment the counter array $Count_h[]$, where $h = tid \mod H$. The bounds for the support of an itemset x are:

$$\sigma(x)^{\text{upper}} = \sum_{h=0}^{H-1} \min_{item \in x} (Count_h[item])$$

$$\sigma(x)^{\text{lower}} = \sum_{h=0}^{H-1} \max\left(0, n_h - \sum_{\text{item} \in \text{x}} (n_h - Count_h[item])\right)$$

where n_h is the total number of transactions associated with the h^{th} hash value.

Consider the same example discussed above, i.e. 30 transactions including items A and B, where $\sigma(A) = 18$ and $\sigma(B) = 18$. Let $H = 3$. Therefore $n_h = 10$, for each $h = 0,1,2$. Suppose that $Count_0[A] = 8$, $Count_0[B] = 7$, $Count_1[A] = 4$, $Count_1[B] = 5$, $Count_2[A] = 6$, and $Count_2[B] = 6$. Using the same notation previously introduced we obtain:

```
              h=0                          h=1                          h=2
    Best case    Worst case      Best case    Worst case      Best case    Worst case
A:  XXXXXXXX..   XXXXXXXX..  A:  XXXX......   XXXX......  A:  XXXXXX....   XXXXXX....
B:  XXXXXXX...   ...XXXXXXX  B:  XXXXX.....   .....XXXXX  B:  XXXXXX....   ....XXXXXX
AB: XXXXXXX...   ...XXXXX..  AB: XXXX......   ..........  AB: XXXXXX....   ....XX....
supp    7            5       supp    4            0       supp    6            2
```

Each pair of columns, which corresponds to a distinct $h = 0, 1, 2$, represents the transactions having a tid mapped into the corresponding location by the hash function. Note that the lower and upper bounds for $\sigma(\{A,B\})$ are, respectively, $5 + 0 + 2 = 7$ and $7 + 4 + 6 = 17$. Note that these two bounds are stricter than 8 and 18, i.e., the ones obtained for $H = 1$.

Both lower bound and upper bound computations can be extended recursively to larger itemsets. This is possible since the reasoning previously explained still holds if we considers the occurrences of itemsets instead of those of single items.

The lower bound computed in this way will be often equal to zero in sparse datasets. Conversely, on dense datasets this method did prove to be effective in narrowing the two bounds.

5. Experimental evaluation

In this section we study the behavior of the proposed method. We run the $\mathtt{AP}_{\mathrm{stream}}$ algorithm on several datasets using different parameters. The goal of these tests is to understand how similarities of the results vary as the stream length increases, how the hash based pruning is effective, and, in general, how dataset peculiarities and invocation parameters affect the accuracy of the results. Furthermore, we want to study how execution time evolves when the stream length increases.

5.1. Assessing accuracy

The method we are proposing yields approximate results. In particular $\mathtt{AP}_{\mathrm{stream}}$ computes itemset supports which may be slightly different from the exact ones. Thus the result set may miss some frequent itemset (false negatives), or include some infrequent itemset (false positives).

5.1.1. Similarity measure

In order to evaluate the accuracy of the results, we need a measure of similarity between two pattern sets. A widely used one has been introduced in [30], and is based on support difference.

Definition 6. (Similarity) Let A and B respectively be the reference (correct) result set and the approximate result set. $sup_A(x) \in [0,1]$ and $sup_B(y) \in [0,1]$, where $x \in A$ and $y \in B$, correspond to the relative support found in A and B respectively. Note that since B corresponds to the frequent itemsets found by the approximate algorithm under observation, $A - B$ thus corresponds to the set of *false negatives*, while $B - A$ are the *false positives*.

The Similarity is thus computed as

$$\mathtt{Sim}_\alpha(\mathtt{A},\mathtt{B}) = \frac{\sum_{\mathtt{x}\in\mathtt{A}\cap\mathtt{B}} \mathtt{max}\{0,\ \ 1-\alpha * |\mathtt{sup}_\mathtt{A}(\mathtt{x}) - \mathtt{sup}_\mathtt{B}(\mathtt{x})|\}}{|\mathtt{A}\cup\mathtt{B}|}$$

where $\alpha \geqslant 1$ is a scaling parameter, which increase the effect of the support dissimilarity. Moreover, $\frac{1}{\alpha}$ indicates the maximum allowable error on (relative) itemset supports. We will use the notation $Sim()$ to indicate the default case for α, i.e. $\alpha = 1$.

This measure of similarity is thus the sum of at most $|A \cap B|$ values in the range $[0,1]$, divided by $|A \cup B|$. Since $|A \cap B| \leqslant |A \cup B|$, similarity lies in $[0,1]$ too.

When an itemset appears in both sets and the difference between the two supports is greater than $\frac{1}{\alpha}$, it does not improve similarity, otherwise similarity is increased according to the scaled difference. If $\alpha = 20$, then the maximum allowable error in the relative support is $1/20 = 0.05 = 5\%$. Supposing that the support difference for a particular itemset is 4%, the numerator of the similarity measure will be increased by a small quantity: $1 - (20 * 0.04) = 0.2$. When α is 1 (default value), only itemsets whose support difference is at most 100% contribute to increase similarity. On the other hand, when we set α to a very high value, only itemsets with a very similar supports in both the approximate and reference sets will contribute to increase the similarity measure.

It is worth noting that the presence of several false positives and negatives in the approximate result set B contributes to reduce our similarity measure, since this entails an increase in $A \cup B$ (the denominator of the Sim_α formula) with respect to $A \cap B$. Moreover, if an itemset has an actual support which is slightly less than $minsup$ but the approximate support (sup_B) is slightly greater than $minsup$, similarity is decreased even if the computed support was almost correct.

Two more classical result approximation measures are Precision and Recall, both originally introduced in the information retrieval context. The Precision is defined as the fraction of patterns contained in the solution that are actually frequent, i.e., it is the probability that a generic returned pattern will be actually frequent. The Recall is defined as the fraction of the total number of frequent pattern that are contained in the solution, i.e., it is the probability that a generic frequent pattern will be found by the algorithm. Both, however, does not consider the correctness of the support, but only the presence in the result set. This may be misleading, in particular when using a high minimum support threshold. On the other hand, a high similarity value ensure high Precision, high Recall, and limited differences between the actual support values and the discovered ones.

5.1.2. Average support range

When bounds on the support of each itemset are available, an intrinsic measure of the correctness of the approximation is the average width of the interval between the upper bound and the lower bound [34].

Definition 7. (Average support range) Let B be the approximate result set, $sup(x)$ the exact support for itemset x and $sup(x)^{\text{lower}}$ and $sup(x)^{\text{upper}}$ the lower and upper bounds on $sup(x)$, respectively. The average support range is thus defined as:

$$ASR(B) = \frac{1}{|B|} \sum_{x \in B} sup(x)^{\text{upper}} - sup(x)^{\text{lower}}$$

Note that, while this definition can be used for every approximate algorithm, how to compute $sup(x)^{\text{lower}}$ and $sup(x)^{\text{upper}}$ is algorithm specific.

5.2. Experimental data

We performed several tests using both real world datasets, mainly from the FIMI'03 contest [18], and synthetic datasets generated using the IBM generator. We randomly shuffled each dataset and used the resulting datasets as input streams.

Table 3 illustrates these datasets along with their cardinality. The datasets having the name starting with T are synthetic datasets, which mimic the behavior of market basket transactions. The sparse dataset family T20I8N5k has transactions composed, on average, of 20 items, chosen from 5000 distinct items, and includes maximal itemsets whose average length is 8. The dataset family T30I30N1k was generated with the parameters briefly indicated in its name. It is a moderately dense dataset, since more than 10,000 frequent itemsets can be extracted even with a minimum support of 30%. A description of all other datasets can be found in [18]. Kosarak and Retail are really sparse datasets, whereas all other the real world datasets used in experimental evaluation are dense. Table 3 also indicates, for each dataset, a short acronym that will be used in our charts for referring to it.

5.3. Experimental Results

For each dataset and several minimum support thresholds, we computed the exact reference solutions by using `DCI` [27], the same FIM algorithm used as a building block for both $\texttt{AP}_{\text{Interp}}$ and $\texttt{AP}_{\text{stream}}$. Then we ran $\texttt{AP}_{\text{stream}}$ for different values of available memory and number of hash entries.

The first test is focused on catching the effect of used memory on the behavior of the algorithm, when the block of transactions processed at a time is sized dynamically according to the available resources.

Table 3
Datasets used in experimental evaluation

Dataset	Reference	#Trans.
accidents	A	340183
kosarak	K	990002
retail	R	88162
pumbs	P	49046
pumbs-star	PS	49046
connect	C	67557
T20I8N5k	S2..6	77302..3189338
T25I20N5k	S7..11	89611..1433580
T30I30N1k	D1..D9	50000..3189338

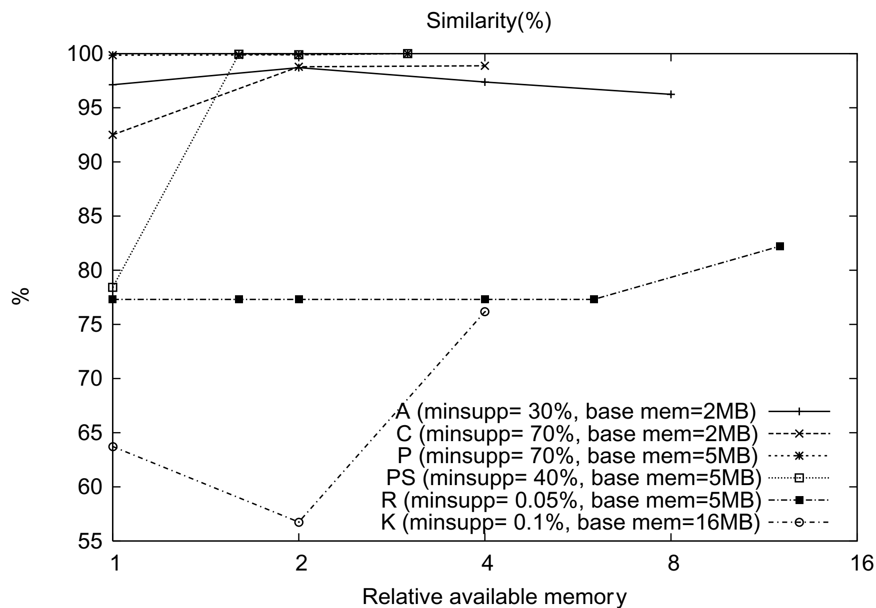

Fig. 4. Similarity as a function of available memory.

In this case data are buffered as long as all the item counters, and the representation of the transactions included in the current block fit the available memory. Note that the size of all frequent itemsets, mined either locally or globally, is not considered in our resource evaluation, since they can be offloaded to disk if needed. The second test is somehow related to the previous one. In this case the amount of required memory is variable, since we determine a-priori the number of transactions to include in a single block, independently of the stream content. The typical use case for $\texttt{AP}_{\text{stream}}$ matches the first test: the user chooses the support, while the other parameters are chosen adaptively, depending on the available system memory and data peculiarities. The second test, with this adaptive behavior disabled, has been inserted for the sake of completeness. Since the datasets used in the tests are quite different, in both cases we used really different ranges of parameters. Therefore, in order to fit all the datasets in the same plot, the number reported in the horizontal axis are relative quantities, corresponding to the block sizes actually used in each test. These relative quantities used in the chart are obtained by dividing the memory/block size used in the specific test by the smallest one for that dataset. For example, the series 50 KB, 100 KB, 400 KB thus becomes 1, 2, 8.

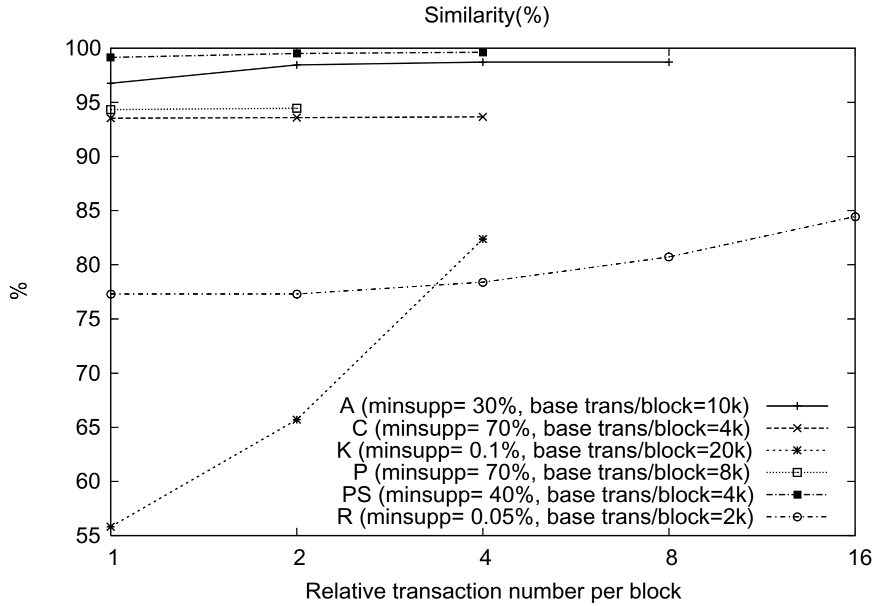

Fig. 5. Similarity as a function of the number of transactions per block.

The plot in Fig. 4 shows the results obtained in the fixed memory case, while the plot in Fig. 5 corresponds to the case when the number of transactions per block is fixed. The relative quantities reported in both plots refer to different base values of either memory or transactions per blocks. These values are reported in the legend of each plot. In general when we increase the number of transactions processed at a time, either statically or dynamically on the basis of the memory available, we also improve the results similarity. Nevertheless the variation is in most cases small, and sometimes there is also a slightly negative trend, caused by the data dependant relationship between used memory and transactions per block. Indeed, a different amount of available memory entails a different division of the stream into blocks, having different sizes and starting points. Occasionally, this could worsen the similarity, in spite of a larger amount of available memory, as in the case of dataset K in the plot in Fig. 4. In our test we noted that choosing an excessively low amount of available memory for some datasets leads to performance degradation, and sometimes also to similarity degradation. The plot in Fig. 7 shows the effectiveness of the hash-based bounds on reducing the Average Support Range (zero corresponds to an exact result). As expected, the improvement is evident only on more dense datasets.

The last batch of tests makes use of a family of synthetic datasets, with homogeneous distribution parameters and varying lengths. Each dataset is obtained from the larger dataset of the series by truncating it to simulate streams with different lengths. For each truncated dataset we computed the exact result set, used as reference value in computing the similarity of the corresponding approximate result obtained by $\texttt{AP}_{\mathrm{stream}}$. The chart in Fig. 7 plots both similarity and ASR as the stream length increases. We can see that similarity remains almost the same, whereas the ASR decreases when an increasing portion of the stream is processed. Finally, the plot in Fig. 8 shows the evolution of execution time as the stream length increases. The execution time increases linearly with the length of the stream. Hence, the average time per transaction is constant if we fix the dataset and the execution parameters.

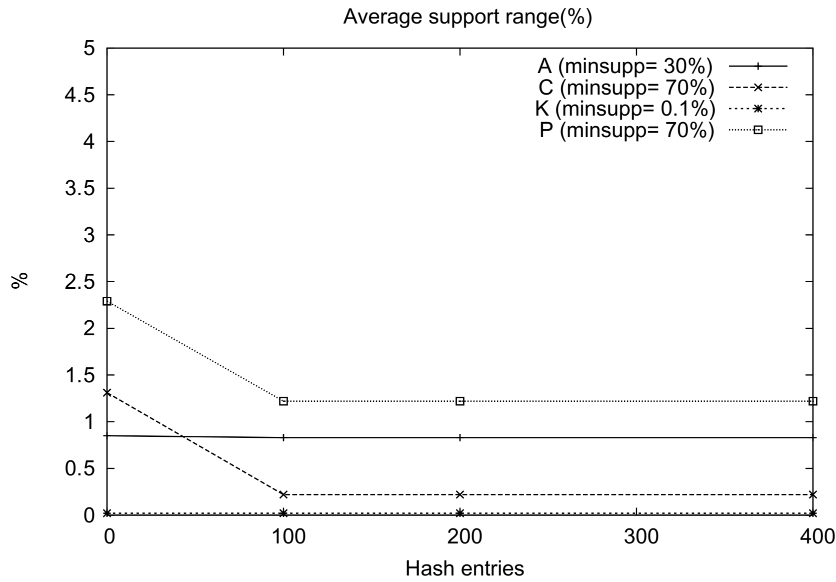

Fig. 6. ASR as a function of the number of hash entries.

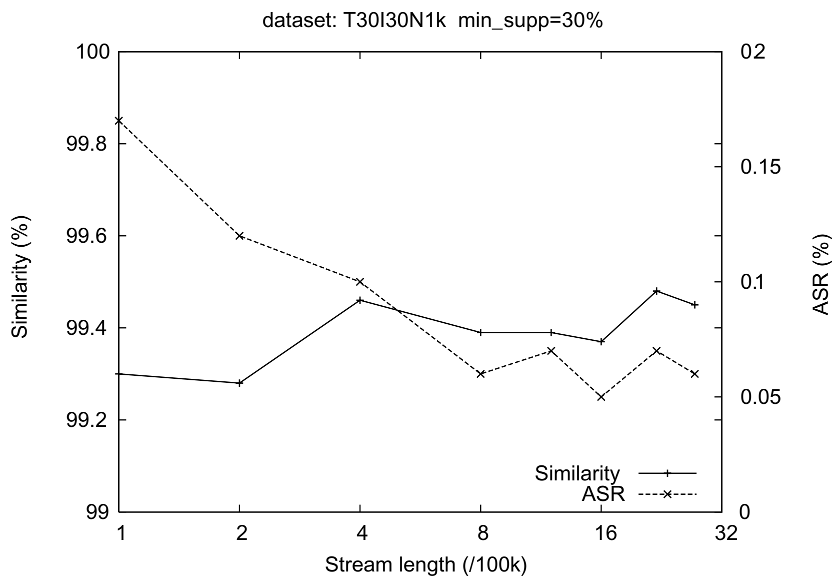

Fig. 7. Similarity and Average Support Range as a function of different stream lengths.

6. Related works

The Association Rule Mining (ARM) in transactional databases has been introduced in [2] and is one of the most popular topics in the KDD field [14,15]. The Frequent Itemset Mining (FIM) is the most computationally expensive phase of ARM. Most FIM algorithms are based on the `Apriori` [4] algorithm, which restricts the search to itemsets whose subsets are all frequent. `Apriori` is a level-wise algorithm, since it examine the k-patterns only when all the frequent patterns of length $k-1$ have been

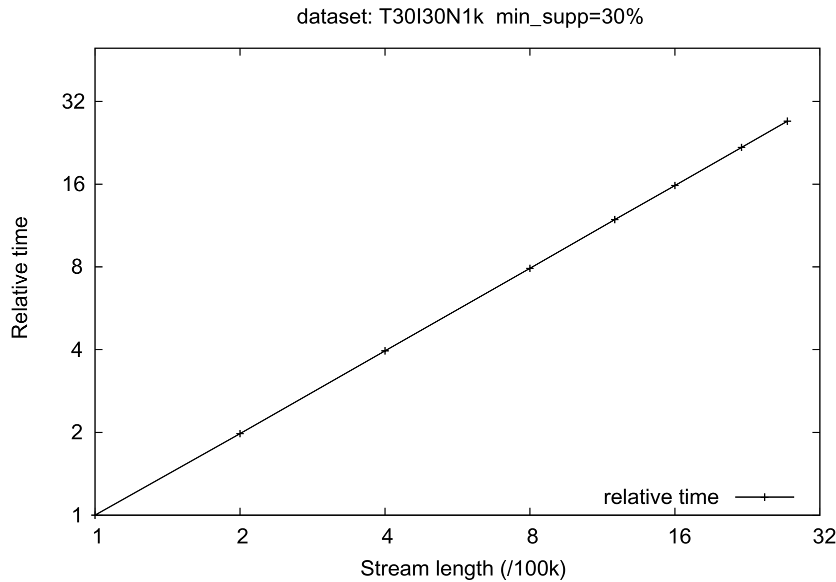

Fig. 8. Execution time as a function of different stream lengths.

discovered. Several other algorithms based on the Apriori principle have been proposed. Some use the same level-wise approach, but introduce efficient optimizations, like a hybrid count/intersection support computation [27], or the reduction of the number of candidates using a hash based technique [29]. Others use a depth-first approach, either class based [38] or projection based [1,20]. Others again use completely different approaches, based on multiple independent computations on smaller part of the dataset, like [31], or incremental computation on an adaptive sample of the data [13,30,32,35]. Parallel (PDM) and distributed (DDM) data-mining are a natural evolution of data-mining technologies, motivated by the need of scalable and high performance systems. A number of parallel algorithms for solving the FIM have been proposed in the last years [3,19]. Most of them can be considered parallelizations of the well-known Apriori algorithm.

Zaki authored a good survey on ARM algorithms and relative parallelization schemas [37]. Agrawal et al. [3] proposed a broad taxonomy of the parallelization strategies that can be adopted for Apriori on distributed-memory architectures. The described approaches constitute a wide spectrum of tradeoffs between computation, communication, memory usage, synchronization, and the use of problem-specific information. The Count Distribution (CD) approach adopts the data-parallel paradigm, according to which the input transaction database is statically partitioned among the processing nodes, while the candidate set C_k is replicated. Count Distribution is an algorithm that can be realized in a distributed setting, since it based on a partitioned dataset, and also because the amount of information exchanged between nodes is limited. The other two methods proposed by Agrawal et al., Data and Candidate Distribution, require moving the dataset. Unfortunately in a distributed environment such dataset is usually already partitioned and distributed on distinct sites, and cannot be moved for several reasons, for example due to the low latency/bandwidth network that connects the sites.

Several DDM FIM algorithms have been proposed, aimed at reducing the amount of communications involved in the Count Distribution method. FDM [9] constitutes an attempt to reduce the amount of communication entailed in the sum-reduction of the local counters in the CD parallelization of the Apriori algorithm. Schuster and Wolff [33] then introduced DDM, whose aim is to reduce the number

of messages exchanged by FDM that, in presence of non-homogeneity of database partitions, quickly becomes similar to the number of messages exchanged by CD. The basic idea of DDM is to verify that an itemset is frequent before collecting its support from every party. The same authors extend the idea of DDM to a dynamic large scale P2P environment [36], i.e., a system based on utilizing free computational/storage resources on non-dedicated machines, where nodes can suddenly depart/join along with the associated database, thus modifying the global result of the computation.

The exact discovery of frequent items in a stream of items may be a highly memory intensive problem [8]. Several relaxed versions of this problem exist, and some interesting ones were introduced in [8,12,23]. The techniques used for solving this family of problems can be classified into two large categories: count-based techniques [12,23–25], sketch-based techniques [8,10,11,24]. The first ones monitor a limited set of potentially "interesting" items, using a counter for each one of them. In this case an error arises when an item is erroneously kept out of the set or inserted too late. The second family provides frequency estimation for every item by using a hash indexed vector of counters. In this case the risk of completely missing the occurrences of an item is avoided, at the cost of looser guarantees on the computed frequencies.

The FIM problem on stream of transactions poses additional memory and computational issues due to the exponential growth of solution size with respect to the corresponding problem on streams of items. Two representative approximate algorithms are derived respectively from LOSSY COUNT [24] and FREQUENT [12,23]. The first one is presented in [24], and is an almost straightforward extension of LOSSY COUNT. The second one is presented in [22], and, even if based on FREQUENT, is significantly different from it, since a property that ensures the correctness in the item case is no longer valid for itemsets. Both algorithms are affected by the issues previously described in the discussion of `Stream Partition`, i.e., they do not consider the possible support count that a pattern could have, even if it has been reported as infrequent. LOSSY COUNT maintains the obvious upper bound that we also used, but no lower bound is exploited.

7. Extensions

The proposed interpolation framework for frequent pattern mining is based on the merge of partial results, using interpolation to replace missing data. The framework was originally proposed for distributed datasets [34], and, in this paper, has been extended to stream datasets. Thanks to the generality of the proposed approach, it can be easily extended also to other, more challenging, cases, like Frequent Sequences, Frequent Closed Itemsets, and settings involving multiple distributed streams. Interestingly, the proposed stream algorithm can be applied, with little modifications, also to a mobile agent setting. In particular it corresponds to the simple case of a single agent that traverses multiple repositories in sequence, carrying partial results along with the code. Thus we plan to investigate the use of this framework in more intricate scenarios, involving largely distributed datasets and several cooperating mobile agents.

In this section we only discuss some of the extensions indicated above, namely the distributed/stream FSM (Frequent Sequence Mining) problem and the FIM problem for multiple distributed streams.

7.1. *Frequent Sequence Mining on distributed/stream data*

The methods presented for frequent itemset extraction can easily be extended to another kind of frequent patterns: the frequent sequences. This only involves minor modifications of the algorithms:

replacing the interpolation formula with one suitable for sequences, and the FIM algorithm with a FSM algorithm. CCSM [28] is an efficient level-wise FSM algorithm, able to handle time constraints, and producing an ordered set of frequent sequences. CCSM is a suitable FSM candidate to be inserted in our distributed and stream framework. Indeed, since CCSM visits level-wise the search space, it extracts the sequences ordered by length. This feature allows $\text{AP}_{\text{stream}}$ and $\text{AP}_{\text{Interp}}$ to merge on-the-fly the sequence patterns as they arrive. Furthermore the on-the-fly merge reduces both memory requirement and computational cost.

As the overall framework remains exactly the same, all the improvements and limits that we have explained for frequent itemsets are still valid. The only problems are those originated by the intrinsic difference between frequent itemset and frequent sequences, which make the result of FSM potentially larger and more likely to be affected by combinatorial explosion.

7.2. *Frequent Itemset Mining on distributed stream data*

The proposed merge/interpolation framework can be extended seamlessly to manage distributed streams in several ways. The most straightforward one is based on the composition of $\text{AP}_{\text{Interp}}$, followed by $\text{AP}_{\text{stream}}$. Each slave is responsible for extracting frequent itemsets from its local streams. The results of each processed block are sent to the master and merged, first among them by using $\text{AP}_{\text{Interp}}$, and then with the past combined results by using $\text{AP}_{\text{stream}}$. The schema on the left of Fig. 9 illustrates this framework. $Res_{\text{node,i}}$ is the FIM result on the i^{th} block of the $node$ stream, whereas Res_i is the result of the merge of all local i^{th} results, and $Hist_Res_i$ is the historical global result, i.e., from the beginning to the i^{th} block.

A first improvement on this base idea could be the replacement of the two cascaded merge phases, one distribution related and the other stream related, with a single one. This would allow for better accuracy of results and stricter bounds, thanks to the reduction of cumulated errors. Clearly, the recount step, used in $\text{AP}_{\text{stream}}$ for assessing the support of recently infrequent itemsets that were frequent in past data, is impossible in both cases. Since the merge is performed in the master node, only the received locally frequent patterns are available. However, this step proved to be effective in our preliminary tests on $\text{AP}_{\text{stream}}$, particularly for dense datasets.

In order to introduce the local recount phase, it is necessary to move the stream merge phase to the slave nodes. In this way, recent data are still available in the reception buffer, and can be used to improve the results. Each slave node then sends its local results, related to the whole history of its streams, to the master node that simply merges them like in $\text{AP}_{\text{Interp}}$. Since these results are sent each time a block is processed, it would be advisable to send only the differences in the results related to the last processed block. This involves rethinking the central merge phase, but in our opinion it should yield better results. The schema on the right of Fig. 9 illustrates this framework. The stream of result generated by each instance of DCI is directly processed by $\text{AP}_{\text{stream}}$, yielding $Hist_Res_{node,i}$, i.e. the results on the whole $node$ stream at time i. $\text{AP}_{\text{Interp}}$ collects these results and outputs the final result $Hist_Res_i$.

The last aspect to consider is synchronization. Each stream evolves, potentially at a different rate with respect to other streams. This means that when the stream reception buffer of a node is full other nodes could be still collecting data. Thus, the collect and merge framework should allow for asynchronous and incremental result merge, with some kind of forced periodical synchronization, if needed. In this case, like in $\text{AP}_{\text{Interp}}$, we are considering a straightforward way of collecting and merging the local results. However, when the number of distributed streams is really high, a better solution is possible. The nodes can be organized in a hierarchy, where the master exchanges messages only with the first level, and intermediate nodes encapsulate their child nodes, returning the result of the merge to the parent node.

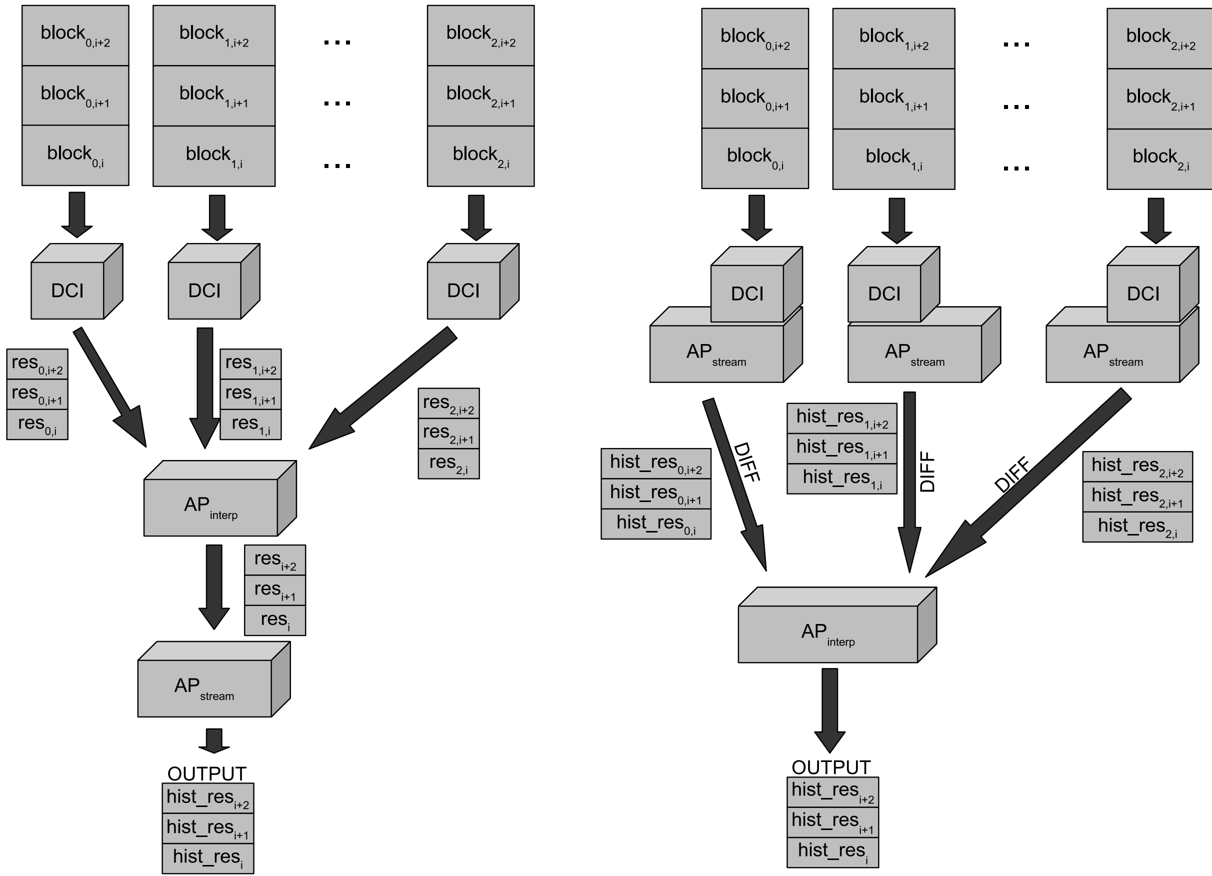

Fig. 9. Distributed stream mining framework. On the left distributed merge followed by stream merge, on the right local stream merge followed by distributed merge.

7.3. Time granularity

The method proposed in this paper yields the most recent solution to the frequent pattern problem in a landmark setting, that is, the returned frequent patterns are referred to the whole stream. While this can be satisfactory in several cases, sometimes the user may be interested in limiting the query time interval or in comparing the solution for different time intervals to discover changes. Our algorithm can be straightforwardly adapted to these time constrained queries, since the merge of local results can be postponed, to enforce the user supplied time constraints. This technique is described in full details in [17]. Here we summarize its main aspects and explain how to integrate our algorithm in a tilted-time window framework.

7.3.1. Tilted-time windows

The users are often interested in analyzing recent data at a finer granularity than past data. The design of tilted-time windows allows for storing in a memory-efficient way the summaries needed to answer queries on long term data, and fine granularity on more recent data.

Figure 10 shows a tilted-time window based on commonly used time intervals: last 4 quarter of an hour, last 24 hours, last 31 days, last 12 months, last years, last 2 years, last 4 years. If we keep track of

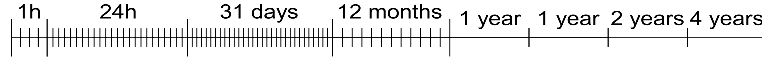

Fig. 10. Natural tilted-time windows.

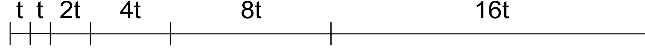

Fig. 11. Logarithmic tilted-time windows.

the support of a pattern for each time interval, we can use such information in order to answer the user query. It should be noted that only 78 counters are used to represent the past 4 year with high granularity on recent data, and a few more counters would allow extending the larger time window to over 100 years. If the available memory is a critical factor, logarithmic tilted-time windows can be used. In this case the size of every window is larger than the more recent one by a fixed factor. Figure 11 shows a logarithmic tilted-time window corresponding to a factor 2. If the time unit t is still a quarter of an hour, the first two intervals on the left represent the last two quarters, the following one the last half-hour and so on. In this case the same 4 years period would require only $\lceil \log_2(4 \times 24 \times 365) \rceil + 1 \approx 19$, which is far less than the number of quarters contained in the same period.

When a time unit elapses, the most recent counter is shifted and replaced by the new support, the previous one is shifted too and so on, summing the supports when needed (e.g., 24 hours make a day). Tilted-time windows can be efficiently updated, if we use some extra memory to store the counters that will replace the current ones while they are incremented. Indeed, the amortized time is O(1) for each pattern and, in the logarithmic case, only one extra counter is needed for each counter to be maintained.

7.3.2. $\mathtt{AP}_{\mathrm{stream}}$ *and tilted-time windows*

Simply merging the set of frequent itemsets for different time intervals, as highlighted in the `Stream Partition` case, leads to an approximation of the support. This is due to the possible occurrence of patterns in intervals where they are not frequent. To address this issue the authors of [17] are forced to maintain also several infrequent patterns, in a number increasing with the required maximum error on the support ε. Since the approach they propose is roughly comparable to a reduction of the minimum support during local computation, the time needed to process each batch can be unreasonable for dense datasets. Even moderately sparse datasets with long transactions may be critical, due to the reduction of the minimum support to ε.

Thus, we propose to avoid the support reduction and to use the interpolation based merge proposed in $\mathtt{AP}_{\mathrm{stream}}$, instead of simply summing the supports when the counters are shifted. In this case the user will not be able to specify a maximal error bound. However, $\mathtt{AP}_{\mathrm{stream}}$ will determine the error bounds on computed patterns, and it will also be able to deal with lower support level, and more complex datasets than the algorithm proposed in [17].

7.3.3. *Dealing with concept drift*

In case the models built on old data become inaccurate, due to a data distribution change, using tilted-time windows can help to avoid the effects of concept drift. Since the pattern frequencies are maintained at different time granularities, we can simply decide to ignore the summaries of older data when they are no longer representatives, that is, when the knowledge they provide is not compatible with current data. However this approach requires being able to decide which part of past data is useful and which is not, and sometimes this is not easily decidable.

A simpler approach consists in gradually decreasing the importance of past data [17], using a fading factor ϕ, applied each time a counter is shifted or merged. Obviously also the window sizes, which corresponds to the supports of the empty pattern, are "faded", so the definition of frequent pattern is still consistent. The main drawback of this approach is that it is not reversible. Hence, it is impossible to apply a different fading factor to past data. However, if we apply the fading factor to the already summarized windows instead of at batch level, we can avoid this issue.

8. Conclusions

In this paper we have discussed $\mathtt{AP}_{\mathrm{stream}}$, a new algorithm for approximate frequent itemset mining on streams. $\mathtt{AP}_{\mathrm{stream}}$ exploits a novel interpolation method to infer the unknown past counts of some itemsets, which are frequents only on recent data. Since the support values computed by the algorithm are approximate, we have also proposed a method for establishing a pair of upper and lower bounds for each interpolated value. These bounds are computed using the knowledge of subpattern frequencies in past data, and of a hash based compressed representation of past data.

Experimental tests shows that the solution produced by $\mathtt{AP}_{\mathrm{stream}}$ is a good approximation of the exact global result. The comparisons with exact results consider both the set of itemsets found and their support. The metric used in order to assess the quality of the algorithm output is the similarity measure introduced in [30], used along with the novel false positive aware similarity proposed in [34]. The interpolation works particularly well for dense dataset, achieving a similarity close to 100% in the best case. The adaptive behavior of $\mathtt{AP}_{\mathrm{stream}}$ allows us to limit the amount of used memory. As expected, we have found that a larger amount of available memory corresponds to a more accurate result. Furthermore, as the length of the processed stream increases, the similarity with the exact result remains almost the same. At the same time, we have observed a decrease in the average difference between upper and lower bounds, which is an intrinsic measure of result accuracy. Finally the time needed to process a block of transactions does not depend on the stream length, hence the total execution time is linear with respect to the stream length.

Acknowledgements

This work was partially supported by the PRIN'04 Research Project entitled "GeoPKDD – Geographic Privacy-aware Knowledge Discovery and Delivery". The datasets used during the experimental evaluation are some of those used for the FIMI'03 (Frequent Itemset Mining Implementations) contest [18]. We thanks the owners of this data and people who made them available in current format. In particular Karolien Geurts [16] for *Accidents*, Ferenc Bodon for *Kosark*, Tom Brijs [6] for *Retail* and Roberto Bayardo for the conversion of *UCI* datasets. Other datasets were generated using the publicly available synthetic data generator code from the IBM Almaden Quest data mining project.

References

[1] R. Agarwal, C. Aggarwal and V.V.V. Prasad, A tree projection algorithm for generation of frequent itemsets, *Journal of Parallel and Distributed Computing* **61**(3) (2001).

[2] R. Agrawal, T. Imielinski and A. Swamim, *Mining Association Rules Between Sets of Items in Large Databases*, In Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data, 1993.

[3] R. Agrawal and J.C. Shafer, Parallel mining of association rules, *IEEE Transactions on Knowledge and Data Engineering* **8**(6) (1996).
[4] R. Agrawal and R. Srikant, *Fast Algorithms for Mining Association Rules in Large Databases*, In Proceedings of the 20th International Conference on Very Large Data Bases, 1994.
[5] B. Babcock, S. Babu, M. Datar, R. Motwani and J. Widom, *Models and Issues in Data Stream Systems*, In Proceedings of the 21st ACM SIGMOD Symposium on Principles of Database Systems, 2002. 23.
[6] T. Brijs, G. Swinnen, K. Vanhoof and G. Wets, *Using Association Rules for Product Assortment Decisions: A Case Study*, In Proceedings of the 5th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 1999.
[7] J.H. Chang and W.S. Lee, *Finding Recent Frequent Itemsets Adaptively Over Online Data Streams*, In Proceedings of the 9th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2003.
[8] M. Charikar, K. Chen and M. Farach-Colton, Finding frequent items in data streams. In Proceedings of the 29th International Colloquium on Automata, Languages and Programming, 2002.
[9] D.W. Cheung, J. Han, V.T. Ng, A.W. Fu and Y. Fu, *A Fast Distributed Algorithm for Mining Association Rules*, In Proceedings of the 4th International Conference on Parallel and Distributed Information Systems, 1996.
[10] G. Cormode and S. Muthukrishnan, *What's Hot and what's Not: Tracking Most Frequent Items Dynamically*, In Proceedings of the 22nd ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems, 2003.
[11] G. Cormode and S. Muthukrishnan, An improved data stream summary: the count-min sketch and its applications, *Journal of Algorithms* **55**(1) (2005).
[12] E.D. Demaine, A. Lóopez-Ortiz and J.I. Munro, *Frequency Estimation of Internet Packet Streams with Limited Space*, In ESA '02: Proceedings of the 10th Annual European Symposium on Algorithms, 2002.
[13] C. Domingo, R. Gavalda and O. Watanabe, Adaptive sampling methods for scaling up knowledge discovery algorithms, *Data Mining and Knowledge Discovery* **6**(2) (2002).
[14] U.M. Fayyad, G. Piatetsky-Shapiro, P. Smyth and R. Uthurusamy, eds, *Advances in Knowledge Discovery and Data Mining*, AAAI Press, 1998.
[15] V. Ganti, J. Gehrke and R. Ramakrishnan, Mining very large databases, *IEEE Computer* **32**(8) (1999).
[16] K. Geurts, G. Wets, T. Brijs and K. Vanhoof, *Profiling High Frequency Accident Locations Using Association Rules*, In Proceedings of the 82nd Annual Transportation Research Board, 2003.
[17] C. Giannella, J. Han, J. Pei, X. Yan and P.S. Yu, *Mining Frequent Patterns in Data Streams at Multiple Time Granularities*, AAAI/MIT Press, 2003. 24.
[18] B. Goethals and M.J. Zaki, eds, *FIMI '03, Frequent Itemset Mining Implementations*, Proceedings of the ICDM 2003 Workshop on Frequent Itemset Mining Implementations, Vol. 90 of CEUR Workshop Proceedings. CEUR-WS.org, 2003.
[19] E.-H.S. Han, G. Karypis and V. Kumar, Scalable parallel data mining for association rules, *In IEEE Transaction on Knowledge and Data Engineering* (2000).
[20] J. Han, J. Pei and Y. Yin, *Mining Frequent Patterns Without Candidate Generation*, In Proceedings of the ACM SIGMOD International Conference on Management of Data, 2000.
[21] J.D. Holt and S.M. Chung, Mining association rules using inverted hashing and pruning, *Information Processing Letters* **83**(4) (2002).
[22] R. Jin and G.G. Agrawal, *An Algorithm for In-Core Frequent Itemset Mining on Streaming Data*, In Fifth IEEE International Conference on Data Mining, 2005.
[23] R.M. Karp, S. Shenker and C.H. Papadimitriou, A simple algorithm for finding frequent elements in streams and bags, *ACM Transactions on Database Systems* **28**(1) (2003).
[24] G. Manku and R. Motwani, *Approximate Frequency Counts Over Data Streams*, In In Proceedings of the 28th International Conference on Very Large Data Bases, 2002.
[25] J. Misra and D. Gries, Finding repeated elements, *Science of Computer Programming* **2**(2) (1982).
[26] A. Mueller, Fast sequential and parallel algorithms for association rules mining: A comparison. Technical Report CS-TR-3515, University of Maryland, 1995.
[27] S. Orlando, P. Palmerini, R. Perego and F. Silvestri, *Adaptive and Resource-Aware Mining of Frequent Sets*, In Proceedings of the 2002 IEEE International Conference on Data Mining, 2002.
[28] S. Orlando, R. Perego and C. Silvestri, *A New Algorithm for Gap Constrained Sequence Mining*, In Proceedings of ACM Symposim on Applied Computing – Data Mining Track, 2004.
[29] J.S. Park, M.S. Chen and P.S. Yu, *An Effective Hash Based Algorithm for Mining Association Rules*, In Proceedings of 1995 ACM SIGMOD International Conference on Management of Data, 1995.
[30] S. Parthasarathy, *Efficient Progressive Sampling for Association Rules*, In Proceedings of the 2002 IEEE International Conference on Data Mining, 2002, 25.
[31] A. Savasere, E. Omiecinski and S.B. Navathe, *An Efficient Algorithm for Mining Association Rules in Large Databases*, In Proceedings of 21th International Conference on Very Large Data Bases, 1995.
[32] T. Scheffer and S. Wrobel, Finding the most interesting patterns in a database quickly by using sequential sampling, *Journal of Machine Learning Research* **3** (2003).

[33] A. Schuster and R. Wolffer, *Communication Efficient Distributed Mining of Association Rules*, In Proceedings of the 2001 ACM SIGMOD International Conference on Management of Data, 2001.
[34] C. Silvestri and S. Orlando, *Distributed Approximate Mining of Frequent Patterns*, In Proceedings of ACM Symposim on Applied Computing – Data Mining track, 2005.
[35] H. Toivonen, *Sampling Large Databases for Association Rules*, In In Proceedings of the 1996 International Conference on Very Large Data Bases. Morgan Kaufman, 1996.
[36] R. Wolff and A. Schuster, *Mining Association Rules in Peer-to-Peer Systems*, In The Third IEEE International Conference on Data Mining, 2003.
[37] M.J. Zaki, Parallel and distributed association mining: A survey, *IEEE Concurrency* **7**(4) (1999).
[38] M.J. Zaki, Scalable algorithms for association mining, *IEEE Transactions on Knowledge and Data Engineering* **12**(3) (2000), 26.