

An Efficient Algorithm for Mining Frequent Sequence with Constraint Programming

John O.R. Aoga(✉)¹, Tias Guns², and Pierre Schaus¹

¹ UCLouvain, ICTEAM (Belgium)

{john.aoga,pierre.schaus}@uclouvain.be

² KU Leuven, DTAI Research group (Belgium)

tias.guns@cs.kuleuven.be

Abstract. The main advantage of Constraint Programming (CP) approaches for sequential pattern mining (SPM) is their modularity, which includes the ability to add new constraints (regular expressions, length restrictions, etc). The current best CP approach for SPM uses a global constraint (module) that computes the projected database and enforces the minimum frequency; it does this with a filtering algorithm similar to the PrefixSpan method. However, the resulting system is not as scalable as some of the most advanced mining systems like Zaki’s cSPADE. We show how, using techniques from both data mining and CP, one can use a generic constraint solver and yet outperform existing specialized systems. This is mainly due to two improvements in the module that computes the projected frequencies: first, computing the projected database can be sped up by pre-computing the positions at which a symbol can become unsupported by a sequence, thereby avoiding to scan the full sequence each time; and second by taking inspiration from the *trailing* used in CP solvers to devise a backtracking-aware data structure that allows fast incremental storing and restoring of the projected database. Detailed experiments show how this approach outperforms existing CP as well as specialized systems for SPM, and that the gain in efficiency translates directly into increased efficiency for other settings such as mining with regular expressions.

1 Introduction

Sequence mining is a widely studied problem concerned with discovering subsequences in a dataset of given sequences, where each (sub) sequence is an ordered list of symbols. It has applications ranging from web usage mining, text mining, biological sequence analysis and human mobility mining [7]. We focus on the problem of finding patterns in sequences of individual symbols, which is the most commonly used setting in those applications.

In recent years, constraint programming (CP) has been proposed as a general framework for pattern mining [4,3,8,5]. The main benefit of CP-based approaches over dedicated algorithms is that it is *modular*. In a CP framework, a problem is expressed as a set of constraints that the solutions must satisfy. Each such a constraint can be seen as a *module*, and can range from being as simple as ensuring

that a subsequence does not contain a certain symbol at a certain position, up to computing the frequency of a pattern in a database. This modularity allows for flexibility, in that certain constraints such as symbol restrictions, length, regular expressions etc can easily be added and removed to existing problems. Another advantage is that improving the efficiency of one constraint will improve the efficiency of all problems involving this constraint.

However, this increased flexibility can come at a cost. Negrevergne et al. [8] have shown that a fine-grained modular approach to sequence mining can support any type of constraints, including gap and span constraints and any quality function beyond frequency, but that this is not competitive with state-of-the-art specialized methods. On the other hand, they showed that by using a global constraint (a module) that computes the pseudo-projection of the sequences in the database similar to PrefixSpan [10], this overhead can be reduced. Kemmar et al. [5,6] propose to use a single global constraint for pseudo-projection as well as frequency counting over all sequences. This approach is much more efficient than the one of [8] that uses many reified constraints. These CP-based methods obtain reasonable performance, especially for mining under regular expressions. While they improve scalability compared to each-other, they are not on par with some of the best specialized systems such as Zaki’s cSpade [18]. In this work, we show for the first time that a generic CP system with a custom global constraint can outperform existing specialised systems including Zaki’s.

The global constraint improves on earlier global constraints for sequence mining by combining ideas from both pattern mining and constraint programming as follows: first, we improve the efficiency of computing the projected database and the projected frequency using last-position lists, similar to the LAPIN algorithm [16] but within a PrefixSpan approach. Second, we take into account not just the efficiency of computing the projected database, but also that of storing and restoring it during depth-first search. For this we use the *trailing* mechanism from CP solvers to avoid unnecessary copying of the pseudo-projection data structure. Such an approach is in fact applicable to any depth-first algorithm in pattern mining and beyond.

By combining the right ingredients from both research communities in a novel way, we end up with an elegant algorithm for the projected frequency computation. When added as a module to a generic CP solver, the resulting system improves both on previous CP-based sequence miners as well as state-of-the-art specialized systems. Furthermore, we show that by improving this one module, these improvements directly translate to other problems using this module, such as regular-expression based sequence mining.

2 Related works

We review specialized methods as well as CP-based approaches. A more thorough review of algorithmic developments is given in [7].

Specialized methods. Introduced by Srikant and Agrawal [1], GSP was the first approach to extract sequential patterns from a sequential database. Many

works have improved on this apriori-based method, typically employing depth-first search. A seminal work is that of PrefixSpan [10]. A prefix in this context is a sequential pattern that can only be extended by appending symbols to it. Given a prefix, one can compute the *projected database* of all suffixes of the sequences that have the prefix as a subsequence. This projected database can then be used to compute the frequency of the prefix and of all its 1-extensions (projected frequency). A main innovation in PrefixSpan is the use of a *pseudo-projected* database: instead of copying the entire (projected) database, one only has to maintain pointers to the position in each sequence where the prefix matched.

Alternative methods such as SPADE [18] and SPAM [2] use a vertical representation of the database, having for each symbol a list of sequence identifiers and positions at which that symbol appears.

Yang et al. have shown [17] that algorithms with either data representation can be improved by precomputing the last position of each symbol in a sequence. This can avoid scanning the projected database, as often the reason for scanning is to know whether a symbol still appears in the projected sequence.

The standard sequence mining settings have been extended in a number of directions, including user-defined constraints on length or on the gap or span of a sequence such as in the cSPADE algorithm [18], closed patterns [15] and algorithms that can handle regular expression constraints on the patterns such as SMA [14]. These constraints are typically hard-coded in the algorithms.

CP-based approaches for SPM. CP-based approaches for sequence mining are gaining interest in the CP community. Early work has focused on fixed-length sequences with wildcards [3]. More generally, [8] proposed two approaches: a full decomposition of the problem in terms of constraints and an approach using a global constraint to construct the pseudo-projected database similar to PrefixSpan. It uses one such constraint for each sequence. Kemmar et al [6] propose to gather all these constraints into a unique global constraint to reduce the overhead of the multiple constraints. They further showed how the constraint can be modified to take a maximal gap constraint into account [5].

3 Sequential Pattern Mining Background

This section introduces the necessary concepts and definitions of sequence mining and constraint programming.

3.1 Sequence Mining Background

Let $I = \{s_1, \dots, s_N\}$ be a set of N symbols. In the remaining of the paper when there is no ambiguity a symbol is simply denoted by its identifier i with $i \in \{1, \dots, N\}$.

Definition 1. *Sequence and sequence database.* A sequence $s = \langle s_1 s_2 \dots s_n \rangle$ over I is an ordered list of (potentially repeating) symbols s_j , $j \in [1, n]$ with $\#s = n$ the length of the sequence s . A set of tuples (sid, s) where sid is a sequence identifier and s a sequence, is called sequence database (*SDB*).

sid	sequence	lastPosList	lastPosMap
sid ₁	⟨ABCBC⟩	[(C,5),(B,4),(A,1)]	{A→1, B→4, C→5,D→0}
sid ₂	⟨BABC⟩	[(C,4),(B,3),(A,2)]	{A→2, B→3, C→4,D→0}
sid ₃	⟨AB⟩	[(B,2),(A,1)]	{A→1, B→2, C→0,D→0}
sid ₄	⟨BCD⟩	[(D,3),(C,2),(B,1)]	{A→0, B→1, C→2,D→3}

Table 1: A sequence database SDB_1 and list of last positions.

1) SDB, 2) lastPosList, 3) lastPosMap

Example 1. Table. 1 shows an example SDB_1 over symbols $I = \{A, B, C, D\}$. For the sequence $s = \langle BABC \rangle$: $\#s = 4$ and $s_1 = B, s_2 = A, s_3 = B, s_4 = C$.

Definition 2. Sub-sequence (\preceq), super-sequence. A sequence $\alpha = \langle \alpha_1 \dots \alpha_m \rangle$ is called a sub-sequence of $s = \langle s_1 s_2 \dots s_n \rangle$ and s is a super-sequence of α iff (i) $m \leq n$ and (ii) for all $i \in [1, m]$ there exist integers j_i s.t. $1 \leq j_1 \leq \dots \leq j_m \leq n$, such that $\alpha_i = s_{j_i}$.

Example 2. For instance $\langle BD \rangle$ is a sub-sequence of $\langle BCCD \rangle$, and inversely $\langle BCCD \rangle$ is the super-sequence of $\langle BD \rangle$: $\langle BD \rangle \preceq \langle BCCD \rangle$.

Definition 3. Cover, Support, Pattern, Frequent Pattern. The cover of sequence p in SDB, denoted by $cover_{SDB}(p)$, is the subset of sequences in SDB that are a super-sequence of p , i.e. $cover_{SDB}(p) = \{(sid, s) \in SDB \mid p \preceq s\}$. The support of p in SDB, denoted by $sup_{SDB}(p)$, is the number of super-sequences of p in SDB: $sup_{SDB}(p) = \#cover_{SDB}(p)$. Any sequence p over symbols in I can be a pattern, and we call a pattern frequent iff $sup_{SDB}(p) \geq \theta$, where θ is a given minimum support threshold.

Example 3. Assume that $p = \langle BC \rangle$ and $\theta = 2$, $cover_{SDB_1}(p) = \{(sid_1, \langle ABCBC \rangle), (sid_2, \langle BABC \rangle), (sid_4, \langle BCD \rangle)\}$ and hence $sup_{SDB_1}(p) = 3$. Hence, p is a frequent pattern for that given threshold.

The sequential pattern mining (SPM) problem, first introduced by Agrawal and Srikant [1], is the following:

Definition 4. Sequential Pattern Mining (SPM). Given an minimum support threshold θ and a sequence database SDB, the SPM problem is to find all patterns p such that $sup_{SDB}(p) \geq \theta$.

Our method uses the idea of a *prefix* and *prefix-projected* database for enumerating the frequent patterns. These concepts were first introduced in the seminal paper that presented the *PrefixSpan* algorithm [10].

Definition 5. Prefix, prefix-projected database Let $\alpha = \langle \alpha_1 \dots \alpha_m \rangle$ be a pattern. If a sequence $\beta = \langle \beta_1 \dots \beta_n \rangle$ is a super-sequence of α : $\alpha \preceq \beta$, then the prefix of β w.r.t. α is the smallest prefix of β that is still a super-sequence of α : $\langle \beta_1 \dots \beta_j \rangle$ s.t. $\alpha \preceq \langle \beta_1 \dots \beta_j \rangle$ and $\nexists j' < j : \alpha \preceq \langle \beta_1 \dots \beta_{j'} \rangle$. The sequence

$\langle \beta_{j+1} \dots \beta_n \rangle$ is called the *suffix* and it represents the *prefix-projection* obtained by projecting the prefix away. A *prefix-projected database* of a pattern α , denoted by $SDB|_{\alpha}$, is the set of *prefix-projections* of all sequences in SDB that are *super-sequences* of α .

Example 4. In SDB_1 , assume $\alpha = \langle A \rangle$, then $SDB_1|_{\alpha} = \{(sid_1, \langle BCBC \rangle), (sid_2, \langle BC \rangle), (sid_3, \langle B \rangle)\}$.

We say that the *prefix-projected frequency* of the symbols I in a prefix-projected database is the number of sequences in which these symbols appear. For $SDB_1|_{\langle A \rangle}$ the prefix-projected frequencies are $A : 0, B : 3, C : 2, D : 0$.

The PrefixSpan algorithm solves the SPM problem by starting from the empty pattern and extending this pattern using depth-first search. At each step it extends a pattern by a symbol and projects the database accordingly. The appended symbol is removed on backtrack. It hence grows the pattern incrementally, which is why it is called a *pattern-growth* method. A frequent pattern in the projected database is also frequent in the original database.

There are two important considerations for the efficiency of the method. The first is that one does not have to consider during search any symbol that is not frequent in the prefix-projected database. The second is that of *pseudo-projection*: to store the prefix-projected database during the depth-first search, it is not necessary to store (and later restore) an entire copy of the projected database. Instead, one only has to store for each sequence the pointer to the position j that marks the end of the prefix in that sequence (remember, the *prefix* of α in β is the smallest prefix $\langle \beta_1 \dots \beta_j \rangle \succeq \alpha$).

Example 5. The projected database $SDB_1|_{\alpha} = \{(sid_1, \langle BCBC \rangle), (sid_2, \langle BC \rangle), (sid_3, \langle B \rangle)\}$ can be represented as a pseudo-projected database as follows: $\{(sid_1, 2), (sid_2, 3), (sid_3, 2)\}$.

3.2 Constraint Programming Background

CP is a powerful declarative paradigm to solve combinatorial satisfaction and optimization problems (see, e.g., [12]). A CP problem (V, D, C) is defined by a set of variables V with their respective domains D (the values that can be assigned to a variable), and a set of constraints C on these variables. A solution of a CP problem is an assignment of the variables to a value from its domain, such that all constraints are satisfied.

At its core, CP solvers are depth-first search algorithms that iterate between *searching* over unassigned variables and *propagating* constraints. Propagation is the act of letting the constraints in C remove unfeasible values from the domains of its variables. This is repeated until *fixed-point*, that is, no more constraint can remove any unfeasible values. Then, a search exploration step is taken by choosing an unassigned variable and assigning it to a value from its current domain, after which propagation is executed again.

Example 6. Let there be 2 variables x, y with domains $D(x) = \{1, 2, 3\}, D(y) = \{3, 4, 5\}$. Then constraint $x + y \geq 5$ can derive during propagation that $1 \notin D(x)$ because the lowest value y can take is 3 and hence $x \geq 5 - \min(D(y)) \geq 5 - 3 \geq 2$.

Constraints and global constraints Many different constraints and their propagation algorithms have been investigated in the CP community. This includes logical and arithmetic ones like the above, up to constraints for enforcing regular expressions or graph theoretic properties. A constraint that enforces some non-trivial or application-dependent property is often called a *global constraint*. For example, [8] introduced a global constraint for the pseudo-projection of a single sequence, and [5] for the entire projected frequency subproblem.

State restoration in CP In any depth-first solver, there must be some mechanism to store and restore some *state*, such that computations can be performed incrementally and intermediate values can be stored. In most of the CP solvers³ a general mechanism, called *trailing* is used for storing and restoring the state (on backtrack) [13]. Externally, the CP solvers typically expose some "reversible" objects whose values are automatically stored and restored on the trail when they change. The most important example are the domains of CP variables. Hence, for a variable the domain modifications (*assign*, *removeValue*) are automatically reversible operations. A CP solver also exposes reversible versions of primitive types such as integers and sets for use within constraint propagators. They are typically used to store incremental computations. CP solvers consist of an efficient implementation of the DFS backtracking algorithm, as well as many constraints that can be called by the fix-point algorithm. The modularity of constraint solvers stems from this ability to add any set of constraints to the fix-point algorithm.

4 Global constraints for projected frequency

We first introduce the basic CP model of frequent sequence mining introduced in [8] and extended in [6]. Then, we present how we improve the computation of the pseudo-projection, followed by the projected frequency counting and pruning.

4.1 Existing methods [8,6]

As explained before, a constraint model consists of variables, domains and constraints. The CP model will be such that a single solution corresponds to a frequent sequence, meaning that all sequences can be extracted by enumerating all solutions.

Let L be an upper bound on the pattern length, e.g. the length of the longest sequence in the database. The variables used to represent the unknown pattern P is modeled as an array of L integer variables $P = [P_1, P_2, \dots, P_L]$. Each variable has an initial domain $\{0, \dots, N\}$, corresponding to all possible symbols identifiers and augmented with an additional identifier 0. The symbol with identifier 0 represents ϵ , the empty symbol. It will be used to denote the end of the sequence in P , using a trailing suffix of such 0's.

³ One notable exception is the Gecode copy-based solver.

Listing 1.1: PrefixProjection(SDB, P, i, θ)

```

1 // pre: variables  $\langle P_1, \dots, P_i \rangle$  are bound,  $SDB$  is given
2 //  $P_i$  is the new instantiated variable since previous call.
3 if ( $P_i == 0$ ) {
4   foreach ( $j \in \{i+1, \dots, L\}$ ) {  $P_j.assign(0)$  }
5 } else if ( $i \geq 2$ ) {
6   projFreqs = ProjectAndGetFreqs( $SDB, P_i, \theta$ )
7   foreach ( $j \in \{i+1, \dots, L\}$ )
8     foreach ( $a \in D(P_j)$ )
9       if ( $a \neq 0$  and projFreqs[a] <  $\theta$ ) {  $P_j.removeValue(a)$  }
10 }

```

Definition 6. A CP model over P represents the frequent sequence mining problem with threshold θ , iff the following three conditions are satisfied by every valid assignment to P :

1. $P_1 \neq 0$
2. $\forall i \in \{2, \dots, L-1\} : P_i = 0 \Rightarrow P_{i+1} = 0$
3. $\#\{(sid, s) \in SDB \langle P_1 \dots P_j \rangle \preceq s\} \geq \theta, j = \max(\{i \in \{1 \dots L\} | P_i \neq 0\})$.

The first requirement states that the sequence may not start with the empty symbol, e.g. no empty sequence. The second requirement enforces that the pattern is in a canonical form such that after the empty symbol, all other symbols are the empty symbol too. Hence, a sequence of length $l < L$ is represented by l non-zero symbols, followed by $L - l$ zero symbols. The last requirement states that the frequency of the non-zero part of the pattern must be above the threshold θ .

Prefix projection global constraint Initial work [8] proposed to decompose these three conditions into separate constraints, including a dedicated global constraint for the inclusion relation $\langle P_1 \dots P_j \rangle \preceq s$ for each sequence separately. It used the pseudo-projection technique of PrefixSpan for this, with the projected frequency enforced on each symbol in separate constraints.

Kemmar et al. [6] extended this idea by encapsulating the filtering of all three conditions into one single (global) constraint called PrefixProjection. It also uses the pseudo-projection idea of PrefixSpan, but over the entire database. The propagation algorithm for this constraint, as executed when the next unassigned variable P_i is assigned during search, is given in Listing 1.1.

An initial assumption is that the database SDB does not contain any infrequent symbols, which is a simple preprocessing step. The code is divided in three parts: (i) if P_i is assigned to 0 the remaining P_k with $k > i$ is assigned to 0; else (ii) from the second position onwards (remember that the first position can take any symbol and be guaranteed to be frequent as every symbol is known to be frequent), the projected database and the projected frequency of each symbol

Listing 1.2: ProjectAndGetFreqs(SDB,a, θ)

```
1  $PSDB_i = \emptyset$ 
2 foreach (sid, start)  $\in PSDB_{i-1}$  {
3   s = SDB[sid]; pos = start
4   while (pos < #s and a  $\neq$  s[pos]) { pos = pos + 1 }
5   if (pos < #s) {  $PSDB_i = PSDB_i \cup \{(sid, pos)\}$  }
6 }
7 projFreqs[a]=0  $\forall a \in \{1, \dots, N\}$ 
8 if ( $\#PSDB_i \geq \theta$ ) {
9   foreach (sid, start)  $\in PSDB_i$  {
10    s = SDB[sid]; existsSymbol[b] = false  $\forall b \in \{1, \dots, N\}$ 
11    foreach (i  $\in \{start, \dots, \#s\}$ ) {
12      if (!existsSymbol[s[i]]) {
13        projFreqs[s[i]] = projFreqs[s[i]]+1
14        existsSymbol[s[i]] = true
15      }
16    } } }
17 return projFreqs
```

is computed; and (iii) all symbols that have a projected frequency below the threshold are removed from the domain of the subsequent pattern variables.

The algorithm for computing the (pseudo) projected database and the projected frequencies of the symbols is given in Listing 1.2. It operates as follows with a the new symbol appended to the prefix of assigned variables since previous call. The first loop at line 2 attempts to discover for each sequence s in the projected database if it can be a sub-sequence of the extended prefix. If yes, this sequence is added to the next projected database at line 5. The second loop at line 9 computes the frequency of each symbol occurring in the projected database but counting it at most once per sequence.

4.2 Improving propagation

Although being the state-of-art approach for solving SPM with CP, the filtering algorithm of Kemmar et al [5] presents room for improvement. We identify four weaknesses and propose solutions to them.

Weakness 1. Databases with long sequences will have a large upper-bound L . For such databases, removing infrequent symbols from all remaining pattern variables P in the loop defined at line 7 of Listing 1.1 can take time. This is not only the case for doing the action, but also for restoring the domains on backtracking. On the other hand, only the next pattern variable P_{i+1} will be considered during search, and in most cases a pattern will never actually be of length L , so all subsequent domain changes are unnecessary. This weakness is a peculiarity of using a fixed-length array P to represent a variable-length

sequence. Mining algorithms typically have a variable length representation of the pattern, and hence only look one position ahead. In our propagator we only remove values from the domain of P_{i+1} .

Weakness 2. When computing the projected frequencies of the symbols, one has to scan each sequence from its current pseudo-projection pointer *start* till the end of the sequence. This can be time consuming in case of many repetitions of only a few symbols for example. Thanks to the *lastPosList* defined next, it is possible to visit only the last position of each symbol occurring after *start*. This idea was first introduced in [17] and exploited in the LAPIN family of algorithms.

Definition 7. (Last position list). For a current sequence s , *lastPosList* is a sequence of pairs (symbol, pos) giving for each symbol that occurs in s its last position: $pos = \max\{p \leq \#s : s[p] = \text{symbol}\}$. The sequence is of length m , the number of distinct symbols in s . This sequence is decreasing according to positions: $lastPosList[i].pos > lastPosList[i + 1].pos \forall i \in \{1, \dots, m - 1\}$.

Example 7. Table 1 shows the *lastPosList* sequences for SDB_1 . We consider the sequence with sid_1 and a prefix $\langle A \rangle$. The computation of the frequencies starts at position 2, remaining suffix is $\langle BCBC \rangle$. Instead of visiting all the 4 positions of this suffix, only the last two can be visited thanks to the information contained in $lastPosList[sid_1]$. Indeed according to $lastPosList[sid_1][1]$ the maximum last position is 5 (corresponding to the last C). Then according to $lastPosList[sid_1][2]$ the second maximum last position is 4 (corresponding to the last position of symbol B). The third maximum last position is 1 for symbol A . Since this position is smaller than 2 (our initial start), we can stop.

Weakness 3. Related to weakness 2, line 4 in Listing 1.2 finds the new position (pos_s) of a in $SDB[sid]$. This code is executed even if the new symbol no longer appears in that sequence. Currently, the code has to loop over the entire sequence until it reaches the end before discovering this.

Assume that the current position in the sequence s is already larger than the position of the last occurrence of a . Then we immediately know this sequence cannot be part of the projected database. To verify this in $O(1)$ time, we use a *lastPosMap* as follows:

Definition 8. (Last position map of symbols). For a given sequence s with id sid , $lastPosMap[sid]$ is a map such that $lastPosMap[sid][i]$ is the last position of symbol i in the sequence s . In case the symbol i is not present: $lastPosMap[sid][i] = 0$ (positions are assumed to start at index 1).

Example 8. Table 1 shows the *lastPosMap* arrays next to SDB_1 . For instance for sid_2 the last position of symbol C is 4.

Weakness 4. Listing 1.2 creates a new set $PSDB_i$ to represent the projected database. This projected database is computed many times during the search, namely at least once in each node of the search tree (more if there are other




Fig. 1: Reversible vectors technique

constraints in the fixPoint set). This is a source of inefficiency for garbage collected languages such as Java but also for C since it induces many "slow" system calls such as free and malloc leading to fragmentation of the memory. We propose to store and restore the pseudo-projected databases with reversible vectors making use of CP trailing techniques. The idea is to use one and the same array throughout the search in the propagator, and only maintain the relevant start/stop position during search. Each call to propagate will read from the previous start to stop position, and write after the previous stop position plus store the new start/stop position. The projected databases are thus *stacked* in the array along a branch of the search tree. We implement the pseudo-projected database with two reversible vectors: *sids* and *poss* respectively for the sequence ids and the current position in the corresponding sequences. The position ϕ is the start entry (in *sids* and *poss*) of the current projected database, and φ is the size of the projected database. We thus have the current projected database contained in sub-arrays $\text{sids}[\phi, \dots, \phi + \varphi - 1]$ and $\text{poss}[\phi, \dots, \phi + \varphi - 1]$. In order to make the projected database reversible, ϕ and φ are reversible integers. That is on backtrack to an ancestor node those integers retrieve their previous value and entries of *sids* and *poss* starting from ϕ can be reused.

Example 9. Figure 1 is an example using SDB_1 . Initially all the sequences are present $\varphi = 4$ and position is initialized $\phi = 0$. The A-projected database contains sequence 1, 2, 3 at positions 1, 2, 1 with $\phi = 4$ and $\varphi = 3$.

Prefix Projection Incremental Counting propagator (PPIC). Putting all the solutions to the identified weaknesses together, we list the code of the main function of our propagator's in Listing 1.3.

The main loop at line 3 iterates over the previous (*parent*) projected database. In case the sequence at index i in the projected database contains the new

Listing 1.3: ProjectAndGetFreqs(SDB, a , θ , $sids$, $poss$, ϕ , φ)

```
1 projFreqs[b]=0  $\forall b \in \{1, \dots, N\}$ 
2  $i = \phi$ ;  $j = \phi + \varphi$ ;  $sup = 0$ 
3 while ( $i < \phi + \varphi$ ) {
4    $sid = sids[i]$ ;  $pos = poss[i]$ ;  $s = SDB[sid]$ 
5   if ( $lastPosMap[sid][a] - 1 \geq start$ ) {
6     //find the next position of  $a$  in  $s$ 
7     while ( $pos < \#s$  and  $a \neq s[pos]$ ) {  $pos = pos + 1$  }
8     // update projected database
9      $sids[j] = sid$ ;  $poss[j] = pos + 1$ ;  $j = j + 1$ ;  $sup = sup + 1$ 
10    // recompute projected frequencies
11    foreach ( $(symbol, pos_x)$  in  $lastPosList[sid]$ ) {
12      if ( $pos_x \leq pos$ ) { break }
13       $projFreqs[symbol] = projFreqs[symbol] + 1$ 
14    }
15     $i = i + 1$ 
16  }
17  $\phi = \phi + \varphi$ ;  $\varphi = sup$ 
18 return  $projFreqs$ 
```

symbol at a subsequent position larger or equal to $start$, the matching position is searched and added to the new projected database (at index j of reversible vectors $sids$ and $poss$) at line 9. Then the contribution of the sequence to the projected frequencies is computed in the loop at line 11. Only the entries in the $lastPosList$ with position larger than current pos are considered (recall that this list is decreasing according to positions). Finally line 17 updates the reversible integers ϕ and φ to reflect the newly computed projected database. Based on these projected frequencies a filtering similar to the one of Listing 1.1 is achieved except that only the domain of the next variable $D(P_{i+1})$ is filtered according to the solution to Weakness 1.

Prefix Projection Decreasing Counting propagator (PPDC). The key idea of this approach is not to count the projected frequencies from scratch, but rather to *decrement* them. More specifically, when scanning the position of the current symbol at line 7, if pos happens to be the last position of a symbol ($pos == lastPosMap[sid][s[pos]]$) then $projFreqs[s[pos]]$ is decremented. This requires $projFreqs$ to be an array of reversible integers. With this strategy the loop at line 11 disappears, but in case the current sequence is not added to the projected database, the frequencies of all its last symbols occurring after pos must also be decremented. This can be done by adding an **else** block to the **if** defined at line 5 that will iterate over the $lastPosList$ and decrement the symbol frequencies.

Example 10. Assume SDB_1 . The initial projected frequency array is $projFreqs = [A:3, B:4, C:3, D:1]$. Consider now the A -projected database illustrated on

Table 2: Dataset Features. Sparsity is equal to $(\frac{1}{\#SDB} \times \sum \frac{\#s}{\#I/s})$

SDB	$\#SDB$	N	$avg(\#s)$	$avg(\#I/s)$	$max(\#s)$	sparsity	description
BIBLE	36369	13905	21.64	17.85	100	1.2	text
FIFA	20450	2990	36.24	34.74	100	1.2	web click stream
Kosarak	69999	21144	7.98	7.98	796	1.0	web click stream
Leviathan	5834	9025	33.81	26.34	100	1.3	text
PubMed	17237	19931	29.56	24.82	198	1.2	bio-medical text
data200k	200000	26	50.25	18.25	86	2.8	synthetic data
protein	103120	25	482.25	19.93	600	24.2	protein sequences

Fig. 1. The projected frequency array becomes $projFreqs=[A:0, B:3, C:2, D:0]$. The entry at A is decremented three times as pos moved beyond its $lastPos$ for each of the sequences sid_1 , sid_2 and sid_3 . Since sid_4 is removed from the projected database, the frequency of all its last symbols occurring after pos is also decremented, that is for entries B , C and D .

PP-mixed. Both PPIC and PPDC approaches can be of interest depending on the number of removed sequences in the projected database. If the number of sequences removed is large then PPIC is preferable. On the other hand is only a few sequences are removed then PPDC can be more interesting. Inspired from the *reset* idea of [11] the PP-mixed approach dynamically chooses the best strategy: if $projFreqs_{SDB}(a) < \#PSDB_i/2$ (i.e., more than half of sequences will be removed) then PPIC is used otherwise PPDC.

4.3 Constraints of SPM

We implemented common constraints such as minimum and maximum pattern size, symbol inclusion/exclusion, and regular expression constraints. Time constraints (maxgap, mingap, maxspan, etc) are outside the scope of this work: they change the definition of what a valid prefix is, and hence require changing the propagator (as in [5]).

5 Experiments

In this section, we report our experimental results on the performance of our approaches with six real-life datasets⁴ and one synthetic (data200k [14]) with various characteristics shown in Table. 2. Sparsity, representing the average of the number of symbols that appear in each sequence, is a good indicator of how sparse or dense a dataset is.

Our work is implemented in Scala in Oscala solver [9] and run under JVM with maximum memory set to 8GB. All our software, datasets and results are available online as open source in order to make this research reproducible (<http://sites.uclouvain.be/cp4dm/spm/>).

⁴ <http://www.philippe-fournier-viger.com/spmf/>

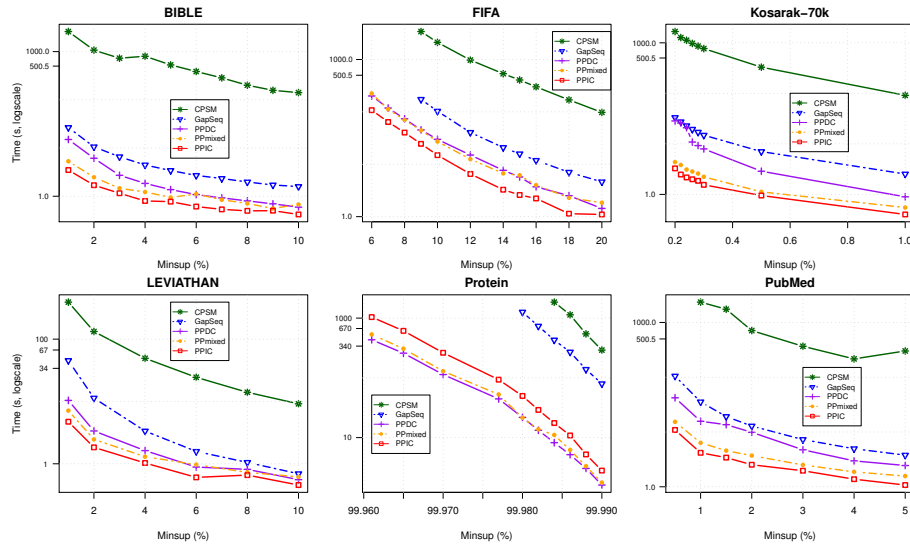


Fig. 2: CPU times for PPIC, PPDC, PPMIXED and GAP-SEQ for several minsup (missing points indicate a timeout)

We used a machine with a 2.7GHz Intel core i5 processor and 8GB of RAM with Linux 3.19.0-32-generic 64 bits distribution Mint 17.3. Execution time limit is set to 3600 seconds (1 hour). Our proposals are compared, first, with CPSM⁵[8] and GAP-SEQ⁶[5], the recently CP-based approaches including Gap constraint and the previous version of GAP-SEQ, PP⁷[6] without Gap but with regular expression constraint. Second, we made comparison with CSPADE⁸[18], PrefixSpan [10]⁹ and SPMF¹⁰.

PPIC vs PPDC vs PPMixed. The CPU time of PPIC, PPDC and PPMIXED models are shown in Fig. 2. PPIC is more efficient than PPDC in 80% of datasets. This is essentially because in many cases at the beginning of mining, there are many unsupported sequences for which the symbol counters must be decremented (compared to not having to increase the counters in PPIC). For instance with BIBLE SDB and $minsup = 10\%$ PPDC need to see 21,979,585 symbols to be complete while only 15,916,652 is needed for PPIC. Unsurprisingly, PPMIXED is between these approaches.

Our proposals vs Gap-Seq (CP method). Fig. 2 confirms CPSM is outperformed by Gap-Seq which itself improves PP (without gap). We can clearly notice our approaches outperform GAP-SEQ (and hence PP) in all cases. In the case of FIFA SDB, GAP-SEQ reach time limit when $minsup \leq 9\%$. PPIC is very effective in large and dense datasets regarding of CPU-times.

⁵ <https://dtai.cs.kuleuven.be/CP4IM/cpsm/>
⁶ <https://sites.google.com/site/cp4spm/>
⁷ <https://sites.google.com/site/prefixprojection4cp/>
⁸ <http://www.cs.rpi.edu/~zaki/www-new/pmwiki.php/Software>
⁹ <http://illimine.cs.uiuc.edu/software/>
¹⁰ <http://www.philippe-fournier-viger.com/spmf/index.php?link=download.php>

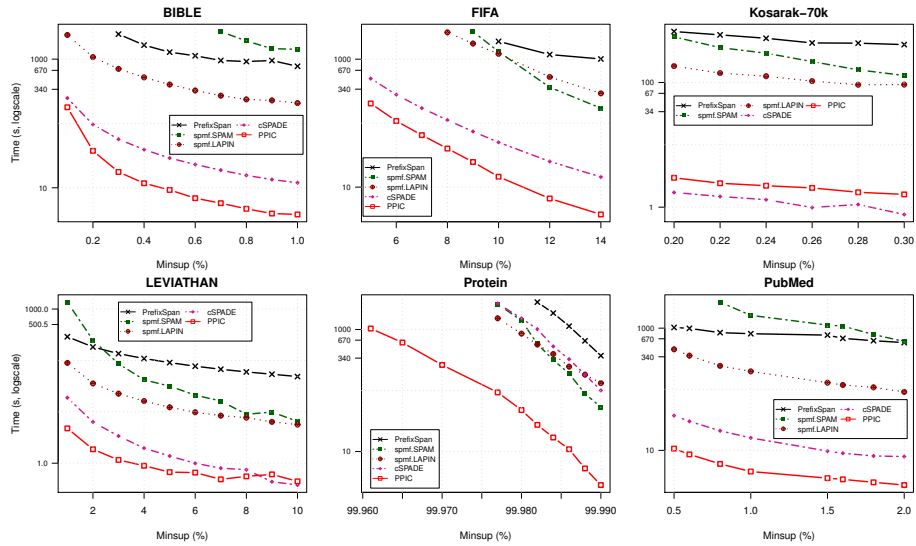


Fig. 3: CPU times for PPIC,PPDC,PPmixed and cSPADE for several *minsup*




Fig. 4: Incremental impact of our solutions to the different weaknesses

Comparison with specialized algorithms. Our third experience is the comparison with specialized algorithms. As we can see in the Fig. 3, we perform better on 84% of the datasets. However, cSPADE is still the most efficient for Kosarak. In fact, Kosarak doesn't contain any symbol repetition in its sequences. So it is a bad case for prefix-projection-based algorithms which need to scan all the positions. On the contrary, with protein dataset (the sparse one) cSPADE requires much more CPU time. The SPMF implementation of SPAM, PrefixSpan and LAPIN appears to be consistently slower than cSPADE but there is no clear domination among these.

Impact of the improvements. Figure 4 shows the incremental impact of our proposed solutions to the weaknesses defined in Section 4.2, starting from reversible vectors (fix of weakness 4) up to all our proposed modifications. Fix 1 has limited effect, while adding fix 3 is data dependent but adding fix2 always improves further.




Fig. 5: Handling of different additional constraints

Handling different additional constraints. In order to illustrate the modularity of our approach we compare with a number of user-defined constraints that can be added as additional modules without changing the main propagator (Fig. 5). (a) We compared PPIC and PP (unfortunately the GAP-SEQ tool does not support a regular expression command-line argument) under various size constraints on the protein dataset with $minsup = 99.984$. (b,c) We also selected data200k adding a regular expression constraint $RE10 = A * B(B|C)D * EF * (G|H)I*$ and $RE14 = A * (Q|BS * (B|C))D * E(I|S) * (F|H)G * R$ [14]. The last experiment reported on Fig.5d consists in combining size and symbols constraints on the protein dataset: only sequential patterns that contain VALINE and GLYCINE twice and ASPARATE and SERINE once are valid. PPIC under constraints still dominates PP.

6 Conclusion

This work improved the existing CP-based sequential pattern mining approaches [8,5] up to the point that it also surpasses specialized mining systems in terms of efficiency. To do so, we combined and adapted a number of ideas from both the sequence mining literature and the constraint programming literature; correspondingly last-position information [16] and reversible data-structures for storing and restoring state during backtracking search. We introduced the PREFIXPROJECTION-INC (PPIC) global constraint and two variants proposing different strategies to compute the projected frequencies: from scratch, by decreasing the counters, or a mix of both. These can be plugged in as modules in a CP solver. These constraints are implemented in Scala and made available in the generic Oscar solver. Furthermore, the approach is compatible with a number of constraints including size and regular expression constraints. There are other constraints which change the subsequence relation and which would hence require hardcoding changes in the propagator (gap [5], span, etc). We think many of our improvements can be applied to such settings as well.

Our work shows that generic CP solvers can indeed be used as framework to build scalable mining algorithms, not just for generic yet less scalable systems as was done for itemset mining [4]. Furthermore, advanced data-structures for backtracking search, such as trailing and reversible vectors, can also be used in non-CP algorithms. This appears to be an understudied aspect of backtracking algorithms in pattern mining and data mining in general. We believe there is much more potential for combinations of techniques from data mining and CP.

References

1. Agrawal, R., Srikant, R.: Mining sequential patterns. In: Data Engineering, 1995. Proceedings of the Eleventh International Conference on. pp. 3–14. IEEE (1995)
2. Ayres, J., Flannick, J., Gehrke, J., Yiu, T.: Sequential pattern mining using a bitmap representation. In: ACM SIGKDD. pp. 429–435 (2002)
3. Coquery, E., Jabbour, S., Sais, L., Salhi, Y.: A SAT-based approach for discovering frequent, closed and maximal patterns in a sequence. In: ECAI (2012)
4. Guns, T., Nijssen, S., De Raedt, L.: Itemset mining: A constraint programming perspective. *Artificial Intelligence* 175(12), 1951–1983 (2011)
5. Kemmar, A., Loudni, S., Lebbah, Y., Boizumault, P., Charnois, T.: A global constraint for mining sequential patterns with gap constraint. CPAIOR16 (2015)
6. Kemmar, A., Loudni, S., Lebbah, Y., Boizumault, P., Charnois, T.: Prefix-projection global constraint for sequential pattern mining. In: Principles and Practice of Constraint Programming. Springer (2015)
7. Mabroukeh, N.R., Ezeife, C.I.: A taxonomy of sequential pattern mining algorithms. *ACM Comput. Surv.* 43(1), 3:1–3:41 (Dec 2010)
8. Negrevergne, B., Guns, T.: Constraint-based sequence mining using constraint programming. In: CPAIOR15. Springer (2015)
9. OscaR Team: OscaR: Scala in OR (2012), available from <https://bitbucket.org/oscarlib/oscar>
10. Pei, J., Han, J., Mortazavi-Asl, B., Pinto, H., Chen, Q., Dayal, U., Hsu, M.C.: Prefixspan: Mining sequential patterns efficiently by prefix-projected pattern growth. In: icccn. p. 0215. IEEE (2001)
11. Perez, G., Regin, J.C.: Improving gac-4 for table and mdd constraints. In: Principles and Practice of Constraint Programming. Springer (2014)
12. Rossi, F., Van Beek, P., Walsh, T.: Handbook of CP. Elsevier (2006)
13. Schulte, C., Carlsson, M.: Finite domain constraint programming systems. Handbook of constraint programming pp. 495–526 (2006)
14. Trasarti, R., Bonchi, F., Goethals, B.: Sequence mining automata: A new technique for mining frequent sequences under regular expressions. In: Data Mining, 2008. ICDM'08. Eighth IEEE International Conference on. pp. 1061–1066. IEEE (2008)
15. Yan, X., Han, J., Afshar, R.: Clospan: Mining closed sequential patterns in large datasets. In: In SDM. pp. 166–177. SIAM (2003)
16. Yang, Z., Kitsuregawa, M.: LAPIN-SPAM: An Improved Algorithm for Mining Sequential Pattern. In: International Conference on Data Engineering (2005)
17. Yang, Z., Wang, Y., Kitsuregawa, M.: LAPIN: effective sequential pattern mining algorithms by last position induction for dense databases. In: DAFSAA. pp. 1020–1023 (2007)
18. Zaki, M.J.: Sequence mining in categorical domains: incorporating constraints. In: Proceedings of the ninth international conference on Information and knowledge management. pp. 422–429. ACM (2000)