

Flow-Based Propagators for the SEQUENCE and Related Global Constraints¹

Michael Maher² and Nina Narodytska³ and Claude-Guy Quimper⁴ and Toby Walsh⁵

Abstract. We propose new filtering algorithms for the SEQUENCE constraint and several extensions which are based on network flows. Our propagator for the SEQUENCE constraint enforces domain consistency in $O(n^2)$ time down a branch of the search tree. This improves upon the best existing domain consistency algorithm by a factor of $O(\log n)$. The flows used in these algorithms are derived from a linear program. Some of them differ from the flows used to propagate global constraints like GCC since the domains of the variables are encoded as costs on the edges rather than capacities. Such flows are efficient for maintaining bounds consistency over large domains and may be useful for other global arithmetic constraints.

1 Introduction

Graph based algorithms play a very important role in constraint programming, especially within propagators for global constraints. For example, Regin’s propagator for the ALLDIFFERENT constraint is based on a perfect matching algorithm [14], whilst his propagator for the GCC constraint is based on a network flow algorithm [15]. Both these graph algorithms are derived from the bipartite value graph, in which nodes represent variables and values, and edges represent domains. For example, the GCC propagator finds a flow in such a graph in which each unit of flow represents the assignment of a particular value to a variable. In this paper, we identify a new way to build graph based propagators for global constraints: we convert the global constraint into a linear program and then convert this into a network flow. These encodings contain several novelties. For example, variables domain bounds can be encoded as costs along the edges. We apply this approach to the SEQUENCE family of constraints. Our results widen the class of global constraints which can be propagated using flow-based algorithms. We conjecture that these methods will be useful to propagate other global constraints.

2 Background

A constraint satisfaction problem (CSP) consists of a set of variables, each with a finite domain of values, and a set of constraints specifying allowed combinations of values for subsets of variables. We use capital letters for variables (e.g. X , Y and S), and lower case for values (e.g. d and d_i). A solution is an assignment of values to the

variables satisfying the constraints. Constraint solvers typically explore partial assignments enforcing a local consistency property using either specialized or general purpose propagation algorithms. A *support* for a constraint C is a tuple that assigns a value to each variable from its domain which satisfies C . A *bounds support* is a tuple that assigns a value to each variable which is between the maximum and minimum in its domain which satisfies C . A constraint is *domain consistent (DC)* iff for each variable X_i , every value in the domain of X_i belongs to a support. A constraint is *bounds consistent (BC)* iff for each variable X_i , there is a bounds support for the maximum and minimum value in its domain. A CSP is DC/BC iff each constraint is DC/BC. A constraint is *monotone* iff there exists a total ordering \prec of the domain values such that for any two values v, w if $v \prec w$ then v is substitutable for w in any support for C .

We also give some background on flows. A *flow network* is a weighted directed graph $G = (V, E)$ where each edge e has a capacity between non-negative integers $l(e)$ and $u(e)$, and an integer cost $w(e)$. A *feasible flow* in a flow network between a source (s) and a sink (t), (s, t) -flow, is a function $f : E \rightarrow \mathbb{Z}^+$ that satisfies two conditions: $f(e) \in [l(e), u(e)]$, $\forall e \in E$ and the *flow conservation* law that ensures that the amount of incoming flow should be equal to the amount of outgoing flow for all nodes except the source and the sink. The *value* of a (s, t) -flow is the amount of flow leaving the sink s . The *cost* of a flow f is $w(f) = \sum_{e \in E} w(e)f(e)$. A *minimum cost flow* is a feasible flow with the minimum cost. The Ford-Fulkerson algorithm can find a feasible flow in $O(\phi(f)|E|)$ time. If $w(e) \in \mathbb{Z}$, $\forall e \in E$, then a minimum cost feasible flow can be found using the successive shortest path algorithm in $O(\phi(f)SPP)$ time, where SPP is the complexity of finding a shortest path in the residual graph. Given a (s, t) -flow f in $G(V, E)$, the *residual graph* G_f is the directed graph (V, E_f) , where $E_f = \{e \text{ with cost } w(e) \text{ and capacity } 0..(u(e) - f(e)) \mid e = (u, v) \in E, f(e) < u(e)\} \cup \{e \text{ with cost } -w(e) \text{ and capacity } 0..(f(e) - l(e)) \mid e = (u, v) \in E, l(e) < f(e)\}$. There are other asymptotically faster but more complex algorithms for finding either feasible or minimum-cost flows [2].

In our flow-based encodings, a consistency check will correspond to finding a feasible or minimum cost flow. To enforce DC, we therefore need an algorithm that, given a minimum cost flow of cost $w(f)$ and an edge e checks if an extra unit flow can be pushed (or removed) through the edge e and the cost of the resulting flow is less than or equal to a given threshold T . We use the residual graph to construct such an algorithm. Suppose we need to check if an extra unit flow can be pushed through an edge $e = (u, v)$. Let $e' = (u, v)$ be the corresponding arc in the residual graph. If $w(e) = 0$, $\forall e \in E$, then it is sufficient to compute strongly connected components (SCC) in the residual graph. An extra unit flow can be pushed through an edge e

¹ NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council.

² NICTA and UNSW, Sydney, Australia

³ NICTA and UNSW, Sydney, Australia

⁴ Ecole Polytechnique de Montreal, Montreal, Canada

⁵ NICTA and UNSW, Sydney, Australia

iff both ends of the edge e' are in the same strongly connected component. If $w(e) \in \mathbb{Z}, \forall e \in E$, the shortest path p between v and u in the residual graph has to be computed. The minimal cost of pushing an extra unit flow through an edge e equals $w(f) + w(p) + w(e)$. If $w(f) + w(p) + w(e) > T$, then we cannot push an extra unit through e . Similarly, we can check if we can remove a unit flow through an edge.

3 The SEQUENCE Constraint

The SEQUENCE constraint was introduced by Beldiceanu and Contejean [5]. It constrains the number of values taken from a given set in any sequence of k variables. It is useful in staff rostering to specify, for example, that every employee has at least 2 days off in any 7 day period. Another application is sequencing cars along a production line (prob001 in CSPLib). It can specify, for example, that at most 1 in 3 cars along the production line has a sun-roof. The SEQUENCE constraint can be defined in terms of a conjunction of AMONG constraints. $\text{AMONG}(l, u, [X_1, \dots, X_k], v)$ holds iff $l \leq |\{i | X_i \in v\}| \leq u$. That is, between l and u of the k variables take values in v . The AMONG constraint can be encoded by channelling into 0/1 variables using $Y_i \leftrightarrow (X_i \in v)$ and $l \leq \sum_{i=1}^k Y_i \leq u$. Since the constraint graph of this encoding is Berge-acyclic, this does not hinder propagation. Consequently, we will simplify notation and consider AMONG (and SEQUENCE) on 0/1 variables and $v = \{1\}$. If $l = 0$, AMONG is an ATMOST constraint. ATMOST is *monotone* since, given a support, we also have support for any larger value [6]. The SEQUENCE constraint is a conjunction of overlapping AMONG constraints. More precisely, $\text{SEQUENCE}(l, u, k, [X_1, \dots, X_n], v)$ holds iff for $1 \leq i \leq n - k + 1$, $\text{AMONG}(l, u, [X_i, \dots, X_{i+k-1}], v)$ holds. A sequence like X_i, \dots, X_{i+k-1} is a *window*. It is easy to see that this decomposition hinders propagation. If $l = 0$, SEQUENCE is an ATMOSTSEQ constraint. Decomposition in this case does not hinder propagation. Enforcing DC on the decomposition of an ATMOSTSEQ constraint is equivalent to enforcing DC on the ATMOSTSEQ constraint [6].

Several filtering algorithms exist for SEQUENCE and related constraints. Regin and Puget proposed a filtering algorithm for the Global Sequencing constraint (GSC) that combines a SEQUENCE and a global cardinality constraint (GCC) [17]. Beldiceanu and Carlsson suggested a greedy filtering algorithm for the CARDPATH constraint that can be used to propagate the SEQUENCE constraint, but this may hinder propagation [3]. Regin decomposed GSC into a set of variable disjoint AMONG and GCC constraints [16] but this decomposition also hinders propagation. Bessiere *et al.* [6] encoded SEQUENCE using a SLIDE constraint, and give a domain consistency propagator that runs in $O(nd^{k-1})$ time. van Hoeve *et al.* [13] proposed two filtering algorithms that establish domain consistency. The first is based on an encoding into a REGULAR constraint and runs in $O(n2^k)$ time, whilst the second is based on cumulative sums and runs in $O(n^3)$ time. Finally, Brand *et al.* [9] studied a number of different encodings of the SEQUENCE constraint. Their asymptotically fastest encoding is based on separation theory and enforces domain consistency in $O(n^2 \log n)$ time down the whole branch of a search tree. One of our contributions is to improve on this bound.

4 Flow-based Propagator for the SEQUENCE Constraint

We will convert the SEQUENCE constraint to a flow by means of a linear program (LP). We shall use $\text{SEQUENCE}(l, u, 3, [X_1, \dots, X_6], v)$

as a running example. We can formulate this constraint simply and directly as an integer linear program:

$$\begin{aligned} l &\leq X_1 + X_2 + X_3 \leq u, \\ l &\leq X_2 + X_3 + X_4 \leq u, \\ l &\leq X_3 + X_4 + X_5 \leq u, \\ l &\leq X_4 + X_5 + X_6 \leq u \end{aligned}$$

where $X_i \in \{0, 1\}$. By introducing surplus/slack variables, Y_i and Z_i , we convert this to a set of equalities:

$$\begin{aligned} X_1 + X_2 + X_3 - Y_1 &= l, & X_1 + X_2 + X_3 + Z_1 &= u, \\ X_2 + X_3 + X_4 - Y_2 &= l, & X_2 + X_3 + X_4 + Z_2 &= u, \\ X_3 + X_4 + X_5 - Y_3 &= l, & X_3 + X_4 + X_5 + Z_3 &= u, \\ X_4 + X_5 + X_6 - Y_4 &= l, & X_4 + X_5 + X_6 + Z_4 &= u \end{aligned}$$

where $Y_i, Z_i \geq 0$. In matrix form, this is:

$$\begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} X_1 \\ \vdots \\ X_6 \\ Y_1 \\ Z_1 \\ \vdots \\ Y_4 \\ Z_4 \end{pmatrix} = \begin{pmatrix} l \\ u \\ l \\ u \\ l \\ u \\ l \\ u \end{pmatrix}$$

This matrix has the *consecutive ones* property for columns: each column has a block of consecutive 1's or -1's and the remaining elements are 0's. Consequently, we can apply the method of Veinott and Wagner [1] (also described in Application 9.6 of [2]) to simplify the problem. We create a zero last row and subtract the i th row from $i + 1$ th row for $i = 1$ to $2n$. These operations do not change the set of solutions. This gives:

$$A\vec{X} = \vec{b},$$

where

$$A = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 1 & 0 & 0 & 0 & -1 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & -1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & -1 & -1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix},$$

$$\vec{X} = (X_1, \dots, X_6, Y_1, Z_1, \dots, Y_4, Z_4)^T,$$

$$\vec{b} = (l, u - l, l - u, u - l, l - u, u - l, l - u, u - l, -u)^T$$

This matrix has a single 1 and -1 in each column. Hence, it describes a network flow problem [2] on a graph $G = (V, E)$ (that is, it is a network matrix). Each row in the matrix corresponds to a node in V and each column corresponds to an edge in E . Down each column, there is a single row i equal to 1 and a single row j equal to -1 corresponding to an edge $(i, j) \in E$ in the graph. We include a source node s and a sink node t in V . Let b be the vector on the right hand side of the equation. If b_i is positive, then there is an edge $(s, i) \in E$ that carries exactly b_i amount of flow. If b_i is negative, there is an edge $(i, t) \in E$ that carries exactly $|b_i|$ amount of flow. The bounds on the variables, which are not expressed in the matrix, are represented as bounds on the capacity of the corresponding edges.

The graph for the set of equations in the example is given in Figure 1. A flow of value $4u - 3l$ in the graph corresponds to a solution. If a feasible flow sends a unit flow through the edge labeled with X_i then $X_i = 1$ in the solution; otherwise $X_i = 0$. Each even numbered

$$A\vec{X} = \vec{b},$$

where

$$A = [A1 \mid A2 \mid -A2]$$

$$A1 = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & -1 & -1 & -1 \end{pmatrix}$$

$$A2 = \begin{pmatrix} -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix}$$

$$\vec{X} = (X_1, \dots, X_6, Y_1, Z_1, \dots, Y_4, Z_4, Q_1, P_1, \dots, Q_4, P_4)^T$$

$$\vec{b} = (l, u - l, l - u, u - l, l - u, u - l, l - u, u - l, -u)^T$$

The flow graph $G = (V, E)$ for the transformed system is presented in Figure 2. Dashed edges have cost 1, while other edges have cost 0. The minimal cost flow in the graph corresponds to a minimal cost solution to the system of equations.

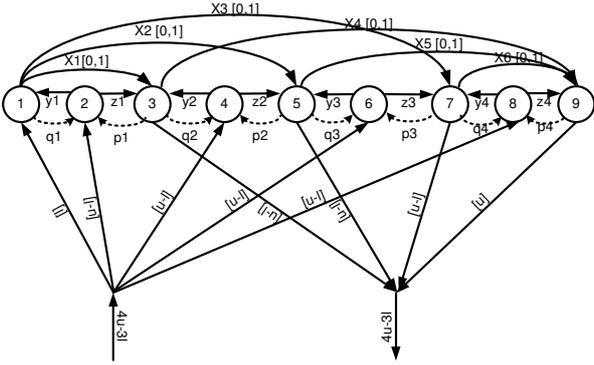


Figure 2. A flow graph for $\text{SOFTSEQUENCE}(l, u, 3, t, [X_1, \dots, X_6])$

Theorem 2 For any constraint $\text{SOFTSEQUENCE}(l, u, k, t, [X_1, \dots, X_n], v)$, there is an equivalent network flow graph. There is a one-to-one correspondence between solutions of the constraint and feasible flows of cost less than or equal to t .

Thus, if the minimal cost flow is greater than $\max(\text{dom}(T))$, then the SOFTSEQUENCE constraint is inconsistent. The minimal cost flow can be found in $O(|V||E| \log \log U \log |V|C) = O(n^2 \log n \log \log u)$ time [2]. Consider the edge (u, v) in the residual graph associated to variable X_i and let $k_{(u,v)}$ be its residual cost. If the flow corresponds to an assignment with $X_i = 0$, pushing a unit of flow on (u, v) results in a solution with $X_i = 1$. Symmetrically, if the flow corresponds to an assignment with $X_i = 1$, pushing a unit of flow on (u, v) results in a solution with $X_i = 0$. If the shortest path in the residual graph between v and u is $k_{(v,u)}$, then the shortest cycle that contains (u, v) has length $k_{(u,v)} + k_{(v,u)}$. Pushing a unit of flow through this cycle results in a flow of cost

$c + k_{(u,v)} + k_{(v,u)}$ which is the minimum-cost flow that contains the edge (u, v) . If $c + k_{(u,v)} + k_{(v,u)} > \max(\text{dom}(T))$, then no flows containing the edge (u, v) exist with a cost smaller or equal to $\max(\text{dom}(T))$. The variable X_i must therefore be fixed to the value taken in the current flow. Following Equation 1, the cost of the variable T must be no smaller than the cost of the solution. To enforce bounds consistency on the cost variable, we increase the lower bound of $\text{dom}(T)$ to the cost of the minimum flow in the graph G .

To enforce DC on the X variables efficiently we can use an all pairs shortest path algorithm on the residual graph. This takes $O(n^2 \log n)$ time using Johnson's algorithm [10]. This gives an $O(n^2 \log n \log \log u)$ time complexity to enforce DC on the SOFTSEQUENCE constraint. The penalty variables used for SOFTSEQUENCE arise directly out of the problem description and occur naturally in the LP formulation. We could also view them as arising through the methodology of [20], where edges with costs are added to the network graph for the hard constraint to represent the softened constraint.

5.1 Soft ATMOSTSEQ Constraint

In many cases, we have only upper bounds and not lower bounds on the frequency of the occurrence of values (i.e. $l = 0$). For instance, this is the case in car sequencing problems. This can be used to simplify propagation. For example, there is a simple propagator to enforce DC on the soft ATMOSTSEQ constraint in just $O(n^2 k)$ time down a branch of the search tree. Consider the assignment which assigns each X_i the smallest value in its domain. Due to monotonicity of the ATMOSTSEQ constraint any other solution X'_i will be greater or equal to this minimal assignment: $X_i \leq X'_i, i = 1, \dots, n$. The violation measure is a monotonically non-decreasing function of the X_i . Consequently, the violation cost for any other solution is greater or equal to the violation cost of this minimal assignment. Hence, if the violation cost for the minimal assignment is greater than the upper bound on the cost variable then the constraint is inconsistent. To enforce DC on soft ATMOSTSEQ , we can use the failed literal test. If a value is pruned from the domain of X_i , then it takes $O(k)$ time to update the cost value of the minimal assignment and $O(nk)$ time to perform the failed literal test for n Boolean variables. Hence, the total time complexity is $O(n^2 k)$ down a branch of the search tree.

6 Generalized SEQUENCE Constraint

To model real world problems, we may want to have different size or positioned windows. For example, the window size in a rostering problem may depend on whether it includes a weekend or not. An extension of the SEQUENCE constraint proposed in [13] is that each AMONG constraint can have different parameters (start position, l , u , and k). More precisely, $\text{GEN-SEQUENCE}(\vec{p}_1, \dots, \vec{p}_m, [X_1, X_2, \dots, X_n], v)$ holds iff $\text{AMONG}(l_i, u_i, k_i, [X_{s_i}, \dots, X_{s_i+k_i-1}], v)$ for $1 \leq i \leq m$ where $\vec{p}_i = \langle l_i, u_i, k_i, s_i \rangle$. Whilst the methods in Section 4 easily extend to allow different bounds l and u for each window, dealing with different windows is more difficult. In general, the matrix now does not have the consecutive ones property. It may be possible to re-order the windows to achieve the consecutive ones property. If such a re-ordering exists, it can be found and performed in $O(m + n + r)$ time, where r is the number of non-zero entries in the matrix [8]. Even when re-ordering cannot achieve the consecutive ones property there may, nevertheless, be an equivalent network matrix. Bixby

and Cunningham [7] give a procedure⁶ to find an equivalent network matrix, when it exists, in $O(mr)$ time. In these cases, the method in Section 4 can be applied to propagate the GEN-SEQUENCE constraint in $O(n^2)$ time down the branch of a search tree.

Not all GEN-SEQUENCE constraints can be expressed as network flows. Consider the GEN-SEQUENCE constraint with $n = 5$, identical upper and lower bounds (l and u), and 4 windows: [1,5], [2,4], [3,5], and [1,3]. We can express it as an integer linear program:

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 \\ -1 & -1 & -1 & -1 & -1 \\ 0 & 1 & 1 & 1 & 0 \\ 0 & -1 & -1 & -1 & 0 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & -1 & -1 & -1 \\ 1 & 1 & 1 & 0 & 0 \\ -1 & -1 & -1 & 0 & 0 \end{pmatrix} \begin{pmatrix} X_1 \\ X_2 \\ X_3 \\ X_4 \\ X_5 \end{pmatrix} \geq \begin{pmatrix} l \\ -u \\ l \\ -u \\ l \\ -u \\ l \\ -u \end{pmatrix} \quad (2)$$

Applying the test described in Section 20.1 of [18] to Example 2, we find that the matrix of this problem is not equivalent to any network matrix.

However, all GEN-SEQUENCE constraint matrices satisfy a weaker property: total unimodularity. A matrix is *totally unimodular* iff every square non-singular submatrix has a determinant of $+1$ or -1 . The advantage of this property is that any totally unimodular system of inequalities with integral constants is solvable in \mathbb{Z} iff it is solvable in \mathbb{R} .

Theorem 3 *The matrix of the inequalities associated with GEN-SEQUENCE constraint is totally unimodular.*

In practice, only integral values for the bounds l_i and u_i are used. Thus the consistency of a GEN-SEQUENCE constraint can be determined via linear programming techniques in $O(n^{3.5} \log u)$ time. Using the failed literal test, we can enforce DC at a cost of $O(n^{5.5} \log u)$ down the branch of a search tree for any GEN-SEQUENCE constraint. This is too expensive to be practical. We can, instead, exploit the fact that the matrix for each GEN-SEQUENCE constraint has the consecutive ones property for rows (before the introduction of slack/surplus variables). Corresponding to the row transformation for matrices with consecutive ones for columns is a change-of-variables transformation into variable $S_j = \sum_{i=1}^j X_i$ for matrices with consecutive ones for rows. This gives the dual of a network matrix. This is the basis of an encoding of SEQUENCE in [9] (denoted there CD). Consequently that encoding extends to GEN-SEQUENCE. Adapting the analysis in [9] to GEN-SEQUENCE, we can enforce DC in $O(nm + n^2 \log n)$ time down the branch of a search tree.

In summary, for a compilation cost of $O(mr)$, we can enforce DC on a GEN-SEQUENCE constraint in $O(n^2)$ down the branch of a search tree, when it has a flow representation, and in $O(nm + n^2 \log n)$ when it does not.

7 SLIDINGSUM Constraint

The SLIDINGSUM constraint [4] is a generalization of the SEQUENCE constraint from Boolean to integer variables, which we extend to allow arbitrary windows. SLIDINGSUM ($[X_1, \dots, X_n], [\vec{p}_1, \dots, \vec{p}_m]$) holds iff $l_i \leq \sum_{j=s_i}^{s_i+k_i-1} X_j \leq u_i$ holds where $\vec{p}_i = \langle l_i, u_i, k_i, s_i \rangle$ is, as with the generalized SEQUENCE, a window. The constraint can be expressed as a linear program \mathcal{P} called the *primal* where W is a matrix encoding the inequalities and the bounds on each variable are given by

$a_i \leq X_i \leq b_i$. Since the constraint represents a satisfaction problem, we minimize the constant 0.

$$\left. \begin{array}{l} \min 0 \\ \left[\begin{array}{c} W \\ -W \\ I \\ -I \end{array} \right] X \geq \left[\begin{array}{c} l \\ -u \\ a \\ -b \end{array} \right] \end{array} \right\} \mathcal{P} \quad (3)$$

The dual \mathcal{D} is however an optimization problem.

$$\left. \begin{array}{l} \min [-l \quad u \quad -a \quad b] Y \\ [W^T \quad -W^T \quad I \quad -I] Y = 0 \\ Y \geq 0 \end{array} \right\} \mathcal{D} \quad (4)$$

Von Neumann's Strong Duality Theorem states that if the primal and the dual problems are feasible, then they have the same objective value. Moreover, if the primal is unsatisfiable, the dual is unbounded. The SLIDINGSUM constraint is thus satisfiable if the objective function of the dual problem is zero. It is unsatisfiable if it tends to negative infinity.

Note that the matrix W^T has the consecutive ones property on the columns. The dual problem can thus be converted to a network flow using the same transformation as with the SEQUENCE constraint. Consider the dual LP of our running example 2:

$$\begin{array}{l} \text{Minimize} \\ -\sum_{i=1}^4 l_i Y_i + \sum_{i=1}^4 u_i Y_{4+i} - \sum_{i=1}^5 a_i Y_{8+i} + \sum_{i=1}^5 b_i Y_{13+i} \\ \text{subject to:} \end{array}$$

$$A\vec{Y} = \vec{b},$$

where

$$A = \begin{pmatrix} 1 & 0 & 0 & 1 & -1 & 0 & 0 & -1 & 1 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & -1 & -1 & 0 & -1 & 0 & 1 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & -1 & -1 & -1 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ 1 & 1 & 1 & 0 & -1 & -1 & -1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\ 1 & 0 & 1 & 0 & -1 & 0 & -1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix},$$

$$\vec{Y} = (Y_1, \dots, Y_{18})^T,$$

$$\vec{b} = (0, \dots, 0)^T.$$

Our usual transformation will turn this into a network flow problem:

$$\begin{array}{l} \text{Minimize} \\ -\sum_{i=1}^4 l_i Y_i + \sum_{i=1}^4 u_i Y_{4+i} - \sum_{i=1}^5 a_i Y_{8+i} + \sum_{i=1}^5 b_i Y_{13+i} \end{array}$$

subject to

$$A\vec{Y} = \vec{0},$$

where

$$A = \begin{pmatrix} 1 & 0 & 0 & 1 & -1 & 0 & 0 & -1 & 1 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & -1 & 0 & 0 & -1 & 1 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & -1 & 0 & 0 & -1 & 1 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 & 1 \\ 0 & -1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 & 0 & 1 \\ -1 & 0 & -1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 1 \end{pmatrix},$$

$$\vec{Y} = (Y_1, \dots, Y_{18})^T.$$

The flow associated with this example is given in Figure 3. There are $n + 1$ nodes labelled from 1 to $n + 1$ where node i is connected

⁶ Another procedure is given in [18].

to node $i + 1$ with an edge of cost $-a_i$ and node $i + 1$ is connected to node i with an edge of cost b_i . For each window \vec{p}_i , we have an edge from s_i to $s_i + k_i$ with cost $-l_i$ and an edge from $s_i + k_i$ to s_i with cost u_i . All nodes have a null supply and a null demand. A flow is therefore simply a circulation i.e., an amount of flow pushed on the cycles of the graph.

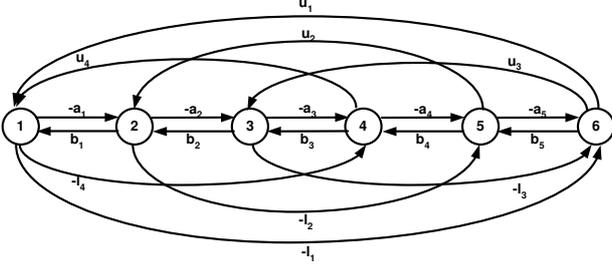


Figure 3. Network flow associated to the SLIDINGSUM constraint posted on the running example.

Theorem 4 *The SLIDINGSUM constraint is satisfiable if and only if there are no negative cycles in the flow graph associated with the dual linear program.*

Proof: If there is a negative cycle in the graph, then we can push an infinite amount of flow resulting in a cost infinitely small. Hence the dual problem is unbounded, and the primal is unsatisfiable. Suppose that there are no negative cycles in the graph. Pushing any amount of flow over a cycle of positive cost results in a flow of cost greater than zero. Such a flow is not optimal since the null flow has a smaller objective value. Pushing any amount of flow over a null cycle does not change the objective value. Therefore the null flow is an optimal solution and since this solution is bounded, then the primal is satisfiable. Note that the objective value of the dual (zero) is in this case equal to the objective value of the primal. \diamond

The flow graph has $O(n)$ nodes and $O(n + m)$ edges. Testing whether there is a negative cycle takes $O(n^2 + nm)$ time using the Bellman-Ford algorithm. We can use this consistency test to construct an efficient bounds consistency propagator. We find for each variable X_i the smallest (largest) value in its domain such that assigning this value to X_i does not create a negative cycle. We compute the shortest path between all pairs of nodes. Johnson's algorithm solves the all-pair shortest path problem in $O(|V|^2 \log |V| + |V||E|)$ time which in our case gives $O(n^2 \log n + nm)$ time. Suppose that the shortest path between i and $i + 1$ has length $s(i, i + 1)$, then for the constraint to be satisfiable, we need $b_i + s(i, i + 1) \geq 0$. Since b_i is a value potentially taken by X_i , we need to have $X_i \geq -s(i, i + 1)$. We therefore assign $\min(\text{dom}(X_i)) \leftarrow \max(\min(\text{dom}(X_i)), -s(i, i + 1))$. Similarly, let the length of the shortest path between $i + 1$ and i be $s(i + 1, i)$. For the constraint to be satisfiable, we need $s(i + 1, i) - a_i \geq 0$. Since a_i is a value potentially taken by X_i , we have $X_i \leq s(i + 1, i)$. We assign $\max(X_i) \leftarrow \min(\max(X_i), s(i + 1, i))$. It is not hard to prove this is sound and complete, removing all values that cause negative cycles. Following [9], we can make the propagator incremental using the algorithm by Cotton and Maler [11] to maintain the shortest path between $|P|$ pairs of nodes in $O(|E| + |V| \log |V| + |P|)$ time upon edge reduction. Each time a lower bound a_i is increased or an

upper bound b_i is decreased, the shortest paths can be recomputed in $O(m + n \log n)$ time.

8 Soft SLIDINGSUM Constraint

The soft SLIDINGSUM constraint is an extension of the SLIDINGSUM constraint. The soft SLIDINGSUM $([X_1, \dots, X_n], [\vec{p}_1, \dots, \vec{p}_m], T)$ introduces a violation variable T and is defined as follow.

$$T \geq \sum_{i=1}^m \max(l_i - \sum_{j=s_i}^{s_i+k_i-1} X_j, \sum_{j=s_i}^{s_i+k_i-1} X_j - u_i, 0) \quad (5)$$

To express the soft SLIDINGSUM constraint as a linear program, we introduce penalty variables for each inequality associated to the hard SLIDINGSUM, namely, Q_i and P_i , $i = 1, \dots, m$ and minimize the sum of penalty variables:

$$\min \sum_{i=1}^m Q_i + P_i \quad (6)$$

$$\sum_{j \in s_i}^{s_i+k_i-1} X_j + Q_i \geq l_i \quad \forall 1 \leq i \leq m \quad (7)$$

$$\sum_{j \in s_i}^{s_i+k_i-1} -X_j + P_i \geq -u_i \quad \forall 1 \leq i \leq m \quad (8)$$

$$X_i \geq a_i, -X_i \geq -b_i, Q_i \geq 0, P_i \geq 0 \quad (9)$$

Rewriting system (6)– (9) in the matrix form, we obtain primal linear program \mathcal{P} :

$$\left. \begin{array}{l} \min e^T Q + e^T P \\ \left[\begin{array}{ccc} W & I_m & 0 \\ -W & 0 & I_m \\ I_n & 0 & 0 \\ -I_n & 0 & 0 \\ 0 & I_m & 0 \\ 0 & 0 & I_m \end{array} \right] \begin{bmatrix} X \\ Q \\ P \end{bmatrix} \geq \left[\begin{array}{c} l \\ -u \\ a \\ -b \\ 0 \\ 0 \end{array} \right] \end{array} \right\} \mathcal{P} \quad (10)$$

where I_n is the $n \times n$ identity matrix, l and u are the vectors containing the m values l_i and u_i , a and b are the vectors containing the n lower and upper bounds a_i and b_i , and e is the vector of dimension m with all components set to one.

The dual problem \mathcal{D} corresponding to the primal problem \mathcal{P} (system (10)) is

$$\left. \begin{array}{l} \min [-l \quad u \quad -a \quad b \quad 0 \quad 0]^T Y \\ \left[\begin{array}{cccccc} W^T & -W^T & I_n & -I_n & 0 & 0 \\ I_m & 0 & 0 & 0 & I_m & 0 \\ 0 & I_m & 0 & 0 & 0 & I_m \end{array} \right] Y = \left[\begin{array}{c} 0 \\ e \\ e \end{array} \right] \\ Y \geq 0 \end{array} \right\} \mathcal{D} \quad (11)$$

where Y is a vector of $4n + 2m$ dual variables.

The dual problem \mathcal{D} can be transformed using row operations to obtain the consecutive ones property on the columns of the matrix. Note that W^T already has the consecutive ones property. For each of the first m columns, one needs to obtain ones between the last entry in W^T set to one and the identity matrix under W^T . This is done

by selecting the row in the identity matrix whose corresponding column is set to one and adding this row to every row above until the consecutive ones property is reached on this column. The principle applies to the m following columns except that the last m equations are negated to obtain columns with negative ones. The following $2n$ columns already had the consecutive ones property and remain unchanged during the transformation. The last $2m$ columns are modified but still satisfy the consecutive ones property. Using the same technique for the SEQUENCE constraint, we obtain a system that can be solved using a network flow algorithm.

Theorem 5 *There is a one-to-one correspondence between solutions of the soft SLIDINGSUM constraint and feasible flows of cost less than or equal to the upper bound of T .*

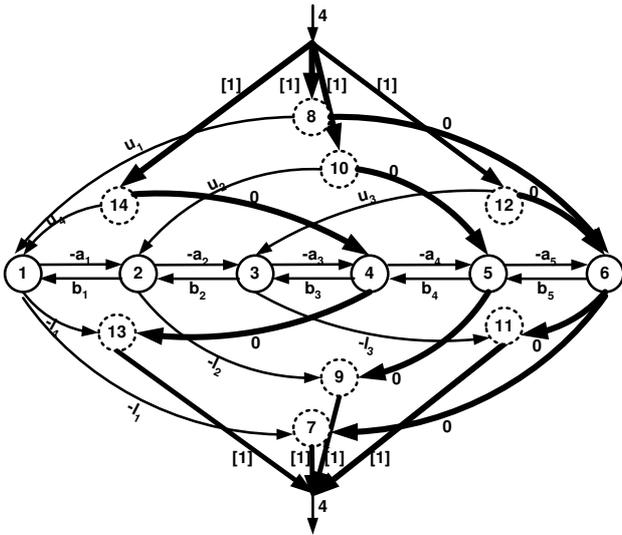


Figure 4. Network flow associated with the soft SLIDINGSUM constraint posted on the running example. The edge capacities are written in square brackets $[]$ to differentiate them from the edge costs. Bold edges show a possible flow in the network.

Figure 8 shows the flow graph for the soft version of the running example in Section 7. Note that the flow graphs for the hard and soft SLIDINGSUM constraints have a very similar structure. Consider the flow graph for the SLIDINGSUM constraint in the running example (Figure 3). It includes arcs that correspond to original variables X and are labeled $-a_i$ or b_i , $i = 1, \dots, n$ and arcs that correspond to linear inequalities and are labeled $-l_i$ or u_i , $i = 1, \dots, m$. The flow graph for the soft SLIDINGSUM constraint contains the same arcs for variables, however, each of the inequality arcs is split into two arcs by introducing a node with unit demand or supply⁷. This leads to a difference between the two flow graphs: the flow graph for the SLIDINGSUM constraint has zero flow circulation, while the flow graph for the soft SLIDINGSUM constraint contains a flow of value m . Note that the capacity of each edge connecting either the source or the sink to other nodes is exactly 1. However, the capacities of the other edges in the flow network are not bounded. Therefore, these edges can carry several units of flow in a feasible minimum cost flow, which makes a flow-based BC propagator for the

⁷ Note that, in contrast to the soft SEQUENCE constraint, this flow graph is not obtained by the methodology of [20].

soft SLIDINGSUM constraint more computationally expensive compared to the hard case. The flow graph for the soft SLIDINGSUM constraint has $O(n + m)$ nodes and $O(n + m)$ edges. The minimal cost flow can be found in $O(|V| \log |E| (|E| + |V| \log |V|)) = O((n + m)^2 \log^2(n + m))$ time [2]. The BC filtering algorithm for the soft SLIDINGSUM constraint works exactly the same as for the soft SEQUENCE constraint (Section 5), except that finding all pairs of shortest paths is replaced with finding all pairs of minimal cost flows. Hence, the total time complexity of the flow-based BC filtering algorithm is $O(n(n + m)^2 \log^2(n + m))$.

9 Cyclic SEQUENCE constraint

In rostering problems, we may wish to produce a cyclic schedule which can be repeated, say, every four weeks. We therefore consider a cyclic version of the SEQUENCE constraint. More precisely, $CYCLICSEQUENCE(l, u, k, [X_1, \dots, X_n], v)$ ensures that between l and u variables in X_i to $X_{1+(i+k-1 \bmod n)}$ takes values in the set v for $1 \leq i \leq n$.

The cyclic SEQUENCE constraint can be expressed with a linear program. Consider, for example, the primal linear system for $CYCLICSEQUENCE(l, u, 2, [X_1, \dots, X_3], v)$.

$$\begin{pmatrix} 1 & 1 & 0 \\ -1 & -1 & 0 \\ 0 & 1 & 1 \\ 0 & -1 & -1 \\ 1 & 0 & 1 \\ -1 & 0 & -1 \end{pmatrix} \begin{pmatrix} X_1 \\ X_2 \\ X_3 \end{pmatrix} \geq \begin{pmatrix} l \\ l \\ -u \\ -u \\ l \\ -u \end{pmatrix} \quad (12)$$

Unfortunately, the matrix at the left-hand side of system (12) is not totally unimodular, because it contains a submatrix with the determinant equal to 2.

$$\begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{pmatrix}$$

Hence, methods employing a network flow or its dual, as in Sections 4 and 6, are not directly applicable to Cyclic SEQUENCE.

10 Experimental Results

To evaluate performance of our filtering algorithms we carried out a series of experiments on random problems. The experimental setup is similar to that in [9]. The first set of experiments compares performance of the flow-based propagator FB on single instance of the SEQUENCE constraint against the $HPRS$ propagator (the third propagator in [13]), the CS encoding of [9], and the $AMONG$ decomposition AD of the SEQUENCE constraint. The second set of experiments compares the flow-based propagator FB_S for the $SOFTSEQUENCE$ constraint and its decomposition into soft $AMONG$ constraints. Experiments were run with ILOG 6.1 on an Intel Xeon 4 CPU, 2.0 Ghz, 4G RAM. Boost graph library version 1.34.1 was used to implement the flow-based algorithms.

10.1 The SEQUENCE constraint

For each possible combination of $n \in \{500, 1000, 2000, 3000, 4000, 5000\}$, $k \in \{5, 15, 50\}$, $\Delta = u - l \in \{1, 5\}$, we generated twenty instances with random lower bounds in the interval $(0, k - \Delta)$. We used random value and variable ordering and a time out of 300 sec. We used

the Ford-Fulkerson algorithm to find a maximum flow. Results for different values of Δ are presented in Tables 1- 2 and Figure 5. First of all, we notice that the *CS* encoding is the best on hard instances ($\Delta = 1$) and the *AD* decomposition is the fastest on easy instances ($\Delta = 5$). This result was first observed in [9]. The *FB* propagator is not the fastest one but has the most robust performance. It is sensitive only to the value of n and not to other parameters, like the length of the window(k) or hardness of the problem(Δ). As can be seen from Figure 5, the *FB* propagator scales better than the other propagators with the size of the problem. It appears to grow linearly with the number of variables, while the *HPRS* propagator displays quadratic growth.

n	k	<i>AD</i>	<i>CS</i>	<i>HPRS</i>	<i>FB</i>
500	7	8 / 2.13	20 / 0.13	20 / 0.35	20 / 0.30
	15	6 / 0.01	20 / 0.09	20 / 0.30	20 / 0.29
	50	2 / 0.02	20 / 0.07	20 / 0.26	20 / 0.28
1000	7	4 / 0.01	20 / 0.71	20 / 2.36	20 / 1.18
	15	2 / 0.59	20 / 0.38	20 / 2.06	20 / 1.17
	50	1 / 0	20 / 0.28	20 / 1.48	20 / 1.14
2000	7	4 / 0.04	20 / 4.25	20 / 18.52	20 / 4.76
	15	0 / 0	20 / 1.84	20 / 15.19	20 / 4.56
	50	1 / 0	20 / 1.16	20 / 13.24	20 / 4.42
3000	7	3 / 0.07	20 / 15.14	20 / 64.04	20 / 10.44
	15	1 / 0	20 / 5.49	20 / 51.04	20 / 11.90
	50	0 / 0	20 / 2.61	20 / 35.48	20 / 10.12
4000	7	3 / 0.12	20 / 30.87	20 / 132.73	20 / 23.25
	15	0 / 0	20 / 14.44	20 / 123.60	20 / 18.61
	50	1 / 0	20 / 4.78	20 / 93.98	20 / 18.97
5000	7	1 / 0	20 / 64.05	15 / 262.17	20 / 36.09
	15	0 / 0	20 / 24.46	17 / 211.17	20 / 34.59
	50	0 / 0	20 / 8.24	19 / 146.63	20 / 31.66
TOTALS					
solved/total		37 / 360	360 / 360	351 / 360	360 / 360
avg tm for solved		0.517	9.943	60.973	11.874
avg bt for solved		17761	429	0	0

Table 1. Randomly generated instances with a single SEQUENCE constraint and $\Delta = 1$. Number of instances solved in 300 sec / average time to solve.

n	k	<i>AD</i>	<i>CS</i>	<i>HPRS</i>	<i>FB</i>
500	7	20 / 0.01	20 / 0.58	20 / 0.15	20 / 0.44
	15	20 / 0.01	20 / 0.69	20 / 0.25	20 / 0.44
	50	18 / 0.02	20 / 0.20	20 / 0.37	20 / 0.42
1000	7	20 / 0.03	20 / 4.33	20 / 0.99	20 / 1.70
	15	20 / 0.03	20 / 4.68	20 / 1.83	20 / 1.70
	50	10 / 0.05	20 / 1.24	20 / 2.73	20 / 1.69
2000	7	20 / 0.07	20 / 32.41	20 / 7.19	20 / 6.62
	15	20 / 0.07	20 / 39.71	20 / 14.89	20 / 6.63
	50	5 / 5.19	20 / 9.52	20 / 13.71	20 / 6.94
3000	7	20 / 0.14	20 / 104.68	20 / 23.85	20 / 14.96
	15	20 / 0.16	20 / 125.11	20 / 44.67	20 / 15.21
	50	5 / 0.29	20 / 22.73	20 / 66.61	20 / 14.61
4000	7	20 / 0.25	17 / 251.56	20 / 55.70	20 / 29.34
	15	20 / 0.22	5 / 179.41	20 / 112.99	20 / 26.99
	50	9 / 0.34	20 / 50.52	17 / 141.25	20 / 26.67
5000	7	20 / 0.36	0 / 0	20 / 109.18	20 / 46.42
	15	20 / 0.36	6 / 160.99	17 / 215.97	20 / 45.97
	50	9 / 0.48	20 / 108.34	11 / 210.53	20 / 44.88
TOTALS					
solved/total		296 / 360	308 / 360	345 / 360	360 / 360
avg tm for solved		0.236	52.708	50.698	16.200
avg bt for solved		888	1053	0	0

Table 2. Randomly generated instances with a single SEQUENCE constraint and $\Delta = 5$. Number of instances solved in 300 sec / average time to solve.

n	k	$\Delta = 1$		$\Delta = 5$	
		<i>AD_S</i>	<i>FB_S</i>	<i>AD_S</i>	<i>FB_S</i>
50	7	6 / 19.30	7 / 27.91	20 / 0.01	20 / 2.17
	15	8 / 36.07	13 / 20.41	11 / 49.49	10 / 30.51
	25	6 / 0.73	10 / 23.27	10 / 6.40	10 / 7.41
100	7	1 / 0	3 / 7.56	19 / 10.50	18 / 16.51
	15	0 / 0	5 / 6.90	3 / 0.01	3 / 7.20
	25	0 / 0	5 / 4.96	5 / 19.07	5 / 23.99
TOTALS					
solved/total		21 / 120	43 / 120	68 / 120	66 / 120
avg tm for solved		19.463	18.034	13.286	13.051
avg bt for solved		245245	343	147434	128

Table 3. Randomly generated instances with 4 soft SEQUENCES. Number of instances solved in 300 sec / average time to solve.

10.2 The Soft SEQUENCE constraint

We evaluated performance of the soft SEQUENCE constraint on random problems. For each possible combination of $n \in \{50, 100\}$, $k \in \{5, 15, 25\}$, $\Delta = \{1, 5\}$ and $m \in \{4\}$ (where m is the number of SEQUENCE constraints), we generated twenty random instances. All variables had domains of size 5. An instance was obtained by selecting random lower bounds in the interval $(0, k - \Delta)$. We excluded instances where $\sum_{i=1}^m l_i \geq k$ to avoid unsatisfiable instances. We used a random variable and value ordering, and a time-out of 300 sec. All SEQUENCE constraints were enforced on disjoint sets of cardinality one. Instances with this set of parameters are hard instances for SEQUENCE propagators [9]. To relax these instances, we allow to violate the SEQUENCE constraint with a cost that has to be less than or equal to 15% of the length of the sequence. Experimental results are presented in Table 3. As can be seen from the table, the *FB_S* algorithms is competitive with the decomposition into soft AMONG constraints on easy problems and outperforms the decomposition on hard problems.

We observed that the *FB_S* propagator is very slow for the soft SEQUENCE constraint. Note that the number of backtracks of *FB_S* is three order of magnitude smaller compared to *AD_S*. We profiled the algorithm and found that it spends most of the time performing the all pairs shortest path algorithm. Unfortunately, this is difficult to compute incrementally because the residual graph can be different on every invocation of the propagator.

11 Conclusion

We have proposed new filtering algorithms for the SEQUENCE constraint and several extensions including the soft SEQUENCE and generalized SEQUENCE constraints which are based on network flows. Our propagator for the SEQUENCE constraint enforces domain consistency in $O(n^2)$ time down a branch of the search tree. This improves upon the best existing domain consistency algorithm by a factor of $O(\log n)$. We also introduced a soft version of the SEQUENCE constraint and propose an $O(n^2 \log n \log \log u)$ time domain consistency algorithm based on minimum cost network flows. These algorithms are derived from linear programs which represent a network flow. They differ from the flows used to propagate global constraints like GCC since the domains of the variables are encoded as costs on the edges rather than capacities. Such flows are efficient for maintaining bounds consistency over large domains. Experimental results demonstrate that the *FB* filtering algorithm is more robust than existing propagators. We conjecture that similar flow based propagators derived from linear programs may be useful for other global arithmetic constraints.

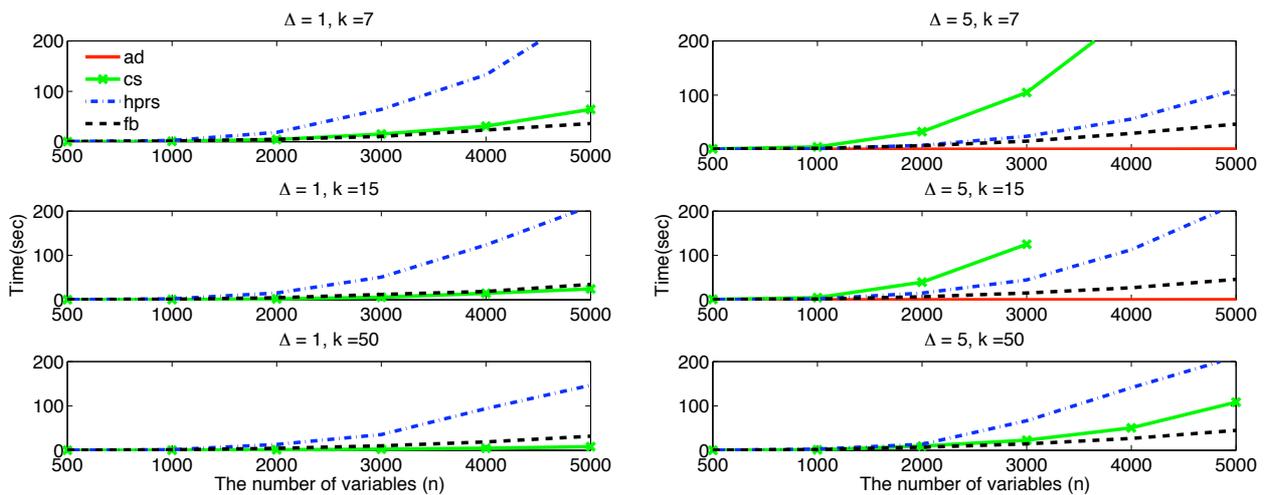


Figure 5. Randomly generated instances with a single SEQUENCE constraints for different combinations of Δ and k .

REFERENCES

- [1] Jr. A.F. Veinott and H.M. Wagner, 'Optimal capacity scheduling I', *Operations Research*, **10**(4), 518–532, (1962).
- [2] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin, *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall, 1993.
- [3] N. Beldiceanu and M. Carlsson, 'Revisiting the cardinality operator and introducing cardinality-path constraint family', in *Proc. of the Int. Conf. on Logic Programming*, pp. 59–73, (2001).
- [4] Nicolas Beldiceanu, 'Global constraint catalog', T-2005-08, SICS Technical Report.
- [5] Nicolas Beldiceanu and Evelyne Contejean, 'Introducing global constraints in CHIP', *Mathematical and Computer Modelling*, **12**, 97–123, (1994).
- [6] Christian Bessiere, Emmanuel Hebrard, Brahim Hnich, Zeynep Kiziltan, and Toby Walsh, 'The slide meta-constraint', Technical report, (2007).
- [7] R.E. Bixby and W.H. Cunningham, 'Converting linear programs to network problems', *Mathematics of Operations Research*, **5**, 321–357, (1980).
- [8] K.S. Booth and G.S. Lueker, 'Testing for the consecutive ones property, interval graphs and graph planarity using PQ-tree algorithms', *Journal of Computer and Systems Sciences*, **13**, 335–379, (1976).
- [9] S. Brand, N. Narodytska, C.-G. Quimper, P. Stuckey, and T. Walsh, 'Encodings of the sequence constraint', in *Proc. of the 13th Int. Conf. on Principles and Practices of Constraint Programming*, (2007).
- [10] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein, *Introduction to Algorithms, Second Edition*, The MIT Press, 2001.
- [11] Scott Cotton and Oded Maler, 'Fast and flexible difference constraint propagation for DPLL(T)', in *Proc. of Theory and Applications of Satisfiability Testing (SAT-2006)*, pp. 170–183, (2006).
- [12] Andrew V. Goldberg and Satish Rao, 'Beyond the flow decomposition barrier', *J. ACM*, **45**, 753–782, (1998).
- [13] Willem-Jan van Hoeve, Gilles Pesant, Louis-Martin Rousseau, and Ashish Sabharwal, 'Revisiting the sequence constraint', in *Proc. of the 12th Int. Conf. on Principles and Practice of Constraint Programming*, pp. 620–634, (2006).
- [14] Jean-Charles Régin, 'A filtering algorithm for constraints of difference in csps', in *Proc. of the 12th National Conf. on AI (AAAI'94)*, volume 1, pp. 362–367, (1994).
- [15] Jean-Charles Régin, 'Generalized arc consistency for global cardinality constraint', in *Proc. of the 12th National Conf. on AI (AAAI'96)*, pp. 209–215, (1996).
- [16] Jean-Charles Régin, 'Combination of among and cardinality constraints', in *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, Second International Conference*, volume 3524, pp. 288–303. Springer, (2005).
- [17] Jean-Charles Régin and Jean-Francois Puget, 'A filtering algorithm for global sequencing constraints', in *Proc. of the 3th Int. Conf. on Principles and Practice of Constraint Programming*, pp. 32–46, (1997).
- [18] Alexander Schrijver, *Theory of linear and integer programming*, John Wiley & Sons, Inc., 1986.
- [19] Christine Solnon, Van Dat Cung, Alain Nguyen, and Christian Artigues, 'The car sequencing problem: overview of state-of-the-art methods and industrial case-study of the ROADEF'2005 challenge problem', *European Journal of Operational Research (EJOR)*, (January 2008). In press.
- [20] Willem Jan van Hoeve, Gilles Pesant, and Louis-Martin Rousseau, 'On global warming: Flow-based soft global constraints', *J. Heuristics*, **12**(4-5), 347–373, (2006).