

# MDD Propagation for Sequence Constraints

**David Bergman**

*School of Business, University of Connecticut  
2100 Hillside Road, Unit 1041, Storrs, CT 06260*

DAVID.BERGMAN@BUSINESS.UCONN.EDU

**Andre A. Cire**

**Willem-Jan van Hoeve**

*Tepper School of Business, Carnegie Mellon University  
5000 Forbes Avenue, Pittsburgh, PA 15213 USA*

ACIRE@ANDREW.CMU.EDU

VANHOEVE@ANDREW.CMU.EDU

## Abstract

We study propagation for the SEQUENCE constraint in the context of constraint programming based on limited-width MDDs. Our first contribution is proving that establishing MDD-consistency for SEQUENCE is NP-hard. Yet, we also show that this task is fixed parameter tractable with respect to the length of the sub-sequences. In addition, we propose a partial filtering algorithm that relies on a specific decomposition of the constraint and a novel extension of MDD filtering to node domains. We experimentally evaluate the performance of our proposed filtering algorithm, and demonstrate that the strength of the MDD propagation increases as the maximum width is increased. In particular, MDD propagation can outperform conventional domain propagation for SEQUENCE by reducing the search tree size and solving time by several orders of magnitude. Similar improvements are observed with respect to the current best MDD approach that applies the decomposition of SEQUENCE into AMONG constraints.

## 1. Introduction

The central inference process of constraint programming is constraint propagation (Rossi et al., 2006; Dechter, 2003; Apt, 2003). While traditional constraint processing techniques were designed for explicitly defined relations of small arity, state-of-the-art constraint programming solvers apply specialized constraint propagation algorithms for global constraints of any arity, often based on efficient combinatorial methods such as network flows (van Hoeve & Katriel, 2006; Régin, 2011).

Conventional constraint propagation algorithms (or domain filtering algorithms) operate on individual constraints of a given problem. Their role is to identify and remove values in the variable domains that are inconsistent with respect to the constraint under consideration. Whenever the domain of a variable is updated (i.e., a value is removed), the constraints in which this variable appears can be reconsidered for inspection. This cascading process of propagating the changes in variable domains through the constraints continues until a fixed point is reached. Most constraint programming solvers assume that the variable domains are finite, which ensures termination of the constraint propagation process. Note that constraint propagation in itself may not be sufficient to determine the resolution of a given problem. Therefore, constraint propagation is normally applied at each search state in a systematic search process.

A major benefit of propagating variable domains is that it can be implemented efficiently in many cases. However, an inherent weakness of domain propagation is that it implicitly represents the Cartesian product of the variable domains as potential solution space. By communicating only domain changes, this limits the amount of information shared between constraints.

To address this shortcoming of domain propagation, Andersen, Hadzic, Hooker, and Tiedemann (2007) proposed the use of multi-valued decision diagrams (MDDs) as an alternative to variable domains in the context of constraint propagation. MDDs are directed acyclic layered graphs that can, in principle, compactly represent all solutions to a combinatorial problem (Wegener, 2000). Andersen et al. (2007) showed that MDDs of limited width can provide a much stronger relaxation of the solution space than the traditional Cartesian product of the variable domains, and as a consequence MDDs allow to represent and communicate more refined information between constraints. By propagating MDDs rather than variable domains, huge reductions in search tree size and computation time can be realized (Andersen et al., 2007; Hadzic et al., 2008a, 2008b; Hoda et al., 2010; Cire & van Hoeve, 2012, 2013).

MDDs can be used to represent individual (global) constraints, subsets of constraints, or all constraints in a given problem. When representing individual constraints, as in (Hawkins et al., 2005; Cheng & Yap, 2008; Hadzic et al., 2009), the higher-level information carried by the MDD is lost when projecting this down to the variable domains for the traditional domain propagation. The highest potential for MDD propagation instead appears to be in representing specific subsets of constraints within the same MDD. That is, for a given set of constraints, we create and maintain one single limited-width MDD, which is then propagated through this constraint set. Since an MDD is defined with respect to a fixed variable ordering, it is most useful to select a subset of constraints compatible with this ordering. When applied in this way, MDD propagation can be implemented in parallel to the existing domain propagation in constraint programming systems, thus complementing and potentially strengthening the domain propagation process. For example, Cire and van Hoeve (2013) introduced MDD propagation for a subset of constraints representing disjunctive scheduling problems. They embedded this as a custom global constraint in the ILOG CP Optimizer constraint programming solver, which greatly improved the performance.

## Methodology

Constraint propagation based on limited-width MDDs amounts to *MDD filtering* and *MDD refinement*. The role of an MDD filtering algorithm is to remove provably inconsistent arcs from the MDD (Hadzic et al., 2008b; Hoda et al., 2010). An MDD refinement algorithm on the other hand, aims at splitting nodes in the MDD to more accurately reflect the solution space (Hadzic et al., 2008a). In order to make this approach scalable and efficient, refinement algorithms must ensure that the MDD remains within a given maximum size (typically by restricting its maximum width—the number of nodes on any layer). By increasing this maximum width, the MDD relaxation can be strengthened to any desired level. That is, a maximum width of 1 would correspond to the traditional Cartesian product of the variable domains, while an infinite maximum width would correspond to an exact MDD representing all solutions. However, increasing the size of the MDD immediately impacts

the computation time, and one typically needs to balance the trade-off between the strength of the MDD and the associated computation time.

In order to characterize the outcome of an MDD filtering algorithm, the notion of *MDD consistency* was introduced by Andersen et al. (2007), similar to domain consistency in finite-domain constraint programming: Given an MDD, a constraint is MDD consistent if all arcs in the MDD belong to at least one solution to the constraint. As a consequence of the richer data structure that an MDD represents, establishing MDD consistency may be more difficult than establishing domain consistency. For example, Andersen et al. (2007) show that establishing MDD consistency on the ALLDIFFERENT constraint is NP-hard, while establishing traditional domain consistency can be done in polynomial time (Régin, 1994).

### Contributions

The main focus of this paper is the SEQUENCE constraint, that is defined as a specific conjunction of AMONG constraints, where an AMONG constraint restricts the occurrence of a set of values for a sequence of variables to be within a lower and upper bound (Beldiceanu & Contejean, 1994). The SEQUENCE constraint finds applications in, e.g., car sequencing and employee scheduling problems (Régin & Puget, 1997; van Hoesve et al., 2009). It is known that classical domain consistency can be established for SEQUENCE in polynomial time (van Hoesve et al., 2006, 2009; Brand et al., 2007; Maher et al., 2008; Downing et al., 2012). Furthermore, (Hoda et al., 2010) present an MDD filtering algorithm for AMONG constraints establishing MDD consistency in polynomial time. However, it remained an open question whether or not MDD consistency for SEQUENCE can be established in polynomial time as well.

In this work, we answer that question negatively and our first contribution is showing that establishing MDD consistency on the SEQUENCE constraint is NP-hard. This is an important result from the perspective of MDD-based constraint programming. Namely, of all global constraints, the SEQUENCE constraint has perhaps the most suitable combinatorial structure for an MDD approach; it has a prescribed variable ordering, it combines sub-constraints on contiguous variables, and existing approaches can handle this constraint fully by using bounds reasoning only.

As our second contribution, we show that establishing MDD consistency on the SEQUENCE is fixed parameter tractable with respect to the lengths of the sub-sequences (the AMONG constraints), provided that the MDD follows the order of the SEQUENCE constraint. The proof is constructive, and follows from a generic algorithm to filter one MDD with another.

The third contribution is a partial MDD propagation algorithm for SEQUENCE, that does not necessarily establish MDD consistency. It relies on the decomposition of SEQUENCE into ‘cumulative sums’, and a new extension of MDD filtering to the information that is stored at its nodes.

Our last contribution is an experimental evaluation of our proposed partial MDD propagation algorithm. We evaluate the strength of our algorithm for MDDs of various maximum widths, and compare the performance with existing domain propagators for SEQUENCE. We also compare our algorithm with the currently best known MDD approach that uses the

natural decomposition of SEQUENCE into AMONG constraints (Hoda et al., 2010). Our experiments demonstrate that MDD propagation can outperform domain propagation for SEQUENCE by reducing the search tree size, and solving time, by several orders of magnitude. Similar results are observed with respect to MDD propagation of AMONG constraints. Our results thus provide further evidence for the power of MDD propagation in the context of constraint programming.

The remainder of this paper is structured as follows. In Section 2, we provide the necessary definitions of MDD-based constraint programming and the SEQUENCE constraint. In Section 3, we present the proof that establishing MDD consistency on SEQUENCE is NP-hard. Section 4 describes that establishing MDD consistency is fixed parameter tractable. In Section 5, the partial MDD filtering algorithm is presented. Section 6 shows the experimental results. We present final conclusions in Section 7.

## 2. Definitions

We first recall some basic definitions of MDD-based constraint programming, following (Andersen et al., 2007; Hoda et al., 2010). In this work, an *ordered Multivalued Decision Diagram (MDD)* is a directed acyclic graph whose nodes are partitioned into  $n+1$  (possibly empty) subsets or *layers*  $L_1, \dots, L_{n+1}$ , where the layers  $L_1, \dots, L_n$  correspond respectively to variables  $x_1, \dots, x_n$ .  $L_1$  contains a single *root* node  $r$ , and  $L_{n+1}$  contains a single *terminal* node  $t$ . For a node  $u$  in the MDD, we let  $L(u)$  denote the index of its layer. The *width* of an MDD is the maximum number of nodes in a layer, or  $\max_{i=1}^n \{|L_i|\}$ . In MDD-based CP, the MDDs typically have a given fixed maximum width.

All arcs of the MDD are directed from an upper to a lower layer; that is, from a node in some  $L_i$  to a node in some  $L_j$  with  $i < j$ . For our purposes it is convenient to assume (without loss of generality) that each arc connects two adjacent layers. Each arc out of layer  $L_i$  is labeled with an element of the domain  $D(x_i)$  of  $x_i$ . For an arc  $a$ , we refer to the label it represents as  $\ell(a)$ . An element in  $D(x_i)$  appears at most once as a label on the arcs out of a given node  $u \in L_i$ . The set  $A(u, v)$  of arcs from node  $u$  to node  $v$  may contain multiple arcs, and we denote each with its label. Let  $A^{in}(u)$  denote the set of arcs coming into node  $u$ .

An arc with label  $v$  leaving a node in layer  $i$  represents an assignment  $x_i = v$ . Each path in the MDD from  $r$  to  $t$  can be denoted by the arc labels  $v_1, \dots, v_n$  on the path and is identified with the solution  $(x_1, \dots, x_n) = (v_1, \dots, v_n)$ . For an MDD  $M$ , we denote the set of associated solutions by  $\text{Sol}(M)$ . For two nodes  $i, j$  in  $M$ , we let  $M_{i,j}$  denote set of paths from  $i$  to  $j$  in  $M$ . For an arc  $a$  in  $M$ , we let  $M|_a$  denote the set of  $r$ - $t$  paths in  $M$  that contain  $a$ .

A path  $v_1, \dots, v_n$  is *feasible* for a given constraint  $C$  if setting  $(x_1, \dots, x_n) = (v_1, \dots, v_n)$  satisfies  $C$ . Constraint  $C$  is feasible on an MDD if the MDD contains a feasible path for  $C$ .

A constraint  $C$  is called *MDD consistent* on a given MDD if every arc of the MDD lies on some feasible path. Thus MDD consistency is achieved when all redundant arcs (i.e., arcs on no feasible path) have been removed. We also say that such MDD is MDD consistent with respect to  $C$ . Domain consistency for  $C$  is equivalent to MDD consistency on an MDD of width one that represents the variable domains. That is, it is equivalent

to MDD consistency on an MDD in which each layer  $L_i$  contains a single node  $s_i$ , and  $A(s_i, s_{i+1}) = D(x_i)$  for  $i = 1, \dots, n$ .

Lastly, we formally recall the definitions of AMONG (Beldiceanu & Contejean, 1994), SEQUENCE (Beldiceanu & Contejean, 1994), and GEN-SEQUENCE (van Hove et al., 2009) constraints. The AMONG constraint counts the number of variables that are assigned to a value in a given set  $S$ , and ensures that this number is between a given lower and upper bound:

**Definition 1** *Let  $X$  be a set of variables,  $l, u$  integer numbers such that  $0 \leq l \leq u \leq |X|$ , and  $S \subset \cup_{x \in X} D(x)$  a subset of domain values. Then we define  $\text{AMONG}(X, l, u, S)$  as*

$$l \leq \sum_{x \in X} (x \in S) \leq u.$$

Note that the expression  $(x \in S)$  is evaluated as a binary value, i.e., resulting in 1 if  $x \in S$  and 0 if  $x \notin S$ . The SEQUENCE constraint is the conjunction of a given AMONG constraint applied to every sub-sequence of length  $q$  over a sequence of  $n$  variables:

**Definition 2** *Let  $X$  be an ordered set of  $n$  variables,  $q, l, u$  integer numbers such that  $0 \leq q \leq n$ ,  $0 \leq l \leq u \leq q$ , and  $S \subset \cup_{x \in X} D(x)$  a subset of domain values. Then*

$$\text{SEQUENCE}(X, q, l, u, S) = \bigwedge_{i=1}^{n-q+1} \text{AMONG}(s_i, l, u, S),$$

where  $s_i$  represents the sub-sequence  $x_i, \dots, x_{i+q-1}$ .

Finally, the *generalized* SEQUENCE constraint extends the SEQUENCE constraint by allowing the AMONG constraints to be specified with different lower and upper bounds, and sub-sequence length:

**Definition 3** *Let  $X$  be an ordered set of  $n$  variables,  $k$  a natural number,  $\vec{s}, \vec{l}, \vec{u}$  vectors of length  $k$  such that  $s_i$  is a sub-sequence of  $X$ ,  $l_i, u_i \in \mathbb{N}$ ,  $0 \leq l_i \leq u_i \leq n$  for  $i = 1, 2, \dots, k$ , and  $S \subset \cup_{x \in X} D(x)$  a subset of domain values. Then*

$$\text{GEN-SEQUENCE}(X, \vec{s}, \vec{l}, \vec{u}, S) = \bigwedge_{i=1}^k \text{AMONG}(s_i, l_i, u_i, S).$$

### 3. MDD Consistency for Sequence is NP-Hard

As stated before, the only known non-trivial NP-hardness result for a global constraint in the context of MDD-based constraint programming is that of Andersen et al. (2007) for the ALLDIFFERENT constraint. A challenge in determining whether a global constraint can be made MDD consistent in polynomial time is that this must be guaranteed for *any* given MDD. That is, in addition to the combinatorics of the global constraint itself, the shape of the MDD adds another layer of complexity to establishing MDD consistency. For proving NP-hardness, a particular difficulty is making sure that in the reduction, the MDD remains of polynomial size. For SEQUENCE constraints, so far it was unknown whether a polynomial-time MDD consistency algorithm exists. In this section we answer that question negatively and prove the following result.

**Theorem 1** *Establishing MDD consistency for SEQUENCE on an arbitrary MDD is NP-hard even if the MDD follows the variable ordering of the SEQUENCE constraint.*

*Proof.* The proof is by reduction from 3-SAT, a classical NP-complete problem (Garey & Johnson, 1979). We will show that an instance of 3-SAT is satisfied if and only if a particular SEQUENCE constraint on a particular MDD  $M$  of polynomial size has a solution. Therefore, establishing MDD consistency for SEQUENCE on an arbitrary MDD is at least as hard as 3-SAT.

Consider a 3-SAT instance on  $n$  variables  $x_1, \dots, x_n$ , consisting of  $m$  clauses  $c_1, \dots, c_m$ . We first construct an MDD that represents the basic structure of the 3-SAT formula (see Example 1 after this proof for an illustration). We introduce binary variables  $y_{i,j}$  and  $\bar{y}_{i,j}$  representing the literals  $x_j$  and  $\bar{x}_j$  per clause  $c_i$ , for  $i = 1, \dots, m$  and  $j = 1, \dots, n$  ( $x_j$  and  $\bar{x}_j$  may or may not exist in  $c_i$ ). We order these variables as a sequence  $Y$ , first by the index of the clauses, then by the index of the variables, and then by  $y_{i,j}, \bar{y}_{i,j}$  for clause  $c_i$  and variable  $x_j$ . That is, we have  $Y = y_{1,1}, \bar{y}_{1,1}, y_{1,2}, \bar{y}_{1,2}, \dots, y_{1,n}, \bar{y}_{1,n}, \dots, y_{m,1}, \bar{y}_{m,1}, \dots, y_{m,n}, \bar{y}_{m,n}$ . We construct an MDD  $M$  as a layered graph, where the  $k$ -th layer corresponds to the  $k$ -th variable in the sequence  $Y$ .

A clause  $c_i$  is represented by  $2n$  consecutive layers corresponding to  $y_{i,1}, \dots, \bar{y}_{i,n}$ . In such part of the MDD, we identify precisely those paths that lead to a solution satisfying the clause. The basis for this is a ‘diamond’ structure for each pair of literals  $(y_{i,j}, \bar{y}_{i,j})$ , that assigns either  $(0, 1)$  or  $(1, 0)$  to this pair. If a variable does not appear in a clause, we represent it using such a diamond in the part of the MDD representing that clause, thus ensuring that the variable can take any assignment with respect to this clause. For the variables that do appear in the clause, we will explicitly list out all allowed combinations.

More precisely, for clause  $c_i$ , we first define a local root node  $r_i$  representing layer  $L(y_{i,1})$ , and we set  $\text{tag}(r_i) = \text{‘unsat’}$ . For each node  $u$  in layer  $L(y_{i,j})$  (for  $j = 1, \dots, n$ ), we do the following. If variable  $x_j$  does not appear in  $c_i$ , or if  $\text{tag}(u)$  is ‘sat’, we create two nodes  $v, v'$  in  $L(\bar{y}_{i,j})$ , one single node  $w$  in  $L(y_{i,j+1})$ , and arcs  $(u, v)$  with label 1,  $(u, v')$  with label 0,  $(v, w)$  with label 0, and  $(v', w)$  with label 1. This corresponds to the ‘diamond’ structure. We set  $\text{tag}(w) = \text{tag}(u)$ . Otherwise (i.e.,  $\text{tag}(u)$  is ‘unsat’ and  $y_{i,j}$  appears in  $c_i$ ), we create two nodes  $v, v'$  in  $L(\bar{y}_{i,j})$ , two nodes  $w, w'$  in  $L(y_{i,j+1})$ , and arcs  $(u, v)$  with label 1,  $(u, v')$  with label 0,  $(v, w)$  with label 0, and  $(v', w')$  with label 1. If  $c_i$  contains as literal  $y_{i,j}$ , we set  $\text{tag}(w) = \text{‘sat’}$  and  $\text{tag}(w') = \text{‘unsat’}$ . Otherwise ( $c_i$  contains  $\bar{y}_{i,j}$ ), we set  $\text{tag}(w) = \text{‘unsat’}$  and  $\text{tag}(w') = \text{‘sat’}$ .

This procedure will be initialized by a single root node  $r$  representing  $L(y_{11})$ . We iteratively append the MDDs of two consecutive clauses  $c_i$  and  $c_{i+1}$  by merging the nodes in the last layer of  $c_i$  that are marked ‘sat’ into a single node, and let this node be the local root for  $c_{i+1}$ . We finalize the procedure by merging all nodes in the last layer that are marked ‘sat’ into the single terminal node  $t$ . By construction, we ensure that only one of  $y_{ij}$  and  $\bar{y}_{ij}$  can be set to 1. Furthermore, the variable assignment corresponding to each path between layers  $L(y_{i,1})$  and  $L(y_{i+1,1})$  will satisfy clause  $c_i$ , and exactly  $n$  literals are chosen accordingly on each such path.

We next need to ensure that for a feasible path in the MDD, each variable  $x_j$  will correspond to the same literal  $y_{i,j}$  or  $\bar{y}_{i,j}$  in each clause  $c_i$ . To this end, we impose the




Figure 1: The MDD corresponding to Example 1.

constraint

$$\text{SEQUENCE}(Y, q = 2n, l = n, u = n, S = \{1\}) \quad (1)$$

on the MDD  $M$  described above. If the sub-sequence of length  $2n$  starts from a positive literal  $y_{i,j}$ , by definition there are exactly  $n$  variables that take value 1. If the sub-sequence starts from a negative literal  $\bar{y}_{i,j}$  instead, the last variable in the sequence corresponds to the value  $x_j$  in the next clause  $c_{i+1}$ , i.e.,  $y_{i+1,j}$ . Observe that all variables except for the first and the last in this sequence will take value 1 already  $n - 1$  times. Therefore, of the first and the last variable in the sequence (which represent  $x_j$  and its complement  $\bar{x}_j$  in any order), only one can take the value 1. That is,  $x_j$  must take the same value in clause  $c_i$  and  $c_{i+1}$ . Since this hold for all sub-sequences, all variables  $x_j$  must take the same value in all clauses.

The MDD  $M$  contains  $2mn + 1$  layers, while each layer contains at most six nodes. Therefore, it is of polynomial size (in the size of the 3-SAT instance), and the overall construction needs polynomial time.  $\square$




Figure 2: The exact MDD for the SEQUENCE constraint of Example 2.

**Example 1** Consider the 3-SAT instance on four Boolean variables  $x_1, x_2, x_3, x_4$  with clauses  $c_1 = (x_1 \vee \bar{x}_3 \vee x_4)$  and  $c_2 = (x_2 \vee x_3 \vee \bar{x}_4)$ . The corresponding MDD used in the reduction is given in Figure 1.

#### 4. MDD Consistency for Sequence is Fixed Parameter Tractable

In this section we show that establishing MDD consistency for SEQUENCE on an arbitrary MDD is fixed parameter tractable, with respect to the length of the sub-sequences  $q$ . It was already shown in (van Hoeve et al., 2006, 2009) that an exact MDD for the SEQUENCE constraint exists with  $O(n2^q)$  nodes (i.e., the ‘unfolded’ automaton of the REGULAR constraint).

**Example 2** Consider the constraint  $\text{SEQUENCE}(X, q = 3, l = 1, u = 2, S = \{1\})$  where  $X = \{x_1, x_2, \dots, x_6\}$  is an ordered set of binary variables. The corresponding exact MDD, following the order of  $X$ , is presented in Figure 2. For convenience, each node in the MDD is labeled with the last  $q - 1$  labels that represent the sub-sequence up to that node (starting  $q - 1$  layers up). For example, the second node in the third layer represents decisions  $x_1 = 0$  and  $x_2 = 1$ , corresponding to sub-sequence 01. To construct the next layer, we either append a 0 or a 1 to this sub-sequence, leading to nodes labeled 10 and 011, respectively. Note that from nodes labeled 00 we must take an arc with label 1, because  $l = 1$ . Similarly for nodes labeled 11. After  $q$  layers, all possible sub-sequences have been created (maximally  $O(2^{q-1})$ ), which thus defines the width of the subsequent layers.



However, since we are given an arbitrary MDD, and not necessarily an exact MDD, we need some additional steps to exploit this connection. For this we apply a generic approach that will not only show fixed parameter tractability for SEQUENCE, but in fact can be applied to determine whether MDD consistency is tractable for any constraint.

Our goal is to establish MDD consistency on a given MDD  $M$  with respect to another MDD  $M'$  on the same set of variables. This is compatible with our earlier definitions since  $M'$  can be interpreted to define a constraint. Namely, by definition,  $M$  is MDD consistent with respect to  $M'$  if every arc in  $M$  belongs to a path for which a corresponding path exists in  $M'$ . More formally, for every  $a$  in  $M$ ,  $\text{Sol}(M|_a) \cap \text{Sol}(M') \neq \emptyset$ . For our purposes, we assume that  $M$  and  $M'$  follow the same variable ordering.

Our algorithm is based on a top-down pass and a bottom-up pass of  $M$  with  $M'$ .

#### 4.1 Top-Down Phase

We present our top-down traversal of  $M$  in Algorithm 1. Let  $r, r'$  be the roots of  $M, M'$ , respectively. The algorithm proceeds by labeling each node  $u \in M, u \neq r$  with a *top-down state*  $s^+(u)$  which represents the following:

$$s^+(u) = \{u' \in M' \mid L(u) = L(u'), \text{Sol}(M_{r,u}) \cap \text{Sol}(M'_{r',u'}) \neq \emptyset\} \quad (2)$$

That is, a node  $u' \in M'$  is in  $s^+(u)$  if there are partial solutions represented by paths in  $M$  ending at  $u$  that are also partial solutions represented by paths in  $M'$  ending at  $u'$ .

---

#### Algorithm 1 TopDownFilter( $M, M'$ )

---

- 1: let  $\tilde{s}^+(r) = \{r'\}$  and  $\tilde{s}^+(u) = \emptyset$  for all other nodes  $u \in M$
  - 2: **for**  $j = 1, \dots, n$  **do**
  - 3:   **for all**  $u \in S_j$  **do**
  - 4:     **for all**  $a = (u, v) \in M$  **do**
  - 5:       **if**  $\exists u' \in \tilde{s}^+(u)$  with arc  $a' = (u', v') \in M'$  with  $d(a) = d(a')$  **then**
  - 6:           $\tilde{s}^+(v) \leftarrow \tilde{s}^+(v) \cup \{v'\}$
  - 7:       **else**
  - 8:          delete  $a$  from  $M$
- 

We will show that the labels  $\tilde{s}^+(u)$  created via Algorithm 1 coincide with the top-down states. The algorithm proceeds as follows. We begin by labeling the root node of  $M$  with the root node of  $M'$ . Then, having created the labels for the nodes in layers  $1, \dots, j$ , we create the labels of nodes in layer  $j + 1$  by examining the nodes in layer  $j$ . For every node  $u$  in layer  $j$  and each node in  $\tilde{s}^+(u)$ , we check if there are outgoing arcs  $a \in M$  and  $a' \in M'$  directed out of these two nodes, respectively, with the same arc-domain. If so, we add the node pointed to by  $a'$  to the top-down state of the node pointed to by  $a$ . If no such arc exists, we can infer that  $\text{Sol}(M|_a) \cap \text{Sol}(M') = \emptyset$  because no path going through  $a$  will represent a solution that is in the set of solutions in  $M'$  so that  $a$  can be filtered.

**Theorem 2** For  $u \in M, u \neq r$ ,  $\tilde{s}^+(u) = s^+(u)$ ; i.e., Algorithm 1 correctly calculates the top-down states of  $M$  and therefore deletes arcs from  $M$  only when  $a$  has no support in  $M'$ .

*Proof.* First, by construction, node  $u' \in \tilde{s}^+(u)$  only if  $L(u) = L(u')$ . Therefore we need only check condition 2 of (2).

We proceed by induction on  $k \leq n$  and show that for any node  $u \in M$  with  $L(u) = k$ , the set  $\tilde{s}^+(u)$  determined by Algorithm 1 correctly calculates the top-down states in (2); i.e.,  $\tilde{s}^+(u) = s^+(u)$ .

For  $k = 1$  the only node in  $L_1$  is  $r$ , the root of  $M$ .  $\tilde{s}^+(r) = \{r'\}$ , the root node of  $M'$ .

Take any node  $u \in L_2$  in  $M$ . A node  $u' \in M'$  is in  $\tilde{s}^+(u)$  if and only if for some arc  $a = (r, u)$ , there is an arc  $a' \in M'$  directed out of  $r'$  to some node  $u' \in M'$  with  $d(a') = d(a)$ . This happens if and only if  $d(a) \in \text{Sol}(M_{r,u})$  and  $d(a') \in \text{Sol}(M'_{r',u'})$ . Therefore,  $d(a) = d(a') \in \text{Sol}(M_{r,u}) \cap \text{Sol}(M'_{r',u'})$ , so that  $u' \in \tilde{s}^+(u)$  if and only if  $\text{Sol}(M_{r,u}) \cap \text{Sol}(M'_{r',u'}) \neq \emptyset$ .

Now, by induction, fix  $k < n$  and we suppose that for all  $u \in M$  with  $L(u) = j \leq k$ ,  $\tilde{s}^+(u) = s^+(u)$ .

Consider any node  $v \in M$  in layer  $k + 1$ . We need to show that  $v' \in \tilde{s}^+(v)$  if and only if  $\text{Sol}(M_{r,v}) \cap \text{Sol}(M'_{r',v'}) \neq \emptyset$ .

We begin with the case that  $v' \notin \tilde{s}^+(v)$ . By way of contradiction, suppose that  $v' \in s^+(v)$ . Then there exists some partial solution  $x'$  on the first  $k$  variables with  $x' \in \text{Sol}(M_{r,v}) \cap \text{Sol}(M_{r,v})$ . This solution corresponds to some path  $p \in M$  and some path  $p' \in M'$  ending at  $v, v'$  respectively. Consider nodes  $u \in p, u' \in p'$  with  $L(u) = L(u') = k$ , which must exist, and let  $d'$  be the domain value of the  $k^{\text{th}}$  variable in  $x'$  so that one can write  $x' = (x'', d')$ . Also, we let  $a_u, a_{u'}$  be the arcs in  $M, M'$  that are directed out of  $u, u'$  that have arc-domain  $d'$ .

By induction, since  $x'' \in \text{Sol}(M_{r,u}) \cap \text{Sol}(M'_{r',u'})$ , the intersection is non-empty and so  $u' \in \tilde{s}^+(u)$ . When Algorithm 1 examines  $u$  and arc  $a_u$ , it will find that  $a_{u'}$  satisfies the condition in line 5 and so will add  $v'$  to  $\tilde{s}^+(v)$ . Since nothing is ever deleted from  $\tilde{s}^+(v)$ , this yields a contradiction, as desired, so that if  $v' \notin \tilde{s}^+(v)$  then  $v' \notin s^+(v)$ .

Now consider the case when  $v' \in \tilde{s}^+(v)$ . This happens if and only if during the execution of Algorithm 1 the condition in line 5 is satisfied by some nodes  $u \in M, u' \in M'$ , and arcs  $a, a'$  (with common arc-domain  $d'$ ). This implies that  $u' \in \tilde{s}^+(u) = s^+(u)$ , by induction. We therefore have  $u' \in s^+(u)$  so that  $\text{Sol}(M_{r,u}) \cap \text{Sol}(M'_{r',u'}) \neq \emptyset$ . Take any partial solution  $x'$  on the first  $k - 1$  variables that is in this common solution set.

$a = (u, v) \in M$ , which is directed out of  $u$ , implies that  $(x', d(a)) = (x', d') \in \text{Sol}(M_{r,v})$ .  $a' = (u', v') \in M'$ , which is directed out of  $u'$ , implies that  $(x', d(a')) = (x', d') \in \text{Sol}(M'_{r',v'})$ . Therefore,  $(x', d') \in \text{Sol}(M_{r,v}) \cap \text{Sol}(M'_{r',v'})$  and so the common solution set is non-empty, as desired, finishing the proof.  $\square$

Note that if an arc  $a = (u, v)$  is deleted and  $u$  no longer has any outgoing arcs, we can in addition delete  $u$  from  $M$  and recursively delete any arcs directed at  $u$ . However, this still need not be done if we allow *dangling* nodes—nodes with no incoming or outgoing arcs.

## 4.2 Bottom-Up Phase

As for the top-down phase, we can filter  $M$  with  $M'$  with a single bottom-up traversal of  $M$ , as presented in Algorithm 2. We proceed nearly identically as in Algorithm 1 except that we start with the terminal node of  $M$  and proceed upwards through  $M$  labeling each node with a *bottom-up state*  $s^-(u)$  which represents the following:

$$s^-(u) = \{u' \in M' \mid L(u) = L(u'), \text{Sol}(M_{u,t}) \cap \text{Sol}(M'_{u',t'}) \neq \emptyset\}$$

---

### Algorithm 2 BottomUpFilter( $M, M'$ )

---

- 1: let  $s^-(t) = \{t'\}$  and  $s^-(u) = \emptyset$  for all other nodes  $u \in M$
  - 2: **for**  $j = n + 1, \dots, 2$  **do**
  - 3:   **for all**  $u \in S_j$  **do**
  - 4:     **for all**  $a = (v, u) \in M$  **do**
  - 5:       **if**  $\exists u' \in s^-(u)$  with arc  $a' = (v', u') \in M'$  with  $d(a) = d(a')$  **then**
  - 6:           $s^-(v) \leftarrow s^-(v) \cup \{v'\}$
  - 7:       **else**
  - 8:          delete  $a$  from  $M$
- 

The proof of correctness of Algorithm 2 follows similarly to the proof of Algorithm 1 and is omitted here. We use the same notation for the labels determined by the algorithm as in the definition of the bottom-up states in the algorithm.

## 4.3 Establishing MDD Consistency

We next combine Algorithms 1 and 2 to establish MDD consistency of  $M$  with respect to  $M'$ . We remark that individually these algorithms are not idempotent—applying the algorithms multiple times may cause more arcs to be filtered. However, by combining the state information from the two algorithms, we can establish MDD consistency in one top-down and bottom-up pass.

---

### Algorithm 3 MDD-Consistency( $M, M'$ )

---

- 1: TopDownFilter( $M, M'$ )
  - 2: BottomUpFilter( $M, M'$ )
  - 3: **for all**  $a = (u, v) \in M$  **do**
  - 4:   **if**  $\nexists a' = (u', v') \in M'$  with  $u' \in s^+(u)$  and  $v' \in s^-(v)$  **then**
  - 5:     delete  $a$  from  $M$
- 

The combined approach is described in Algorithm 3. We first apply the algorithms TopDownFilter and BottomUpFilter. Then, for each arc  $a = (u, v)$  in  $M$ , we determine whether a corresponding arc  $a'$  exists in  $M'$  by combining all nodes in the top-down state  $s^+(u)$  and the bottom-up state  $s^-(v)$ . If no such arc exists, we delete  $a$  from  $M$ .

For an MDD  $M$  with node set  $V$  and arc set  $A$ , we let the *size* of  $M$  be  $|M| := |V| + |A|$ . Furthermore, we let  $w(M)$  denote the width of  $M$ .

**Theorem 3** *Algorithm 3 establishes MDD-consistency on  $M$  with respect to  $M'$  in  $O(|M| \cdot (w(M'))^2)$  time.*

*Proof.* Take any solution  $x'$  that uses arc  $a = (u, v)$ . Let  $L(u) = j$  so that  $x'$  can be split into 3 sub-solutions  $x' = (x^1, d(a), x^2)$  where  $x^1$  represents the first  $j - 1$  values in  $x'$ ,  $d(a)$  is the  $j^{\text{th}}$  value in  $x'$ , and  $x^2$  represents the final  $n - j$  values in  $x'$ .

This solution is also represented by a path in  $M'$  if and only if there exists arc  $a' = (u', v') \in M$  with  $L(u') = L(u)$  and  $L(v') = L(v)$  for which  $d(a') = d(a)$ ,  $x^1 \in \text{Sol}(M'_{r',u'})$  and  $x^2 \in \text{Sol}(M'_{v',t'})$ . Algorithm 3 checks for every solution in  $\text{Sol}(M|a)$  this condition through the top-down and bottom-up states.

The time complexity for TopDownFilter and BottomUpFilter is  $O(|M| \cdot w(M'))$ , since we evaluate each arc in  $M$  exactly once, and each such evaluation may involve up to  $w(M')$  nodes in  $M'$ . Lines 3-5 in Algorithm 3 evaluate each arc in  $M$ , and the comparison of all nodes in  $s^+(u)$  with those in  $s^-(v)$  (line 4) takes at most  $(w(M'))^2$  steps.  $\square$

We remark that MDD-consistency does not ensure that each solution in  $\text{Sol}(M)$  is represented by some path in  $M'$ . This would be the effect of intersecting  $M$  and  $M'$ . MDD-consistency merely establishes that each arc in  $M$  belongs to some solution is also in  $\text{Sol}(M')$ . Intersecting  $M$  and  $M'$  would establish the strongest form of consistency between the two MDDs as possible. It creates a new MDD  $\tilde{M}$  which contains exactly the set of solutions  $\text{Sol}(M) \cap \text{Sol}(M')$ . The limitation in employing MDD intersections, however, is that the width of  $\tilde{M}$  may be as large as the product of the widths of  $M$  and  $M'$  so intersecting  $M$  with multiple MDDs will, in general, increase the size of the MDD exponentially.

#### 4.4 Application to Sequence

We next apply Theorem 3 to the SEQUENCE constraint.

**Corollary 1** *Let  $X$  be an ordered sequence of variables,  $C = \text{SEQUENCE}(X, q, l, u, S)$  a sequence constraint, and  $M$  an arbitrary MDD following the variable ordering of  $X$ . Establishing MDD consistency for  $C$  on  $M$  is fixed parameter tractable with respect to parameter  $q$ .*

*Proof.* We know from (van Hoeve et al., 2006, 2009) that there exists an exact MDD  $M'$  of size  $O(n2^{q-1})$  that represents  $C$ . The result then follows immediately from applying Theorem 3.  $\square$

We note that Theorem 3 can also be applied to obtain the tractability of establishing MDD consistency on other constraints. Consider for example the constraint AMONG( $x_1, x_2, \dots, x_n, l, u, S$ ). For any variable ordering, we can construct an exact MDD in a top-down procedure by associating with each node  $v$  the number of variables taking a value in  $S$  along the path from  $r$  to  $v$ , representing the ‘length’ of that path. Nodes with the same length are equivalent and can be merged. Because the largest layer has at most  $u$  different path lengths, the exact MDD has size  $O(nu)$ , and by Theorem 3 establishing MDD consistency is tractable for AMONG. Indeed, Hoda et al. (2010) also showed that MDD consistency can be established for this constraint, with quadratic time complexity.

The converse of Theorem 3 does not hold: there exist constraints for which MDD consistency can be established in polynomial time on any given MDD, while a minimal

reduced exact MDD has exponential size. As a specific example, consider linear inequality constraints of the form  $\sum_{i=1}^n a_i x_i \geq b$  where  $x_i$  is an integer variable,  $a_i$  is a constant, for  $i = 1, \dots, n$ , and  $b$  is a constant. MDD consistency can be established for such constraints in linear time, for any given MDD, by computing the shortest  $r$ - $t$  path (relative to the coefficients  $a_i$ ) for each arc (Andersen et al., 2007). However, Hosaka et al. (1997) provide the following explicit linear inequality. For  $k$  even and  $n = k^2$ , consider  $\sum_{1 \leq i, j \leq k} a_{ij} x_{ij} \geq k(2^{2k} - 1)/2$ , where  $x_{ij}$  is a binary variable, and  $a_{ij} = 2^{i-1} + 2^{k+j-1}$ , for  $1 \leq i, j \leq k$ . They show that, for any variable order, the size of the reduced ordered BDD for this inequality is bounded from below by  $\Omega(2^{\sqrt{n}/2})$ .

## 5. Partial MDD Filtering for Sequence

In many practical situations the value of  $q$  will lead to prohibitively large exact MDDs for establishing MDD consistency, which limits the applicability of Corollary 1. Therefore we next explore a more practical partial filtering algorithm that is polynomial also in  $q$ .

One immediate approach is to propagate the SEQUENCE constraint in MDDs through its natural decomposition into AMONG constraints, and apply the MDD filtering algorithms for AMONG proposed by Hoda et al. (2010). However, it is well-known that for classical constraint propagation based on variable domains, the AMONG decomposition can be substantially improved by a dedicated domain filtering algorithm for SEQUENCE (van Hove et al., 2006, 2009; Brand et al., 2007; Maher et al., 2008). Therefore, our goal in this section is to provide MDD filtering for SEQUENCE that can be stronger in practice than MDD filtering for the AMONG decomposition, and stronger than domain filtering for SEQUENCE. In what follows, we assume that the MDD at hand respects the ordering of the variables in the SEQUENCE constraint.

### 5.1 Cumulative Sums Encoding

Our proposed algorithm extends the original domain consistency filtering algorithm for SEQUENCE by van Hove et al. (2006) to MDDs, following the ‘cumulative sums’ encoding as proposed by Brand et al. (2007). This representation takes the following form. For a sequence of variables  $X = x_1, x_2, \dots, x_n$ , and a constraint  $\text{SEQUENCE}(X, q, l, u, S)$ , we first introduce variables  $y_0, y_1, \dots, y_n$ , with respective initial domains  $D(y_i) = [0, i]$  for  $i = 1, \dots, n$ . These variables represent the cumulative sums of  $X$ , i.e.,  $y_i$  represents  $\sum_{j=1}^i (x_j \in S)$  for  $i = 1, \dots, n$ . We now rewrite the SEQUENCE constraint as the following system of constraints:

$$y_i = y_{i-1} + \delta_S(x_i) \quad \forall i \in \{1, \dots, n\}, \tag{3}$$

$$y_{i+q} - y_i \geq l \quad \forall i \in \{0, \dots, n - q\}, \tag{4}$$

$$y_{i+q} - y_i \leq u \quad \forall i \in \{0, \dots, n - q\}, \tag{5}$$

where  $\delta_S : X \rightarrow \{0, 1\}$  is the indicator function for the set  $S$ , i.e.,  $\delta_S(x) = 1$  if  $x \in S$  and  $\delta_S(x) = 0$  if  $x \notin S$ . Brand et al. (2007) show that establishing singleton bounds consistency on this system suffices to establish domain consistency for the original SEQUENCE constraint.

In order to apply similar reasoning in the context of MDDs, the crucial observation is that the domains of the variables  $y_0, \dots, y_n$  can be naturally represented at the *nodes* of the

MDD. In other words, a node  $v$  in layer  $L_i$  represents the domain of  $y_{i-1}$ , restricted to the solution space formed by all  $r$ - $t$  paths containing  $v$ . Let us denote this information for each node  $v$  explicitly as the interval  $[\text{lb}(v), \text{ub}(v)]$ , and we will refer to it as the ‘node domain’ of  $v$ . Following the approach of Hoda et al. (2010), we can compute this information in linear time by one top-down pass, by using equation (3), as follows:

$$\begin{aligned} \text{lb}(v) &= \min_{(u,v) \in A^{\text{in}}(v)} \{ \text{lb}(u) + \delta_S(\ell(u, v)) \}, \\ \text{ub}(v) &= \max_{(u,v) \in A^{\text{in}}(v)} \{ \text{ub}(u) + \delta_S(\ell(u, v)) \}, \end{aligned} \quad (6)$$

for all nodes  $v \neq r$ , while  $[\text{lb}(r), \text{ub}(r)] = [0, 0]$ .

As the individual AMONG constraints are now posted as  $y_{i+q} - y_i \geq l$  and  $y_{i+q} - y_i \leq u$ , we also need to compute for a node  $v$  in layer  $L_{i+1}$  all its ancestors from layer  $L_i$ . This can be done by maintaining a vector  $\mathcal{A}_v$  of length  $q + 1$  for each node  $v$ , where  $\mathcal{A}_v[i]$  represents the set of ancestor nodes of  $v$  at the  $i$ -th layer above  $v$ , for  $i = 0, \dots, q$ . We initialize  $\mathcal{A}_r = [r, \emptyset, \dots, \emptyset]$ , and apply the recursion

$$\begin{aligned} \mathcal{A}_v[i] &= \cup_{(u,v) \in A^{\text{in}}(v)} \mathcal{A}_u[i-1] \quad \text{for } i = 1, 2, \dots, q, \\ \mathcal{A}_v[0] &= \{v\}. \end{aligned}$$

The resulting top-down pass itself takes linear time (in the size of the MDD), while a direct implementation of the recursive step for each node takes  $O(qW^2)$  operations, where  $W$  is the maximum width of the MDD. Now, the relevant ancestor nodes for a node  $v$  in layer  $L_{i+q}$  are stored in  $\mathcal{A}_v[q]$ , a subset of layer  $L_i$ .

We similarly compute all descendant nodes of  $v$  in a vector  $\mathcal{D}_v$  of length  $q + 1$ , such that  $\mathcal{D}_v[i]$  contains all descendants of  $v$  in the  $i$ -th layer below  $v$ , for  $i = 0, 1, \dots, q$ . We initialize  $\mathcal{D}_1 = [\mathbf{1}, \emptyset, \dots, \emptyset]$ .

Alternatively, we can approximate the information from the ancestor and descendant nodes by only maintaining a lower and upper bound on the node domains of  $\mathcal{A}_v$ , resp.,  $\mathcal{D}_v$ , resulting in a recursive step of  $O(qW)$  operations per node.

## 5.2 Processing the Constraints

We next process each of the constraints (3), (4), and (5) in turn to remove provably inconsistent arcs, while at the same time we filter the node information.

Starting with the ternary constraints of type (3), we remove an arc  $(u, v)$  if  $\text{lb}(u) + \delta_S(\ell(u, v)) > \text{ub}(v)$ . Updating  $[\text{lb}(v), \text{ub}(v)]$  for a node  $v$  is done similar to the rules (6) above:

$$\begin{aligned} \text{lb}(v) &= \max \{ \text{lb}(v), \min_{(u,v) \in A^{\text{in}}(v)} \{ \text{lb}(u) + \delta_S(\ell(u, v)) \} \}, \\ \text{ub}(v) &= \min \{ \text{ub}(v), \min_{(u,v) \in A^{\text{in}}(v)} \{ \text{ub}(u) + \delta_S(\ell(u, v)) \} \}, \end{aligned} \quad (7)$$

In fact, the resulting algorithm is a special case of the MDD consistency equality propagator of Hadzic et al. (2008a), and we thus inherit the MDD consistency for our ternary constraints.

Next, we process the constraints (4) and (5) for a node  $v$  in layer  $L_{i+1}$  ( $i = 0, \dots, n$ ). Recall that the relevant ancestors from  $L_{i+1-q}$  are  $\mathcal{A}_v[q]$ , while its relevant descendants




Figure 3: MDD propagation for the constraint  $\text{SEQUENCE}(X, q = 3, l = 1, u = 2, S = \{1\})$  of Example 3.

from  $L_{i+1+q}$  are  $\mathcal{D}_v[q]$ . The variable corresponding to node  $v$  is  $y_i$ , and it participates in four constraints:

$$\begin{aligned}
 y_i &\geq l + y_{i-q}, \\
 y_i &\leq u + y_{i-q}, \\
 y_i &\leq y_{i+q} - l, \\
 y_i &\geq y_{i+q} - u.
 \end{aligned} \tag{8}$$

Observe that we can apply these constraints to filter *only* the node domain  $[\text{lb}(v), \text{ub}(v)]$  corresponding to  $y_i$ . Namely, the node domains corresponding to the other variables  $y_{i-q}$  and  $y_{i+q}$  may find support from nodes in layer  $L_{i+1}$  other than  $v$ . We update  $\text{lb}(v)$  and  $\text{ub}(v)$  according to equations (8):

$$\begin{aligned}
 \text{lb}(v) &= \max\left\{ \text{lb}(v), \quad l + \min_{u \in \mathcal{A}_v[q]} \text{lb}(u), \quad \min_{w \in \mathcal{D}_v[q]} \text{lb}(w) - u \right\}, \\
 \text{ub}(v) &= \min\left\{ \text{ub}(v), \quad u + \max_{u \in \mathcal{A}_v[q]} \text{ub}(u), \quad \max_{w \in \mathcal{D}_v[q]} \text{ub}(w) - l \right\}.
 \end{aligned} \tag{9}$$

The resulting algorithm is a specific instance of the generic MDD consistent binary constraint propagator presented by Hoda et al. (2010), and again we inherit the MDD consistency for these constraints. We can process the constraints in linear time (in the size of the MDD) by a top-down and bottom-up pass through the MDD.

**Example 3** Consider the constraint  $\text{SEQUENCE}(X, q = 3, l = 1, u = 2, S = \{1\})$  with the ordered sequence of binary variables  $X = \{x_1, x_2, x_3, x_4, x_5\}$ . Assume we are given the MDD in Figure 3.a. In Figure b. we show the node domains that result from processing rules (6). Figure c. shows the resulting MDD after processing the constraints via the rules (7) and (9). For example, consider the middle node in the fourth layer, corresponding to variable

$y_3$ . Let this node be  $v$ . It has initial domain  $[0, 2]$ , and  $\mathcal{A}_v[q]$  only contains the root node, which has domain  $[0, 0]$ . Since  $l = 1$ , we can reduce the domain of  $v$  to  $[1, 2]$ . We can next consider the arcs into  $v$ , and conclude that value 1 in its domain is not supported. This further reduces the domain of  $v$  to  $[2, 2]$ , and allows to eliminate one incoming arc (from the first node of the previous layer).

The resulting MDD in Figure c. reflects all possible deductions that can be made by our partial algorithm. We have not established MDD consistency however, as witnessed by the infeasible path  $(1, 1, 0, 0, 0)$ .

Observe that our proposed algorithm can be applied immediately to the more general GEN-SEQUENCE constraints in which each AMONG constraint has its individual  $l, u$  and  $q$ , as the cumulative sums encoding can be adjusted in a straightforward manner to represent these different values.

## 6. Computational Results

The purpose of our computational results is to evaluate empirically the strength of the ‘partial filtering’ MDD propagator from Section 5. We perform three main comparisons. First, we want to assess the impact of increasing the maximum width of the MDD on the filtering. Second, we want to compare the MDD propagation with the classical domain propagation for SEQUENCE. In particular, we wish to evaluate the computational overhead of MDD propagation relative to domain propagation, and to what extent MDD propagation can outperform domain propagation. Third, we compare the filtering strength of our MDD propagator for SEQUENCE to the filtering strength of the MDD propagators for the individual AMONG constraints, being the best MDD approach for SEQUENCE so far (Hoda et al., 2010).

We have implemented our MDD propagator for SEQUENCE as a global constraint in IBM ILOG CPLEX CP Optimizer 12.4, using the C++ interface. For the MDD propagator for AMONG, we apply the code of (Hoda et al., 2010). For the domain propagation, we implemented two models. The first uses the domain consistent propagator for SEQUENCE from (van Hoeve et al., 2009), running in  $O(n^3)$  time. The currently fastest domain consistent propagator for SEQUENCE runs in  $O(n^2)$  time (Maher et al., 2008; Downing et al., 2012), but unfortunately the code of (Downing et al., 2012) is not publicly available. To control for this, we applied as second domain propagation model the decomposition into cumulative sums, which uses no explicit global constraint for SEQUENCE. Propagating this decomposition takes  $O(n^2)$  in the worst case, as it considers  $O(n)$  variables and constraints while the variable domains contain up to  $n$  elements. We note that for almost all test instances, the cumulative sums encoding established domain consistency. As an additional advantage, the cumulative sums encoding permits a more insightful comparison with our MDD propagator, since both are based on the cumulative sums decomposition.

Because single SEQUENCE constraints can be solved in polynomial time, we consider instances with multiple SEQUENCE constraints in our experiments. We assume that these are defined on the same ordered set of variables. To measure the impact of the different propagation methods correctly, all approaches apply the same fixed search strategy, i.e., following the given ordering of the variables, with a lexicographic value ordering heuristic.



For each method, we measure the number of backtracks from a failed search state as well as the solving time. All experiments are performed using a 2.33GHz Intel Xeon machine.

### 6.1 Systems of Sequence Constraints

We first consider systems of multiple independent SEQUENCE constraints. We generate instances with  $n = 50$  variables each having domain  $\{0, 1, \dots, 10\}$ , and 5 SEQUENCE constraints. For each SEQUENCE constraint, we set the length of sub-sequence uniform randomly between  $[5, n/2)$  as

$$q = (\text{rand()} \% ((n/2) - 5)) + 5.$$

Here, `rand()` refers to the standard C++ random number generator, i.e., `rand() % k` selects a number in the range  $[0, k - 1]$ . Without the minimum length of 5, many of the instances would be very easy to solve by either method. We next define the difference between  $l$  and  $u$  as  $\Delta := (\text{rand()} \% q)$ , and set

$$\begin{aligned} l &:= (\text{rand()} \% (q - \Delta)), \\ u &:= l + \Delta. \end{aligned}$$

Lastly, we define the set of values  $S$  by first defining its cardinality as  $(\text{rand()} \% 11) + 1$ , and then selecting that many values uniformly at random from  $\{0, 1, \dots, 10\}$ . We generated 250 such instances in total. All instances will be made available via <http://www.andrew.cmu.edu/user/vanhoeve/mdd/>.




Figure 4: Performance comparison of domain and MDD propagators for the SEQUENCE constraint. Each data point reflects the total number of instances that are solved by a particular method within the corresponding time limit.





Figure 5: Comparing domain and MDD propagation for SEQUENCE constraints. Each data point reflects the number of backtracks (a.) resp. solving time in seconds (b.) for a specific instance, when solved with the best domain propagator (cumulative sums encoding) and the MDD propagator with maximum width 32. Instances for which either method needed 0 backtracks (a.) or less than 0.01 seconds (b.) are excluded. Here, TO stands for ‘timeout’ and represents that the specific instance could not be solved within 1,800s (Fig. b.). In Figure a., these instances are labeled separately by TO (at tick-mark  $10^8$ ); note that the reported number of backtracks after 1,800 seconds may be much less than  $10^8$  for these instances. All reported instances with fewer than  $10^8$  backtracks were solved within the time limit.

We solve each instance using the domain consistency propagator for SEQUENCE, the cumulative sums encoding (domain propagation), and the MDD propagator with maximum widths 2, 4, 8, 16, 32, 64, 128. Each method is given a maximum time limit of 1,800 seconds per instance.

We compare the performance of domain propagation and MDD propagation in Figure 4. In this figure, we report for each given time point how many instances could be solved within that time by a specific method. Observe that the cumulative sums domain propagation, although not guaranteed to establish domain consistency, outperforms the dedicated SEQUENCE propagator. Also, MDD propagation with maximum width 2 can already substantially outperform domain propagation. We can further observe that larger maximum widths require more time for the MDDs to be processed, but in the end it does allow to solve more instances: maximum MDD width 128 permits to solve all 250 instances within the given time limit, whereas domain propagation can solve 230 (SEQUENCE), respectively 232 (cumulative sums), instances.



a. Survival function with respect to backtracks      b. Survival function with respect to solving time

Figure 6: Evaluating the impact of increased width for MDD propagation via survival function plots with respect to search backtracks (a.) and solving time (b.). Both plots are in log-log scale. Each data point reflects the percentage of instances that require at least that many backtracks (a.) resp. seconds (b.) to be solved by a particular method.

To illustrate the difference between domain and MDD propagation in more detail, Figure 5 presents scatter plots comparing domain propagation (cumulative sums) with MDD propagation (maximum width 32). This comparison is particularly meaningful because both propagation methods rely on the cumulative sums representation. For each instance, Figure 5.a depicts the number of backtracks while Figure 5.b depicts the solving time of both methods. The instances that were not solved within the time limit are collected under ‘TO’ (time out) for that method. Figure 5.a demonstrates that MDD propagation can lead to dramatic search tree reductions, by several orders of magnitude. Naturally, the MDD propagation comes with a computational cost, but Figure 5.b shows that for almost all instances (especially the harder ones), the search tree reductions correspond to faster solving times, again often several orders of magnitude.

We next evaluate the impact of increasing maximum widths of the MDD propagator. In Figure 6, we present for each method the ‘survival function’ with respect to the number of backtracks (a.) and solving time (b.). Formally, when applied to combinatorial backtrack search algorithms, the survival function represents the probability of a run taking more than  $x$  backtracks (Gomes et al., 2005). In our case, we approximate this function by taking the proportion of instances that need at least  $x$  backtracks (Figure 6.a), respectively seconds (Figure 6.b). Observe that these are log-log plots. With respect to the search tree size, Figure 6.a clearly shows the strengthening of the MDD propagation when the maximum width is increased. In particular, the domain propagation reflects the linear behavior over

| Requirement  | SEQUENCE( $X, q, l, u, S$ )              |
|--|--|
| At least 20 work shifts every 28 days:             | SEQUENCE( $X, 28, 20, 28, \{D, E, N\}$ ) |
| At least 4 off-days every 14 days:                 | SEQUENCE( $X, 14, 4, 14, \{O\}$ )        |
| At least 4 off-days every 14 days:                 | SEQUENCE( $X, 14, 4, 14, \{O\}$ )        |
| Between 1 and 4 night shifts every 14 days:        | SEQUENCE( $X, 14, 1, 4, \{N\}$ )         |
| Between 4 and 8 evening shifts every 14 days:      | SEQUENCE( $X, 14, 4, 8, \{E\}$ )         |
| Nights shifts cannot appear on consecutive days:   | SEQUENCE( $X, 2, 0, 1, \{N\}$ )          |
| Between 2 and 4 evening/night shifts every 7 days: | SEQUENCE( $X, 7, 2, 4, \{E, N\}$ )       |
| At most 6 work shifts every 7 days:                | SEQUENCE( $X, 7, 0, 6, \{D, E, N\}$ )    |

Table 1: Nurse rostering problem specification. Variable set  $X$  represents the shifts to be assigned over a sequence of days. The possible shifts are day (D), evening (E), night (N), and day off (O).

several orders of magnitude that is typical for heavy-tailed runtime distributions. Naturally, similar behavior is present for the MDD propagation, but in a much weaker form for increasing maximum MDD widths. The associated solving times are presented in Figure 6.b. It reflects similar behavior, but also takes into account the initial computational overhead of MDD propagation.

## 6.2 Nurse Rostering Instances

We next consider a more structured problem class inspired by nurse rostering problems. The problem is to design a work schedule for a nurse over a given horizon of  $n$  days. On each day, a nurse can either work a day shift (D), evening shift (E), night shift (N), or have a day off (O). We introduce a variable  $x_i$  for each day  $i = 1, \dots, n$ , with domain  $D(x_i) = \{O, D, E, N\}$  representing the shift. We impose the eight SEQUENCE constraints modeling the requirements listed in Table 1.

By the combinatorial nature of this problem, the size of the CP search tree turns out to be largely independent on the length of the time horizon, when a lexicographic search (by increasing day  $i$ ) is applied. We however do consider instances with various time horizons ( $n = 40, 60, 80, 100$ ), to address potential scaling issues.

The results are presented in Table 2. The columns for ‘Domain Sequence’ show the total number of backtracks (BT) and solving time in seconds (CPU) for the domain consistent SEQUENCE propagator. Similarly, the columns for ‘Domain Cumul. Sums’ show this information for the cumulative sums domain propagation. The subsequent columns show these numbers for the MDD propagator, for MDDs of maximum width 1, 2, 4, and 8. Note that propagating an MDD of width 1 corresponds to domain propagation, and indeed the associated number of backtracks is equivalent to the domain propagator of the cumulative sums. As a first observation, a maximum width of 2 already reduces the number of backtracks by a factor 8.3. For maximum width of 8 the MDD propagation even allows to solve the problem without search. The computation times are correspondingly reduced, e.g., from 157s (resp. 96s) for the domain propagators to 0.10s for the MDD propagator (width 8) for

| $n$ | Domain Sequence |        | Domain Cumul. Sums |       | MDD Width 1 |        | MDD Width 2 |       | MDD Width 4 |      | MDD Width 8 |      |
|-----|-----------------|--------|--------------------|-------|-------------|--------|-------------|-------|-------------|------|-------------|------|
|     | BT              | CPU    | BT                 | CPU   | BT          | CPU    | BT          | CPU   | BT          | CPU  | BT          | CPU  |
| 40  | 438,059         | 43.83  | 438,059            | 32.26 | 438,059     | 54.27  | 52,443      | 12.92 | 439         | 0.44 | 0           | 0.02 |
| 60  | 438,059         | 78.26  | 438,059            | 53.40 | 438,059     | 80.36  | 52,443      | 18.36 | 439         | 0.68 | 0           | 0.04 |
| 80  | 438,059         | 124.81 | 438,059            | 71.33 | 438,059     | 106.81 | 52,443      | 28.58 | 439         | 0.94 | 0           | 0.06 |
| 100 | 438,059         | 157.75 | 438,059            | 96.27 | 438,059     | 135.37 | 52,443      | 37.76 | 439         | 1.22 | 0           | 0.10 |

Table 2: Comparing domain propagation and the MDD propagation for SEQUENCE on nurse rostering instances. Here,  $n$  stands for the number of variables, BT for the number of backtracks, and CPU for solving time in seconds.

the instance with  $n = 100$ . Lastly, we can observe that in this case MDD propagation does not suffer from scaling issues when compared to domain propagation.

As a final remark, we also attempted to solve these nurse rostering instances using the SEQUENCE domain propagator of CP Optimizer (`IloSequence`). It was able to solve the instance with  $n = 40$  in 1,150 seconds, but none of the others instances were solved within the time limit of 1,800 seconds.

### 6.3 Comparing MDD Filtering for Sequence and Among

In our last experiment, we compare our SEQUENCE MDD propagator to the MDD propagator for AMONG constraints by Hoda et al. (2010). Our main goal is to determine whether a large MDD is by itself sufficient to solve these problem (irrespective of propagating AMONG or a cumulative sums decomposition), or whether the additional information obtained by our SEQUENCE propagator makes the difference.

We apply both methods, MDD propagation for SEQUENCE and MDD propagation for AMONG, to the data set of Section 6.1 containing 250 instances. The time limit is again 1,800 seconds, and we run the propagators with maximum MDD widths 2, 8, 32, and 128.

We first compare the performance of the MDD propagators for AMONG and SEQUENCE in Figure 7. The figure depicts the number of instances that can be solved within a given time limit for the various methods. The plot indicates that the AMONG propagators are much weaker than the SEQUENCE propagator, and moreover that larger maximum widths alone do not suffice: using the SEQUENCE propagator with maximum width 2 outperforms the AMONG propagators for all maximum widths up to 128.

The scatter plot in Figure 8 compares the MDD propagators for AMONG and SEQUENCE in more detail, for widths 2, 8, 32, and 128 (instances that take 0 backtracks, resp. less than 0.01 seconds, for either method are discarded from Figure 8.a, resp. 8.b). For smaller widths, there are several instances that the AMONG propagator can solve faster, but the relative strength of the SEQUENCE propagator increases with larger widths. For width 128, the SEQUENCE propagator can achieve orders of magnitude smaller search trees and solving time than the AMONG propagators, which again demonstrates the advantage of MDD propagation for SEQUENCE when compared to the AMONG decomposition.





Figure 7: Performance comparison of MDD propagation for SEQUENCE and AMONG for various maximum widths. Each data point reflects the total number of instances that are solved by a particular method within the corresponding time limit.



a. Number of backtracks

b. Solving time

Figure 8: Evaluating MDD propagation for SEQUENCE and AMONG for various maximum widths via scatter plots with respect to search backtracks (a.) and solving time (b.). Both plots are in log-log scale and follow the same format as Figure 5.

## 7. Conclusion

Constraint propagation with limited-width MDDs has recently been shown to be a powerful alternative to the conventional propagation of variable domains in constraint programming. In this work, we have studied MDD propagation for the SEQUENCE constraint, which appears in, e.g., rostering and scheduling applications. We have first proved that establishing MDD consistency for SEQUENCE is NP-hard. However, we have also shown that this task is fixed parameter tractable with respect to the length of the sub-sequences defined by the constraint, provided that the MDD follows the variable ordering provided by the constraint. We then proposed a practical MDD propagation algorithm for sequence that is also polynomial in the length of the sub-sequences, which is based on a cumulative decomposition. We provided extensive experimental results comparing our MDD propagator for SEQUENCE to domain propagators for SEQUENCE as well as existing MDD propagator for AMONG. Our computational experiments have shown that our MDD propagator for SEQUENCE can outperform domain propagators by orders of magnitude in terms of search tree size and solving time. Similar results were obtained when compared to the existing MDD propagator for AMONG, which demonstrates that in practice a large MDD alone is not sufficient to solve these problems; specific MDD propagators for global constraints such as SEQUENCE can lead to orders of magnitude speedups.

## References

- Andersen, H., Hadzic, T., Hooker, J., & Tiedemann, P. (2007). A Constraint Store Based on Multivalued Decision Diagrams. In *Proceedings of CP*, Vol. 4741 of *LNCS*, pp. 118–132. Springer.
- Apt, K. R. (2003). *Principles of Constraint Programming*. Cambridge University Press.
- Beldiceanu, N., & Contejean, E. (1994). Introducing global constraints in CHIP. *Journal of Mathematical and Computer Modelling*, 20(12), 97–123.
- Brand, S., Narodytska, N., Quimper, C., Stuckey, P., & Walsh, T. (2007). Encodings of the sequence constraint. In *Proceedings of CP*, Vol. 4741 of *LNCS*, pp. 210–224. Springer.
- Cheng, K., & Yap, R. (2008). Maintaining Generalized Arc Consistency on Ad Hoc r-Ary Constraints. In *Proceedings of CP*, Vol. 5202 of *LNCS*, pp. 509–523. Springer.
- Cire, A. A., & van Hoes, W.-J. (2012). MDD Propagation for Disjunctive Scheduling. In *Proceedings of ICAPS*, pp. 11–19. AAAI Press.
- Cire, A. A., & van Hoes, W.-J. (2013). Multivalued Decision Diagrams for Sequencing Problems. *Operations Research*, to appear.
- Dechter, R. (2003). *Constraint Processing*. Morgan Kaufmann.
- Downing, N., Feydy, T., & Stuckey, P. (2012). Explaining flow-based propagation. In *Proceedings of CPAIOR*, Vol. 7298 of *LNCS*, pp. 146–162. Springer.
- Garey, M., & Johnson, D. (1979). *Computers and Intractability - A Guide to the Theory of NP-Completeness*. Freeman.
- Gomes, C. P., Fernández, C., Selman, B., & Bessière, C. (2005). Statistical Regimes Across Constrainedness Regions. *Constraints*, 10(4), 317–337.

- Hadzic, T., Hooker, J., O’Sullivan, B., & Tiedemann, P. (2008a). Approximate Compilation of Constraints into Multivalued Decision Diagrams. In *Proceedings of CP*, Vol. 5202 of *LNCS*, pp. 448–462. Springer.
- Hadzic, T., Hooker, J., & Tiedemann, P. (2008b). Propagating Separable Equalities in an MDD Store. In *Proceedings of CPAIOR*, Vol. 5015 of *LNCS*, pp. 318–322. Springer.
- Hadzic, T., O’Mahony, E., O’Sullivan, B., & Sellmann, M. (2009). Enhanced Inference for the Market Split Problem. In *Proceedings of ICTAI*, pp. 716–723. IEEE.
- Hawkins, P., Lagoon, V., & Stuckey, P. (2005). Solving Set Constraint Satisfaction Problems Using ROBDDs. *JAIR*, 24(1), 109–156.
- Hoda, S., van Hoeve, W.-J., & Hooker, J. (2010). A Systematic Approach to MDD-Based Constraint Programming. In *Proceedings of CP*, Vol. 6308 of *LNCS*, pp. 266–280. Springer.
- Hosaka, K., Takenaga, Y., Kaneda, T., & Yajima, S. (1997). Size of ordered binary decision diagrams representing threshold functions. *Theoretical Computer Science*, 180, 47–60.
- Maher, M., Narodytska, N., Quimper, C.-G., & Walsh, T. (2008). Flow-Based Propagators for the SEQUENCE and Related Global Constraints. In *Proceedings of CP*, Vol. 5202 of *LNCS*, pp. 159–174. Springer.
- Régin, J.-C. (1994). A Filtering Algorithm for Constraints of Difference in CSPs. In *Proceedings of AAAI*, Vol. 1, pp. 362–367. AAAI Press.
- Régin, J.-C. (2011). Global Constraints: A Survey. In Van Hentenryck, P., & Milano, M. (Eds.), *Hybrid Optimization*, pp. 63–134. Springer.
- Régin, J.-C., & Puget, J.-F. (1997). A Filtering Algorithm for Global Sequencing Constraints. In *Proceedings of CP*, Vol. 1330 of *LNCS*, pp. 32–46. Springer.
- Rossi, F., van Beek, P., & Walsh, T. (Eds.). (2006). *Handbook of Constraint Programming*. Elsevier.
- van Hoeve, W.-J., & Katriel, I. (2006). Global Constraints. In Rossi, F., van Beek, P., & Walsh, T. (Eds.), *Handbook of Constraint Programming*, chap. 6. Elsevier.
- van Hoeve, W.-J., Pesant, G., Rousseau, L.-M., & Sabharwal, A. (2006). Revisiting the Sequence Constraint. In *Proceedings of CP*, Vol. 4204 of *LNCS*, pp. 620–634. Springer.
- van Hoeve, W.-J., Pesant, G., Rousseau, L.-M., & Sabharwal, A. (2009). New Filtering Algorithms for Combinations of Among Constraints. *Constraints*, 14, 273–292.
- Wegener, I. (2000). *Branching Programs and Binary Decision Diagrams: Theory and Applications*. SIAM monographs on discrete mathematics and applications. Society for Industrial and Applied Mathematics.