

Compact-Table: Efficiently Filtering Table Constraints with Reversible Sparse Bit-Sets

Jordan Demeulenaere¹, Renaud Hartert¹, Christophe Lecoutre²,
Guillaume Perez³, Laurent Perron⁴, Jean-Charles Régin³, Pierre Schaus¹

¹UCLouvain, Belgium ²CRIL, Univ. Artois & CNRS, F-62300 Lens, France
³University of Nice, France ⁴Google, France,

Abstract. In this paper, we describe Compact-Table (CT), a bitwise algorithm to enforce Generalized Arc Consistency (GAC) on table constraints. Although this algorithm is the default propagator for table constraints in or-tools and OspaR, two publicly available CP solvers, it has never been described so far. Importantly, CT has been recently improved further with the introduction of residues, resetting operations and a data-structure called reversible sparse bit-set, used to maintain tables of supports (following the idea of tabular reduction): tuples are invalidated incrementally on value removals by means of bit-set operations. The experimentation that we have conducted with OspaR shows that CT outperforms state-of-the-art algorithms STR2, STR3, GAC4R, MDD4R and AC5-TC on standard benchmarks.

1 Introduction

Table constraints, also called extension(al) constraints, explicitly express the allowed combinations of values for the variables they involve as sequences of tuples, which are called tables. Table constraints can theoretically encode any kind of constraints and are among the most useful ones in Constraint Programming (CP). Indeed, they are often required when modeling combinatorial problems in many application fields. The design of filtering algorithms for such constraints has generated a lot of research effort, see [1,23,21,10,33,17,20,12,30].

Over the last decade, many developments have thus been achieved for enforcing the well-known property called Generalized Arc Consistency (GAC) on binary and/or non-binary extensionally defined constraints. Among successful techniques, we find:

- bitwise operations that allow performing parallel operations on bit vectors. Already exploited during the 70's [32,27], they have been applied more recently to the enforcement of arc consistency on binary constraints [3,22].
- residual supports (residues) that store the last found supports of each value. Initially introduced for ensuring optimal complexity [2], they have been shown efficient in practice [18,24,19] when used as simple sentinels.
- tabular reduction, which is a technique that dynamically maintains the tables of supports. Based on the structure of sparse sets [4,16], variants of Simple Tabular Reduction (STR) have been proved to be quite competitive [33,17,20].

- resetting operations that saves substantial computing efforts in some particular situations. They have been successfully applied to the algorithm GAC4 [30].

In this paper, we introduce a very efficient GAC algorithm¹ for table constraints that combines the use of bitwise operations, residual supports, tabular reduction, and resetting operations. It is called Compact-Table (CT), and originates from or-tools, the Google solver that won the latest MiniZinc Challenges. It is important to note that or-tools developers prefer to focus on highly-optimized implementations of a few important (global) constraints instead of having many of them. Through the years, CT has reached a good level of maturity because it has been continuously improved and extended with many cutting edge ideas such as those introduced earlier. Unfortunately, the core algorithm of CT has not been described in the literature so far² and is thus seldom used as a reference for practical comparisons. The first version of CT implemented in or-tools, with a bit-set representation of tables, dates back to 2012, whereas the version of CT presented in this paper is exactly the last one implemented in OscaR [29].

Outline. After presenting related works in Section 2, we introduce some technical background in Section 3. Then, we recall in Section 4 usual state restoration mechanisms implemented in CP solvers, and describe reversible sparse bit-sets in Section 5. In Section 6, we describe our algorithm CT. Before concluding, we present in Section 7 the results of an experimentation we have conducted with CT and its contenders on a large variety of benchmarks.

2 Related Work

Propagators for table constraints are filtering procedures used to enforce GAC. Given the importance of table constraints, it is not surprising that much research has been carried out in order to find efficient propagators. This section briefly describes the most efficient ones.

Generic Algorithms. On the one hand, GAC3 is a classical general-purpose GAC algorithm [25] for non-binary constraints. Each call to this algorithm for a constraint requires testing if each value is still supported by a valid tuple accepted by the constraint. Several improvements to fasten the search for a support gave birth to variants such as GAC2001 [2] and GAC3^{rm} [19]. Unfortunately, the worst-case time complexity of all these algorithms grows exponentially with the arity of the constraints. On the other hand, GAC4 [28] is a value-based algorithm, meaning here that for each value, it maintains a set of valid tuples supporting it. Each time a value is removed, all supporting tuples are removed from the associated sets, which allows us to identify values without any more supports.

¹ We are aware of an independent work [34] on a similar topic, but hadn't the opportunity of reading it at the time of writing our paper.

² Note that some parts of this paper were published in a Master Thesis report [7].

GAC4R is a recent improvement of GAC4 [30], which recomputes the sets of supporting tuples from scratch (referred to as a resetting operations) when it appears to be less costly than updating them based on the removed values.

AC5 Instantiations. In [12], Mairy et al. introduce several instantiations of the generic AC5 algorithm for table constraints, the best of them being AC5-TCOptSparse. This algorithm shares some similarities with GAC4 since it pre-computes lists of supporting tuples which allows us to retrieve efficiently new supports by iterating over these lists. Note that a reversible integer, i.e., an integer storage location with a facility to restore its successive values, is used to indicate the current position of a support in each list. This algorithm is implemented in Comet, and has been shown to be efficient on ternary and quaternary constraints.

Simple Tabular Reduction. STR1 [33] and STR2 [17] are algorithms that globally enforce GAC by traversing the constraint tables while dynamically maintaining them: each call to the algorithm for a constraint removes the invalid tuples from its table. The improvements brought in STR2 avoid unnecessary operations by considering only relevant subsets of variables when checking the validity of a tuple, and collecting supported values. Contrary to its predecessors, STR3 [20] is a fine-grained (or value-based) algorithm. For each value, it initially computes a static array of tuples supporting it, and keeps a reversible integer *curr* that indicates the position of the last valid tuple in the array. STR3 also maintains the set of valid tuples. STR3 is shown to be complementary to STR2, being more efficient when the tables are not reduced drastically during search.

Compressed Representations. Other algorithms gamble on the compression of tables to reduce the time needed to ensure GAC. The most promising data structure allowing a more compact representation is the Multi-valued Decision Diagram (MDD) [31], but note that the order of variables used to build an MDD may significantly impact its size. Two notable algorithms using MDDs as main data structure are *mddc* [6] and MDD4R [30]. The former does not modify the decision diagram and performs a depth-first search of the MDD during propagation to detect which parts of the MDD are consistent or not. MDD4R dynamically maintains the MDD by deleting nodes and edges that do not belong to a solution. Each value is matched with its corresponding edges in the MDD, so, when a value has none of its edges present in the MDD, it can be removed. On the other hand, some other forms of compression have been studied from the concepts of compressed tuples [14,35], short supports [13] and sliced tables [11].

3 Technical Background

A *constraint network* (CN) N is composed of a set of n variables and a set of e constraints. Each *variable* x has an associated domain, denoted by $dom(x)$,

that contains the finite set of values that can be assigned to it. Each *constraint* c involves an ordered set of variables, called the *scope* of c and denoted by $scp(c)$, and is semantically defined by a *relation*, denoted by $rel(c)$, which contains the set of tuples allowed for the variables involved in c . The arity of a constraint c is $|scp(c)|$, i.e., the number of variables involved in c . A (positive) *table constraint* c is a constraint such that $rel(c)$ is defined explicitly by listing the tuples that are allowed by c .

Example 1. The constraint $x \neq y$ with $x \in \{1, 2, 3\}$ and $y \in \{1, 2\}$ can be alternatively defined by the table constraint c such that $scp(c) = \{x, y\}$ and $rel(c) = \{(1, 2), (2, 1), (3, 1), (3, 2)\}$. We also write:

$$\langle x, y \rangle \in T \quad \text{with} \quad T = \langle (1, 2), (2, 1), (3, 1), (3, 2) \rangle$$

Let $\tau = (a_1, a_2, \dots, a_r)$ be a tuple of values associated with an ordered set of variables $X = \{x_1, x_2, \dots, x_r\}$. The i th value of τ is denoted by $\tau[i]$ or $\tau[x_i]$. The tuple τ is valid iff $\forall i \in 1..r, \tau[i] \in dom(x_i)$. An r -tuple τ is a *support* on the r -ary constraint c iff τ is a valid tuple that is allowed by c . If τ is a support on a constraint c involving a variable x and such that $\tau[x] = a$, we say that τ is a *support for* (x, a) on c . Generalized Arc Consistency (GAC) is a well-known domain-filtering consistency defined as follows:

Definition 1. A constraint c is generalized arc consistent (GAC) iff $\forall x \in scp(c), \forall a \in dom(x)$, there exists at least one support for (x, a) on c . A CN N is GAC iff every constraint of N is GAC.

Enforcing GAC is the task of removing from domains all values that have no support on a constraint. Many algorithms have been devised for establishing GAC according to the nature of the constraints. For table constraints, STR [33] is such an algorithm: it removes invalid tuples during search of supports using a sparse set data structure which separates valid tuples from invalid ones. This method of seeking supports improves search time by avoiding redundant tests on invalid tuples that have already been detected as invalid during previous GAC enforcements. STR2 [17], an optimization of STR, limits some operations concerning the validity of tuples and the identification of supports, through the introduction of two sets called S^{sup} and S^{val} (described later in Section 6).

4 Reversible Objects and Implementation Details

Trail and Timestamping. The issue of storing related states of the solving process is essential in CP. In many solvers³, a general mechanism is used for doing and undoing (on backtrack) the current state. This mechanism is called a trail and it was first introduced in [9] for implementing non-deterministic search. A trail is a stack of pairs $(location, value)$ where *location* stands for any piece

³ One notable exception is Gecode, a copy-based solver.

of memory (e.g., a variable), which can be restored when backtracking. Typically, at each search node encountered during the solving process, the constraint propagation algorithm is executed. A same filtering procedure (propagator) can thus be executed several times at a given node. Consequently, if one is interested in storing some information concerning a filtering procedure, the value of a same memory location can be changed several times. However, stamping that is part of the "folklore" of programming [15] can be used to avoid storing a same memory location on the trail more than once per search node. The idea behind timestamping is that only the final state of a memory location is relevant for its restoration on backtrack. The trail contains a general time counter that is incremented at each search node, and a timestamp is attached to each memory location indicating the time at which its last storage on the trail happened. If a memory location changes and its timestamp matches the current time of the trail then there is no need to store it again. CP solvers generally expose some "reversible" objects to the users using this trail+timestamping mechanism. The most basic one is the reversible version of primitive types such as `int` or `long` values. In the following, we denote by `rint` and `rlong` the reversible versions of `int` and `long` primitive types.

Reversible Sparse Sets. Reversible primitive types can be used to implement more complex data structures such as reversible sets. It was shown in [16] how to implement a reversible set using a single `rint` that represents the current size (limit) of the set. In this structure, which is called reversible sparse set, an array of size n is used to store the permutation from 0 to $n - 1$. All values in this permutation array at indices smaller than or equal to a variable *limit* are considered as part of the set, while the others are considered as removed. When iterating on current values of the set (with decreasing indices from *limit* to 0), the value at the current index can be removed in $O(1)$ by just swapping it with the value stored at *limit* and decrementing *limit*. Making a sparse set reversible just requires managing a single `rint` for *limit*. On backtrack, when the limit is restored, all concerned removed values are restored in $O(1)$.

Domains and Deltas. In Oskar [29], the implementation of domains relies on reversible sparse sets. One advantage is that one can easily retrieve the set of values removed from a domain between any two calls to a given filtering procedure. All we need to store in the filtering procedure is the last size of the domain. The delta set (set of values removed between the two calls) is composed of all the values located between the current size and the last recorded size. More details on this cheap mechanism to retrieve the delta sets can be found in [16].

5 Reversible Sparse Bit-Sets

This section describes the class `RSparseBitSet` that is the main data structure for our algorithm to maintain the supports. In what follows, when we refer to an array t , $t[0]$ denotes the first element (indexing starts at 0) and $t.length$ the number of its cells (size). Also, 0^k will stand for a sequence of k bits set to 0.

Algorithm 1: Class RSparseBitSet

```
1 words: array of rlong // words.length = p
2 index: array of int // index.length = p
3 limit: rint
4 mask: array of long // mask.length = p

5 Method isEmpty(): Boolean
6 | return limit = -1

7 Method clearMask()
8 | foreach i from 0 to limit do
9 | | offset ← index[i]
10 | | mask[offset] ← 064

11 Method reverseMask()
12 | foreach i from 0 to limit do
13 | | offset ← index[i]
14 | | mask[offset] ← ~mask[offset] // bitwise NOT

15 Method addToMask(m: array of long)
16 | foreach i from 0 to limit do
17 | | offset ← index[i]
18 | | mask[offset] ← mask[offset] | m[offset] // bitwise OR

19 Method intersectWithMask()
20 | foreach i from limit downto 0 do
21 | | offset ← index[i]
22 | | w ← words[offset] & mask[offset] // bitwise AND
23 | | words[offset] ← w
24 | | if w = 064 then
25 | | | index[i] ← index[limit]
26 | | | index[limit] ← offset
27 | | | limit ← limit - 1

28 Method intersectIndex(m: array of long): int
29 | /* Post: returns the index of a word where the bit-set
30 | intersects with m, -1 otherwise */
31 | foreach i from 0 to limit do
32 | | offset ← index[i]
33 | | if words[offset] & m[offset] ≠ 064 then
34 | | | return offset
35 | return -1
```

The class `RSparseBitSet`, which encapsulates four fields and 6 methods, is given in Algorithm 1. One important field is `words`, an array of p 64-bit words (actually, reversible long integers), which defines the current value of the bit-set:

the i th bit of the j th word is 1 iff the $(j - 1) \times 64 + i$ th element of the (initial) set is present. Initially, all words in this array have all their bits at 1, except for the last word that may involve a suffix of bits at 0. For example, if we want to handle a set initially containing 82 elements, then we build an array with $p = \lceil 82/64 \rceil = 2$ words that initially looks like:

```
words: 11111111111111111111111111111111 11111111111111111111111111111111
       11111111111111111100000000000000 00000000000000000000000000000000
```

Because, in our context, only non-zero words (words having at least one bit set to 1) are relevant when processing operations on the bit-set, we rely on the sparse-set technique by managing in an array `index` the indices of all words: the indices of all non-zero words are in `index` at positions less than or equal to the value of a variable `limit`, and the indices of all zero-words are in `index` at positions strictly greater than `limit`. For our example, we initially have:

```
words: 11111111111111111111111111111111 11111111111111111111111111111111
       11111111111111111100000000000000 00000000000000000000000000000000
index: 0 1
limit : 1
```

If we suppose now that the 66 first elements of our set above are removed, we obtain:

```
words: 00000000000000000000000000000000 00000000000000000000000000000000
       00111111111111111100000000000000 00000000000000000000000000000000
index: 1 0
limit: 0
```

The class invariant describing the state of a reversible sparse bit-set is the following:

- `index` is a permutation of $[0, \dots, p - 1]$, and
- $\text{words}[\text{index}[i]] \neq 0^{64} \Leftrightarrow i \leq \text{limit}, \forall i \in 0..p - 1$

Note that the reversible nature of our object comes from 1) an array of reversible long (denoted `rlong`) (instead of simple longs) to store the bit words, and 2) the reversible prefix size of non-zero words by using a reversible int (`rint`).

A `RSparseBitSet` also contains a local temporary array, called `mask`. Is used to collect elements with Method `addToMask()`, and can be cleared and reversed too. A `RSparseBitSet` can only be modified by means of the method `intersectWithMask()` which is an operation used to intersect with the elements collected in `mask`. An illustration of the usage of these methods is given in next example.

Example 2. Figure 1 illustrates the use of Methods `addToMask()` and `intersectWithMask()`. We assume that the current state of the bit-set is given by the value of `words`, and that `clearMask()` has been called such that `mask` is initially

words	1	0	1	0	1	1	1	1
addToMask	1	1	1	0	1	0	0	0
addToMask	0	0	0	1	0	0	0	1
mask	1	1	1	1	1	0	0	1
intersectWithMask	1	0	1	0	1	0	0	1

Fig. 1: RSparseBitSet example

empty. Then two bit-sets are collected in **mask** by calling `addToMask()`. The value of **mask** is represented after these two operations. Finally `intersectWithMask()` is executed and the new value of the bit-set **words** is given at the last row of Figure 1.

We now describe the implementation of the methods in `RSparseBitSet`. Method `isEmpty()` simply checks if the number of non-zero words is different from zero (if the limit is set to -1, it means that all words are non-zero). Method `clearMask()` sets to 0 all words of **mask** corresponding to non-zero words of **words**, whereas Method `reverseMask()` reverses all words of **mask**. Method `addToMask()` applies a word by word logical bit-wise *or* operation. Once again, notice that this operation is only applied to words of **mask** corresponding to non-zero words of **words**. Method `intersectMask()` considers each non-zero word of **words** in turn and replaces it by its intersection with the corresponding word of **mask**. In case the resulting new word is zero, it is swapped with the last non-zero word and the value of **limit** is decremented. Finally, Method `intersectIndex()` checks if a given bit-set (array of longs) intersects with the current bit-set: it returns the index of the first word where an intersection can be proved, -1 otherwise.

6 Compact-Table (CT) Algorithm

As STR2 and STR3, Compact-Table (CT) is a GAC algorithm that dynamically maintains the set of valid supports regarding the current domain of each variable. The main difference is that CT is based on an object `RSparseBitSet`. In this set, each tuple is indexed by the order it appears in the initial table. Invalid tuples are removed during the initialization as well as values that are not supported by any tuple. The class `ConstraintCT`, Algorithm 2, allows us to implement any positive table constraint c while running the CT algorithm to enforce GAC.

6.1 Fields

As fields of Class `ConstraintCT`, we first find **scp** for representing the scope of c and **currTable** for representing the current table of c by means of a reversible sparse bit-set. If $\langle \tau_0, \tau_1, \dots, \tau_{t-1} \rangle$ is the initial table of c , then **currTable** is a `RSparseBitSet` object (of initial size t) such that the value i is contained (is set to 1) in the bit-set if and only if the i th tuple is valid:

$$i \in \text{currTable} \Leftrightarrow \forall x \in \text{scp}(c), \tau_i[x] \in \text{dom}(x)$$

We also have three fields S^{val} , S^{sup} and `lastSizes` in the spirit of STR2. The set S^{val} contains variables whose domains have been reduced since the previous invocation of the filtering algorithm on c . To set up S^{val} , we need to record the domain size of each modified variable x right after the execution of CT on c : this value is recorded in `lastSizes[x]`. The set S^{sup} contains unfixed variables (from the scope of the constraint c) whose domains contain each at least one value for which a support must be found. These two sets allow us to restrict loops on variables to relevant ones.

We also have a field `supports` containing static data. During the set up of the table constraint c , CT also computes a static array of words `supports[x, a]`, seen as a bit-set, for each variable-value pair (x, a) where $x \in \text{scp}(c) \wedge a \in \text{dom}(x)$: the bit at position i in the bit-set is 1 if and only if the tuple τ_i in the initial table of c is a support for (x, a) .

T	x	y	z
0	a	a	a
1	a	a	b
2	a	b	c
3	b	a	a
	a	c	b
4	a	b	b
5	b	a	b
6	b	b	a
7	b	b	b

(a) The indexed tuples

<code>currTable</code>	1	1	1	1	1	1	1	1
<code>supports[x, a]</code>	1	1	1	0	1	0	0	0
<code>supports[x, b]</code>	0	0	0	1	0	1	1	1
<code>supports[y, a]</code>	1	1	0	1	0	1	0	0
<code>supports[y, b]</code>	0	0	1	0	1	0	1	1
<code>supports[y, d]</code>	0	0	0	0	0	0	0	0
<code>supports[z, a]</code>	1	0	0	1	0	0	1	0
<code>supports[z, b]</code>	0	1	0	0	1	1	0	1
<code>supports[z, c]</code>	0	0	1	0	0	0	0	0

(b) The corresponding bit-sets

Fig. 2: Illustration of the data structures after the initialization of $\langle x, y, z \rangle \in T$. The tuple (a, c, b) will not be indexed and d will be removed from $\text{dom}(y)$.

Example 3. Figure 2 shows an illustration of the content of those bit-sets after the initialization of the following table constraint $\langle x, y, z \rangle \in T$, with:

- $\text{dom}(x) = \{a, b\}$, $\text{dom}(y) = \{a, b, d\}$, $\text{dom}(z) = \{a, b, c\}$
- $T = \langle (a, a, a), (a, a, b), (a, b, c), (b, a, a), (a, c, b), (a, b, b), (b, a, b), (b, b, a), (b, b, b) \rangle$

The tuple (a, c, b) is initially invalid because $c \notin \text{dom}(y)$, and thus will not be indexed. Value d will be removed from $\text{dom}(y)$ given that it is not supported by any tuple.

Finally, we have an array `residues` such that for each variable-value pair (x, a) , `residues[x, a]` denotes the index of the word where a support was found for (x, a) the last time one was sought for.

Algorithm 2: Class ConstraintCT

```
1 scp: array of variables // Scope
2 currTable: RSparseBitSet // Current table
3  $S^{\text{val}}, S^{\text{sup}}$  // Temporary sets of variables
4 lastSizes // lastSizes[x] is the last size of the domain of x
5 supports // supports[x,a] is the bit-set of supports for (x,a)
6 residues // residues[x,a] is the last found support for (x,a)

7 Method updateTable()
8   foreach variable  $x \in S^{\text{val}}$  do
9     currTable.clearMask()
10    if  $|\Delta_x| < |dom(x)|$  then // Incremental update
11      foreach value  $a \in \Delta_x$  do
12        | currTable.addToMask(supports[x,a])
13      currTable.reverseMask()
14    else // Reset-based update
15      foreach value  $a \in dom(x)$  do
16        | currTable.addToMask(supports[x,a])
17    currTable.intersectWithMask()
18    if currTable.isEmpty() then
19      | break

20 Method filterDomains()
21   foreach variable  $x \in S^{\text{sup}}$  do
22     foreach value  $a \in dom(x)$  do
23       index  $\leftarrow$  residues[x,a]
24       if currTable.words[index] & supports(x,a)[index] =  $0^{64}$  then
25         | index  $\leftarrow$  currTable.intersectIndex(supports[x,a])
26         | if index  $\neq -1$  then
27           | residues[x,a]  $\leftarrow$  index
28         else
29           |  $dom(x) \leftarrow dom(x) \setminus \{a\}$ 
30     lastSize[x]  $\leftarrow |dom(x)|$ 

31 Method enforceGAC()
32    $S^{\text{val}} \leftarrow \{x \in scp : |dom(x)| \neq lastSize[x]\}$ 
33   foreach variable  $x \in S^{\text{val}}$  do
34     | lastSize[x]  $\leftarrow |dom(x)|$ 
35    $S^{\text{sup}} \leftarrow \{x \in scp : |dom(x)| > 1\}$ 
36   updateTable()
37   if currTable.isEmpty() then
38     | return Backtrack
39   filterDomains()
```

6.2 Methods

The main method in `ConstraintCT` is `enforceGAC()`. After the initialization of the sets \mathbf{S}^{val} and \mathbf{S}^{sup} , CT updates `currTable` to filter out (indices of) tuples that are no more supports, and then considers each variable-value pair to check whether these values still have a support.

Updating the Current Table For each variable $x \in \mathbf{S}^{\text{val}}$, i.e., each variable x whose domain has changed since the last time the filtering algorithm was called, `updateTable()` performs some operations. This method assumes an access to the set of values Δ_x removed from $\text{dom}(x)$ since the last call to `enforceGAC()`. There are two ways of updating `currTable`, either incrementally or from scratch after resetting. Note that the idea of using resets has been proposed in [30] and successfully applied to GAC4 and MDD4, with the practical interest of saving computational effort in some precise contexts. This is the strategy implemented in `updateTable()`, by considering a reset-based computation when the size of the domain is smaller than the number of deleted values.

In case of an incremental update (line 10), the union of the tuples to be removed is collected by calling `addToMask()` for each bit-set (of supports) corresponding to removed values, whereas in case of a reset-based update (line 14), we perform the union of the tuples to be kept. To get a mask ready to apply, we just need to reverse it when it has been built from removed values. Finally, the (indexes of) tuples of `currTable` not contained in the mask, built from x , are directly removed by means of `intersectWithMask()`. When there are no more tuples in the current table, a failure is detected, and `updateTable()` is stopped by means of a loop break.

Filtering of Domains Values are removed from the domain of some variables during the search of a solution, which can lead to inconsistent values in the domain of other variables. As `currTable` is a reversible and dynamically maintained structure, the value of some bits changes from 1 to 0 when tuples become invalid (or from 0 to 1 when the search backtracks). On the contrary, the `supports` bit-sets are only computed at the creation of the constraint and are not maintained during search. It follows from the definition of those bit-sets that (x, a) has a valid support if and only if

$$(\text{currTable} \cap \text{supports}[x, a]) \neq \emptyset \quad (1)$$

Therefore, each time a tuple becomes invalid, the constraint must check this condition for every variable value pair (x, a) such that $a \in \text{dom}(x)$, and remove a from $\text{dom}(x)$ if the condition is not satisfied any more. This operation is efficiently implemented in `filterDomains()` with the help of residues and the method `intersectIndex()`.

Example 4. The same set of tuples as in Example 3 is considered. Suppose now that a was removed from $\text{dom}(x)$ (by another constraint) after the initialization. Given that the domain of x is reduced, when `updateTable()` is called by

enforceGAC(), all tuples supporting a (because $\Delta_x = \{a\}$) will be invalidated. Figure 3a illustrates the intermediary bit-sets used to compute the new value currTable^{out} from currTable^{in} and $\text{supports}[x, a]$. Then filterDomains() computes for each variable-value pair (x_i, a_i) (with $x_i \in \mathbf{S}^{sup}$ and $a_i \in \text{dom}(x)$) the intersection of its associated set of supports with currTable as shown in Figure 3b. Given that the intersection for $\text{supports}[z, c]$ and currTable is empty, c is removed from $\text{dom}(z)$.




Fig. 3: Illustration of enforceGAC() after the removal of a from $\text{dom}(x)$.

6.3 Improvements

The algorithm in Section 6.2 can be improved to avoid unnecessary computations in some cases.

Filtering out bounded variables. The initialization of \mathbf{S}^{val} at line 32 can be only performed from unbound variables (and the last assigned variable), instead of the whole scope. We can maintain them in a reversible sparse set.

Last modified variable. It is not necessary to attempt to filter values out from the domain of a variable x if this was the only modified variable since the last call to enforceGAC(). Indeed, when updateTable() is executed, the new state of currTable will be computed from $\text{dom}(x)$ or Δ_x only. Because every value of x had a support in currTable the last time the propagator was called, we can omit filtering $\text{dom}(x)$ by initially removing x from \mathbf{S}^{sup} .

7 Experiments

We experimented CT on 1,621 CSP instances involving (positive) table constraints (15GB of uncompressed files in format XCSP 2.1). This corresponds to a large variety of instances, taken from 37 series. For guiding search, we used binary branching with *domain over degree* as variable ordering heuristic and




Fig. 4: Performance Profile

min value as value ordering heuristic. A timeout of 1,000 seconds was used for each instance. The tested GAC algorithms are CT, STR2 [17], STR3 [20], GAC4 [28,30], GAC4R [30], MDDR [30] and AC5TCRecomp [26]. All scripts, codes and benchmarks allowing to reproduce our experiments are available at <https://bitbucket.org/pschaus/xp-table>. The experiments were run on a 32-core machine (1400MHz cpu) with 100GB using Java(TM) SE Runtime Environment (build 1.8.0_60-b27) with 10GB of memory allocated (-Xmx option).

Performance Profiles. Let $t_{i,s}$ represent the time obtained with filtering algorithm $s \in S$ on instance $i \in I$. The performance ratio is defined as follows: $r_{i,s} = \frac{t_{i,s}}{\min\{t_{i,s} | s \in S\}}$. A ratio $r_{i,s} = 1$ means that s was the fastest on instance i . The performance profile [8] is a cumulative distribution function of the performance of s (speedup) compared to other algorithms: $\rho_s(\tau) = \frac{1}{|I|} \times |\{i \in I | r_{i,s} \leq \tau\}|$.

Our results are visually aggregated to form a performance profile in Figure 4 generated by means of the online tool [5] <http://sites.uclouvain.be/performance-profile>. Note that we filtered out the instances that i) could not be solved within 1,000 seconds by all algorithms ii) were solved in less than 2 seconds by the slowest algorithm, and iii) required less than 500 backtracks. The final set of instances used to build the profile is composed of 227 instances. An interactive performance profile is also available at <https://www.info.ucl.ac.be/~pschaus/assets/publi/performance-profile-ct> to let the interested reader deactivate some family of instances to analyze the results more closely.

Table 1 reports the speedup statistics of CT over the other algorithms. A first observation is that CT is the fastest algorithm on 94.47% of the instances. Among all tested algorithms, AC5TCRecomp obtains the worse results. Then it is not clear which one among STR2, STR3, GAC4 and GAC4R is the second

Speedup	STR2	STR3	GAC4	GAC4R	MDD4R	AC5-TC	Best2
geometric mean	5.09	4.03	7.05	6.15	6.57	19.22	2.75
min	0.76	1.09	0.92	1.13	0.13	1.05	0.13
max	88.58	51.04	173.24	208.52	50.84	1850.82	15.99
st.dev.	10.64	4.36	19.67	18.57	9.46	134.13	2.87

Table 1: Speedup analysis of CT over the other algorithms. Column 'Best2' corresponds to a virtual second best solver (minimum time of all algorithms except CT).




Fig. 5: Performance Profiles with dynamic (CT), recomputation (CTR) and incremental (CTI) strategies.

best algorithm. Based on the geometric mean speedup, STR3 seems to be the second best algorithm followed by STR2, GAC4R and MDD4R. Importantly, one can observe that the geometric mean speedup of CT over the best of the other algorithms is about 2.75.

Impact of Resetting Operations. In Algorithm 2, the choice of being incremental or not, when updating `currTable`, depends on the size of several sets and is thus dynamic. We propose to analyze two variants of Algorithm 2 when this choice is static:

- Full incremental (CTI): only the body of the 'if' at line 10 is executed (deltas are systematically used).
- Full re-computation (CTR) : only the body of the 'else' at line 14 is executed (domains are systematically used).

The performance profiles with these two variants are given in Figure 5, and the speedup table of the static versions over the dynamic one is given in Table 2.

Speedup	CTI	CTR	Best
geometric mean	1.04	1.34	0.99
min	0.44	0.53	0.44
max	3.23	4.39	1.96
st.dev.	0.38	0.65	0.27

Table 2: Speedup analysis of the two static variants over CT.

As can be seen from both the performance profiles and the speedup table, the dynamic version using the resetting operations dominates the static ones. The geometric mean speedup is around 4% over CTI and 34% over CTR.

Contradiction with Previous Results. In [26], AC5TCRecomp was presented as being competitive with STR2. When we analyzed the code⁴ of STR2 used in [26], it appeared that STR2 was implemented in Comet using built-in sets (triggering the garbage collection of Comet). We thus believe that the results and conclusions in [26] may over-penalize the performance of STR2. Our results also somehow contradict the results in [30] where STR3 and STR2 were dominated by MDD4R and GAC4R. When analyzing the performance of the implementation of STR2 and STR3 used in [30] with or-tools, it appears that it is not as competitive as that in AbsCon (sometimes slower by a factor of 3). The results presented in [30] may thus also over-penalize STR2 and STR3.

One additional contribution of this work is a fined-tuned implementation of the best filtering algorithms for table constraints. The implementation of these algorithms in OscaR was optimized, and checked to be close in performance to the ones by the original authors. For CT, STR2 and STR3, a comparison was made with AbsCon, and for CT, MDD4R and GAC4R, a comparison was made with or-tools. Our implementation required a development effort of 10 man-months in order to obtain an efficient implementation of each algorithm. It involved the expertise of several OscaR developers and a deep analysis of the existing implementations in AbsCon and or-tools. The implementation of all algorithms used in this paper is open-source and part of OscaR release 3.1.0.

8 Conclusion

In this paper, we have shown that Compact-Table (CT) is a robust algorithm that clearly dominates state-of-the-art propagators for table constraints. CT benefits from well-tried techniques: bitwise operations, residual supports, tabular reduction and resetting operations. We believe that CT can be easily implemented using the reversible sparse bit-set data structure.

⁴ available at <http://becool.info.ucl.ac.be>

References

1. C. Bessiere and J.-C. Régin. Arc consistency for general constraint networks: preliminary results. In *Proceedings of IJCAI'97*, pages 398–404, 1997.
2. C. Bessiere, J.-C. Régin, R. Yap, and Y. Zhang. An optimal coarse-grained arc consistency algorithm. *Artificial Intelligence*, 165(2):165–185, 2005.
3. C. Bliet. Wordwise algorithms and improved heuristics for solving hard constraint satisfaction problems. Technical Report 12-96-R045, ERCIM, 1996.
4. P. Briggs and L. Torczon. An efficient representation for sparse sets. *ACM Letters on Programming Languages and Systems*, 2(1-4):59–69, 1993.
5. Sascha Van Cauwelaert, Michele Lombardi, and Pierre Schaus. A visual web tool to perform what-if analysis of optimization approaches. Technical report, UCLouvain, 2016.
6. K. Cheng and R. Yap. An MDD-based generalized arc consistency algorithm for positive and negative table constraints and some global constraints. *Constraints*, 15(2):265–304, 2010.
7. J. Demeulenaere. Efficient algorithms for table constraints. Technical report, Master Thesis, under the supervision of P. Schaus, UCLouvain, 2015.
8. E.D. Dolan and J.J. Moré. Benchmarking optimization software with performance profiles. *Mathematical programming*, 91(2):201–213, 2002.
9. R.W. Floyd. Nondeterministic algorithms. *Journal of the ACM*, 14(4):636–644, 1967.
10. I.P. Gent, C. Jefferson, I. Miguel, and P. Nightingale. Data structures for generalised arc consistency for extensional constraints. In *Proceedings of AAAI'07*, pages 191–197, 2007.
11. N. Gharbi, F. Hemery, C. Lecoutre, and O. Roussel. Sliced table constraints: Combining compression and tabular reduction. In *Proceedings of CPAIOR'14*, pages 120–135, 2014.
12. P. Van Hentenryck J.-B. Mairiy and Y. Deville. Optimal and efficient filtering algorithms for table constraints. *Constraints*, 19(1):77–120, 2014.
13. C. Jefferson and P. Nightingale. Extending simple tabular reduction with short supports. In *Proceedings of IJCAI'13*, pages 573–579, 2013.
14. G. Katsirelos and T. Walsh. A compression algorithm for large arity extensional constraints. In *Proceedings of CP'07*, pages 379–393, 2007.
15. D.E. Knuth. *The Art of Computer: Combinatorial Algorithms*, volume 4. Addison-Wesley, 2015.
16. V. le Clément de Saint-Marcq, P. Schaus, C. Solnon, and C. Lecoutre. Sparse-sets for domain implementation. In *Proceeding of TRICS'13*, pages 1–10, 2013.
17. C. Lecoutre. STR2: Optimized simple tabular reduction for table constraints. *Constraints*, 16(4):341–371, 2011.
18. C. Lecoutre, F. Boussemart, and F. Hemery. Exploiting multidirectionality in coarse-grained arc consistency algorithms. In *Proceedings of CP'03*, pages 480–494, 2003.
19. C. Lecoutre and F. Hemery. A study of residual supports in arc consistency. In *Proceedings of IJCAI'07*, pages 125–130, 2007.
20. C. Lecoutre, C. Likitvivatanavong, and R. Yap. STR3: A path-optimal filtering algorithm for table constraints. *Artificial Intelligence*, 220:1–27, 2015.
21. C. Lecoutre and R. Szymanek. Generalized arc consistency for positive table constraints. In *Proceedings of CP'06*, pages 284–298, 2006.

22. C. Lecoutre and J. Vion. Enforcing arc consistency using bitwise operations. *Constraint Programming Letters*, 2:21–35, 2008.
23. O. Lhomme and J.-C. Régin. A fast arc consistency algorithm for n-ary constraints. In *Proceedings of AAAI'05*, pages 405–410, 2005.
24. C. Likitvivanavong, Y. Zhang, J. Bowen, and E.C. Freuder. Arc consistency in MAC: a new perspective. In *Proceedings of CPAI'04 workshop held with CP'04*, pages 93–107, 2004.
25. A.K. Mackworth. Consistency in networks of relations. *Artificial Intelligence*, 8(1):99–118, 1977.
26. J.-B. Mairy, P. van Hentenryck, and Y. Deville. An optimal filtering algorithm for table constraints. In *Proceedings of CP'12*, pages 496–511, 2012.
27. J.J. McGregor. Relational consistency algorithms and their application in finding subgraph and graph isomorphisms. *Information Sciences*, 19:229–250, 1979.
28. R. Mohr and G. Masini. Good old discrete relaxation. In *Proceedings of ECAI'88*, pages 651–656, 1988.
29. Oskar Team. Oskar: Scala in OR, 2012. Available from <https://bitbucket.org/oscarlib/oscar>.
30. G. Perez and J.-C. Régin. Improving GAC-4 for Table and MDD constraints. In *Proceedings of CP'14*, pages 606–621, 2014.
31. A. Srinivasan, T. Kam, S. Malik, and R.K. Brayton. Algorithms for discrete function manipulation. In *Proceedings of ICCAD'90*, pages 92–95, 1990.
32. J.R. Ullmann. An algorithm for subgraph isomorphism. *Journal of the ACM*, 23(1):31–42, 1976.
33. J.R. Ullmann. Partition search for non-binary constraint satisfaction. *Information Science*, 177:3639–3678, 2007.
34. R. Wang, W. Xia, R. Yap, and Z. Li. Optimizing simple table reduction with bitwise representation. In *Proceedings of IJCAI'16*, 2016.
35. W. Xia and R. Yap. Optimizing STR algorithms with tuple compression. In *Proceedings of CP'13*, pages 724–732, 2013.