

Arc consistency for general constraint networks: preliminary results

Christian Bessière
LIRMM-CNRS (UMR 5506)
161 rue Ada
34392 Montpellier cedex 5, France
Email: bessiere@lirmm.fr

Jean-Charles Régin
ILOG S.A.
9 rue de Verdun
94253 Gentilly Cedex, France
Email: regin@ilog.fr

Abstract

Constraint networks are used more and more to solve combinatorial problems in real-life applications. Much activity is concentrated on improving the efficiency of finding a solution in a constraint network (the constraint satisfaction problem, CSP). Particularly, arc consistency caught many researchers' attention, involving the discovery of a large number of algorithms. And, for the last two years, it has been shown that maintaining arc consistency during search is a worthwhile approach. However, results on CSPs and on arc consistency are almost always limited to binary constraint networks. The CSP is no longer an academic problem, and it is time to deal with non-binary CSPs, as widely required in real world constraint solvers. This paper proposes a general schema to implement arc consistency on constraints of any arity when no specific algorithm is known. A first instantiation of the schema is presented here, which deals with constraints given by a predicate, by the set of forbidden combinations of values, or by the set of allowed ones.

1 Introduction

1.1 Overview

Constraint satisfaction problems (CSPs) occur widely in artificial intelligence. They are used more and more in real-life applications, such as resource allocation, crew scheduling, time tabling, frequency allocation, car sequencing, etc. The CSP involves finding a solution in a constraint network, i.e. finding values for problem variables subject to constraints.

The general task of solving a CSP being NP-hard, many researchers have concentrated on improving the efficiency of finding a solution in a constraint network. Particularly, arc consistency caught many researchers' attention, involving the discovery of a large number of

algorithms [Mackworth, 1977a], [Mohr and Henderson, 1986], [Bessière, 1994], [Bessière *et al.*, 1995]. And, recently, the value of these studies on arc consistency increased since it has been shown that maintaining arc consistency during search is definitely a worthwhile approach when solving large and hard problems [Sabin and Freuder, 1994], [Bessière *et al.*, 1995], [Bessière and Régin, 1996], [Grant and Smith, 1996].

However, results about CSP solving and about arc consistency are almost always limited to *binary* constraint networks, justified by the fact that any non-binary constraint network can be translated into an equivalent binary one with additional variables [Rossi *et al.*, 1990], [Dechter, 1990]. But, in practical cases, it is often inconceivable to translate a non-binary constraint into an equivalent set of binary ones with the technique cited above because of the underlying computational and memory cost. When a constraint is *representable* [Montanari, 1974]), it is sometimes possible to replace it by a set of binary constraints without introducing new variables and without generating explicitly the set of allowed tuples of each binary constraint. For example, a constraint enforcing that the three variables x_1 , x_2 , and x_3 must all take different values can equivalently be replaced by the three binary constraints $x_1 \neq x_2$, $x_1 \neq x_3$, and $x_2 \neq x_3$. But, as it has been shown in [Régin, 1994], non-binary constraints lose a part of their semantics when encoded into a set of binary ones in this way. This leads, for example, to less pruning for arc consistency algorithms handling them.

Hence, if we no longer want to consider the CSP as an academic problem we must be able to deal with any kind of constraints, and so, with non-binary constraints, as widely required in real-world constraint solvers.

1.2 Previous work

In the constraint satisfaction community, the number of works involving non-binary constraint networks is rather small. For the particular case of arc consistency, we know two algorithms capable of achieving it. Mackworth [Mackworth, 1977b] proposed the algorithm CN, which

is a kind of generalization of AC-3 to non-binary constraints. As AC-3, that algorithm has a bad worst-case time complexity ($O(er^2d^{r+1})$, with e the number of constraints in the network, r the maximal arity of the constraints, and d the size of the largest domain). Mohr and Massini [Mohr and Masini, 1988] proposed GAC4. It is based on the same idea as AC-4 (computing the number of supports for each value in each domain and removing those with this number equal to zero). Thus, it gets rid of the huge worst-case time complexity of CN (GAC4 is in $O(ed^r)$), but has the same drawbacks as AC-4: space complexity (because of the lists of supported values), and average time complexity. The respective drawbacks of CN and GAC4 are even more important than on their binary versions, so that CN can only be applied on ternary constraints and very small domains, and GAC4 on very tight constraints, where the number of allowed tuples of values is very small.

Finally, we can point out the work of van Beek and Dechter [van Beek and Dechter, 1995], who proposed another definition for arc consistency on non-binary constraint networks, namely *relational arc consistency*. This definition is much stronger than the classical one since for each constraint it requires global consistency on the underlying subnetwork (i.e. the network involving the variables of the given constraint, and all the “smaller” constraints defined on some of these variables).

1.3 Our purpose

Our general aim is to propose a new schema in order to perform arc consistency on any real world constraint network with a reasonably low time and space complexity. In industrial applications, a constraint can be given in many different forms: by the set of allowed tuples in extension (generally when the constraint is very tight), by the set of forbidden tuples in extension (when the constraint is very loose), by a conjunctive constraint, by an arithmetic relation, or by any predicate for which no particular semantics is known (data base query, user’s context-dependent constraint, etc.). Thus, a suitable arc consistency algorithm should be able to efficiently handle any of these constraints.

The particular aim of this paper is to present the first steps of our work. First, we will present our general schema for arc consistency on non-binary constraint networks. This schema is based on the AC-7 schema given in [Bessi re *et al.*, 1995]. It makes use of the “current support” idea, and of “multidirectionality” (the generalization of bidirectionality to non-binary constraints) in order to save as many constraint checks as possible. Second, we will instantiate this schema with some of the most frequently occurring forms of constraint representations.

Constraints defined by a predicate for which no partic-

ular semantics is known are one of the most important forms that have to be handled. Indeed, if we except the predicates for which specific algorithms are known (arithmetic relations [Van Hentenryck *et al.*, 1992], cardinality constraints [R egin, 1996], etc.), the only available algorithms for non-binary constraints are CN and GAC4. GAC4 deals only with constraints given in extension, and then is not capable of handling a predicate, even if, theoretically, it is always possible to build the whole set of allowed tuples by checking all the combinations of values for the variables involved in the constraint. It would be impracticable, at least because of space requirements. Moreover, asking the value of a predicate for a particular combination of values (a constraint check) can be very costly in practice, when the answer is contained in a database, or requires heavy computation. This definitively eliminates CN, with its huge time complexity, doing and doing again many times the same constraint checks. Therefore, we will propose an instantiation of the schema that will efficiently handle predicates, minimizing the number of constraint checks. However, this schema is certainly not competitive on predicates for which a specific algorithm is available.

Afterwards, we will instantiate the schema for constraints defined by a set of forbidden tuples. This kind of representation is of practical interest when constraints are very loose.

Finally, we will show that the schema can also deal with constraints defined by the list of allowed tuples given in extension. It leads to an improvement on GAC4, which was already written for that kind of constraint.

2 Preliminaries

Constraint network. A finite *constraint network* \mathcal{P} is defined as a set of n variables $X = \{x_1, \dots, x_n\}$, a set of current *domains* $\mathcal{D} = \{D(x_1), \dots, D(x_n)\}$ where $D(x_i)$ is the finite set of possible *values* for variable x_i , and a set \mathcal{C} of *constraints* between variables. We introduce the particular notation $\mathcal{D}_0 = \{D_0(x_1), \dots, D_0(x_n)\}$ to represent the set of initial domains of \mathcal{P} . Indeed, we consider that any constraint network \mathcal{P} can be associated with an initial domain \mathcal{D}_0 (containing \mathcal{D}), on which constraint definitions were stated. A total ordering $<_d$ can be defined on $D(x_i), \forall x_i \in X$, without loss of generality. **Constraints.** A constraint C on the ordered set of variables $X(C) = (x_{i_1}, \dots, x_{i_r})$ is a subset of the Cartesian product $D_0(x_{i_1}) \times \dots \times D_0(x_{i_r})$ that specifies the *allowed* combinations of values for the variables $x_{i_1} \times \dots \times x_{i_r}$. An element of $D_0(x_{i_1}) \times \dots \times D_0(x_{i_r})$ is called a *tuple on* $X(C)$. Two tuples τ and τ' on $X(C)$ can be ordered by the natural lexicographic order $<_{lo}$ in which $\tau <_{lo} \tau'$ iff $\exists k/\tau[1..k-1] = \tau'[1..k-1]$ and $\tau[k] <_d \tau'[k]$ ($\tau[1..k]$ being the prefix of size k of τ , and $\tau[k]$ the k^{th} value of τ). The tuples of

$D_0(x_{i_1}) \times \dots \times D_0(x_{i_r})$ not allowed by C are called the *forbidden* tuples of C . Verifying whether a given tuple τ is allowed by C or not is called a *constraint check*. $|X(C)|$ is the *arity* of C .

A constraint C involving the subset of variables $X(C) = (x_{i_1}, \dots, x_{i_r})$ can be defined by the set of allowed tuples (resp. the set of forbidden tuples) given in extension when the constraint is tight (resp. is loose), or by an arithmetic relation. More generally, it can be represented by any Boolean function f_C defined on $D_0(x_{i_1}) \times \dots \times D_0(x_{i_r})$.

Solutions. A *solution* of a constraint network is an instantiation of the variables such that all the constraints are satisfied.

Notation. A value a for a variable x is often denoted by (x, a) . $\text{var}(C, i)$ represents the i^{th} variable of $X(C)$, while $\text{index}(C, x)$ is the position of variable x in $X(C)$.

Arc consistency. Let $\mathcal{P} = (X, \mathcal{D}, \mathcal{C})$ be a constraint network, C a constraint in \mathcal{C} .

A tuple τ of $X(C)$ is *valid* if $\forall (x, a) \in \tau, a \in D(x)$; otherwise, it is *rejected*.

A value $a \in D(x)$ is *consistent with* C iff $x \notin X(C)$, or $\exists \tau$ allowed by C , such that $a = \tau[\text{index}(C, x)]$ and τ is valid. (τ is then called a *support* for (x, a) on C .)

C is *arc consistent* iff $\forall x_i \in X(C), D(x_i) \neq \emptyset$ and $\forall a \in D(x_i), a$ is consistent with C .

\mathcal{P} is *arc consistent* iff all the constraints of \mathcal{C} are arc consistent.

We achieve arc consistency in \mathcal{P} by removing every value not consistent with at least one constraint in \mathcal{C} .

3 A general schema for arc consistency

As we pointed out in Section 1, our general aim is, on the one hand, to provide a schema sufficiently general to be instantiated with any kind of constraint. On the other hand, this schema must be powerful enough to avoid as many constraint checks as possible when achieving arc consistency.

We did not choose an AC-3 like or an AC-4 like schema for efficiency reasons. To be reasonably efficient, the schema has to be based on AC-6 or AC-7. They are both based on the search of a single support for each value, the worst-case time complexity being optimal in both algorithms. The difference between them is whether or not they deal with bidirectionality [Bessièrè *et al.*, 1995].

In a non-binary constraint network, any constraint is multidirectional, as any constraint is bidirectional in a binary constraint network. *Multidirectionality*, indeed, is the fact that for any constraint C , a tuple τ on $X(C)$ is a support for the value $\tau[\text{index}(C, x)]$ (x being a variable involved in C) iff $\forall y \in X(C)$, τ is a support for $\tau[\text{index}(C, y)]$. We say that an algorithm “deals with” multidirectionality iff it never checks whether a tuple is

Algorithm 1: function propagation

```

propagation (in  $C$ : constraint; in  $x$ : variable;
             in  $a$ : value;
             in out deletionStream: list): boolean
    for each  $\tau \in S_C(x, a)$  do
1   for each  $(z, c) \in \tau$  do remove  $\tau$  from  $S_C(z, c)$ ;
2   for each  $(y, b) \in S(\tau)$  do
       remove  $(y, b)$  from  $S(\tau)$ ;
       if  $b \in D(y)$  then
3          $\sigma \leftarrow \text{seekInferableSupport}(C, y, b)$ ;
4         if  $\sigma \neq \text{nil}$  then
           add  $(y, b)$  in  $S(\sigma)$ ;
           else
5              $\sigma \leftarrow \text{seekNextSupport}(C, y, b, \text{last}_C(y, b))$ ;
6             if  $\sigma \neq \text{nil}$  then
               add  $(y, b)$  in  $S(\sigma)$ ;
                $\text{last}_C(y, b) \leftarrow \sigma$ ;
7             for  $k$  from 1 to  $|X(C)|$  do
                 add  $\sigma$  in  $S_C(\text{var}(C, k), \sigma[k])$ ;
           else
8             remove  $b$  from  $D(y)$ ;
           if  $D(y) = \emptyset$  then return false;
9             add  $(y, b)$  in deletionStream;

return true;

```

a support for a value when it has already been checked for another value, and never looks for a support for a value on a constraint C when a tuple supporting this value has already been checked.

And, dealing with multidirectionality appears to be promising on non-binary constraints. A tuple τ , indeed, allowed by a constraint C , can support $|X(C)|$ values. Thus the possible savings are much more important than on binary constraints. AC-7 could save d^2 constraint checks thanks to bidirectionality, where d is the domain size of the variables involved (see [Bessièrè *et al.*, 1995]), while d^r constraint checks can be saved thanks to multidirectionality on an r -ary constraint.

The framework we propose is then an AC-7 like schema in which search for support (function `seekNextSupport`) is instantiated differently depending on the type of the constraint involved. This schema (named `GAC-schema`) is able to handle any type of constraint, as soon as the corresponding function `seekNextSupport` is available. Before giving the `seekNextSupport` function associated with the types of constraints we will study in this paper (in the next section), let us describe the function `propagation` (see Alg. 1). The program including it must create and initialize the data structures (S_C , S , and last_C), and call `propagation(C, x, a, deletionStream)` for each constraint C involving x , each time a value (x, a) is removed from $D(x)$, in order to propagate the consequences of this deletion.

S_C , S , and $last_C$ must be initialized in a way such that:

- $S_C(x, a)$ contains all the allowed tuples τ that are the current support for some value, and such that $\tau[index(C, x)] = a$.
- $S(\tau)$ contains all values for which τ is the current support.
- $last_C(y, b)$ is the last tuple returned by `seekNextSupport` as a support for (y, b) if `seekNextSupport(C, y, b, -)` has already been called; `nil` otherwise. $last_C(y, b)$ is perhaps no longer a support in C for (y, b) (if it is no longer valid and `seekInferableSupport` has found a new support for (y, b)). It is just here to give the point where `seekNextSupport` will have to restart the search for a support for (y, b) on C at the next call. There is an underlying ordering on the tuples, which is proper to `seekNextSupport`.

Property 1 $\forall C \in \mathcal{C}, \forall \text{allowed tuple } \tau : S(\tau) \neq \emptyset \Leftrightarrow \forall x \in X(C), \tau \in S_C(x, \tau[index(C, x)])$.

Property 2 *Given any constraint C of arity r , the space complexity of the data structure of GAC-schema for C is $O(r^2d)$.*

Proof. Each value has at most one support, then $\sum |S(\tau)| \leq rd$ and there are at most rd tuples in memory. Each value is associated with one $last_C$, so the $last_C$ data structure requires rd pointers. A tuple contains r elements, thus the set of all tuples that support at least one value can be represented in $O(r^2d)$. The number of elements that belong to S_C is bounded above by r^2d by property 1 and because $\sum |S(\tau)| \leq rd$. Since an element of an S_C list corresponds to a pointer to a tuple, the space complexity of the data structures of GAC-schema for C is $O(r^2d)$. ■

Each time a value (x, a) is removed from $D(x)$, we must propagate this deletion to each constraint C involving x . So, for all the values (y, b) that were supported by a tuple containing a in position $index(C, x)$, we must find another support (line 2). To take multidirectionality into account, we first check (lines 3 and 4) if there exists a valid tuple containing (y, b) that is already the current support for another value (function `seekInferableSupport`, Alg. 2). If not, the function `seekNextSupport` will look for another support for (y, b) , starting the search from $last_C(y, b)$ (line 5). If a new support σ is found for (y, b) on C , (y, b) is recorded as being currently supported by σ (line 6), and for each value contained in σ , we store the fact that σ is the current support for some value (line 7); otherwise, (y, b) has to be removed (lines 8 to 9).

4 The main types of constraint representations

In this section, we will present the instantiation of the schema in order to deal with three types of constraints.

Algorithm 2: function `seekInferableSupport`

```

seekInferableSupport (in  $C$ : constraint; in  $y$ : variable;
                      in  $b$ : value): tuple
while  $S_C(y, b) \neq \emptyset$  do
   $\sigma \leftarrow first(S_C(y, b))$ ;
  if  $\exists k/\sigma[k] \notin D(var(C, k))$  then remove  $\sigma$  from  $S_C(y, b)$ ;
  else return  $\sigma$  /*  $\sigma$  is a support */;
return nil ;

```

Constraints given by a predicate for which we do not know specific algorithms is the most important type. Indeed, there does not exist any tool to process a predicate in reasonable time and space. Afterwards, we will present a type of constraint which can be viewed as a special case of predicate constraints: constraints given by the set of forbidden tuples. Finally, we will give the `seekNextSupport` function capable of processing constraints given by the set of allowed tuples.

4.1 Predicates

When a constraint C is defined by a predicate for which no particular semantics is known, it is necessary to define an ordering on the tuples that will be followed when looking for a support for a value (y, b) on C . By means of the function `nextTuple` (Alg. 3), the lexicographic order \prec_{lo} defined in Section 2 is used to examine the valid tuples on $X(C)$. Given any tuple τ on $X(C)$, and any index $k \leq |X(C)|$, `nextTuple(C, y, b, τ , k)` will return the smallest valid tuple σ greater than τ w.r.t. \prec_{lo} , such that the value of y in σ is b , and such that the prefixes of size k of σ and τ are not equal. The second parameter returned by `nextTuple` is the smallest index k' on which σ and τ have different values.

When the function `seekNextSupport` (see Alg. 4) is called to search for a new support for (y, b) on C , we know from `seekInferableSupport` that there does not exist any tuple already checked that supports (y, b) on C . However, we have to take care to avoid checking a tuple τ which has already been unsuccessfully checked for another value of another variable of $X(C)$. We name *candidate* a valid tuple which has never been checked and thus, which could be a support for (y, b) , as op-

Algorithm 3: function `nextTuple`

```

nextTuple (in  $C$ : constraint; in  $y$ : variable; in  $b$ : value;
           in  $\tau$ : tuple; in  $k$ : index): (tuple, index)
/* returns  $(\sigma, k')$  where:
  •  $\sigma$  is the smallest valid tuple such that  $\sigma[index(C, y)] = b$ ,
     $\tau \prec_{lo} \sigma$  and  $\sigma[1..k] \neq \tau[1..k]$ 
  •  $k'$  is such that  $\sigma[1..k' - 1] = \tau[1..k' - 1]$  and  $\sigma[k'] \neq \tau[k']$ 
otherwise returns  $(nil, -)$  */

```

Note that k' will always be less than or equal to k .

Algorithm 4: function `seekNextSupport`

```

seekNextSupport (in  $C$ : constraint; in  $y$ : variable;
                 in  $b$ : value; in  $\tau$ : tuple): tuple
1 if  $\tau \neq \text{nil}$  then
  |  $(\sigma, \text{dummy}) \leftarrow \text{nextTuple}(C, y, b, \tau, |X(C)|)$ ;
  else
    | for  $i$  from 1 to  $|X(C)|$  do
      | |  $\sigma[i] \leftarrow \text{first}(\text{var}(C, i))$ ;
2 |  $\sigma[\text{index}(C, y)] \leftarrow b$ ;
3  $\sigma \leftarrow \text{seekCandidateTuple}(C, y, b, \sigma, 1)$ ;
   $\text{found} \leftarrow \text{false}$ ;
4 while  $(\sigma \neq \text{nil})$  and  $(\text{not found})$  do
5 | if  $f_C(\sigma)$  then  $\text{found} \leftarrow \text{true}$ ;
  | else
    | |  $(\sigma, k) \leftarrow \text{nextTuple}(C, y, b, \sigma, |X(C)|)$ ;
    | |  $\sigma \leftarrow \text{seekCandidateTuple}(C, y, b, \sigma, k)$ ;
6 | return  $\sigma$ ;

```

posed to the valid tuples that have already been checked not to be allowed by C (when we look for support for another value). Dealing with multidirectionality implies only checking the candidates. This can be done by means of the function `seekCandidateTuple` (see Alg. 5).

`seekCandidateTuple`(C, y, b, σ, k) returns the smallest (w.r.t. \prec_{lo}) candidate greater than or equal to σ , where σ is valid and $\sigma[1..k-1]$ has been verified to be a possible prefix for a candidate. For each index from k to $|X(C)|$, `seekCandidateTuple` verifies whether σ is greater than $\text{last}_C(\text{var}(C, k), \sigma[k])$, (denoted by λ) (lines 2 to 4). If σ is smaller than λ , the search for a candidate has to jump forward: either to the smallest valid tuple following σ with a prefix different from $\sigma[1..k]$ (if σ and λ were diverging before k) (line 6), or directly to the valid tuple following λ (if σ and λ were diverging after k) (line 8). When we jump forward to the next valid tuple greater than σ or λ , some values before index k may have changed. Then, the value k goes back to the smallest index where the value of σ has changed (lines 7 and 9) to keep the property of line 1. When k reaches $|X(C)| + 1$, σ is a candidate and is returned.

The function `seekNextSupport`(C, y, b, τ) returns the smallest tuple greater than τ which is checked to be allowed by C . From line 1 to line 2, `seekNextSupport` assigns to σ the smallest valid tuple following τ (depending on whether τ is `nil` or not). Afterwards, σ is assigned to the smallest candidate (line 3). The search for a support for (y, b) on C is done in lines 4 to 6: we check σ and jump to the next candidate until $f_C(\sigma)$ returns `true` (line 5).

A sketch of proof.

We will simply show here that `seekCandidateTuple` cannot miss any candidates when jumping forward in lines 6 and 8.

line 6: suppose there is a candidate σ' between σ and the tuple returned by `nextTuple`(C, y, b, σ, k). This means

Algorithm 5: function `seekCandidateTuple`

```

seekCandidateTuple (in  $C$ : constraint; in  $y$ : variable;
                   in  $b$ : value;
                   in  $\sigma$ : tuple; in  $k$ : index): tuple
  while  $(\sigma \neq \text{nil})$  and  $(k \leq |X(C)|)$  do
1 | /*  $\sigma$  is candidate till index  $k-1$  */;
  | if  $\text{last}_C(\text{var}(C, k), \sigma[k]) \neq \text{nil}$  then
2 | |  $\lambda \leftarrow \text{last}_C(\text{var}(C, k), \sigma[k])$ ;
  | |  $\text{split} \leftarrow 1$ ;
  | | while  $\sigma[\text{split}] = \lambda[\text{split}]$  do
3 | | |  $\text{split} \leftarrow \text{split} + 1$ ;
  | | if  $\sigma[\text{split}] < \lambda[\text{split}]$  then
4 | | | if  $\text{split} < k$  then
5 | | | |  $(\sigma, k') \leftarrow \text{nextTuple}(C, y, b, \sigma, k)$ ;
6 | | | |  $k \leftarrow k' - 1$ ;
  | | | else
7 | | | |  $(\sigma, k') \leftarrow \text{nextTuple}(C, y, b, \lambda, |X(C)|)$ ;
8 | | | |  $k \leftarrow \min(k, k' - 1)$ ;
9 | |  $k \leftarrow k + 1$ ;
  return  $\sigma$ ;

```

that $\sigma'[1..k] = \sigma[1..k]$; otherwise σ' would be the tuple returned by `nextTuple`. So, σ' is smaller than λ because of lines 4 and 5. $\sigma'[k]$ being equal to $\sigma[k]$, σ' cannot be a candidate because $\lambda = \text{last}_C(\text{var}(C, k), \sigma[k])$ (see the definition of last_C).

line 8: suppose there is a candidate σ' between σ and λ . Then, $\sigma'[1..k]$ is equal to $\sigma[1..k]$ since $\sigma[1..k] = \lambda[1..k]$ (lines 3 to 5). Once again that is impossible because $\lambda = \text{last}_C(\text{var}(C, k), \sigma[k])$. Finally, σ' cannot be equal to λ since λ is no longer a valid tuple. (Otherwise `seekInferableSupport` would have found it to support (y, b) .) ■

4.2 Constraints given in extension**Negative constraints**

In this case, the constraint C is given in extension by the set of forbidden tuples, denoted by $\overline{T}(C)$. Even without loss of generality, we can assume that the constraint is very loose. (Otherwise, the space in memory needed prohibits us from using that representation.) Then $\overline{T}(C)$ contains few elements with regard to the set of tuples on $|X(C)|$. So, if we use the previous method with the predicate $f_C(\sigma)$ such that $f_C(\sigma) \Leftrightarrow \sigma \notin \overline{T}(C)$, then only a few constraint checks will be needed to find a new support or to prove there is none, because almost all valid tuples are candidates.

Such a predicate $f_C(\sigma)$ can be efficiently implemented by using a method like hashing [Sedgewick, 1990]. Hashing, indeed, permits us to find whether an element belongs to a set S , with an $O(|S|)$ space complexity, and an average time complexity close to a constant.

Positive constraints

In this case, the set of allowed tuples, denoted by $T(C)$, is explicitly given. Thus, for one constraint, the space

Algorithm 6: function seekNextSupport

```

seekNextSupport (in  $C$ : constraint; in  $y$ : variable;
                in  $b$ : value; in  $dummy$ ): tuple
while  $elt(y, b) \neq nil$  do
   $\sigma \leftarrow Tuple(elt(y, b))$ ;
  if  $isValid(\sigma)$  then return  $\sigma$ ;
   $elt(y, b) \leftarrow next(elt(y, b))$ ;
return nil ;

```

complexity will depend on the size of $T(C)$ similarly to GAC4. So, at first glance, it seems that the space complexity of GAC4 cannot be improved. But, in practice, we may have a problem in which some of the constraints appear many times involving different variables. For instance, in configuration problems, a configurator has to choose the most appropriate components among a catalog of predefined components and arrange them according to constraints that can be assembly rules, performance, etc. Most of the constraints involved with the same kind of components are repeated and are given by their common set of allowed tuples.

For a constraint C , the representation of the data structures of GAC4 corresponds to a particular representation of $T(C)$. So, if two constraints are given by $T(C)$, then GAC4 needs two distinct data structures, each with a $O(|T(C)|)$ space complexity. Hence, for p repeated constraints the space complexity of GAC4 will be $O(p \cdot |T(C)|)$. The algorithm we presented only requires an explicit representation of S and S_C lists, and $last_C$ pointers, for each constraint. This leads to a $O(r^2d)$ space complexity per constraint, r being the arity of the constraint involved (see Property 2). If the representation of $T(C)$ can be shared by p repeated constraints, then the global space complexity will be $O(|T(C)| + p \cdot r^2d)$, and a factor p will be gained with regard to GAC4.

Such an algorithm can be defined by representing the common set $T(C)$ as GAC4 does for one constraint and by using for each repeated constraint and each value (y, b) a pointer that indicates the last element reached in $T(C)$. We will denote it by $elt(y, b)$. The algorithm does no longer compute the tuples and test them; it only looks for a support directly in $T(C)$. The new algorithm is given by the function seekNextSupport (Alg. 6) which replaces the previous one. The function $next(elt(y, b))$ gives the next tuple in $T(C)$ containing b for variable y .

5 Space and time analysis

Let C be a constraint of arity r with a tightness t (proportion of forbidden tuples), and assume that all domains have a size d . Then, if a pointer is represented by 4 bytes, a good estimation of the memory required is:

- $16r \cdot |\{tuples\ on\ X(C)\ allowed\ by\ C\}| = 16rd^r(1 - t)$ Bytes for GAC4
- $24r^2d$ Bytes for the GAC-schema, according to property 2 (C being a predicate¹)

The table below gives some results about the space requirements of GAC4 and the GAC-schema. The sizes are given in Megabytes. Only the configurations associated with bold numbers can be used in practice. (“-” means that more than 256 Mb are needed.) Bear in mind that these space requirements are given for only **one** constraint.

t	0.001	0.02	0.2	0.5	0.8	0.98	0.999
$r = 8, d = 10$							
GAC-s	0.015 for any tightness						
GAC4	-	-	-	-	-	256	12.8
$r = 5, d = 20$							
GAC-s	0.012 for any tightness						
GAC4	256	251	205	128	51	5.1	0.26
$r = 4, d = 50$							
GAC-s	0.019 for any tightness						
GAC4	-	-	-	200	80	8	0.4

Finally, to be exhaustive on complexities, we report the worst-case time complexities of the different versions of the GAC-schema presented in this paper. Let C be a constraint of arity r . If C is a predicate, the GAC-schema has a time complexity bounded above by $O(d^r)$. This is a gain of r^2d with respect to the $O(r^2d^{r+1})$ of CN. If C is a negative constraint, the complexity remains very close to $O(d^r)$ (see subsection 4.2). If C is a positive constraint, the worst-case time complexity of the GAC-schema is $O(|T(C)|)$. This is the same as the one of GAC4.

6 A cryptogram as an example

In this section we briefly present a very small example (a cryptogram) to show some of the advantages of dealing with non-binary predicates. (This example is implemented in ILOG Solver 4.0, in which the GAC-schema has been inserted.) For a complete presentation of the example, see [ILOG, 1997].

The cryptogram is given in Fig. 1. The problem is to find a one to one mapping from letters $\{A, B, D, E, G, L, N, O, R, T\}$ to $\{0..9\}$ in such a way that the addition obtained by replacing each letter by its associated value is consistent. It is possible to solve this problem without the GAC-schema, by encoding it (for instance) as the following constraint network: a variable for each letter; a domain containing the numbers 0..9 for each variable; and two non-binary constraints for which efficient specific algorithms are known, “ $100,000D + 10,000O + 1,000N + 100A + 10L + D +$

¹For constraints given in extension we should add the size of the constraint representation (see subsection 4.2).

D O N A L D
+ G E R A L D

= R O B E R T

Figure 1: The cryptogram

$100,000G + 10,000E + 1,000R + 100A + 10L + D = 100,000R + 10,000O + 1,000B + 100E + 10R + T$, and “all-different($A, B, D, E, G, L, N, O, R, T$)”. However, with this representation, some of the knowledge we have on the problem is not stated, while it could have been helpful in solving the problem efficiently. We know, for example, that in the rightmost column $D + D$ is necessarily equal to T or $10 + T$ since there is no carry. We know that on the leftmost column R is equal to $D + G$ or $D + G + 1$ since there is no letter on the left of R . More generally, for the third column for example, we know that $(N + R = B) \vee (N + R = 10 + B) \vee (N + R + 1 = B) \vee (N + R + 1 = 10 + B)$. All these constraints (six in the cryptogram of Fig. 1) are predicates, but no already known specific algorithm can handle them (they are not arithmetic relations, but disjunctions of several arithmetic relations).

The table below presents the number of backtracks (#bt) and cpu-time performances of solving the cryptogram with ILOG Solver 4.0 (maintaining arc consistency during search) on a Pentium Pro 200, with or without the **GAC-schema** and the 6 additional constraints, and with or without the use of the minimal domain dynamic variable ordering (DVO).

	without DVO		with DVO	
	#bt	time	#bt	time
without GAC-s	4,612	0.72 sec.	138	0.05 sec.
with GAC-s	61	0.1 sec.	1	0.00 sec.

Of course it is a very small illustration that is given here with that sample problem. In this problem, indeed, the additional constraints are only added to improve the search. They are not necessary to encode the problem: the first two constraints were already ensuring that the solutions of the constraint network were the solutions of the cryptogram.

7 Conclusion

While arc consistency had been widely studied by the CSP community, there did not exist any algorithm that efficiently achieves arc consistency on non-binary constraints. Thus, we presented **GAC-schema**, which is built to take into account the last improvements available on binary constraints (AC-7 like schema). We saw that it is even more important than on binary constraints to use such improvements (e.g. multidirectionality). With specialized instantiations, this schema is capable of efficiently dealing with predicates, positive constraints, or negative constraints. The perspectives of this work are

to propose other instantiations of the schema in order to improve arc consistency processing on some other types of constraint representations frequently occurring in industrial problems.

References

- [Bessière and Régin, 1996] C. Bessière and J.C. Régin. MAC and combined heuristics: two reasons to forsake FC (and CBJ?) on hard problems. In *CP'96*, pages 61–75, Cambridge MA.
- [Bessière *et al.*, 1995] C. Bessière, E.C. Freuder, and J.C. Régin. Using inference to reduce arc consistency computation. In *IJCAI'95*, pages 592–598, Montréal, Canada.
- [Bessière, 1994] C. Bessière. Arc-consistency and arc-consistency again. *Artificial Intelligence*, 65:179–190, 1994.
- [Dechter, 1990] R. Dechter. On the expressiveness of networks with hidden variables. In *AAAI'90*, pages 556–562, Boston MA.
- [Grant and Smith, 1996] S.A. Grant and B.M. Smith. The phase transition behavior of maintaining arc consistency. In *ECAI'96*, pages 175–179, Budapest, Hungary.
- [ILOG, 1997] ILOG. *User's manual*. ILOG Solver, 4.0 edition, 1997.
- [Mackworth, 1977a] A.K. Mackworth. Consistency in networks of relations. *Artificial Intelligence*, 8:99–118, 1977.
- [Mackworth, 1977b] A.K. Mackworth. On reading sketch maps. In *IJCAI'77*, pages 598–606, Cambridge MA.
- [Mohr and Henderson, 1986] R. Mohr and T.C. Henderson. Arc and path consistency revisited. *Artificial Intelligence*, 28:225–233, 1986.
- [Mohr and Masini, 1988] R. Mohr and G. Masini. Good old discrete relaxation. In *ECAI'88*, pages 651–656, Munchen, FRG.
- [Montanari, 1974] U. Montanari. Networks of constraints: Fundamental properties and applications to picture processing. *Information Science*, 7:95–132, 1974.
- [Régin, 1994] J.C. Régin. A filtering algorithm for constraints of difference in CSPs. In *AAAI'94*, pages 362–367, Seattle WA.
- [Régin, 1996] J.C. Régin. Generalized arc consistency for global cardinality constraint. In *AAAI'96*, pages 209–215, Portland OR.
- [Rossi *et al.*, 1990] F. Rossi, C. Petrie, and V. Dhar. On the equivalence of constraint satisfaction problems. In *ECAI'90*, pages 550–556, Stockholm, Sweden.
- [Sabin and Freuder, 1994] D. Sabin and E.C. Freuder. Contradicting conventional wisdom in constraint satisfaction. In *PPCP'94*, Seattle WA.
- [Sedgewick, 1990] R. Sedgewick. *Algorithms in C*. Addison-Wesley Publishing Company, 1990.
- [van Beek and Dechter, 1995] P. van Beek and R. Dechter. On the minimality and global consistency of row-convex constraint networks. *Journal of the ACM*, 42(3):543–561, 1995.
- [Van Hentenryck *et al.*, 1992] P. Van Hentenryck, Y. Deville, and C.M. Teng. A generic arc-consistency algorithm and its specializations. *Artificial Intelligence*, 57:291–321, 1992.