Refining the Basic Constraint Propagation Algorithm

Christian Bessière* **— Jean-Charles Régin****

* *member of the* COCONUT *group*
LIRMM–CNRS (UMR 5506)
161 rue Ada
34392 Montpellier Cedex 5, France
Email: bessiere@lirmm.fr

** *ILOG*
1681, route des Dolines
06560 Valbonne, France
Email: regin@ilog.fr

ABSTRACT. *Constraint propagation is the main feature of any constraint solver. This is thus of prime importance to manage constraint propagation as efficiently as possible, justifying the use of the best algorithms. But the ease of integration is also one of the concerns when implementing an algorithm in a constraint solver. This paper focuses on AC-3, which is the simplest arc consistency algorithm known so far. We propose two refinements that preserve as much as possible the ease of integration into a solver (no heavy data structure to be maintained during search), while giving some noticeable improvements in efficiency. One of the proposed refinements is analytically compared to AC-6, showing interesting properties.*

RÉSUMÉ. *La propagation de contraintes est au coeur de tous les solveur avec contraintes. Il est donc important d'effectuer cette propagation de la façon la plus efficace possible, ce qui justifie l'utilisation des meilleurs algorithmes. Mais la simplicité d'intégration est aussi une des préoccupations quand on implémente un algorithme dans un solveur. Cet article s'intéresse à AC-3, qui est l'algorithme de cohérence d'arc le plus simple. Nous proposons deux améliorations de AC-3 qui préservent autant que possible sa facilité d'intégration (pas de structures de données complexes à maintenir pendant la recherche), tout en apportant un gain sensible en temps. Une des améliorations proposées dispose en plus de bonnes propriétés théoriques.*

KEYWORDS: *constraint satisfaction problems, constraint propagation, arc consistency*

MOTS-CLÉS : *programmation par contraintes, propagation de contraintes, cohérence d'arc*

1. Introduction

Constraint propagation is the basic operation in constraint programming. It is now well-recognized that its extensive use is necessary when we want to efficiently solve hard constraint satisfaction problems. All the constraint solvers use it as a basic step. Thus, each improvement that can be incorporated in a constraint propagation algorithm has an immediate effect on the behavior of the constraint solving engine. In practical applications, many constraints are of well-known types for which specific algorithms are available. These algorithms generally receive a set of removed values for one of the variables involved in the constraint, and propagate these deletions according to the constraint. They are usually as cheap as one can expect in cpu time. This state of things implies that most of the existing solving engines are based on a constraint-oriented or variable-oriented propagation scheme (ILOG Solver, CHOCO, etc.). And AC-3, with its natural both constraint-oriented [MAC 77] and variable-oriented [MCG 79] propagation of the constraints, is the generic constraint propagation algorithm which fits the best this propagation scheme. Its successors, AC-4, AC-6, and AC-7, indeed, were written with a value-oriented propagation. This is one of the reasons why AC-3 is the algorithm which is usually used to propagate those constraints for which nothing special is known about the semantics (and then for which no specific algorithm is available). This algorithm has a second strong advantage when compared to AC-4, AC-6 or AC-7, namely, its independence with regard to specific data structure which should be maintained if used during a search procedure. (And following that, a greater easiness to implement it.) On the contrary, its successors, while more efficient when applied to networks where much propagation occurs ([BES 95, BES 99]), need to maintain some additional data structures.

In this paper, our purpose is to present two new algorithms, `AC2000` and `AC2001`, which are both variable-oriented, and which improve AC-3 in efficiency (both in terms of constraint checks and cpu time). `AC2000`, like AC-3, is free of any data structure to be maintained during search. `AC2001`, at the price of a slight extra data structure (just an integer for each value-constraint pair) reaches an optimal worst-case time complexity. It leads to substantial gains, which are shown both on randomly generated and real-world instances of problems. A comparison with AC-6 shows interesting theoretical properties. Regarding the human cost of their implementation, `AC2000` needs five lines more than the classical AC-3, and `AC2001` needs the management of its additional data structure.

2. Preliminaries

Constraint network. A finite binary *constraint network* $\mathcal{P} = (\mathcal{X}, \mathcal{D}, \mathcal{C})$ is defined as a set of n *variables* $\mathcal{X} = \{X_1, \ldots, X_n\}$, a set of *domains* $\mathcal{D} = \{D(X_1), \ldots, D(X_n)\}$, where $D(X_i)$ is the finite set of possible *values* for variable X_i, and a set $\mathcal{C}$ of e binary *constraints* between pairs of variables. A constraint C_{ij} on the ordered set of variables (X_i, X_j) is a subset of the Cartesian product $D(X_i) \times D(X_j)$ that specifies the *allowed* combinations of values for the variables X_i and X_j. (For each constraint

C_{ij}, a constraint C_{ji} is defined between (X_j, X_i), allowing the same pairs of values in the reverse order.) Verifying whether a given pair (v_i, v_j) is allowed by C_{ij} or not is called a *constraint check*. A *solution* of a constraint network is an instantiation of the variables such that all the constraints are satisfied.

Arc consistency. Let $\mathcal{P} = (\mathcal{X}, \mathcal{D}, \mathcal{C})$ be a constraint network, and C_{ij} a constraint in $\mathcal{C}$. A value $v_i \in D(X_i)$ is *consistent with* C_{ij} iff $\exists v_j \in D(X_j)$ such that $(v_i, v_j) \in C_{ij}$. (v_j is then called a *support* for (X_i, v_i) on C_{ij}.) A value $v_i \in D(X_i)$ is *viable* iff it has support in all $D(X_j)$ such that $C_{ij} \in \mathcal{C}$. $\mathcal{P}$ is *arc consistent* iff all the values in all the domains are viable. We achieve arc consistency in $\mathcal{P}$ by removing every value which is not viable.

Algorithm 1: `Main algorithm`

function `AC` **(in** $\mathcal{X}$: set**): Boolean**
- $Q \leftarrow \emptyset$;
- **for each** $X_i \in X$ **do**
  - **for each** X_j *such that* $C_{ij} \in \mathcal{C}$ **do**
    - **if** *`Revise`*$(X_i, X_j,$ ***false***[1]$)$ **then**
      - **if** $D(X_i) = \emptyset$ **then** return **false** ;
      - $Q \leftarrow Q \cup \{X_i\}$;
- return `Propagation`(Q);

Algorithm 2: `Subprocedures for AC3`

function `Propagation3` **(in** Q: set**): Boolean**
- **while** $Q \neq \emptyset$ **do**
  - pick X_j from Q;
  - **for each** X_i *such that* $C_{ij} \in \mathcal{C}$ **do**
    - **if** *`Revise3`*(X_i, X_j) **then**
      - **if** $D(X_i) = \emptyset$ **then** return **false** ;
      - $Q \leftarrow Q \cup \{X_i\}$;
- return **true** ;

function `Revise3` **(in** X_i, X_j: variable**): Boolean**
- CHANGE $\leftarrow$ **false**;
- **for each** $v_i \in D(X_i)$ **do**
  - **if** $\nexists v_j \in D(X_j) / C_{ij}(v_i, v_j)$ **then**
    - remove v_i from $D(X_i)$;
    - CHANGE $\leftarrow$ **true**;
- return CHANGE ;

Before presenting the algorithms we propose in the next sections, let us briefly recall the AC-3 algorithm. We present it with the structure proposed by McGregor [MCG 79], which is built on a variable-based propagation scheme. This will be necessary in the following sections. The main algorithm (see Algorithm 1) is very close to

1. The third parameter is useless for AC-3 but will be used in the next versions of `Revise`.

the original AC-3, with an initialization phase (lines 1 to 3), a propagation phase (line 4), and with the use of the function Revise(X_i, X_j) that removes from $D(X_i)$ the values without support in $D(X_j)$ (line 2). But instead of handling a queue of the constraints to be propagated, it uses a queue[2] of the variables that have seen their domain modified. When a variable X_j is picked from the queue (line 2 of Propagation3 in Algorithm 2), all the constraints involving X_j are propagated with Revise. This is the only change w.r.t. the Mackworth's version of AC-3.[3] (This algorithm has been presented in [CHM 98] under the name AC-8.)

3. AC2000

If we closely examine the behavior of AC-3, we see that removing a single value v_j from a domain $D(X_j)$ (inside function Revise3) is enough to put X_j in the propagation queue (line 3 of AC in Algorithm 1 and line 5 of Propagation3 in Algorithm 2), and to provoke a call to Revise3(X_i, X_j) for every constraint C_{ij} involving X_j (lines 3 and 4 of Propagation3). Revise3 will look for a support for every value in $D(X_i)$ whereas for some of them v_j was perhaps not even a support. We exploit this remark in AC2000. Instead of looking blindly for a support for a value $v_i \in D(X_i)$ each time $D(X_j)$ is modified, we do that only under some conditions. In addition to the queue Q of variables modified, we use a second data structure, $\Delta(X_j)$,[4] which for each variable X_j contains the values removed since the last propagation of X_j. When a call to Revise2000$(X_i, X_j, lazymode)$ is performed, instead of systematically looking whether a value v_i still has support in $D(X_j)$, we first check that v_i really lost a support, namely one of its supports is in $\Delta(X_j)$ (right side of line 3 of Revise2000 in Algorithm 3). The larger $\Delta(X_j)$ is, the more expensive that process is, and the greater the probability to actually find a support to v_i in this set is. So, we perform this "lazymode" only if $\Delta(X_j)$ is sufficiently smaller than $D(X_j)$. We use a parameter Ratio to decide that. (See line 1 of Propagation2000 in Algorithm 3.) The Boolean lazymode is set to true when the ratio is not reached. Otherwise, lazymode is false, and Revise2000 performs exactly as Revise, going directly to lines 4 to 5 of Revise2000 without running the right side of line 3.

Analysis

Let us first briefly prove AC2000 correctness. Assuming AC-3 is correct, we just have to prove that the lazy mode does not let arc-inconsistent values in the domain. The only way the search for support for a value v_i in $D(X_i)$ is skipped, is

2. We name it a queue, but, as in AC-3, it has to be implemented as a set since in line 3 of Algorithm 1 and in line 5 of Algorithm 2 we add X_i to Q only if it does not already belong to it.
3. McGregor's version of the function Revise differs from Mackworth's one. Algorithm Revise3 is the Mackworth's version [MAC 77].
4. We take this name from [Van 92], where Δ denotes the same thing, but for a different use.

Algorithm 3: `Subprocedures for AC2000`

function `Propagation2000` (**in** Q: set)**: Boolean**
  while $Q \neq \emptyset$ **do**
    pick X_j from Q;
1   lazymode $\leftarrow$ $(|\Delta(X_j)| < \mathsf{Ratio} \cdot |D(X_j)|)$;
    for each X_i *such that* $C_{ij} \in \mathcal{C}$ **do**
      if *Revise2000*$(X_i, X_j, lazymode)$ **then**
        if $D(X_i) = \emptyset$ **then** return **false** ;
        $Q \leftarrow Q \cup \{X_i\}$;
2   reset $\Delta(X_j)$;
  return **true** ;

function `Revise2000` (**in** X_i, X_j: variable;
            in lazymode: **Boolean**)**: Boolean**
  CHANGE $\leftarrow$ **false**;
  for each $v_i \in D(X_i)$ **do**
3   **if** $\neg lazymode$ ***or*** $\exists v_j \in \Delta(X_j) / C_{ij}(v_i, v_j)$ **then**
4   **if** $\nexists v_j \in D(X_j) / C_{ij}(v_i, v_j)$ **then**
        remove v_i from $D(X_i)$;
        add v_i to $\Delta(X_i)$;
5   CHANGE $\leftarrow$ **true**;
  return CHANGE ;

when we could not find a support for v_i in $\Delta(X_j)$ (line 3 of `Revise2000`). Since $\Delta(X_j)$ contains all the values deleted from $D(X_j)$ since its last propagation (line 2 of `Propagation2000`), this means that v_i has exactly the same set of supports as before on C_{ij}. Thus, looking again for a support for v_i is useless. It remains consistent with C_{ij}.

The space complexity of `AC2000` is bounded above by the sizes of Q and Δ. Q is in $O(n)$ and Δ is in $O(nd)$, where d is the size of the largest domain. This gives a $O(nd)$ overall complexity.

The organization of `AC2000` is the same as in AC-3. The main change is in function `Revise2000`, where $\Delta(X_j)$ and $D(X_j)$ are examined instead of only $D(X_j)$. Their total size is bounded above by d. This leads to a worst-case where d^2 checks are performed, as in `Revise3`. Thus, the overall time complexity is in $O(ed^3)$ since `Revise2000` can be called d times per constraint. This is as in AC-3.

4. AC2001

In Section 3, we proposed an algorithm, which, like AC-3, does not need special data structures to be maintained during search. (Except the current domains, which have to be maintained by any search algorithm performing some look-ahead.) During a call to `Revise2000`, we have to look whether v_i has a support in $\Delta(X_j)$, to know whether it lost supports or not. In this last case, we have to look again for a support for

Algorithm 4: `Subprocedures for AC2001`

function `Propagation2001` (**in** Q: set)**: Boolean**
  while $Q \neq \emptyset$ **do**
    pick X_j from Q;
    for each X_i *such that* $C_{ij} \in \mathcal{C}$ **do**
1   **if** `Revise2001`$(X_i, X_j,$ ***true***$)$ **then**
        if $D(X_i) = \emptyset$ **then** return **false** ;
        $Q \leftarrow Q \cup \{X_i\}$;
    reset $\Delta(X_j)$;
  return **true** ;

function `Revise2001` (**in** X_i, X_j: variable;
  in lazymode: **Boolean**)**: Boolean**
  CHANGE $\leftarrow$ **false**;
  for each $v_i \in D(X_i)$ **do**
2   **if** $\neg lazymode$ ***or*** $Last(X_i, v_i, X_j) \in \Delta(X_j)$ **then**
3   **if** $\exists v_j \in D(X_j) / v_j >_d Last(X_i, v_i, X_j) \wedge C_{ij}(v_i, v_j)$
      then $Last(X_i, v_i, X_j) \leftarrow v_j$;
      else
        remove v_i from $D(X_i)$;
        add v_i to $\Delta(X_i)$;
        CHANGE $\leftarrow$ **true**;
  return CHANGE ;

v_i on C_{ij}. If we could remember what was the support found for v_i in $D(X_j)$ the last time we revised C_{ij}, we would just have to check whether this last support belongs to $\Delta(X_j)$ or not. Adding a very light extra data structure to remember the last support of a value on a constraint leads to the algorithm `AC2001` that we present in this section.

Let us store in $Last(X_i, v_i, X_j)$ the value that has been found as a support for v_i at the last call to `Revise2001`$(X_i, X_j, lazymode)$. The function `Revise2001` will always run in the lazy mode (see line 1 of `Propagation2001` in Algorithm 4), except during the initialization phase (line 2 of Algorithm 1), in which the $Last$ structure has not been initialized yet. In line 2 of `Revise2001` (see Algorithm 4) we check whether $Last(X_i, v_i, X_j)$ belongs to $\Delta(X_j)$. If it is the case, we look for a new support for v_i in $D(X_j)$; otherwise nothing is done since $Last(X_i, v_i, X_j)$ is still in $D(x_j)$. In line 3 of `Revise2001` we can see another advantage of storing $Last(X_i, v_i, X_j)$: If the supports are checked in $D(X_j)$ in a given ordering "$<_d$", we know that there isn't any support for v_i before $Last(X_i, v_i, X_j)$ in $D(X_j)$. Thus, we can look for a new support only on the values greater than $Last(X_i, v_i, X_j)$.

Analysis

Proving correctness of `AC2001` can be done very quickly since the framework of the algorithm is very close to `AC2000`. They have exactly the same initialization phase since the lazy mode is set to **false** during this whole step. During the propagation, they diverge in the way they handle this lazy mode. While Revise2000(X_i, X_j, **true**) was looking first in $\Delta(X_j)$ whether v_i had lost a support, Revise2001(X_i, X_j, **true**) just checks whether $Last(X_i, v_i, X_j)$ belongs to $\Delta(X_j)$. But we see that checking that $Last(X_i, v_i, X_j) \in \Delta(X_j)$ is sufficient to ensure that v_i still has a support in $D(X_j)$. And if a search for a new support has to be done, it is sufficient to limit this search to the values of $D(X_j)$ greater than $Last(X_i, v_i, X_j)$ w.r.t. to the ordering $<_d$ used to visit $D(X_j)$. Indeed, the previous call to Revise2001 stopped as soon as the value $Last(X_i, v_i, X_j)$ was found. It was then the smallest support for v_i in $D(X_j)$ w.r.t. $<_d$.

The space complexity of AC2001 is bounded above by the size of Q, Δ, and $Last$. Q is in $O(n)$ and Δ in $O(nd)$. But $Last$ is in $O(ed)$ since each value v_i has a $Last$ pointer for each constraint involving X_i. This gives a $O(ed)$ overall complexity.

As in AC-3 and AC2000, the function Revise2001 can be called d times per constraint in AC2001. But, at each call to Revise2001($X_i, X_j, lazymode$), for each value $v_i \in D(X_i)$, there will be a test on the $Last(X_i, v_i, X_j)$, and a search for support only on the values of $D(X_j)$ greater than $Last(X_i, v_i, X_j)$. Thus, the total work that can be performed for a value v_i over the d possible calls to Revise2001 on a pair (X_i, X_j) is bounded above by d tests on $Last(X_i, v_i, X_j)$ and d constraint checks. The overall time complexity is then bounded above by $d \cdot (d + d) \cdot 2 \cdot e$, which is in $O(ed^2)$. This is optimal [MOH 86]. AC2001 is the first optimal arc consistency algorithm proposed in the literature which is free of any lists of supported values. Indeed, the other optimal algorithms, AC-4, AC-6, AC-7, and AC-Inference all used these lists.

5. Experiments

In the sections above, we presented two refinements of AC-3, namely AC2000 and AC2001. It remains to see whether they are effective in saving constraint checks and/or cpu time when compared to AC-3. AS we said previously, the goal is not to compete with AC6/AC7, which have very subtle data structure for the propagation phase. An improvement (even small) w.r.t. AC-3 would fulfill our expectations.

5.1. *Arc consistency as a preprocessing*

The first set of experiments we performed should permit to see the effect of our refinements when arc consistency is used as a preprocessing (without search). In this case, the chance to have some propagations is very small on real instances. We have to fall in the phase transition of arc consistency (see [GEN 97]). So, we present results for randomly generated instances (those presented in [BES 99]), and for only one real-world instance. For the random instances, we used a model B generator [PRO 96]. The parameters are $< N, D, C/p1, T/p2 >$, where N is the number of

	AC-3		AC2000		AC2001	
	#ccks	time	#ccks	time	#ccks	time
<150, 50, 500, 1250>	100,010	0.04	100,010	0.05	100,010	0.05
<150, 50, 500, 2350>	507,783	0.18	507,327	0.18	487,029	0.17
<150, 50, 500, 2296>	2,860,542	1.06	1,601,732	0.69	688,606	0.35
<50, 50, 1225, 2188>	4,925,403	1.78	3,038,280	1.25	1,147,084	0.62
SCEN#08	4,084,987	1.67	3,919,078	1.65	2,731,158	1.26

Table 1. *Arc consistency results in mean number of constraint checks (#ccks) and mean cpu time in seconds (time) on a PC Pentium II 300MHz (50 instances generated for each random class).*

variables, D the size of the domains, C the number of constraints (their density $p1 = 2C/N \cdot (N-1)$), and T the number of forbidden tuples (their tightness $p2 = T/D^2$). The real-world instance, SCEN#08, is taken from the FullRLFAP archive,[5] which contains instances of radio link frequency assignment problems (RLFAPs). They are described in [CAB 99]. The parameter Ratio used in AC2000 is set to 0.2. Table 1 presents the results for four classes of random instances plus the real-world one.

The upper two are under-constrained (<$150, 50, 500/0.045, 1250/$**0.5**>) and over-constrained (<$150, 50, 500/0.045, 2350/$**0.94**>) problems. They represent cases where there is little or no propagation to reach the arc consistent or arc inconsistent state. This is the best case for AC-3, which performs poorly during propagation. We can see that AC2000 and AC2001 do not suffer from this.

The third and fourth experiments are at the phase transition of arc consistency for sparse (<$150, 50, 500/$**0.045**$, 2296/0.918$>) and dense (<$50, 50, 1225/$**1.0**$, 2188/0.875$>) problems. We can assume there is much propagation on these problems before reaching the arc consistent state. This has a significant impact on the respective efficiencies of the algorithms. The smarter the algorithm is, the lower the number of constraint checks is. AC2001 dominates AC2000, which itself dominates AC-3. And the cpu time follows this trend.

The lower experiment reports the results for the SCEN#08. This is one of the instances in FullRLFAP for which arc consistency is sufficient to detect inconsistency.

5.2. *Maintaining arc consistency during search*

The second set of experiments we present in this section shows the effect of our refinements when arc consistency is maintained during search (MAC algorithm [SAB 94]) to find the first solution. We present results for all the instances contained in the FullRLFAP archive for which more than 2 seconds were necessary to find a solution or prove that none exists. We took again 0.2 for the Ratio in AC2000. It has

5. We thank the Centre d'Electronique de l'Armement (France).

	MAC-3		MAC2000		MAC2001	
	#ccks	time	#ccks	time	#ccks	time
SCEN#01	5,026,208	2.33	4,319,423	2.10	1,983,332	1.66
SCEN#11	77,885,671	39.50	77,431,840	38.22	9,369,298	22.69
GRAPH#09	6,269,218	2.95	5,164,692	2.57	2,127,598	2.04
GRAPH#10	6,790,702	3.04	5,711,923	2.68	2,430,109	1.86
GRAPH#14	5,503,326	2.53	4,253,845	2.14	1,840,886	1.67

Table 2. *Results for search of the first solution with a MAC algorithm in mean number of constraint checks (#ccks) and mean cpu time in seconds (time) on a PC Pentium II 300MHz.*

to be noticed that the original question in these instances is not satisfiability but the search of the "best" solution, following some criteria. It is of course out of the scope of this paper.

From these instances we can see a slight gain for `AC2000` on AC-3. It dominates it on all the instances except SCEN#11. On SCEN#11, it can be noticed that with Ratio = 0.1, `AC2000` performs better than AC-3. A more significant gain can be seen for `AC2001`, with up to 9 times less constraint checks and twice less cpu time on SCEN#11. As for the experiments performed on random instances at the phase transition of arc consistency, this tends to show that the trick of storing the $Last$ data structure significantly pays off. However, we have to keep in mind that we are only comparing algorithms with simple data structures. This prevents them from reaching the efficiency of algorithms using lists of supported values. (E.g., a MAC algorithm using AC-6 for enforcing arc consistency needs only 14.69 seconds to solve the SCEN#11 instance.)

6. `AC2001` vs AC-6

In the previous sections, we proposed two algorithms based on AC-3 to achieve arc consistency on a binary constraint network. `AC2000` is close to AC-3, from which it inherits its $O(ed^3)$ time complexity and its $O(nd)$ space complexity. `AC2001`, thanks to its additional data structure, has an optimal $O(ed^2)$ worst-case time complexity, and an $O(ed)$ space complexity. These are the same characteristics as AC-6.[6] So, we can ask the question: "What are the differences between `AC2001` and AC-6?".

Let us first briefly recall the AC-6 behavior [BES 94]. AC-6 looks for one support (the *first* one or *smallest* one with respect to the ordering $<_d$) for each value (X_i, v_i) on each constraint C_{ij} to prove that (X_i, v_i) is currently viable. When (X_j, v_j) is found as the smallest support for (X_i, v_i) on C_{ij}, (X_i, v_i) is added to $S[X_j, v_j]$, the list

6. We do not speak about AC-7 here, since it is the only one among these algorithms to deal with the bidirectionality of the constraints (namely, the fact that $C_{ij}(v_i, v_j) = C_{ji}(v_j, v_i)$).

of values currently having (X_j, v_j) as smallest support. If (X_j, v_j) is removed from $D(X_j)$, it is added to the $DeletionSet$, which is the stream driving propagations in AC-6. When (X_j, v_j) is picked from the $DeletionSet$, AC-6 looks for the *next* support (i.e., greater than v_j) in $D(X_j)$ for each value (X_i, v_i) in $S[X_j, v_j]$.

To allow a closer comparison, we will suppose in the following that AC-6 data structures are decomposed by variable and constraint. The $DeletionSet$ is split on each variable. This leads to the data structure Δ of AC2001 (since $\cup_{X_i \in \mathcal{X}} \Delta(X_i)$ in AC2001 corresponds to $DeletionSet$ in AC-6). The $S[X_i, v_i]$ lists are split on each constraint C_{ij} involving X_i, leading to a structure $S[X_i, v_i, X_j]$, as in AC-7.

Property 1 *Let $\mathcal{P} = (\mathcal{X}, \mathcal{D}, \mathcal{C})$ be a constraint network. If we suppose AC2001 and AC-6 follow the same ordering of variables and values when looking for supports and propagating deletions, then, enforcing arc consistency on $\mathcal{P}$ with AC2001 requires the same constraint checks as with AC-6.*

Proof. Since they follow the same ordering, both algorithms perform the same constraint checks in the initialization phase: they stop search for support for a value v_i on C_{ij} as soon as the first v_j in $D(X_j)$ compatible with v_i is found, or when $D(X_j)$ is exhausted (then removing v_i). During the propagation phase, both algorithms look for a new support for a value v_i on C_{ij} only when v_i has lost its current support v_j in $D(X_j)$ (i.e., $v_i \in S[X_j, v_j, X_i]$ for AC-6, and $v_j = Last(X_i, v_i, X_j)$ for AC2001). Both algorithms start the search for a new support for v_i at the value in $D(X_j)$ immediately greater than v_j w.r.t. the $D(X_j)$ ordering. Thus, they will find the same new support for v_i on C_{ij}, or will remove v_i, at the same time, and with the same constraint checks. And so on. □

From property 1, we see that the difference between AC2001 and AC-6 cannot be characterized by the number of constraint checks they perform. We will then focus on the way they find which values should look for a new support. For that, both algorithms handle their specific data structures. Let us characterize the number of times each of them checks a structure when a set $\Delta(X_j)$ of deletions is propagated on a given constraint C_{ij}.

Property 2 *Let C_{ij} be a constraint in a network $\mathcal{P} = (\mathcal{X}, \mathcal{D}, \mathcal{C})$. Let $\Delta(X_j)$ be a set of values removed from $D(X_j)$ that have to be propagated on C_{ij}. If,*

$A = |\Delta(X_j)| + \sum_{v_j \in \Delta(X_j)} |S[X_j, v_j, X_i]|$,

$B = |D(X_i)|$, *and*

C = *# checks performed on C_{ij} to propagate $\Delta(X_j)$,*

then, $A + C$ and $B + C$ represent the number of operations AC-6 and AC2001 will respectively perform to propagate $\Delta(X_j)$ on C_{ij}.

Proof. From property 1 we know that AC-6 and AC2001 perform the same constraint checks. The difference is in the process leading to them. AC-6 traverses the $S[X_j, v_j, X_i]$ list for each $v_j \in \Delta(X_j)$ (i.e., A operations), and AC2001 checks whether $Last(X_i, v_i, X_j)$ belongs to $\Delta(X_j)$ for every v_i in $D(X_i)$ (i.e., B operations). □

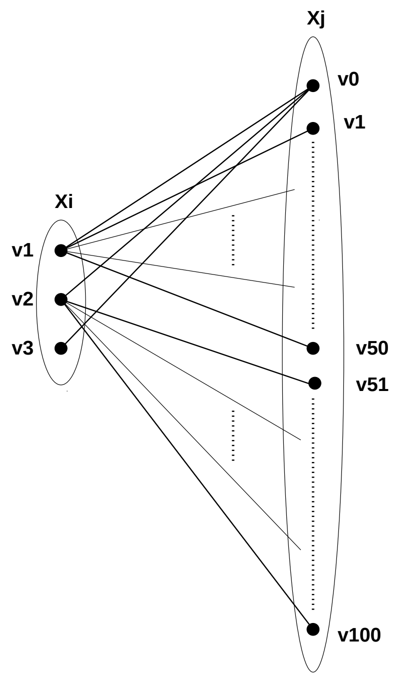

Figure 1. *The constraint example*

We illustrate this on the extreme case presented in Figure 1. In that example, the three values of X_i are all compatible with the first value v_0 of X_j. In addition, (X_i, v_1) is compatible with all the values of X_j from v_1 to v_{50}, and (X_i, v_2) with all the values of X_j from v_{51} to v_{100}. Imagine that for some reason, the value v_3 has been removed from $D(X_i)$ (i.e., $\Delta(X_i) = \{v_3\}$). This leads to $A = 1$, $B = 101$, and $C = 0$, which is a case in which propagating with AC-6 is much better than with `AC2001`, even if none of them needs any constraint check. Indeed, AC-6 just checks that $S[X_i, v_3, X_j]$ is empty,[7] and stops. `AC2001` takes one by one the 101 values v_j of $D(X_j)$ to check that their $Last(X_j, v_j, X_i)$ is not in $\Delta(X_i)$. Imagine now that the values v_1 to v_{100} of $D(X_j)$ have been removed (i.e., $\Delta(X_j) = \{v_1, \ldots, v_{100}\}$). Now, $A = 100$, $B = 3$, and $C = 0$. This means that `AC2001` will clearly outperform AC-6. Indeed, AC-6 will check for all the 100 values v_j in $\Delta(X_j)$ that $S[X_j, v_j, X_i]$ is empty,[8] while `AC2001` just checks that $Last(X_i, v_i, X_j)$ is not in $\Delta(X_j)$ for the 3 values in $D(X_i)$.

7. The only value compatible with (X_i, v_3) is (X_j, v_0), which is currently supported by (X_i, v_1).
8. Indeed, (X_j, v_0) is the current support for the three values in $D(X_i)$ since it is the smallest in $D(X_j)$ and it is compatible with every value in $D(X_i)$.

Discussion

Thanks to property 2, we have characterized the amount of effort necessary to AC-6 and AC2001 to propagate a set $\Delta(X_j)$ of removed values on a constraint C_{ij}. $A - B$ gives us information on which algorithm is the best to propagate $\Delta(X_j)$ on C_{ij}. We can then easily imagine a new algorithm, which would start with an AC-6 behavior on all the constraints, and would switch to AC2001 on a constraint C_{ij} as soon as $A - B$ would be positive on this constraint (and then forget the S lists on C_{ij}). Switching from AC2001 to AC-6 on this constraint is no longer possible because we can deduce in constant time that $Last(X_i, v_i, X_j) = v_j$ when v_i belongs to $S[X_j, v_j, X_i]$, but we cannot obtain cheaply $S[X_j, v_j, X_i]$ from the $Last$ structure. A more elaborated version would maintain the S lists even in the AC2001 behavior (putting v_i in $S[X_j, v_j, X_i]$ each time v_j is found as being the $Last(X_i, v_i, X_j)$). This would permit to switch from AC-6 to AC2001 or the reverse at any time on any constraint in the process of achieving arc consistency. These algorithms are of course far from our initial purpose of emphasizing easiness of implementation since they require the S and $Last$ structures to be maintained during search.

7. Conclusion

We presented AC2000 and AC2001, two refinements in AC-3. The first one improves slightly AC-3 in efficiency (number of constraint checks and cpu time) although it does not need any new data structure to be maintained during search. The second, AC2001, needs an additional data structure, the $Last$ supports, which should be maintained during search. This data structure permits a significant improvement on AC-3, and decreases the worst-case time complexity to the optimal $O(ed^2)$. AC2001 is the first algorithm in the literature achieving optimally arc consistency while being free of any lists of supported values. Its behavior is compared to that of AC-6, making a contribution to the understanding of the different AC algorithms, and opening an opportunity of improvement. This is in the same vein as the work on AC-3 vs AC-4 [WAL 93], which was leading up to AC-6.

Acknowledgements

We thank the anonymous reviewer who suggested an improvement in the presentation of the algorithm AC2000.

8. References

[BES 94] BESSIÈRE C., "Arc-consistency and arc-consistency again", *Artificial Intelligence*, vol. 65, 1994, p. 179-190.

[BES 95] BESSIÈRE C., FREUDER E. C., RÉGIN J. C., "Using inference to reduce arc consistency computation", *Proceedings IJCAI'95*, Montréal, Canada, 1995, p. 592-598.

[BES 99] BESSIÈRE C., FREUDER E., RÉGIN J., "Using constraint metaknowledge to reduce arc consistency computation", *Artificial Intelligence*, vol. 107, 1999, p. 125-148.

[CAB 99] CABON C., DE GIVRY S., LOBJOIS L., SCHIEX T., WARNERS J., "Radio link frequency assignment", *Constraints*, vol. 4, 1999, p. 79–89.

[CHM 98] CHMEISS A., JÉGOU P., "Efficient Path-Consistency Propagation", *International Journal on Artificial Intelligence Tools*, vol. 7, num. 2, 1998, p. 121-142.

[GEN 97] GENT I., MACINTYRE E., PROSSER P., SHAW P., WALSH T., "The constrainedness of arc consistency", *Proceedings CP'97*, Linz, Austria, 1997, p. 327–340.

[MAC 77] MACKWORTH A., "Consistency in Networks of Relations", *Artificial Intelligence*, vol. 8, 1977, p. 99-118.

[MCG 79] MCGREGOR J., "Relational consistency algorithms and their application in finding subgraph and graph isomorphism", *Information Science*, vol. 19, 1979, p. 229–250.

[MOH 86] MOHR R., HENDERSON T., "Arc and Path Consistency Revisited", *Artificial Intelligence*, vol. 28, 1986, p. 225-233.

[PRO 96] PROSSER P., "An empirical study of phase transition in binary constraint satisfaction problems", *Artificial Intelligence*, vol. 81, 1996, p. 81-109.

[SAB 94] SABIN D., FREUDER E., "Contradicting Conventional Wisdom in Constraint Satisfaction", *Proceedings PPCP'94*, Seattle WA, 1994.

[Van 92] VAN HENTENRYCK P., DEVILLE Y., TENG C., "A generic arc-consistency algorithm and its specializations", *Artificial Intelligence*, vol. 57, 1992, p. 291-321.

[WAL 93] WALLACE R., "Why AC-3 is Almost Always Better than AC-4 for Establishing Arc Consistency in CSPs", *Proceedings IJCAI'93*, Chambéry, France, 1993, p. 239–245.