

Optimal Distributed Arc-Consistency

Youssef Hamadi

Thales Research & Technology

Domaine de Corbeville

91404 Orsay Cedex France

Abstract. This paper presents *DisAC-9*, the first optimal distributed algorithm performing the arc-consistency of a constraint network. Our method is optimal according to the number of message passing operations. This algorithm can firstly, give speed-up over the fastest central arc-consistency algorithms; secondly, achieve the fast processing of distributed constraint satisfaction problems (DCSP). Experimental results include classical benchmark and large hard randoms problems. These results allow us to show that the phase transition phenomenon of distributed arc-consistency is closely related to the granularity of the distributed system. The consequences of this analysis are showed to be very important for the future of distributed constraint satisfaction.

Keywords: Parallel relaxation, distributed constraint satisfaction, arc-consistency, distributed problem solving, phase transition, parallel algorithms

1. Introduction

The constraint satisfaction problem (CSP) is a powerful framework for general problem solving. It involves finding a solution to a constraint network, i.e., finding one of d values for n variables subject to constraints that are restrictions on which combinations of values are acceptable. It is widely used in Artificial Intelligence (AI), with applications ranging from machine vision to crew scheduling and many other fields (see [21] for a survey). Since it is an NP-complete task, many filtering techniques have been designed. These methods reduce the problem search space by pruning the set of possible values. Owing to their reasonable complexities, arc-consistency algorithms which detect and eliminate inconsistencies involving all pairs of variables are widely used [19].

In the last few years, the distributed constraint satisfaction (DCSP) paradigm emerges from the necessity of solving distributed AI problems. These distributed problems are more and more crucial because of the wide distribution of information allowed by Internet facilities. As a result, several authors tried to bring classical CSP procedures to the distributed framework [22, 10, 9, 8].

More relevant works on distributed arc-filtering were brought by P. R. Cooper and M. J. Swain who proposed in [4] two massively parallel



© 2001 Kluwer Academic Publishers. Printed in the Netherlands.

(nd processes) versions of AC-4 [14]. More recently, T. Nguyen and Y. Deville [17] distributed AC-4 in a more realistic scale thanks to a granularity varying from 1 to n process(es).

Nevertheless, these works do not address the real difficulty of distributed systems which is the complexity of message passing operations. S. Kasif who has studied the complexity of parallel relaxation [13] concludes that discrete relaxation is inherently a sequential process. This means that due to dependencies between deletions, a distributed/parallel algorithm is likely to perform a lot of messages since it is likely to perform in a sequential fashion.

We have learned from all the previous works and focused our work on message passing reduction. The DisAC-9¹ algorithm we propose here is optimal in the number of message passing operations. This major improvement over the previous results was made possible by the exploitation of the bidirectionality property of constraint relations which allows agents to induce acquaintances deletions. Since bidirectionality is a general property of constraints, DisAC-9 is a general purpose algorithm. Its worst time complexity is $O(n^2d^3)$ which allows us to reach optimality in message operations, nd . According to that, our method can achieve the fast processing of DCSPs and even give major speed-up in the processing of large CSPs.

The paper is organized as follows. We first give background about constraints, arc-consistency and message passing in section 2. Section 3 presents related works. Then we present the centralized algorithm AC-6. Next we introduce our distributed algorithm, DisAC-9. Theoretical analysis follows in section 6. Before concluding, we present experimental results in section 7.

2. Definitions

2.1. CONSTRAINT AND ARC-CONSISTENCY BACKGROUND

A *binary CSP*, $P = (\mathcal{X}, \mathcal{D}, \mathcal{C})$ involves, \mathcal{X} a set of n variables $\{i, j, \dots\}$, $\mathcal{D} = \{D_i, D_j, \dots\}$ where each element is a variable's domain, \mathcal{C} the set of binary constraints $\{C_{ij}, \dots\}$ where C_{ij} is a constraint between i and j . $C_{ij}(a, b) = \text{true}$ means that the association value a for i and b for j is allowed. A *solution* to a constraint satisfaction problem is an assignment of the variables in a way that all the constraints are satisfied. A constraint network (CN) $\mathcal{G} = (\mathcal{X}, \mathcal{C})$ can be associated to any binary CSP.

¹ DisAC-9, because our algorithm is not a distributed version of AC-6, AC-7 or AC-8.

Each element of a domain represents a *value*. A *label* (i, a) is the assignment of the value a to the variable i . It is supported by a value b along C_{ij} *iff* $C_{ij}(a, b) = \text{true}$, b is called a *support* for a along C_{ij} . The value a is also a support for b along C_{ji} . A value a from D_i is *viable* *iff* $\forall C_{ij}$, there exists a support b for a in D_j . The domain \mathcal{D} of a constraint network is *arc-consistent* for this CN if for every variable i , all values in D_i are viable. The maximal arc-consistent domain *Dmax-AC* of a constraint network is defined as the union of all domains included in \mathcal{D} and arc-consistent for this CN. This maximal domain is also *arc-consistent* and is computed by an arc-consistency algorithm.

2.2. DISTRIBUTED CONSTRAINT SATISFACTION BACKGROUND

A *binary distributed CSP*, $P = (\mathcal{X}, \mathcal{D}, \mathcal{C}, \mathcal{A})$ involves a binary CSP, $(\mathcal{X}, \mathcal{D}, \mathcal{C})$ partitioned among \mathcal{A} , a set of p autonomous agents ($1 \leq p \leq n$). Each agent owns the domains/constraints on its variables. An agent that owns the variable i owns D_i and all the binary constraints C_{ij} .

We assume the following communication model [22]. Agents communicate by sending messages. An agent can send messages to other agents *iff* he knows their addresses in the network. For the transmission between any pair of agents, messages are received in the order in which they were sent. The delay in delivering messages is finite.

For readers unfamiliar with exchanges in distributed systems, we present in table I the real data transmission costs for several parallel computers [5]. In local area networks (LAN), these costs are orders of magnitudes larger. These values were computed in an empty communication network, which means that they are lower bounds values. In real systems, the more you send messages, the longer the transmissions.

Table I. The real cost of message passing operations

	iPSC	nCUBE/10	iPSC/2	nCUBE/2	iPSC/860	CM-5
T_s [$\mu\text{s}/\text{msg}$]	4100	400	700	160	160	86
T_b [$\mu\text{s}/\text{byte}$]	2.8	2.6	0.36	0.45	0.36	0.12
T_{fp} [$\mu\text{s}/\text{flop}$]	25	8.3	3.4	0.50	0.033	0.33

The row T_s represents the *start-up* cost of a message, T_b is the cost per byte. We give T_{fp} as a reference point, which is the cost of one floating point operation on each machine. Each message reception incurs a cost, generally this cost is similar to T_s . So the estimated cost for a message m is, $\geq 2 * T_s + |m| * T_b$.

These measures show that message passing is critical. Our work considers for discrete relaxation the reduction of message passing operations.

3. Previous works

We emphasize previous works related to distributed arc-filtering here. We do not present works which use the PRAM machine model like [20] because this model is not representative of the current computing facilities which all use the message passing paradigm. We also do not present works which are not based on optimal central method [18]. In table II we summarize the main features of 3 methods. These methods propose a distributed implementation of the AC-4 sequential algorithm. The *#ps* column gives the granularity, *first step* gives the complexity of the first step of the method since these are two step procedures like the central AC-4. Then, *complexity* gives the method full complexity and *#msg* the number of message passing operations required in the worst case scenario.

A distributed algorithm is *local*, if its messages transmissions are limited, i.e. its transmissions do not involve the whole processes. This definition of locality is sufficient for our needs. For a comprehensive work on locality in distributed systems, see [16]. The column *local* specifies whether a method is local or not.

Table II. Distributed arc-consistency algorithms

Algo.	#ps	first step	complexity	#msg	local
AC Chip	nd	nd	$O(nd)$	-	-
ACP	nd	nd	$O(nd \log_2(nd))$	$n^2 d^2$	no
DisAC-4	$1 \leq p \leq n$	$(n^2 d^2)/p$	$O((n^2 d^2)/p)$	' nd '	no

AC Chip This algorithm presented in [4] defines a VLSI circuit. It uses a very fine granularity and it does not use explicit message passing operations. Both allows a rather good complexity. Nevertheless it would need $O(n^2 d^2)$ initializing steps for loading the network in the digital chip. This work is related to the “hardware for AI” stream which adapt computers architecture to problem solving (e.g. [11]).

ACP This is the software version of the previous algorithm [4]. It uses nd synchronous processes which communicate by message operations.

Each message is sent with $O(\log_2(nd))$ steps and is received by the whole processes. The method is not local.

DisAC-4 This coarse grain method uses p processes [17] sharing an Ethernet network. This allows “efficient” broadcasting of messages between processes. Nevertheless, this brings to an underlying synchronism since “collision networks” cannot carry more than one message in a given time. If we consider a more realistic framework for the method, each broadcast requires p messages which brings to n^2d messages.

4. The arc-consistency algorithm AC-6

We briefly describe here Bessière’s AC-6 algorithm [1]. This method introduces a total ordering between values in each domain, it computes one support (the first one) for each label (i, a) on each constraint C_{ij} . That to make sure of the current viability of (i, a) . When (j, b) is found as the smallest support of (i, a) on C_{ij} , (i, a) is added to S_{jb} the list of values currently having (j, b) as smallest support. Then, if (j, b) is deleted, AC-6 looks for a new support in D_j for all values in S_{jb} . During this search, the method only checks values greater than b . The algorithm incrementally computes a support for each label on each constraint relation in $O(n^2d^2)$. On average cases, it performs fewer checks than AC-4 whose the initialization step globally computes and stores the whole support relations.

In the following, we adapt this incremental behavior to our distributed framework, moreover we drastically improve efficiency by reaching optimality in the message transmissions.

5. The distributed arc-consistency algorithm DisAC-9

The key steps of this new distributed algorithm are the following. According to local knowledge, each agent that owns a set of variables starts computing inconsistent labels. After this computation, *selected labels* are transmitted. This means that only the deleted labels that induce new deletions are sent to an acquaintance. After this initializing step, each agent starts the processing of incoming messages. These messages will modify its local knowledge about related acquaintances. This will allow it to update support information for its labels. Obviously, this can lead to new deletions and maybe to new outgoing messages. Since communications occurs between acquaintances agents, DisAC-9 is local (see section 3).

5.1. MINIMAL MESSAGE PASSING

Constraints bidirectionality allows agents to partially share acquaintances knowledge: since $C_{ij}(a, b) = C_{ji}(b, a)$, an agent owning i but not owning j knows any C_{ji} relations. This knowledge can be used to infer over acquaintances deletions.

More generally, *agent-1* currently deleting (i, a) computes for any acquaintance *agent-2*, labels (j, b) having (i, a) as support. If another local value (i, a') can support (j, b) , it is useless to inform *agent-2* for (i, a) deletion. Otherwise, *agent-2* must delete (j, b) as an outcome of the deletion of (i, a) . It must be informed of the local deletion.

5.2. A COMPLETE EXAMPLE

Figure 1 illustrates this aspect on a particular DCSP. In this problem, two agents share a constraint C_{12} . The constraint relation is represented in the left part of the figure by the micro-structure graph. In this graph, allowed pairs of values are linked. In the right part, we represent the support relation computed by the two agents. An arrow goes from value a to value b if b is the current support for a .




Figure 1. Message minimizing in a DCSP

For *agent-1*, support lists are the following, $S_{2a} = \{(1, a)\}$, $S_{2b} = \{(1, b), (1, c)\}$ and $S_{2c} = \{(1, d)\}$. The second agent, *agent-2* computes $S_{1a} = \{(2, a)\}$, $S_{1b} = \{(2, b)\}$, $S_{1c} = \{(2, c)\}$ and $S_{1d} = \{\}$.

An extra data structure is used during the message minimization process. $first_j[i][a] = b$ tracks that according to C_{ij} , (j, b) is the first label supporting (i, a) in D_j . The *first* table of *agent-1* contains the following values, $first_2[1][a] = a$, $first_2[1][b] = b$, $first_2[1][c] = b$, $first_2[1][d] = c$. For *agent-2*, $first_1[2][a] = a$, $first_1[2][b] = b$ and $first_1[2][c] = c$.

Assuming the removal of $(1, c)$ (due to an incoming event) let us consider the notification of *agent-2* about this deletion. We can just consider in D_2 values greater or equal to $first_2[1][c]$, the first support

for $(1, c)$. Among these values, we must check those possibly supported by $(1, c)$. We find $(2, b)$ and $(2, c)$. For $(2, b)$ we determine $(1, b)$ as a possible support and for $(2, c)$ we find $(1, d)$. According to these local computations, *agent-1* states that any message about $(1, c)$ deletion is useless. Such a message could update the *agent-2* knowledge but it cannot lead to further domain reductions.

Now assuming the loss of $(1, b)$, we detect that $(2, b)$ must be deleted by *agent-2* since there are no more supports in D_1 . Hence, we must send deletion information to *agent-2*. But if we only inform for $(1, b)$ deletion, *agent-2*, by considering S_{1b} will compute $(1, c)$ as the new support for $(2, b)$ since it assumes that $(1, c)$ is still viable. A solution is to inform for all previous deletions. In this example, *agent-1* informs in the same message for deletions of $(1, b)$ and $(1, c)$. As a result *agent-2* will delete $(2, b)$ but it will not inform for this deletion *agent-1* since the agent knows that $(1, b)$ was deleted. It will remove $(2, c)$ from S_{1c} and add it to S_{1d} .

This example shows that by performing more local computations, an agent can prevent outgoing messages. This balance between local work and message transmissions is quite interesting in a distributed environment (see table I).

5.3. ALGORITHM

After this illustration, we present the related algorithm. Our distributed system contains $1 \leq p \leq n$ autonomous agents (cf. 2.2) (plus an extra agent performing termination detection see section 6).

The knowledge of agent k ($1 \leq k \leq p$) is represented/handled by the following data structures/primitives.

Data structures

- *Acquaintances*, agents sharing constraints with k are referenced in this set
 - *localVar*, this set stores the variables for k
 - *localD_i* this set stores the current domain of i ($\forall i \in \text{localVar}$)
 - *localM_i* boolean state vector, keeps tracks of deleted values from the initial domain D_i ($\forall i \in \text{localVar}$)
 - *linkedVar*, this set stores the non local variables which are linked to local variables by a constraint relation
- $$\text{linkedVar} = \{i \in \mathcal{X}, \text{ tq } i \notin \text{localVar} \text{ and } \exists j \in \text{localVar} \text{ tq } C_{ij} \in \mathcal{C}\}$$

- *linkedM_i* boolean state vector, keeps track of external variables deleted labels, ($\forall i \in \text{linkedVar}$)
- A set of lists, $S_{jb} = \{(i, a) | a \in D_i \text{ and } b \text{ is the smallest value in } D_j \text{ supporting } (i, a)\}$
- *first* an integer table, $\text{first}_j[i][a] = b$ tracks that according to C_{ij} , (j, b) is the first label supporting (i, a)
- *localList*, *extList*, these lists store the deleted labels, the second one is used for external propagation
- *sendSet*, vector of sets, stores outcoming deletions informations
- *inform*, boolean table, $\text{inform}_{acc}[i]$ means that *acc* has been selected for *localD_i* transmission

Constant time primitives

- *affected(j)* returns the acquaintances that owns the variable *j*. *owns(A, i)* returns *true* if agent *A* owns the variable *i*
- *higher(D)* returns the last element of the domain *D*. if $D = \emptyset$ returns -1 . *next(a, M)*, *M* is a state vector, returns the smallest indice *ind*, greater than *a* | $M[\text{ind}] = \text{true}$
- *get(S)* returns the first element from *S*. *addTo(S, a)* inserts *a* in *S*. *emptyP(S)* returns *true* if *S* is empty
- *localDUpdate(i, a)* this primitive realizes several important updates. It removes the value *a* from *localD_i* then it stores this deletion in *localM_i* ($\text{localM}_i[a] \leftarrow \text{false}$). After performing these updates, the method checks the size of *localD_i*. If ($\text{localD}_i = \emptyset$), a *stop* message is broadcasted to the whole processes and termination occurs with the detection of problem inconsistency

Message passing primitives

- *getMsg()* blocking instruction, returns the first incoming message
- *broadcast(m, P)* the message *m* is broadcasted to processes in the set *P*, it uses $O(\log_2(|P| + 1))$ operations

Algorithm 1: DisAC-9 main

```

begin
  localList  $\leftarrow \emptyset$ ; extList  $\leftarrow \emptyset$ 
  for each  $i \in localVar$  do
    1  localDi  $\leftarrow D_i$ 
    2  for each  $a \in D_i$  do localMi[a]  $\leftarrow true$ ; Sia  $\leftarrow \emptyset$ 
  for each  $j \in linkedVar$  do
    3  for each  $a \in D_j$  do linkedMj[a]  $\leftarrow true$ ; Sja  $\leftarrow \emptyset$ 
    4  for each  $i \in localVar$  do
      for each  $a \in D_i$  do firstj[i][a]  $\leftarrow 0$ 
  %
  % first step, init. supports
  for each arc (i, j) | i  $\in localVar$  do
    for each  $a \in localD_i$  do
      b  $\leftarrow 0$ 
      5  nextSupport(j, i, a, b, emptySupport)
      if emptySupport then
        6  localDUpdate(i, a)
        7  addTo(localList, (i, a))
        8  addTo(extList, (i, a))
      else
        9  addTo(Sjb, (i, a))
        10 if j  $\in linkedVar$  then firstj[i][a]  $\leftarrow b$ 
  % internal and external propagations
  11 processLists(localList, extList)
  %
  % second step, interactions
  termination  $\leftarrow false$ 
  while !termination do
    12 m  $\leftarrow getMsg()$ 
    switch m do
      13 case stop
        termination  $\leftarrow true$ 
      14 case deletedLabels;set
        while !emptyP(set) do
          15 (j, Mj)  $\leftarrow get(set)$ 
          for each  $a \in D_j$  do
            16 if linkedMj[a]  $\neq M_j[a]$  then
              17 linkedMj[a]  $\leftarrow false$ 
              18 addTo(localList, (j, a))
          19 processLists(localList, extList)
end

```

Each agent starts with a call to the *main* procedure (see algo. 1). This procedure has two steps. In the first one, agent's knowledge allows the filtering of local domains. For each local variable i , each value a in $localD_i$ is checked for viability. To achieve this test, agent looks for support for each related constraint C_{ij} . The *nextSupport* procedure (see algo. 3) looks for the first support according to C_{ij} . If this procedure

Algorithm 2: DisAC-9 processLists

Algorithm: processLists(**in-out:** *localList*, *extList*)

```

begin
1   for each  $j \in \text{Acquaintances}$  do
2       for each  $i \in \text{localVar} \mid \exists k \in \text{linkedVar}$  and owns( $j, k$ ) and  $C_{ik}$  do
3            $\text{inform}_j[i] \leftarrow \text{false}$ 
4       while !emptyP(localList) do
5           ( $j, e$ )  $\leftarrow \text{get}(\text{localList})$ 
6           deletionProcessing( $j, e$ )
7       while !emptyP(extList) do
8           ( $j, e$ )  $\leftarrow \text{get}(\text{extList})$ 
9           selectiveSend( $j, e, \text{sendSet}$ )
10      for each  $j \in \text{Acquaintances}$  do
11          if  $\text{sendSet}_j \neq \emptyset$  then broadcast( $j, \text{deletedLabels} : \text{sendSet}_j$ )
end

```

Algorithm 3: DisAC-9 nextSupport

Algorithm: nextSupport(**in:** j, i, a ; **in-out:** b ; **out:** *emptySupport*)

```

begin
1   if  $b \leq \text{higher}(D_j)$  then
2       if  $j \in \text{localVar}$  then  $\text{stateVector} \leftarrow \text{localM}_j$ 
3       else  $\text{stateVector} \leftarrow \text{linkedM}_j$ 
4       while !stateVector[ $b$ ] do  $b++$ 
5       emptySupport  $\leftarrow \text{false}$ 
6       while !emptySupport and ! $C_{ij}(a, b)$  do
7           if  $b < \text{higher}(D_j)$  then  $b \leftarrow \text{next}(b, \text{stateVector})$ 
8           else emptySupport  $\leftarrow \text{true}$ 
9       else emptySupport  $\leftarrow \text{true}$ 
end

```

Algorithm 4: DisAC-9 deletionProcessing

Algorithm: deletionProcessing(**in:** j, e ; **in-out:** *localList*, *extList*)

```

begin
1   while !emptyP( $S_{j_e}$ ) do
2       ( $i, a$ )  $\leftarrow \text{get}(S_{j_e})$ 
3       if  $a \in \text{localD}_i$  then
4            $c \leftarrow e + 1$ 
5           nextSupport( $j, i, a, c, \text{emptySupport}$ )
6           if emptySupport then
7               localDUpdate( $i, a$ )
8               addTo(localList, ( $i, a$ ))
9               addTo(extList, ( $i, a$ ))
10          else
11              addTo( $S_{j_c}, (i, a)$ )
12              if  $j \in \text{linkedVar}$  then  $\text{first}_j[i][a] \leftarrow c$ 
end

```

Algorithm 5: DisAC-9 selectiveSend

Algorithm: selectiveSend(in : (i, a) ; in-out : $sendSet$)

```

begin
1  for each arc  $(i, j) \mid j \in linkedVar$  do
2     $acc \leftarrow affected(j)$ 
3    if  $!inform_{acc}[i]$  then
4       $b \leftarrow first_j[i][a]$ 
       $endD_j \leftarrow false$ 
       $selected \leftarrow false$ 
5      while  $!selected$  and  $!endD_j$  and  $b \leq higher(D_j)$  do
6        if  $linkedM_j[b]$  and  $C_{ij}(a, b)$  then
           $endD_i \leftarrow false$ 
           $otherSupport \leftarrow false$ 
           $a' \leftarrow 0$ 
7          while  $!otherSupport$  and  $!endD_i$  and  $a' \leq higher(localD_i)$  do
8            if  $localM_i[a']$  and  $C_{ij}(a', b)$  then  $otherSupport \leftarrow true$ 
            else
              if  $a' < higher(localD_i)$  then  $a' \leftarrow next(a', localD_i)$ 
              else  $endD_i \leftarrow true$ 
9          if  $!otherSupport$  then  $selected \leftarrow true$ 
10         if  $b < higher(D_j)$  then  $b \leftarrow next(b, D_j)$ 
          else  $endD_j \leftarrow true$ 
11        if  $selected$  then
12          addTo ( $sendSet_{acc}, (i, localM_i)$ )
13           $inform_{acc}[i] \leftarrow true$ 
end

```

returns with $emptySupport=false$, then b is the first support for (i, a) in D_j . We can add (i, a) to local support list S_{jb} and update $first$ for later use (lines 9 and 10). Otherwise, (i, a) is non-viable. A call to $localDUpdate$ (line 6) removes this label and test for domain wipe-out. Then the deletion is stored in both $localList$ and $extList$ (lines 7,8).

At the end of this initialization step, the *main* procedure processes deletions by calling *processLists* (see line 11). This procedure first, uses *localList* for the local propagation of deletions. This is done by calling the *deletionProcessing* procedure (see algo. 4) for each deleted label (j, e) . This algorithm, removes labels (i, a) from S_{je} , then for each viable one ($a \in localD_i$), a call to *nextSupport* looks for a new support (j, c) greater than (j, e) . If such a new support is found, (i, a) is added to S_{jc} . Otherwise, the procedure updates $localD_i$ and stores (i, a) in both *localList* and *extList* (lines 2,3,4).

After these local propagations, *processLists* considers each locally deleted label from *extList*. For each label (j, e) , a call to *selectiveSend* allows the agent to find which acquaintances to inform.

The *selectiveSend* procedure implements the previously detailed minimizing. While considering the deleted label (i, a) , it finds for each non local variable j such that C_{ij} exists, the acquaintance acc that owns j (see line 2). Then, if acc has not be previously selected for $localD_i$ transmission ($inform_{acc}[i] = false$), it considers labels (j, b) supported by (i, a) . Otherwise it is useless to check $localD_i$ transmission to acc .

We start the exploration of these labels by checking the current support of (i, a) ($first_j[i][a]$) since it is useless to consider previous labels (line 4). Two booleans are used, $endD_j$ is used to stop checking against j values, $selected$ is set to stop checking against acc .

Among j values, we can restrict to viable ones ($linkedM_j[b]=true$) (line 6). Then if $C_{ij}(a, b)$ holds *true*, the procedure looks for another possible support (i, a') . If there is no (i, a') label such that $C_{ij}(a', b)$, acc is selected for message transmission (line 9). We add the pair $(i, localM_i)$ in $sendSet_{acc}$. This inclusion of the current domain of i allows the transmission of previously non informed labels deletions. Of course this can lead to long messages but sizes are not prohibitives comparing to start-up/reception costs² (see table I).

After these computations, *processLists* respectively addresses its *sendSet* set to each acquaintance (line 7).

The second step of the *main* procedure allows agents to revise their knowledge according to incoming messages (line 12). There are two kinds. The *stop* messages which stop the processing. These messages can result from both a domain wipe-out (see *localDUpdate* primitive) or can be addressed by an extra agent, called *System*. This agent is in charge of global termination detection. In our system, we use a global state detection algorithm [2]. So global termination occurs with a domain wipe-out or when the p agents are waiting for a message and when no message transits in the communication network.

The processing of incoming deletions messages (*deletedLabels:set*) starts with line 14. Remember that agents receive a set of pairs $(j, localD_j)$. They temporary store these pairs in (j, M_j) . Then by comparing their local beliefs stored in $linkedM_j$ and the new reported situation they can detect new deletions (line 16). Agents store these new deletions and add them in their own *localList*. After considering the whole incoming *set*, a call to *processLists* allows the local propagation of externals deletions.

² In [7] we present discrete relaxation with both optimal number and optimal size of messages.

6. Analysis

6.1. COMPLEXITIES

The worst time execution of a distributed algorithm occurs when it proceeds with a sequential behavior. With our DisAC-9, this occurs when only one value is deleted at a time. This leads to nd successive deletions between $p = n$ processes. In $O(nd^2)$ our *selectiveSend* procedure allows the detection of the unique outcome of each deletion. While receiving a deletion message each agent updates *linkedM*[j] in $O(d)$, then *processLists* leads to one call to *deletionProcessing*. The S_{je} list contains exactly one label and *nextSupport* uses $O(d)$ steps in search of a new support. So each deletion conducts to $O(nd^2 + d)$ constant time operations. Since there are $O(nd)$ successive deletions, the overall time complexity is $O(n^2d^3)$ with exactly nd point-to-point messages.

With n variables, each one using d values, each agent data structure uses $O(nd)$ space. Since there are n agents in the worst case, the total space requirement is $O(n^2d)$.

6.2. CORRECTNESS AND MESSAGE PASSING OPTIMALITY

We present here the main steps of the complete proofs (see [7] for details).

For correctness we must ensure at the termination,

- (i) $\forall a \in localD_i, (i, a) \in Dmax-AC$, this can be achieved by considering the deletions made by the algorithm.
- (ii) $\forall (i, a) \in Dmax-AC, a \in localD_i$, suppose that a current deletion (j, b) is in $Dmax-AC$, then, an inductive proof which supposes that all the previous deletions are out of $Dmax-AC$, gives (j, b) out of $Dmax-AC$.

For optimality, we must show that,

- (i) each *deletedLabels* message induces at least a label deletion for the receiver. This is exactly what is checked by the *selectiveSend* procedure.
- (ii) each label deletion made during the interactions step is the outcome of an incoming *deletedLabels* message. We can observe that during the second step, deletions are made by *deletionProcessing* which is called by *processLists* which is always called for the processing of a *deletedLabels* message.

6.3. TERMINATION

Termination is detected by the System agent, this agent performs a classical global detection state algorithm [2]. In our system, it detects

a situation where each agent is waiting for an incoming message and no message transits between agents. Domain wipe-out is another case of algorithm termination (see *localDUpdate*).

The cost of this detection is directly related to the distributed algorithm message passing complexity. An optimal method like DisAC-9 has the lowest possible cost (see [7]).

7. Experiments

We have made experiments on an IBM *sp2* supercomputer³ (*C++* language and *MPI* message passing library [15]). We have used two algorithms against DisAC-9. First the sequential AC-6 to assess distribution benefits. Then to strictly evaluate our message minimizing we have built a modified version of our algorithm called DisAC-6⁴ which does not use our minimizing method:

- each deleted label (i, a) is selected for diffusion to *acc* if this agent owns a variable j such that C_{ij} exists
- *sendSet* are composed of (j, a) labels so reception of deletion messages is straight forward ($linkedM[j][a] \leftarrow false$ and *localList* inclusion).

This simplified implementation has a worst time complexity $O(nd(n+d))$ requiring n^2d messages. Since all deleted labels are transmitted to acquaintances, it also allows comparisons with previous works (see section 3). Nevertheless, this new algorithm is local since transmissions occur with respect to the CN. Moreover, it uses incremental computations. According to that, we think that DisAC-6 should outperform previous works.

7.1. HARD RANDOM PROBLEMS

We have generated random problems with 90 variables, domains sizes set to 20, and connectivity p_1 set to 0.2, 0.5 and 0.8. Tightness p_2 varies from 0.1 to 0.9 with a small increment on the phase transition region [6]. For the distributed methods we set $p=90$ which is the finest granularity in our framework since each agent owns one variable and $p=15$ (each agent owns 6 variables). For each connectivity/tightness pair we have generated 100 randoms instances.

³ Message passing latency is $35\mu s$, bandwidth is $90Mo/s$, each CPU is a *PowerPC 120Mhz*.

⁴ DisAC-6 is the first distributed implementation of AC-6.




Figure 2. Median time in seconds (the second figure is a zoomed view)

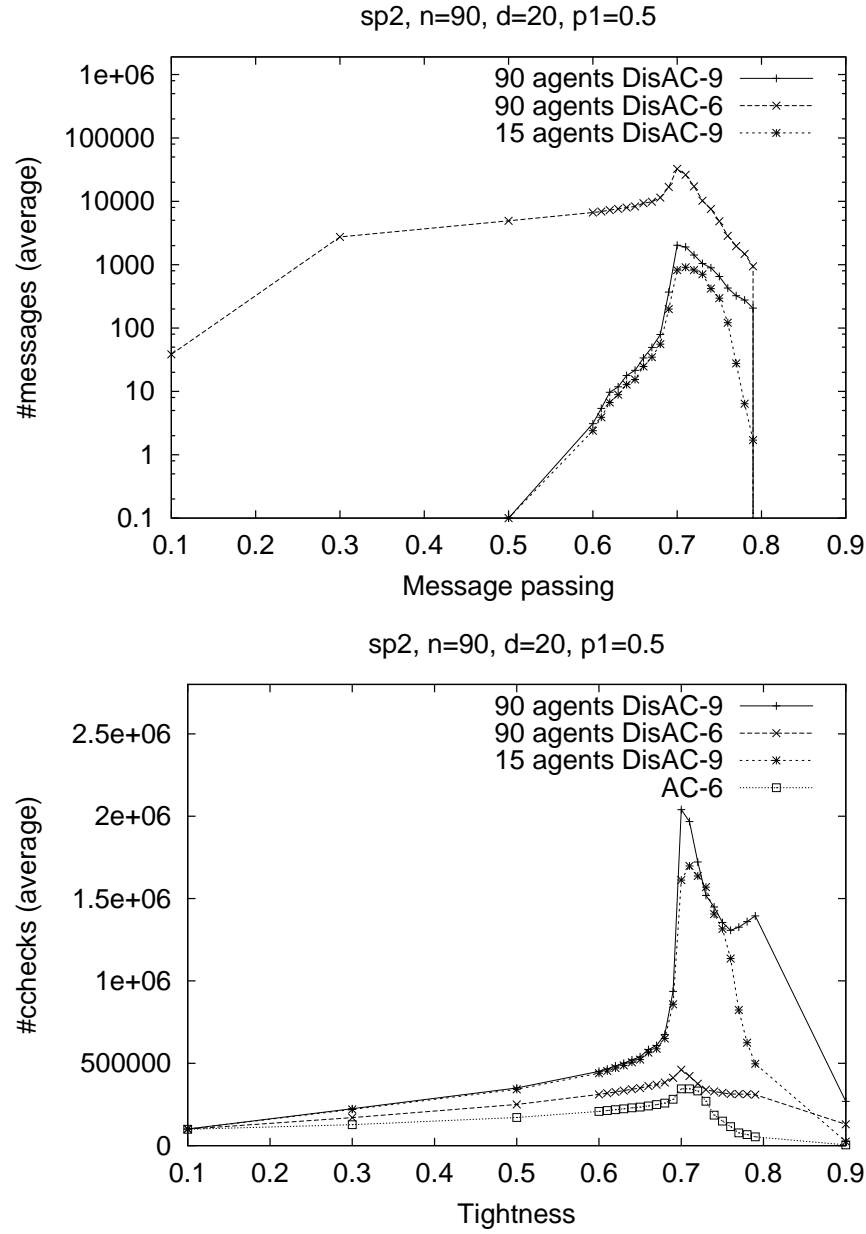


Figure 3. Message passing and constraints checks

We detail results with $p_1 = 0.5$. Figure 2 presents median time in seconds for the three algorithms. These results show the efficiency of DisAC-9. With 90 agents, our procedure complexity peak is 1.06 seconds, while DisAC-6 needs up to 22.23 seconds.

Figure 3, presents the amount of *deletedLabels* messages processed⁵ in the whole system. DisAC-9 process up to 2029 messages while the second method process up to 32493 messages. Interestingly we can observe that in DisAC-9 messages transmissions start at $p_2 = 0.5$. Of course this reduction has a cost in constraint checks (see figure 3), but it is limited to 4.41 times.

AC-6 time peak is 0.6 second while DisAC-9 with 15 agents uses up to 0.17 second. For $p_2 \in \{0.71, 0.72, 0.73\}$, even DisAC-9 with 90 agents is faster than AC-6.

Table III. Random problems computational time peaks

p_1	time			#msg			#ccks		
	0.2	0.5	0.8	0.2	0.5	0.8	0.2	0.5	0.8
DisAC-6 90 ag.	10.52s	22.23s	-	13660.20	32493	-	206243	460272	-
DisAC-9 90 ag.	0.79s	1.06s	1.50s	1663	2029	2431	761619	2.03e6	3.50e6
ratio	13.23	21	-	8.20	16	-	0.27	0.23	-
AC-6	0.28s	0.6s	0.91s	-	-	-	171370	344953	501970
DisAC-9 15 a.	0.09s	0.17s	0.25s	865.4	908.8	823.5	660964	1.69e6	2.60e6
ratio	3.06	3.53	3.64	-	-	-	0.26	0.20	0.19

Table III presents computational time peaks for various connectivities. The first part of the table gives results with $p = 90$. For $p_1 = 0.8$ and $p_2 \geq 0.63$, DisAC-6 exceeds sp2 message passing capacities⁶, so comparisons are not available with this connectivity parameter.

Interestingly, savings are increasing with p_1 . Message passing is reduced up to 16 times, CPU time is reduced up to 21 times. DisAC-9 uses up to 4.41 times more checks to limit messages. We can explain this limited overhead if we remark that useless messages induce local checks for the receiver. These checks are used to update support information.

⁵ We focus on processed messages since when an agent is aware of the instance inconsistency it stops the processing of messages. However when looking at the whole amount of *deletedLabels* messages, the present analysis remains exact (see [7] for the distinction between “processed” and “sent” messages).

⁶ Overflow in message buffering capacities.

The second part of the table compares fifteen DisAC-9 agents against AC-6. Again, speed-up increases with connectivity. It ranges from 3.06 to 3.64.

As shown in [13] discrete relaxation is P-complete. This avoid linear speed-up but our results show that a controlled distributed behavior can significantly enhance the processing of this problem.

7.2. THE ZEBRA PROBLEM

This usual benchmark has 25 variables, 5 labels per variable ($n = 25, d = 5$) and 61 binary constraints. In our experiments we have tried $p \in \{2, 3, 4, 5, 10, 25\}$ which brings 11, 16, 27, 11, 43, 61 inter-agents constraints. For each granularity we have made one hundred runs.

Table IV. The zebra problem

Algorithm	$p =$	1	2	3	4	5	10	25
AC-6	time	0.004s	-	-	-	-	-	-
	#msg	-	-	-	-	-	-	-
	#cks	1161	-	-	-	-	-	-
DisAC-9	time	-	0.0052s	0.07s	0.0104s	0.0102s	0.0176s	0.0524s
	#msg	-	6	17	21.18	19.3	43.56	64.02
	#cks	-	1463	1667.18	1815.58	1391.5	2444.58	3098.82
DisAC-6	time	-	0.008s	0.0092s	0.0096s	0.011s	0.0328s	0.1616s
	#msg	-	8	34	44.42	36.84	135.4	321
	#cks	-	1248	1298.5	1290.06	1251.62	1347.24	1356

Table IV shows our results. Due to the small size of the problem, distributed methods do not provide any speed-up over the central algorithm. But speed-up is not our only goal since we try to propose efficient methods for distributed problem solving. We can observe that the methods are strongly dependent of the number of inter-agents constraints. As expected, DisAC-9 is faster (more than 3 times) than the other algorithm and particularly for finest granularities. Message passing is reduced up to 5 times, checks are increased up to 2.28 times.

We can illustrate here the overhead implied by the distributed framework. DisAC-6 with 2 agents uses on average 1248 constraint checks while AC-6 uses 1161. This overhead comes from delayed informations in the system (see [7]).

7.3. THE PHASE TRANSITION OF DISTRIBUTED ARC-CONSISTENCY

Figure 2 shows some surprising new results. The time peaks of the methods are not located at the same tightness value. For the sequential AC-6, it occurs for $p_2 = 0.71$, for the methods using 90 agents, the time peak is located at $p_2 = 0.69$ while with 15 agents, the peak occurs at $p_2 = 0.7$. Our data results show 0%, 11%, 88% and 100% of inconsistent instances respectively at $p_2 = 0.68, 0.69, 0.7$ and 0.71 .




Figure 4. DCSP AC complexities peaks characterization

To provide a general explanation of this phenomenon, let us consider AC-6. It checks the set of arcs in a sequential way. So, if an instance is inconsistent according to the last arc (i, j) , the detection is very expensive. While with our distributed methods using n agents, the two agents respectively owning i and j rapidly detect inconsistency and immediately broadcast the end of the processing. With intermediate granularities, the previous reasoning occurs. But since the methods own several variables and consider their arcs successively, they are less efficient for these detections.

We characterize these observations in figure 4. The sequential processing peak is located (on average) at 50% of consistent instances [6]. Distributed processing time peak (with $p = n$) appears with the first inconsistent instances. Intermediate granularities peaks appear in between.

The above explains why inconsistent instances are relatively easier for distributed methods (see [7] for a comprehensive analysis). We can now interrogate ourselves on the exploitation of this phenomenon. In

the sequential framework, the understanding of phase transition was used for at least three things. Firstly, it allows researchers to focus their experiments on hard instances; secondly, it is now possible to predict the complexity of a computation by considering some parameters of an instance; finally, it gives rise to new algorithms capable of using some knowledge on these computational cliffs [3, 12].

In the distributed framework, we think that the understanding of phase transition could be even more fruitful. We showed here that some seemingly innocuous changes to the problem granularity can have drastic effects on the overall performance. An instance could meet a computational peak at a particular solving granularity and avoid it by using a different grain.

Our results suggest that a distributed system could use its knowledge about the phase transition structure to dynamically rearrange the distribution of subproblems. This reallocation could allow the system to avoid time peaks. In order to use the system resources efficiently, this structural transformation should be made before solving the problem. A negotiation stage between agents could be used here.

8. Conclusion

We have provided DisAC-9, a coarsely grained distributed algorithm performing the arc-filtering of binary constraint networks. Our method is optimal in message passing operations which are critical in distributed systems. DisAC-9 allows efficient processing of DCSP and with coarser granularity, it can speed-up the central AC-6.

Experiments were made on a state-of-the-art supercomputer. This is not a common framework, particularly if we consider Internet applications for DCSPs. We choose this framework for testing DisAC-9 in a hard way, on a computer with very fast communication. In a LAN, observed speed-up are larger. The reason we found is that message passing performance decrease more rapidly with the amount of exchanged messages.

Surprisingly, the experiments allow a characterization of computational peaks in distributed arc-consistency problems. The hardest instances do not occur at the same value of tightness for CSP and for DCSP. Inconsistent problems are relatively easier for distributed procedures and the location of their time peaks depends of their granularities. As far as we known this is the first identification of the special nature of phase transition in DCSP. This behavior requires a full theoretical study [7]. A possible consequence is to increase efficiency by changing the granularity.

As a practical extension of this work, we are combining it with a distributed search procedure. Similar combinations have brought interesting results in the centralized framework [19] and should be studied for distributed CSPs. This cooperation may prune inconsistent parts of a distributed search space and by the way could save critical message passing operations.

Acknowledgements

The author wishes to thank the French CINES for supercomputing facilities and the anonymous reviewers for their comments. Special thanks are addressed to Philippe Gérard and Simon De Givry.

References

1. Bessière, C.: 1994, 'Arc-consistency and arc-consistency again'. *Artificial Intelligence* **65**(1), 179–190.
2. Chandy, K. M. and L. Lamport: 1985, 'Distributed Snapshots: Determining Global States of Distributed Systems'. *TOCS* **3**(1), 63–75.
3. Clearwater, S. H. and T. Hogg: 1994, 'Exploiting Problem Structure in Genetic Algorithms'. In: *Proceedings of AAAI*, pp. 1310–1315.
4. Cooper, P. R. and M. J. Swain: 1992, 'Arc consistency: Parallelism and Domain Dependence'. *AI* **58**(1–3), 207–235.
5. Culler, D. E., L. T. Liu, R. P. Martin, and C. Yoshikawa: 1996, 'LogP Performance Assessment of Fast Network Interfaces'. *IEEE Micro*.
6. Gent, I., E. M. Intyre, P. Prosser, and T. Walsh: 1997, 'The Constrainedness of Arc Consistency'. In: *Principles and Practice of Constraint Programming*, pp. 327–340.
7. Hamadi, Y.: 1999, 'Traitement des problèmes de satisfaction de contraintes distribués'. Ph.D. thesis, Université Montpellier II. (in french).
8. Hamadi, Y.: 2001a, 'Conflicting Agents in Distributed Search'. Technical Report HPL-2001-222, HP Labs.
9. Hamadi, Y.: 2001b, 'Interleaved Backtracking in Distributed Constraint Networks'. In: IEEE (ed.): *13th International Conference on Tools with Artificial Intelligence*, p. (to appear).
10. Hamadi, Y., C. Bessière, and J. Quinqueton: 1998, 'Backtracking in Distributed Constraint Networks'. In: *ECAI*, pp. 219–223.
11. Hamadi, Y. and D. Merceron: 1997, 'Reconfigurable Architectures: A New Vision for Optimization Problems'. In: G. Smolka (ed.): *Proceedings of the 3rd International Conference on Principle and Practice of Constraint Programming CP97*, Vol. 1330 of *LNCS*. Linz, AUSTRIA, pp. 209–221.
12. Hogg, T.: 1995, 'Exploiting Problem Structure as a Search Heuristic'. Unpublished.
13. Kasif, S.: 1990, 'On the Parallel Complexity of Discrete Relaxation in Constraint Satisfaction Networks'. *AI* **45**(3), 275–286.

14. Mohr, R. and T. C. Henderson: 1986, 'Arc and Path Consistency Revisited'. *Artificial Intelligence* **28**, 225–233.
15. MPIF, M. P. I. F.: 1994, 'MPI: A Message-Passing Interface Standard'. *Int. Journal of Supercomputer Applications* **8**(3/4).
16. Naor, M. and L. Stockmeyer: 1995, 'What Can be Computed Locally ?'. *SIAM J. Comput.* **24**(6), 1259–1277.
17. Nguyen, T. and Y. Deville: 1995, 'A Distributed Arc-Consistency Algorithm'. In: *First Int. Workshop on concurrent Constraint Satisfaction*.
18. Prosser, P., C. Conway, and C. Muller: 1992, 'A distributed constraint maintenance system'. In: *Proc. of the 12th Int. Conf. on AI*. pp. 221–231.
19. Sabin, D. and E. C. Freuder: 1994, 'Contradicting Conventional Wisdom in Constraint Satisfaction'. In: *ECAI*. pp. 125–129.
20. Samal, A. and T. Henderson: 1987, 'Parallel consistent labeling Algorithms'. *Int. J. Parallel Program.* **16**, 341–364.
21. Simonis, H.: 1996, 'A problem classification scheme for finite domain constraint solving'. In: *Proc. of the CP'96 workshop on Constraint Programming Applications: An Inventory and Taxonomy*. pp. 1–26.
22. Yokoo, M. and K. Hirayama: 1996, 'Distributed Breakout Algorithm for Solving Distributed Constraint Satisfaction Problems'. In: *ICMAS*. pp. 401–408.