

Agile Asynchronous Backtracking for Distributed Constraint Satisfaction Problems

Christian Bessiere¹ El Houssine Bouyakhf²

¹LIRMM, University of Montpellier

Montpellier, France

bessiere@lirmm.fr bouyakhf@fsr.ac.ma

Younes Mechqrane² Mohamed Wahbi^{1,2}

²LIMIARF, University Mohammed V-Agdal

Rabat, Morocco

yemechqrane@spsm.ma wahbi@lirmm.fr

Abstract—Asynchronous Backtracking is the standard search procedure for distributed constraint reasoning. It requires a total ordering on the agents. All polynomial space algorithms proposed so far to improve Asynchronous Backtracking by reordering agents during search only allow a limited amount of reordering. In this paper, we propose Agile-ABT, a search procedure that is able to change the ordering of agents more than previous approaches. This is done via the original notion of termination value, a vector of stamps labelling the new orders exchanged by agents during search. In Agile-ABT, agents can reorder themselves as much as they want as long as the termination value decreases as the search progresses. Our experiments show the good performance of Agile-ABT when compared to other dynamic reordering techniques.

Keywords-constraint reasoning, distributed constraint solving

I. INTRODUCTION

Various application problems in distributed artificial intelligence are concerned with finding a consistent combination of agent actions (e.g., distributed resource allocation [5], sensor networks [1]). Such problems can be formalized as Distributed Constraint Satisfaction Problems (DisCSPs). DisCSPs are composed of agents, each owning its local constraint network. Variables in different agents are connected by constraints. Agents must assign values to their variables so that all constraints are satisfied. Several distributed algorithms for solving DisCSPs have been developed, among which Asynchronous Backtracking (ABT) is the central one [10, 2]. ABT is an asynchronous algorithm executed autonomously by each agent in the distributed problem. Agents do not have to wait for decisions of others but they are subject to a total (priority) order on them. Each agent tries to find an assignment satisfying the constraints with what is currently known from higher priority agents. When an agent assigns a value to its variable, the selected value is sent to lower priority agents. When no value is possible for a variable, the inconsistency is reported to higher agents in the form of a nogood. ABT computes a solution (or detects that no solution exists) in a finite time. The total order is static. Now, it is known from centralized CSPs that adapting the order of variables dynamically during search drastically fastens the search procedure.

Asynchronous Weak Commitment (AWC) dynamically reorders agents during search by moving the sender of a nogood higher in the order than the other agents in

the nogood [9]. But AWC requires exponential space for storing nogoods. Silaghi et al. (2001) tried to hybridize ABT with AWC. Abstract agents fulfill the reordering operation to guarantee a finite number of asynchronous reordering operations. In [6], the heuristic of the centralized dynamic backtracking was applied to ABT. However, in both studies, the improvement obtained on ABT was minor. Zivan and Meisels (2006) proposed Dynamic Ordering for Asynchronous Backtracking (ABTDO). When an agent assigns a value to its variable, ABTDO can reorder lower priority agents. A new kind of ordering heuristics for ABTDO is presented in [13]. In the best of those heuristics, the agent that generates a nogood is placed between the last and the second last agents in the nogood if its domain size is smaller than that of the agents it passes on the way up.

In this paper, we propose Agile-ABT, an asynchronous dynamic ordering algorithm that does not follow the standard restrictions in asynchronous backtracking algorithms. The order of agents appearing *before* the agent receiving a backtrack message can be changed with a great freedom while ensuring polynomial space complexity. Furthermore, that agent receiving the backtrack message, called the backtracking *target*, is not necessarily the agent with the lowest priority within the conflicting agents in the current order. The principle of Agile-ABT is built on termination values exchanged by agents during search. A termination value is a tuple of positive integers attached to an order. Each positive integer in the tuple represents the expected current domain size of the agent in that position in the order. Orders are changed by agents without any global control so that the termination value decreases lexicographically as the search progresses. Since, a domain size can never be negative, termination values cannot decrease indefinitely. An agent informs the others of a new order by sending them its new order and its new termination value. When an agent compares two contradictory orders, it keeps the order associated with the smallest termination value.

The rest of the paper is organized as follows. Section II recalls basic definitions. Section III describes the concepts needed to select new orders that decrease the termination value. We give the details of our algorithm in Section IV and we prove it in Section V. An experimental evaluation is given in Section VI. Section VII concludes the paper.

II. PRELIMINARIES

The Distributed Constraint Satisfaction Problem (DisCSP) has been formalized in [10] as a tuple $(\mathcal{A}, \mathcal{X}, \mathcal{D}, \mathcal{C})$, where \mathcal{A} is a set of agents, \mathcal{X} is a set of variables $\{x_1, \dots, x_n\}$, where each variable x_i is controlled by one agent in \mathcal{A} . $\mathcal{D} = \{D_1, \dots, D_n\}$ is a set of domains, where D_i is a finite set of values to which variable x_i may be assigned. The initial domain size of a variable x_i is denoted by d_i^0 . \mathcal{C} is a set of binary constraints that specify the combinations of values allowed for the two variables they involve. A constraint $c_{ik} \in \mathcal{C}$ between two variables x_i and x_k is a subset of the Cartesian product $c_{ik} \subseteq D_i \times D_k$.

For simplicity purposes, we consider a restricted version of DisCSP where each agent controls exactly one variable. We use the terms agent and variable interchangeably and we identify the agent ID with its variable index. All agents maintain their own counter, and increment it whenever they change their value. The current value of the counter tags each generated assignment.

Definition 1: An **assignment** for an agent $A_i \in \mathcal{A}$ is a tuple (x_i, v_i, t_i) , where v_i is a value from the domain of x_i and t_i is the tag value. When comparing two assignments, the most up to date is the one with the highest tag t_i . Two sets of assignments $\{(x_{i_1}, v_{i_1}, t_{i_1}), \dots, (x_{i_k}, v_{i_k}, t_{i_k})\}$ and $\{(x_{j_1}, v_{j_1}, t_{j_1}), \dots, (x_{j_q}, v_{j_q}, t_{j_q})\}$ are **coherent** if every common variable is assigned the same value in both sets.

A_i stores a unique order denoted by o_i . Agents appearing before A_i in o_i are the higher agents (predecessors) denoted by $Pred(A_i)$ and conversely the lower agents (successors) $Succ(A_i)$ are agents appearing after A_i .

Definition 2: The **AgentView** of an agent A_i is an array containing the most up to date assignments of $Pred(A_i)$.

Agents can infer inconsistent sets of assignments, called **nogoods**. A nogood can be represented as an implication. There are clearly many different ways of representing a given nogood as an implication. For example, $\neg[(x_1=v_1) \wedge \dots \wedge (x_k=v_k)]$ is logically equivalent to $[(x_2=v_2) \wedge \dots \wedge (x_k=v_k)] \rightarrow (x_1 \neq v_1)$. When a nogood is represented as an implication, the left hand side (lhs) and the right hand side (rhs) are defined from the position of \rightarrow . A nogood is **compatible** with an order o_i if all agents in $lhs(nogood)$ appear before $rhs(nogood)$ in o_i .

The current domain of x_i is the set of values $v \in D_i$ such that $x_i \neq v$ does not appear in any of the right hand sides of the nogoods stored by A_i . Each agent keeps only one nogood per removed value. The size of the current domain of A_i is denoted by d_i .

III. INTRODUCTORY MATERIAL

Before presenting Agile-ABT, we need to introduce new notions and to present some key subfunctions.

A. Reordering details

In order to allow agents to asynchronously propose new orders, they must be able to coherently decide which order to select. We propose that the priority between the different orders is based on *termination values*. Informally, if $o_i = [A_1, \dots, A_n]$ is the current order known by an agent A_i , then the tuple of domain sizes $[d_1, \dots, d_n]$ is the termination value of o_i on A_i . To build termination values, agents need to exchange *explanations*.

Definition 3: An **explanation** e_j is an expression $lhs(e_j) \rightarrow d_j$, where $lhs(e_j)$ is the conjunction of the left hand sides of all nogoods stored by A_j as justifications of value removals for x_j , and d_j is the number of values not pruned by nogoods in the domain of A_j . d_j is $rhs(e_j)$.

Each time an agent communicates its assignment to other agents (by sending them an **ok?** message, see Section IV), it inserts its explanation in the **ok?** message for allowing other agents to build their termination value.

The variables in the left hand side of an explanation e_j must precede the variable x_j in the order because the assignments of these variables have been used to determine the current domain of x_j . An explanation e_j induces ordering constraints, called *safety conditions* in [3].

Definition 4: A **safety condition** is an assertion $x_k \prec x_j$. Given an explanation e_j , $S(e_j)$ is the set of safety conditions induced by e_j , where $S(e_j) = \{(x_k \prec x_j) \mid x_k \in lhs(e_j)\}$.

An explanation e_j is **compatible** with an order o if all variables in $lhs(e_j)$ appear before x_j in o . Each agent A_i stores a set E_i of explanations sent by other agents. During search, E_i is updated to remove explanations that are no longer valid.

Definition 5: An explanation e_j in E_i is **valid** on agent A_i if it is compatible with the current order o_i and $lhs(e_j)$ is coherent with the AgentView of A_i .

When E_i contains an explanation e_j associated with A_j , A_i uses this explanation to justify the size of the current domain of A_j . Otherwise, A_i assumes that the size of the current domain of A_j is equal to d_j^0 . The termination value depends on the order and the set of explanations.

Definition 6: Let E_i be the set of explanations stored by A_i , o be an order on the agents such that every explanation in E_i is compatible with o , and $o(k)$ be such that $A_{o(k)}$ is the k th agent in o . The **termination value** $TV(E_i, o)$ is the tuple $[tv^1, \dots, tv^n]$, where $tv^k = rhs(e_{o(k)})$ if $e_{o(k)} \in E_i$, otherwise, $tv^k = d_{o(k)}^0$.

In Agile-ABT, an order is always associated with a termination value. When comparing two orders the *strongest* order is that associated with the lexicographically *smallest* termination value. In case of ties, we use the lexicographic order on agents IDs, the smaller being the stronger.

B. The backtracking target

When all the values of an agent A_i are ruled out by nogoods, these nogoods are resolved, producing a new

nogood ng . ng is the conjunction of lhs of all nogoods stored by A_i . If ng is empty, then the inconsistency is proved. Otherwise, one of the conflicting agents must change its value. In standard ABT, the backtracking target (i.e., the agent that must change its value) is the agent with the lowest priority. Agile-ABT overcomes this restriction by allowing A_i to select with great freedom the backtracking target. When a new nogood ng is produced by resolution, the only condition to choose a variable x_k as the backtracking target is to find an order o' such that $TV(up_E_i, o')$ is lexicographically smaller than the termination value associated with the current order of A_i . up_E_i is obtained by updating E_i after placing x_k in $rhs(ng)$.

In [3], Ginsberg and McAllester proposed Partial Order Dynamic Backtracking (POB), a polynomial space algorithm for centralised CSPs that attempted to address the rigidity of dynamic backtracking. POB maintains a set of ordering constraints (aka. safety conditions) on the variables. These ordering constraints imply only a partial order on the variables. This provides flexibility in the reordering of variables in a nogood. Agile-ABT has some similarities with POB because Agile-ABT also maintains a set of safety conditions (induced by explanations). However, the set of safety conditions maintained by Agile-ABT allows more total orderings than the set of safety conditions maintained by POB. In addition, whenever a new nogood is generated by POB, the target of this nogood must be selected such that the safety conditions induced by the new nogood satisfy all existing safety conditions. On the contrary, Agile-ABT allows discarding explanations, and thus, relaxing some of the safety conditions. These two points give Agile-ABT more flexibility in choosing the backtracking target.

Function `updateExplanations` takes as arguments the set E_i , the nogood ng and the variable x_k to place in the rhs of ng . `updateExplanations` removes all explanations that are no longer coherent after placing x_k in the right hand side of ng . It updates the explanation of agent A_k stored in A_i and it returns a set of explanations up_E_i .

```
function updateExplanations( $E_i, ng, x_k$ )
1.  $up\_E_i \leftarrow E_i$ ; setRhs( $ng, x_k$ ) ;
2. remove each  $e_j \in up\_E_i$  such that  $x_k \in lhs(e_j)$  ;
3. if ( $e_k \notin up\_E_i$ ) then setLhs( $e_k, \emptyset$ ); setRhs( $e_k, d_k^0$ ); add  $e_k$  to  $up\_E_i$ ;
4. setLhs( $e'_k, lhs(e_k) \cup lhs(ng)$ ); setRhs( $e'_k, rhs(e_k) - 1$ ) ;
5. replace  $e_k$  by  $e'_k$  ;
6. return  $up\_E_i$ ;
```

This function does not create cycles in the set of safety conditions $S(up_E_i)$ if $S(E_i)$ is acyclic. Indeed, all the explanations added or removed from $S(E_i)$ to obtain $S(up_E_i)$ contain x_k . Hence, if $S(up_E_i)$ contains cycles, all these cycles should contain x_k . However, there does not exist any safety condition of the form $x_k \prec x_j$ in $S(up_E_i)$ because all of these explanations have been removed in line 2. Thus, $S(up_E_i)$ cannot be cyclic. As we will show in

Section IV, the updates performed by A_i ensure that $S(E_i)$ always remains acyclic. As a result, $S(up_E_i)$ is acyclic as well, and it can be represented by a directed acyclic graph $G = (N, U)$ where $N = \{x_1, \dots, x_n\}$. An edge $(j, l) \in U$ if $(x_j \prec x_l) \in S(up_E_i)$. Any topological sort of G is an order compatible with the safety conditions induced by up_E_i .

C. Decreasing termination values

Termination of Agile-ABT is based on the fact that the termination values associated with orders selected by agents decrease as search progresses. To speed up the search, Agile-ABT is written so that agents decrease termination values whenever they can. When an agent resolves its nogoods, it checks whether it can find a new order of agents such that the associated termination value is smaller than that of the current order. If so, the agent will replace its current order and termination value by those just computed, and will inform all other agents.

Assume that after resolving its nogoods, an agent A_i , decides to place x_k in the rhs of the nogood (ng) produced by the resolution and let $up_E_i = \text{updateExplanations}(E_i, ng, x_k)$. The function `computeOrder` takes as parameter the set up_E_i and returns an order up_o compatible with the partial ordering induced by up_E_i . Let G be the acyclic directed graph associated with up_E_i . The function `computeOrder` works by determining, at each iteration p , the set $Roots$ of vertices that have no predecessor. As we aim at minimizing the termination value, function `computeOrder` selects the vertex x_j in $Roots$ that has the smallest domain size. This vertex is placed at the p th position and removed from G . Finally, p is incremented and all outgoing edges from x_j are removed from G .

```
function computeOrder( $up\_E_i$ )
7.  $G = (N, U)$  is the acyclic graph associated to  $up\_E_i$  ;
8.  $p \leftarrow 1$ ;  $o$  is an array of length  $n$  ;
9. while  $G \neq \emptyset$  do
10.    $Roots \leftarrow \{x_j \in N \mid x_j \text{ has no incoming edges}\}$  ;
11.    $o[p] \leftarrow x_j$  such that  $d_j = \min\{d_k \mid x_k \in Roots\}$  ;
12.   remove  $x_j$  from  $G$ ;  $p \leftarrow p + 1$  ;
13. return  $o$ ;
```

Having proposed an algorithm that determines an order with small termination value for a given backtracking target x_k , one needs to know how to choose this variable to obtain an order decreasing more the termination value. The function `chooseVariableOrder` iterates through all variables x_k included in the nogood, computes a new order and termination value with x_k as the target (lines 16–17), and stores the target and the associated order if it is the strongest order found so far (lines 18–19). Finally, the information corresponding to the strongest order is returned.

IV. THE ALGORITHM

Each agent keeps some amount of local information about the global search, namely an `AgentView`, a `NogoodStore`,

```

function chooseVariableOrder ( $E_i$ ,  $ng$ )
14.  $o' \leftarrow o_i$ ;  $TV' \leftarrow TV_i$ ;  $E' \leftarrow nil$ ;  $x' \leftarrow nil$ ;
15. for each  $x_k \in ng$  do
16.    $up\_E_i \leftarrow updateExplanations(E_i, ng, x_k)$  ;
17.    $up\_o \leftarrow computeOrder(up\_E_i)$ ;  $up\_TV \leftarrow TV(up\_E_i, up\_o)$  ;
18.   if ( $up\_TV$  is smaller than  $TV'$ ) then
19.      $x' \leftarrow x_k$ ;  $o' \leftarrow up\_o$ ;  $TV' \leftarrow up\_TV$ ;  $E' \leftarrow up\_E_i$  ;
20. return  $\langle x', o', TV', E' \rangle$ ;

```

a set of explanations (E_i), a current order (o_i) and a termination value (TV_i). Agile-ABT allows the following types of messages (where A_i is the sender):

- **ok?** message is sent by A_i to lower agents to ask whether a chosen value is acceptable. Besides the chosen value, the **ok?** message contains an explanation e_i which communicates the current domain size of A_i . An **ok?** message also contains the current order o_i and the current termination value TV_i stored by A_i .
- **ngd** message is sent by A_i when all its values are ruled out by its NogoodStore. This message contains a nogood, as well as o_i and TV_i .
- **order** message is sent to propose a new order. This message includes the order o_i proposed by A_i accompanied by the termination value TV_i .

Agile-ABT (Figure 1) is executed on every agent A_i . After initialization, each agent assigns a value and informs lower priority agents of its decision (CheckAgentView call, line 22) by sending **ok?** messages. Then, a loop considers the reception of the possible message types. If no message is traveling through the network, the state of quiescence is detected by a specialized algorithm [4], and a global solution is announced. The solution is given by the current variables’ assignments.

When an agent A_i receives a message (of any type), it checks if the order included in the received message is stronger than its current order o_i (CheckOrder call, lines 28, 32 and 34). If it is the case, A_i replaces o_i and TV_i by those newly received (line 42). The nogoods and explanations that are no longer compatible with o_i are removed to ensure that $S(E_i)$ remains acyclic (line 43).

If the message was an **ok?** message, the AgentView of A_i is updated to include the new assignments (UpdateAgentView call, line 29). Beside the assignment of the sender, A_i also takes newer assignments contained in the left hand side of the explanation included in the received **ok?** message to update its AgentView. Afterwards, the nogoods and the explanations that are no longer coherent with AgentView are removed (UpdateAgentView line 53). Then, if the explanation in the received message is valid, A_i updates the set of explanations by storing the newly received explanation. Next, A_i calls the procedure CheckAgentView (line 31).

When receiving an **order** message, A_i processes the new order (CheckOrder) and calls CheckAgentView (line 33).

When A_i receives a **ngd** message, it calls CheckOrder and UpdateAgentView (lines 34 and 35). The nogood

```

procedure Agile-ABT ( )
21.  $t_i \leftarrow 0$ ;  $TV_i \leftarrow [\infty, \infty, \dots, \infty]$ ;  $end \leftarrow \text{false}$ ;  $v_i \leftarrow \text{empty}$ ;
22. CheckAgentView() ;
23. while ( $\neg end$ ) do
24.    $msg \leftarrow \text{getMsg}()$ ;
25.   switch ( $msg.type$ ) do
26.      $ok? : \text{ProcessInfo}(msg)$ ;  $ngd : \text{ResolveConflict}(msg)$ ;
27.      $order : \text{ProcessOrder}(msg)$ ;  $stp : end \leftarrow \text{true}$ ;
procedure ProcessInfo( $msg$ )
28. CheckOrder ( $msg.Order, msg.TV$ ) ;
29. UpdateAgentView ( $msg.Assig \cup \text{lhs}(msg.Exp)$ ) ;
30. if ( $msg.Exp$  is valid) then add ( $msg.Exp, E$ ) ;
31. CheckAgentView() ;
procedure ProcessOrder( $msg$ )
32. CheckOrder ( $msg.Order, msg.TV$ ) ;
33. CheckAgentView() ;
procedure ResolveConflict( $msg$ )
34. CheckOrder ( $msg.Order, msg.TV$ ) ;
35. UpdateAgentView ( $msg.Assig \cup \text{lhs}(msg.Nogood)$ ) ;
36. if (Coherent ( $msg.Nogood, AgentView \cup x_i=v_i$ )) then
37.   if (Compatible ( $msg.Nogood, o_i$ )) then
38.     add ( $msg.Nogood, NogoodStore$ );  $v_i \leftarrow \text{empty}$ ;
39.     CheckAgentView();
40.   else if ( $\text{rhs}(msg.Nogood)=v_i$ ) then
41.     sendMsg:ok?( $v_i, e_i, o_i, TV_i$ ) to  $msg.Sender$  ;
procedure CheckOrder ( $o, TV$ )
42. if ( $o$  is stronger than  $o_i$ ) then  $o_i \leftarrow o$ ;  $TV_i \leftarrow TV$ ;
43. remove nogoods and explanations incompatible with  $o_i$ ;
procedure CheckAgentView()
44. if ( $\neg \text{Consistent}(v_i, AgentView)$ ) then
45.    $v_i \leftarrow \text{ChooseValue}()$  ;
46.   if ( $v_i$ ) then sendMsg:ok?( $v_i, e_i, o_i, TV_i$ ) to  $Succ(A_i)$ ;
47.   else Backtrack();
48. else if ( $o_i$  was modified) then
49.   sendMsg:ok?( $v_i, e_i, o_i, TV_i$ ) to  $Succ(A_i)$ ;
procedure UpdateAgentView( $Assignments$ )
50. for each  $var \in Assignments$  do
51.   if ( $Assignments[var].l > AgentView[var].l$ ) then
52.      $AgentView[var] \leftarrow Assignments[var]$ ;
53. remove nogoods and explanations incoherent with AgentView;
procedure Backtrack()
54.  $ng \leftarrow \text{solve}(NogoodStore)$  ;
55. if ( $ng = \text{empty}$ ) then
56.    $end \leftarrow \text{true}$ ; sendMsg:stp( $system$ );
57.    $\langle x_k, o', TV', E' \rangle \leftarrow \text{chooseVariableOrder}(E_i, ng)$  ;
58. if ( $TV'$  is smaller than  $TV_i$ ) then
59.    $TV_i \leftarrow TV'$ ;  $o_i \leftarrow o'$ ;  $E_i \leftarrow E'$  ;
60.    $\text{setRhs}(ng, x_k)$ ;
61.   sendMsg:ngd( $ng, o_i, TV_i$ ) to  $A_k$  ;
62.   remove  $e_i$  from  $E_i$  ;
63.   broadcastMsg:order( $o_i, TV_i$ ) ;
64. else
65.    $\text{setRhs}(ng, x_k)$ ;
66.   sendMsg:ngd( $ng, o_i, TV_i$ ) to  $A_k$  ;
67.   UpdateAgentView ( $x_k \leftarrow \text{unknown}$ ) ;
68. CheckAgentView() ;
function ChooseValue()
69. for each ( $v \in D_i$  not eliminated by NogoodStore) do
70.   if ( $\exists x_j \in AgentView$  such that  $\neg \text{Consistent}(v, x_j)$ ) then
71.     add ( $x_j=v_j \Rightarrow x_i \neq v$ , NogoodStore) ;
72. if ( $D_i = \emptyset$ ) then return ( $\text{empty}$ );
73. else  $t_i \leftarrow t_i+1$  return ( $v$ );
/*  $v \in D_i$  */

```

Figure 1. The Agile-ABT algorithm.

contained in the message is accepted if it is coherent with the AgentView and the assignment of x_i and compatible with the current order of A_i . Otherwise, the nogood is discarded and an **ok?** message is sent to the sender as in ABT (lines 40 and 41). When the nogood is accepted, it is stored, acting as justification for removing the current value of A_i (line 38). A new value consistent with the AgentView is searched (CheckAgentView call, line 39).

The procedure CheckAgentView checks if the current value v_i is consistent with the AgentView. If v_i is consistent, A_i checks if o_i was modified (line 48). If so, A_i must send its assignment to lower priority agents through **ok?** messages. If v_i is not consistent with its AgentView, A_i tries to find a consistent value (ChooseValue call, line 45). In this process, some values of A_i may appear as inconsistent. In this case, the nogoods justifying their removal are added to the NogoodStore (line 71 of function ChooseValue). If a new consistent value is found, an explanation e_i is built and the new assignment is notified to the lower priority agents of A_i through **ok?** messages (line 46). Otherwise, every value of A_i is forbidden by the NogoodStore and A_i has to backtrack (Backtrack call, line 47).

In procedure Backtrack, A_i resolves its nogoods, deriving a new nogood (ng). If ng is empty, the problem has no solution. A_i terminates execution after sending a **stp** message (line 56). Otherwise, one of the agents included in ng must change its value. The function chooseVariableOrder selects the variable to be changed (x_k) and a new order (o') such that the new termination value TV' is as small as possible. If TV' is smaller than that stored by A_i , the current order and the current termination value are replaced by o' and TV' and A_i updates its explanations by that returned by chooseVariableOrder (line 59). Then, a **ngd** message is sent to the agent A_k owner of x_k (line 61). e_k is removed from E_i since A_k will probably change its explanation after receiving the nogood (line 62). Afterwards, A_i sends an **order** message to all other agents (line 63). When TV' is not smaller than the current termination value, A_i cannot propose a new order and the variable to be changed (x_k) is the variable that has the lowest priority according to the current order of A_i (lines 65 and 66). Next, the assignment of x_k (the target of the backtrack) is removed from the AgentView of A_i (line 67). Finally, the search is continued by calling the procedure CheckAgentView (line 68).

V. CORRECTNESS AND COMPLEXITY

We demonstrate that Agile-ABT is sound, complete and terminates, with a polynomial space complexity.

Theorem 1: Agile-ABT requires $O(nd + n^2)$ space per agent.

Proof: The size of nogoods, explanations, termination values, and orderings, is bounded by n . Now, on each agent, Agile-ABT only stores one nogood per value, one

explanation per agent, one termination value and one ordering. Thus, the space complexity of Agile-ABT is in $O(nd + n^2 + n + n) = O(nd + n^2)$ on each agent. ■

Theorem 2: Agile-ABT is sound.

Proof: (Sketch) When the state of quiescence is reached, all agents necessarily know the order o that is the strongest ever computed. In addition, all agents know the most up to date assignments of all their predecessors in o . Thus, any constraint c_{ik} between agents A_i and A_k has been successfully checked by the agent with lowest priority in o . Otherwise that agent would have tried to change its value and would have either sent an **ok?** or a **ngd** message, breaking the quiescence. ■

Theorem 3: Agile-ABT is complete.

Proof: All nogoods are generated by logical inferences from existing constraints. Therefore, an empty nogood cannot be inferred if a solution exists. ■

The proof of termination is built on lemmas 1 and 2.

Lemma 1: For every agent A_i , while no solution is found and the inconsistency of the problem is not proved, the termination value stored by A_i decreases after a finite amount of time.

Proof: (Sketch) If an agent gets stuck a sufficiently long time with the same termination value, all agents will eventually have that same termination value and the order o to which it was attached. At this point, Agile-ABT works exactly like ABT, which is complete and terminates. Thus, either a solution is found or the first agent in the current order, $A_{o(1)}$, will receive a nogood with empty lhs that prunes one of its remaining values. As soon as $A_{o(1)}$ has sent its new domain size (smaller than tv^1) to the lower agents in o , any backtracking agent will generate a smaller termination value. ■

Lemma 2: For any termination value $TV = [tv^1, \dots, tv^n]$ generated by an agent, we have $tv^j \geq 0, \forall j \in 1..n$

Proof: (Sketch) All explanations e_k stored by an agent A_i have $rhs(e_k) \geq 1$ because it represents the current domain size of A_k . Now, termination values are built with rhs of explanations. The only case where a tv^j can be zero (line 4) is when $A_{o(j)}$ is selected by A_i to be the backtracking target, and in such a case, the explanation $e_{o(j)}$ is removed just after sending the nogood to $A_{o(j)}$ (line 62). Hence, A_i never stores an explanation e_k with $rhs(e_k) = 0$ and cannot produce a termination value with a negative element. ■

Theorem 4: Agile-ABT terminates.

Proof: Direct from Lemmas 1 and 2. ■

VI. EXPERIMENTAL RESULTS

We compared Agile-ABT to ABT, ABTDO, and ABTDO with retroactive heuristics. All experiments were performed on the DisChoco 2.0 [8] platform,¹ in which agents are

¹<http://www.lirmm.fr/coconut/dischoco/>

simulated by Java threads that communicate only through message passing. We evaluate the performance of the algorithms by communication load and computation effort. Communication load is measured by the total number of messages exchanged among agents during algorithm execution ($\#msg$), including termination detection (system messages). Computation effort is measured by an adaptation of the number of non-concurrent constraint checks (generic number of non-concurrent constraint checks $\#gncccs$ [11]).

For ABT, we implemented the standard version where we use counters for tagging assignments. For ABTDO [12], we implemented the best version, using the *nogood-triggered* heuristic where the receiver of a nogood moves the sender to be in front of all other lower priority agents (denoted by ABTDO-ng). For ABTDO with retroactive heuristics [13], we implemented the best version, in which a nogood generator moves itself to be in a higher position between the last and the second last agents in the generated nogood. However, it moves before an agent only if its current domain is smaller than the domain of that agent (denoted by ABTDO-Retro).

Uniform binary random DiscSPs

The algorithms are tested on uniform binary random DiscSPs characterized by $\langle n, d, p_1, p_2 \rangle$, where n is the number of agents/variables, d the number of values per variable, p_1 the ratio of existing binary constraints, and p_2 the constraint tightness defined as the ratio of forbidden value pairs. We solved instances of two classes of problems: sparse problems $\langle 20, 10, 0.2, p_2 \rangle$ and dense problems $\langle 20, 10, 0.7, p_2 \rangle$. We vary the tightness p_2 from 0.1 to 0.9 by steps of 0.1. For each pair of fixed density and tightness (p_1, p_2) we generated 25 instances, solved 4 times each. We report average over the 100 runs.

Figure 2 presents the results on the sparse instances ($p_1 = 0.2$). In term of computational effort ($\#gncccs$) (top of Figure 2), ABT is the less efficient algorithm. ABTDO-ng improves ABT by a large scale and ABTDO-Retro is more efficient than ABTDO-ng. These findings are similar to those reported in [13]. Agile-ABT outperforms all these algorithms, suggesting that on sparse problems, the more sophisticated the algorithm is, the better it is. Regarding the number of exchanged messages ($\#msg$) (bottom of Figure 2), the picture is a bit different. ABTDO-ng and ABTDO-Retro require more messages than ABT. Agile-ABT is the algorithm that requires the smallest number of messages. This is not only because Agile-ABT terminates faster than the other algorithms (see $\#gncccs$). Agile-ABT is more parsimonious than ABTDO algorithms in proposing new orders. Termination values seem to focus changes on those which will pay off.

Figure 3 presents the results on the dense instances ($p_1 = 0.7$). Some differences appear compared to sparse problems. Concerning $\#gncccs$ (top of Figure 3), ABTDO

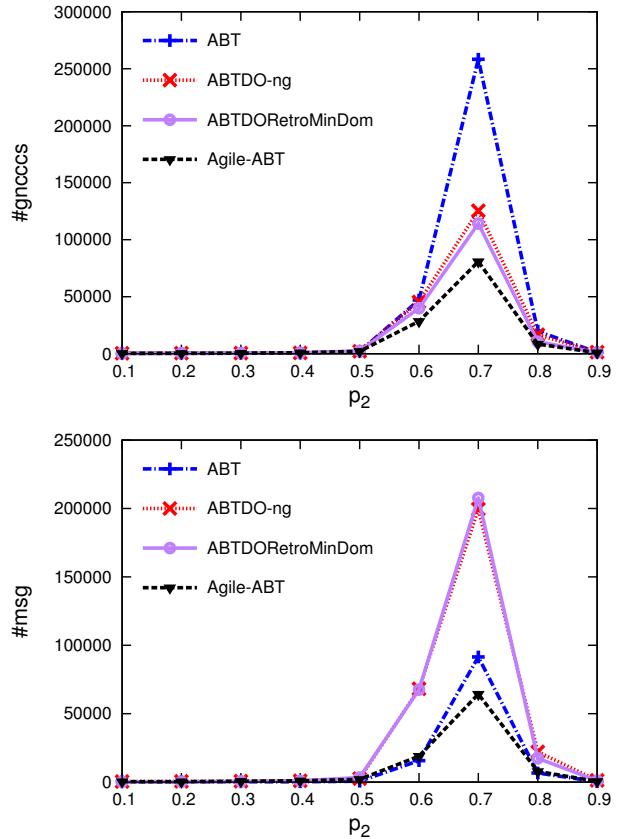


Figure 2. Total $\#msg$ exchanged and $\#gncccs$ performed on sparse problems ($p_1 = 0.2$).

algorithms deteriorate compared to ABT. However, Agile-ABT still outperforms all these algorithms. Regarding communication load ($\#msg$) (bottom of Figure 3), ABTDO-ng and ABTDO-Retro show the same bad performance as in sparse problems. Agile-ABT shows similar communication load as ABT. This confirms its good behavior observed on sparse problems.

Distributed sensor-mobile problems

The distributed sensor-mobile problem [1] is a benchmark based on a real distributed problem. It consists of n sensors that track m mobiles. Each mobile must be tracked by 3 sensors. Each sensor can track at most one mobile. A solution must satisfy visibility and compatibility constraints. The visibility constraint defines the set of sensors that are visible to each mobile. The compatibility constraint defines the compatibility among sensors. We encode SensorDCSP in DiscSP as follows. Each agent represents one mobile. There are three different variables per agent, one for each sensor that we need to allocate to the corresponding mobile. The value domain of each variable is the set of sensors that can detect the corresponding mobile. The intra-agent constraints between the variables of one agent (mobile) specify that the three sensors assigned to the mobile must be

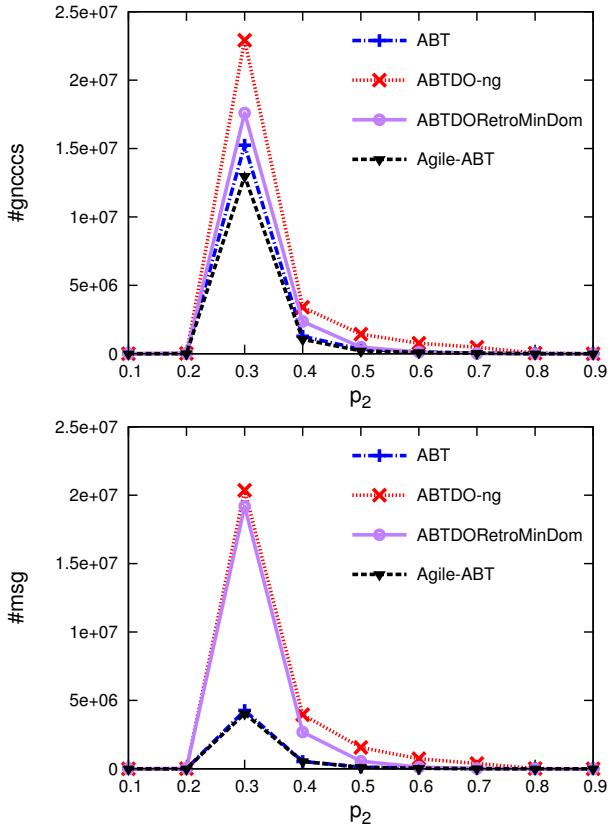


Figure 3. Total $\#msg$ exchanged and $\#gncccs$ performed on dense problems ($p_1 = 0.7$).

distinct and pair-wise compatible. The inter-agent constraints between the variables of different agents specify that a given sensor can be selected by at most one agent. In our implementation of the DisCSP algorithms, this encoding is translated to an equivalent formulation where we have three virtual agents for every real agent, each virtual agent handling a single variable. Problems are characterized by $\langle n, m, p_c, p_v \rangle$, where n is the number of sensors, m is the number of mobiles, p_c is the probability that two sensors are compatible and p_v is the probability that a sensor is visible to a mobile. We present results for class $\langle 25, 5, 0.4, p_v \rangle$ where we vary p_v from 0.1 to 0.9 by steps of 0.1. Again, for each p_v we generated 25 instances, solved 4 times each and averaged over the 100 runs. The results are shown in Figure 4.

When comparing the speed-up of algorithms (top of Figure 4), Agile-ABT is slightly dominated by ABT and ABTDO-ng in the interval [0.3 0.5], while outside of this interval, Agile-ABT outperforms all the algorithms. Nonetheless, the performance of ABT and ABTDO-ng dramatically deteriorate in the interval [0.1 0.3]. Concerning communication load (bottom of Figure 4), as opposed to other dynamic ordering algorithm, Agile-ABT is always better than or as good as standard ABT.

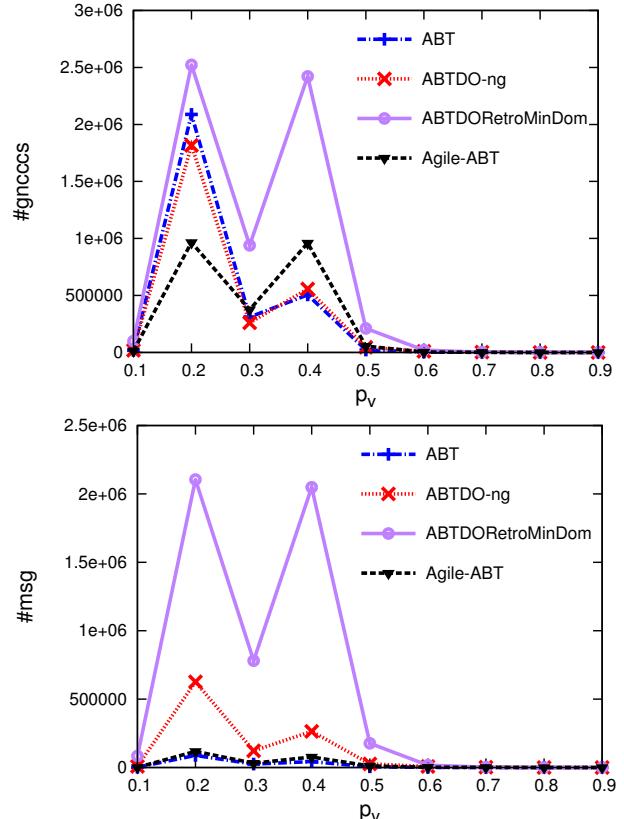


Figure 4. Total $\#msg$ exchanged and $\#gncccs$ performed on sensor-mobile problems.

Discussion

From the experiments above we can conclude that Agile-ABT outperforms other algorithms in term of computation effort ($\#gncccs$) when solving random DisCSP problem. On structured problems (SensorDCSP), our results suggest that Agile-ABT is more robust than other algorithms whose performance is sensitive to the type of problems solved. Concerning communication load ($\#msg$), Agile-ABT is more robust than other versions of ABT with dynamic agent ordering. As opposed to them, it is always better than or as good as standard ABT on difficult problems.

At first sight, Agile-ABT seems to need less messages than other algorithms but these messages are longer than messages sent by other algorithms. One could object that for Agile-ABT, counting the number of exchanged messages is biased. However, counting the number of exchanged messages would be biased only if $\#msg$ was smaller than the number of physically exchanged messages (going out from the network card). Now, in our experiments, they are the same. The International Organization for Standardization (ISO) has designed the Open Systems Interconnection (OSI) model to standardize networking. TCP and UDP are the principal Transport Layer protocols using OSI model. The internet protocols IPv4 (<http://tools.ietf.org/html/rfc791>) and IPv6 (<http://tools.ietf.org/html/rfc2460>) specify the min-

imum datagram size that we are guaranteed to send without fragmentation of a message (in one physical message). This is 568 bytes for IPv4 and 1,272 bytes for IPv6 when using either TCP or UDP (UDP is 8 bytes less than TCP, see RFC-768 –<http://tools.ietf.org/html/rfc768>).

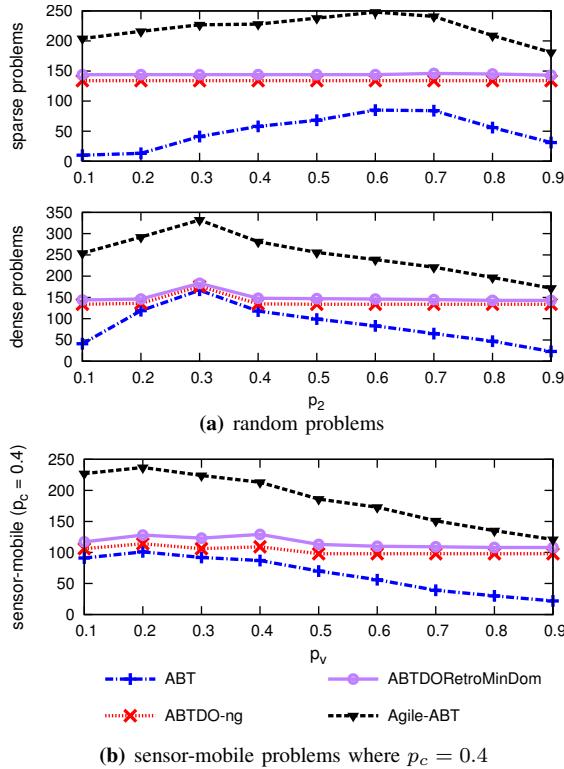


Figure 5. Maximum message size in bytes.

Figure 5 shows the size of the longest message sent by each algorithm on our random and sensor problems. It is clear that Agile-ABT requires lengthy messages compared to other algorithms. However, the longest message sent is always less than 568 bytes (in the worst case it is less than 350, see Figure 5(a)).

VII. CONCLUSION

We have proposed Agile-ABT, an algorithm that is able to change the ordering of agents more agilely than all previous approaches. Thanks to the original concept of termination value, Agile-ABT is able to choose a backtracking target that is not necessarily the agent with the current lowest priority within the conflicting agents. Furthermore, the ordering of agents appearing before the backtracking target can be changed. These interesting features are unusual for an algorithm with polynomial space complexity. Our experiments confirm the significance of these features.

REFERENCES

- [1] R. Bejar, C. Domshlak, C. Fernandez, K. Gomes, B. Krishnamachari, B. Selman, and M. Valls. Sensor networks and distributed CSP: communication, computation and complexity. *Artificial Intelligence*, 161:1-2: 117–148, 2005.
- [2] C. Bessiere, A. Maestre, I. Brito, and P. Meseguer. Asynchronous backtracking without adding links: a new member in the ABT family. *Artificial Intelligence*, 161:1-2:7-24, 2005.
- [3] Matthew L. Ginsberg and David A. McAllester. GSAT and dynamic backtracking. In *KR*, pages 226–237, 1994.
- [4] K. Chandy Mani and L. Lamport. Distributed snapshots: determining global states of distributed systems. *ACM Transactions on Computer Systems*, 3(1):63–75, 1985. ISSN 0734-2071.
- [5] A. Petcu and B. Faltings. A value ordering heuristic for distributed resource allocation. In *Proceeding of CSCLP04*, Lausanne, Switzerland, 2004.
- [6] M. C. Silaghi. Generalized dynamic ordering for asynchronous backtracking on DisCSPs. In *DCR workshop, AAMAS-06*, Hakodate, Japan, 2006.
- [7] M.-C. Silaghi, D. Sam-Haroud, and B. Faltings. Hybridizing ABT and AWC into a polynomial space, complete protocol with reordering. Technical report, 2001.
- [8] M. Wahbi, R. Ezzahir, C. Bessiere, and E. H. Bouyakhf. DisChoco 2: A platform for distributed constraint reasoning. In *Proceedings of the IJCAI'11 workshop on Distributed Constraint Reasoning*, pages 112–121, 2011
- [9] M. Yokoo. Asynchronous weak-commitment search for solving distributed constraint satisfaction problems. In *Proceeding of CP*, pages 88–102, Cassis, France, 1995.
- [10] M. Yokoo, E. H. Durfee, T. Ishida, and K. Kuwabara. Distributed constraint satisfaction problem: Formalization and algorithms. *IEEE Transactions on Knowledge and Data Engineering*, 10:673–685, 1998.
- [11] R. Zivan and A. Meisels. Message delay and DisCSP search algorithms. *Annals of Mathematics and Artificial Intelligence*, 46(4):415–439, 2006. ISSN 1012-2443.
- [12] R. Zivan and A. Meisels. Dynamic ordering for asynchronous backtracking on DisCSPs. *Constraints*, 11(2-3):179–197, 2006. ISSN 1383-7133.
- [13] R. Zivan, M. Zazone, and A. Meisels. Min-domain retroactive ordering for asynchronous backtracking. *Constraints*, 14(2):177–198, 2009. ISSN 1383-7133.