

# Min-Domain ordering for Asynchronous Backtracking\*

Roie Zivan, Moshe Zazone, and Amnon Meisels  
Department of Computer Science,  
Ben-Gurion University of the Negev,  
Beer-Sheva, 84-105, Israel {zivanr,moshezaz,am}@cs.bgu.ac.il

June 18, 2007

## Abstract

Ordering heuristics are a powerful tool in CSP search algorithms. Among the most successful ordering heuristics are heuristics which enforce a *fail first* strategy by using the min-domain property [HE80, BR96, SG98, Dec03]. Ordering heuristics have been introduced recently to *Asynchronous backtracking (ABT)*, for distributed constraints satisfaction (*DisCSP*) [ZM05]. However, the pioneering study of dynamically ordered ABT, *ABT\_DO*, has shown that a straightforward implementation of the min-domain heuristic does not produce the expected improvement over a static ordering. The best ordering heuristic for asynchronous backtracking was found to be the *Nogood-triggered* heuristic.

The present paper proposes an asynchronous dynamic ordering which does not follow the standard restrictions on the position of reordered agents in *ABT\_DO*. Agents can be moved to a position that is higher than that of the target of the backtrack (*culprit*).

Combining the Nogood-triggered heuristic and the min-domain property in this new class of heuristics results in the best performing version of *ABT\_DO*. The new version of retroactively ordered ABT is faster by a large factor than the best form of ABT.

## 1 Introduction

Distributed constraint satisfaction problems (*DisCSPs*) are composed of agents, each holding its local constraints network, that are connected by constraints among variables of different agents. Agents assign values to variables, attempting to generate a locally consistent assignment that is also consistent with all constraints between agents (cf. [Yok00, SGM96]). To achieve this goal, agents check the value assignments to their variables for local consistency and exchange messages with other agents, to check consistency of their proposed assignments against constraints with variables owned by different agents [Yok00, BMBM05].

---

\*Supported by the Lynn and William Frankel center for Computer Sciences and the Paul Ivanier Center for Robotics and Production Management.

A search procedure for a consistent assignment of all agents in a distributed CSP (*DisCSP*), is a distributed algorithm. All agents cooperate in search for a globally consistent solution. The solution involves assignments of all agents to all their variables and exchange of information among all agents, to check the consistency of assignments with constraints among agents.

*Asynchronous Backtracking (ABT)* is one of the most efficient and robust algorithms for solving distributed constraints satisfaction problems. Asynchronous Backtracking was first presented by Yokoo [YDIK98, Yok00] and was developed further and studied in [Ham01, BMM01, SF05, BMBM05]. Agents in the *ABT* algorithms perform assignments asynchronously according to their current view of the system's state. The method performed by each agent is in general simple. Later versions of *ABT* use polynomial space memory and perform dynamic backtracking [BMM01, BMBM05]. The versions of asynchronous backtracking in all of the above studies use a static priority order among all agents.

An asynchronous algorithm with dynamic ordering was proposed by [Yok95], Asynchronous Weak Commitment (*AWC*). According to [Yok00], *AWC* outperforms *ABT* on specific applications (N-queens, Graph-coloring). The heuristic used by *AWC* is very specific. Move any agent that sends back a *Nogood* to be first in the order of all agents [Yok00]. However, in order to be complete, *AWC* uses exponential space for storing *Nogoods* [Yok00]. This can be a problem when solving hard instances of *DisCSPs*.

An attempt to combine *ABT* with *AWC* was reported by [SSHF01]. In order to perform asynchronous finite reordering operations [SSHF01] propose that the reordering operation will be performed by abstract agents. In a later study the exact heuristic of *Dynamic Backtracking* [Gin93] was proposed for *ABT* with dynamic ordering [Sil06]. The results presented in both studies [SSHF01, Sil06] show minor improvements to *ABT* with static ordering.

A general algorithm for dynamic ordering in asynchronous backtracking, *ABT\_DO*, was presented in [ZM05]. The *ABT\_DO* algorithm uses polynomial space, similarly to standard *ABT*. In the *ABT\_DO* algorithm the agents of the *DisCSP* choose orders dynamically and asynchronously. Agents in *ABT\_DO* perform computation according to the current, most updated order they hold. There are three rules on the changes of orderings of agents in *ABT\_DO*. Each agent can change the order of agents who have a lower priority than its own. An agent can propose an order change each time it replaces its assignment. Each order is time-stamped according to the assignments performed by agents [ZM06a].

Ordering heuristics are one of the most powerful tools for increasing the efficiency of standard *CSP* search algorithms [HE80, BR96, Dec03, Wal05]. Heuristics which are successful for sequential CSP algorithms were shown to fail to improve run-time when combined with asynchronous backtracking [ZM06a]. Thus, an investigation of ordering heuristics for asynchronous backtracking on *DisCSPs* is needed.

The results presented in [ZM06a] show that the performance of *ABT\_DO* is highly dependent on the selected heuristic. The classic *Min-Domain* heuristic was implemented by including the current domain size of agents in the messages they send. Surprisingly, this heuristic which in centralized algorithms and in distributed algorithms using a sequential assignment protocol produces a large improvement over static order, was found not to be efficient for Asynchronous Backtracking. A heuristic which achieved a significant improvement was inspired by *Dynamic Backtrack-*

ing [Gin93, Bak94] in which the agent which sends a *Nogood* is advanced in the new order to be immediately after the agent to whom the *Nogood* was sent. The explanation for the success of this heuristic is that it does not cause the removal of relevant *Nogoods* as do other heuristics [ZM06a].

In the present paper we investigate the relation between the success of this heuristic and the *min-domain* heuristic which was found to be successful for sequential assignments (synchronous) algorithms on *DisCSPs* [BM04, MZ07]. We demonstrate the effect of *Nogood* loss as a result of reordering on the failure of the *min domain* heuristic. Removal of *Nogoods* cause the return of values to the agent’s domains. This harms the accuracy of the information that agents hold on the domain size of other agents. On the other hand, the *Nogood-triggered* heuristic of [ZM05] does not lose valid information and moves agents with a potential of having a small domain to a higher position.

In order to maximize the *min-domain* property, a more flexible heuristic is proposed, which violates the restrictions on the ordering of agents in [ZM05]. We study changes of order that move agents to a higher position, replacing agents that were ahead of them including the first agent. This new type of heuristics is termed *Retroactive ordering* and is based on a slightly modified version of *ABT\_DO*.

The results presented in the present paper show that moving the *Nogood* sender as high as possible in the priority order is successful only if the domain size of agents is taken into consideration. Our experiments show that the successful heuristics are those that support a *min-domain* scheme in which agents are moved to a higher position only if their current domain size is smaller than the current domain of agents they are moved in front of. Moving an agent before the agents which are included in the *Nogood* actually *enlarges* its domain. The best heuristic in the present paper is that agents which generate a *Nogood* are placed in the new order between the last and the second last agents in the generated *Nogood*. This heuristic is the asynchronous form of the Min-Domain heuristic. Agents are moved to a higher position only if their domain is smaller than the agents they pass on the way up. Our results on both random *DisCSPs* and on structured *DisCSPs* show that the proposed heuristic improves the best results to date [ZM06a] by a large factor.

Distributed *CSPs* are presented in Section 2. A description of Asynchronous backtracking with dynamic ordering (*ABT\_DO*) and its best heuristic is presented in Section 3. Section 4 presents an investigation of the existing heuristics and offers reasons for their performance in previous papers. Section 5 present the general scheme of retroactive heuristics for *ABT\_DO* and a correctness proof. An extensive experimental evaluation, which compares standard and retroactive heuristics of *ABT\_DO* is in Section 7. The experiments were conducted on randomly generated *DisCSPs* and on Course Scheduling problems. Section 8 presents a discussion of the relation between the experimental results and the *min-domain* heuristic.

## 2 Distributed Constraint Satisfaction

A distributed constraint satisfaction problem - *DisCSP* is composed of a set of  $k$  agents  $A_1, A_2, \dots, A_k$ . Each agent  $A_i$  contains a set of constrained variables  $X_{i_1}, X_{i_2}, \dots, X_{i_{n_i}}$ . Constraints or **relations**  $R$  are subsets of the Cartesian product of the domains of the constrained variables. For a set of constrained variables  $X_{i_k}, X_{j_l}, \dots, X_{m_n}$ , with domains of values for each variable  $D_{i_k}, D_{j_l}, \dots, D_{m_n}$ , the constraint is defined as

$R \subseteq D_{i_k} \times D_{j_l} \times \dots \times D_{m_n}$ . A **binary constraint**  $R_{ij}$  between any two variables  $X_j$  and  $X_i$  is a subset of the Cartesian product of their domains;  $R_{ij} \subseteq D_j \times D_i$ . In a distributed constraint satisfaction problem *DisCSP*, constrained variables can belong to different agents [YDIK98, SGM96]. Each agent has a set of constrained variables, i.e. a *local constraint network*.

An assignment (or a label) is a pair  $\langle var, val \rangle$ , where  $var$  is a variable of some agent and  $val$  is a value from  $var$ 's domain that is assigned to it. A *compound label* (or a partial solution) is a set of assignments of values to a set of variables. A **solution**  $P$  to a *DisCSP* is a compound label that includes all variables of all agents, that satisfies all the constraints. Agents check assignments of values against non-local constraints by communicating with other agents through sending and receiving messages. Agents exchange messages with agents whose assignments may be in conflict [BMBM05]. Agents connected by constraints are therefore called *neighbors*. The ordering of agents is termed *priority*, so that agents that are later in the order are termed "lower priority agents" [Yok00, BMBM05].

The following assumptions are routinely made in studies of *DisCSPs* and are assumed to hold in the present study [Yok00, BMBM05].

1. All agents hold exactly one variable.
2. Messages arrive at their destination in finite time.
3. Messages sent by agent  $A_i$  to agent  $A_j$  are received by  $A_j$  in the order they were sent.

### 3 ABT with Dynamic Ordering

Each agent in *ABT-DO* holds a *Current\_order* which is an ordered list of pairs. Every pair includes the ID of one of the agents and a counter. Each agent can propose a new order for agents that have lower priority (i.e. are in a lower position in the current order), each time it replaces its assignment. This way the sending of an ordering proposal message always coincides with an assignment message (an **ok?** message [?, ZM05]). An agent  $A_i$  can propose an order according to the following rules:

1. Agents with higher priority than  $A_i$  and  $A_i$  itself, do not change priorities in the new order.
2. Agents with lower priority than  $A_i$ , in the current order, can change their priorities in the new order *but not to a higher priority than  $A_i$  itself* (This rule enables a more flexible order than in the centralized case).

The counters attached to each agent ID in the *order* list form a time-stamp. Initially, all time-stamp counters are set to zero and all agents start with the same *Current\_order*. Each agent  $A_i$  that proposes a new order, changes the order of the pairs in its own ordered list and updates the counters as follows:

1. The counters of agents with higher priority than  $A_i$ , according to the *Current\_order*, are not changed.
2. The counter of  $A_i$  is incremented by one.

3. The counters of agents with lower priority than  $A_i$  in the *Current\_order* are set to zero.

In *ABT*, agents send **ok?** messages to their neighbors whenever they perform an assignment. In *ABT\_DO*, an agent can choose to change its *Current\_order* after changing its assignment. If that is the case, besides sending **ok?** messages an agent sends **order** messages to all lower priority agents. The **order** message includes the agent's new *Current\_order*.

An agent which receives an **order** message must determine if the received order is more updated than its own *Current\_order*. It decides by comparing the time-stamps lexicographically. Since orders are changed according to the above rules, every two orders must have a common prefix of agents' IDs. The agent that performs the change does not change its own position and the positions of higher priority agents.

When an agent  $A_i$  receives an order which is more up to date than its *Current\_order*, it replaces its *Current\_order* by the received order. The new order might change the location of the receiving agent with respect to other agents (in the new *Current\_order*). In other words, one of the agents that had higher priority than  $A_i$  according to the old order, now has a lower priority than  $A_i$  or vice versa. Therefore,  $A_i$  rechecks the consistency of its current assignment and the validity of its stored *Nogoods* (the explanations for removing values from its domain [?, ZM05]) according to the new order. If the current assignment is inconsistent according to the new order, the agent makes a new attempt to assign its variable. In *ABT\_DO* agents send **ok?** messages to all constraining agents (i.e. their neighbors in the constraints graph). Although agents might hold in their *Agent\_views* assignments of agents with lower priorities, according to their *Current\_order*, they eliminate values from their domain *only if they violate constraints with higher priority agents*.

A *Nogood* message (i.e. a message carrying a partial assignment which was found to be inconsistent [?, ZM05]) is always checked according to the *Current\_order* of the receiving agent. If the receiving agent is not the lowest priority agent in the *Nogood* according to its *Current\_order*, it sends the *Nogood* to the lowest priority agent and sends an **ok?** message to the sender of the *Nogood*. This is a similar operation to that performed in standard *ABT* for any unaccepted (inconsistent) *Nogood* [BMBM05].

## 4 Investigation of asynchronous heuristics

In this section we offer explanations for the failure of the *min-domain* heuristic and the success of the *Nogood-triggered* heuristic when used in asynchronous backtracking [ZM05, ZM06a]. Consider the example in Figure 1. The agents are ordered by their indices. Each agent has a single variable and three values,  $a$ ,  $b$  and  $c$ , in its domain. The eliminated values are crossed and each points to its eliminating explanation (i.e. the assignment which caused its removal). The circled values represent the current assignments. In this example, agent  $A_5$  has exhausted its domain and must create a *Nogood*. The *Nogood* it generates includes the assignments of  $A_1$  and  $A_2$  therefore the *Nogood* is sent to  $A_2$ . According to the rules of the *ABT\_DO* algorithm agent  $A_2$  can reorder agents  $A_3$ ,  $A_4$  and  $A_5$ . Now, if it will reorder them according to their current domain sizes then  $A_3$  and  $A_4$  will switch places. But, since both of the values eliminated from the domain of  $A_4$  are in conflict with the assignment of  $A_3$  then after




Figure 1: Heuristics example before backtrack.




Figure 2: After reordering using the NG-triggered heuristic.

they change places, these values will be returned to the domain of  $A_4$  and its domain size will be larger than the domain of  $A_3$ .

In contrast, if  $A_2$  reorders according to the *Nogood-triggered* heuristic then the only agent to change places is  $A_5$  which is moved to be after  $A_2$  and before  $A_3$ . Now, after  $A_2$  replaces its assignment we get the situation in Figure 2. We can see that an agent with a small domain was moved forward while the others kept their domain sizes and places.

The example demonstrates why the *min-domain* heuristic fails when used in asynchronous backtracking. In asynchronous backtracking, all agents hold an assignment throughout the search. Conflicts with these assignments effect the size of domains of other agents. For each value which is removed from an agent's domain an explanation *Nogood* is stored. When an agent is moved in front of an agent whose assignment is included in one of its *Nogoods*, this *Nogood* must be eliminated and the corresponding value is returned to the domain. Thus, in contrast to sequential ordering algorithms, in asynchronous backtracking the resulting domain sizes after reordering cannot be anticipated by the ordering agent. The example demonstrates how this phenomena does not affect the *Nogood-triggered* heuristic.

Following the example one can see that the *Nogood-triggered* heuristic is successful because in many cases it moves an agent with a small domain to a higher position.

Only values whose *Nogood* explanation includes the assignment of the culprit agent are returned to the moving agent’s domain. In fact, the agent can be moved up passed the culprit, and as long as it does not pass the second last assignment in the *Nogood* its domain size will stay the same. In Figure 2, Agent  $A_5$  is moved right after agent  $A_2$ . Its domain size is one, since the *Nogoods* of its other two values are valid. If  $A_5$  is moved before  $A_2$  its domain size will stay the same as both eliminating *Nogoods* include only the assignment of  $A_1$ . However, if  $A_5$  will be moved in front of  $A_1$  then all its values will return to its domain. This possibility of moving an agent with a small domain beyond the culprit agent to a higher position is the basic motivation for retroactive ordering.

## 5 Retroactive ordering heuristics for ABT

In contrast to the rules of *ABT\_DO* of the previous section, the present paper proposes a new type of ordering. The new type of ordering can change the order of agents with higher priority than the agent which replaces its assignment. A *retroactive* heuristic would enable moving the *Nogood* sender to a higher position than the *Nogood* receiver. In order to preserve the correctness of the algorithm, agents must be allowed to store *Nogoods*. In order to generate a general scheme for retroactive heuristics, one can define a global space limit for the storage of *Nogoods*. The specific realization is to limit the storage of *Nogoods* that are smaller or equal to some predefined size  $k$ .

The proposed ordering heuristic is triggered by the sending of a *Nogood*. The reordering operation can be generated by either the *Nogood* generator or by the *Nogood* receiver (but not by both). In contrast to [ZM05, ZM06a] we choose the *Nogood* sender to be the one to reorder. This is since the only agent which can lose a relevant *Nogood* as a result of the reordering is the *Nogood* sender (the only one moving to a higher position). Therefore, since it is aware of its own state and the others do not lose information, the *Nogood* sender is the best candidate for selecting the new order.

The new order is selected according to the following rules:

1. The *Nogood* generator can be moved to any position in the new order.
2. If the *Nogood* generator is moved to a position which is before the second last in the *Nogood* (the one before the culprit) all the agents included in the *Nogood* must hold the *Nogood* until the search is terminated.
3. Agents with lower priority than the *Nogood* receiver can change order but not move in front of it (as in standard *ABT\_DO*).

According to the above rules, agents which detect a dead end are moved to a higher position in the priority order. If the length of the created *Nogood* is larger than  $k$ , they can be moved up to the place that is right after the agent which is the last to be included in the *Nogood* according to the current order and is not the culprit (i.e. second last in the *Nogood*).

If the length of the created *Nogood* is smaller or equal to  $k$ , the sending agent can be moved to a position before all the participants in the *Nogood* and the *Nogood* is sent and saved by all of them. In the extreme case where  $k$  is equal to the number of agents in the *DisCSP* (i.e.  $k = N$ ), the *Nogood* sender can always move to be first in the priority order and the resulting algorithm is a generalization of *AWC* [Yok00].

```

when received (ok?, (xj, dj)) do:
1. add (xj, dj) to agent_view;
2. remove inconsistent nogoods;
3. check_agent_view;

when received (order, received_order) do:
1. if (received_order is more updated than Current_order)
2.   Current_order ← received_order;
3.   remove inconsistent nogoods;
4.   check_agent_view;

when received (nogood, xj, nogood)
1. old_value ← current_value
2. if (nogood contains an agent xk
   with lower priority than xi and nogood.size > K)
3.   send (nogood, (xi, nogood)) to xk;
4. else
5.   if (nogood consistent with {Agent_view ∪
   current_assignment} or nogood.size ≤ K)
6.     store nogood;
7.     if (nogood contains an agent xk that is not its neighbor)
8.       request xk to add xi as a neighbor;
9.       add (xk, dk) to agent_view;
10.    if(xi is with lowest priority in nogood)
11.     check_agent_view;
12. if(old_value = current_value)
13.   send (ok?, (xi, current_value)) to xj;

```

Figure 3: Retroactive ABT\_DO algorithm (first part)

Figures 3 and 4 present the code of *Retroactive ABT\_DO*. The difference from standard *ABT\_DO* in the code performed when a *Nogood* is received (Figures 3) derives from the different possible types of *Nogoods*. A *Nogood* smaller or equal to  $k$  is actually a constraint that will be stored by the agent until the search is terminated. In the case of *Nogoods* which are longer than  $k$ , the algorithm treats them as in standard *ABT\_DO* i.e. accepts them only if the receiver is the lowest priority agent in the *Nogood* and the *Nogood* is consistent with the *Agent\_view* and *current\_assignment* of the receiver. In any case of acceptance of a *Nogood*, the agent searches for a new assignment only if it happens to be the lowest priority agent in the *Nogood*. As stated above, our choice is that only the *Nogood* generator is allowed to change order.

Procedure **backtrack** (Figure 4) is largely changed in the retroactive heuristic version of *ABT\_DO*. When an agent creates a *Nogood* it determines whether it is larger than  $k$  or not. If it is larger then a single *Nogood* is sent to the lowest priority agent in the *Nogood* in the same way as in *ABT\_DO*. Consequently, the agent selects a new order in which it puts itself not higher than the second lowest priority agent in the *Nogood*. When the *Nogood* is smaller or equal to  $k$ , if it is the first time this *Nogood* is generated, the *Nogood* is sent to all the agents included in the *Nogood* and the agent moves itself to an unlimited position in the new order (In this case the function **choose\_new\_order** is called with no limitations). In both cases, order messages are sent to all the lower priority agents in the new order. The assignment of the lowest pri-

```

procedure backtrack
1.  $nogood \leftarrow \mathbf{resolve\_inconsistent\_subset}$ ;
2. if ( $nogood$  is empty)
3.   broadcast to other agents that there is no solution;
4.   stop;
5. select  $(x_j, d_j)$  where  $x_j$  has the lowest priority in  $nogood$ ;
6. if ( $nogood.size > K$ )
7.    $Current\_order \leftarrow \mathbf{choose\_new\_order}()$ 
       where  $x_l$  has the second lowest priority in  $nogood$ ;
8.   send (nogood,  $x_i$ ,  $nogood$ ) to  $x_j$ ;
9.   else if ( $is\_new(nogood)$ )
10.   $new\_position \leftarrow unlimited$ 
11.  send (nogood,  $x_i$ ,  $nogood$ ) to all agents in  $nogood$ ;
12.  store sent  $nogood$ ;
13.  $Current\_order \leftarrow \mathbf{choose\_new\_order}(x_i)$ 
14. send (order,  $Current\_order$ ) to lower priority agents;
15. remove  $(x_j, d_j)$  from  $agent\_view$ ;
16. remove all  $nogoods$  containing  $(x_j, d_j)$ ;
17. check\_agent\_view;

```

```

procedure check\_agent\_view
1. if ( $current\_assignment$  is not consistent with all
       higher priority assignments in  $Agent\_view$ )
2.   if (no value in  $D_i$  is consistent with all higher priority
         assignments in  $Agent\_view$ )
3.     backtrack;
4.   else
5.     select  $d \in D_i$  where  $Agent\_view$  and  $d$  are consistent;
6.      $current\_value \leftarrow d$ ;
7.     send (ok?,  $(x_i, d)$ ) to  $neighbors$ ;

```

Figure 4: The Retroactive ABT\_DO algorithm

riority agent in the  $Nogood$  is removed from the  $Agent\_view$ , the relevant  $Nogoods$  are removed and the agent attempts to re-assign its variable by calling **check\\_agent\\_view**.

Procedure **check\\_agent\\_view** (Figure 3) is slightly changed from that of  $ABT\_DO$  [ZM06a] since the change of order in the new scheme is performed by the  $Nogood$  sender and not by its receiver.

## 6 Correctness of Retroactive $ABT\_DO$

In order to prove the correctness of *Retroactive ABT\_DO* we assume the correctness of the standard  $ABT\_DO$  algorithm (see proof in [ZM06a]) and prove that the changes made for retroactive heuristics do not damage its correctness. We first prove the case for no  $Nogood$  storage ( $k = 0$ ):

**theorem 1** *Retroactive ABT\_DO is correct when  $k = 0$ .*

There are two differences between standard  $ABT\_DO$  and *Retroactive ABT\_DO* with  $k = 0$ . First, order is changed whenever a  $Nogood$  is sent and not when an

assignment is replaced. This change does not make a difference in the correctness since when a *Nogood* is sent there are two possible outcomes. Either the *Nogood* receiver replaces its assignment, which makes it effectively the same as in standard *ABT\_DO*, or the *Nogood* is rejected. A rejected *Nogood* can only be caused by a change of assignment either of the receiving agent or of an agent with higher priority. In all of these cases, the most relevant order is determined lexicographically. Ties which could not have been generated in standard *ABT\_DO*, are broken using the agents indexes.

The second change in the code for  $k = 0$  is that in *Retroactive ABT\_DO* a *Nogood* sender can move to a position in front of the agent that receives the *Nogood*. Since the *Nogood* sender is the only agent moving to a higher position, it is the only one that can lose a *Nogood* as a result. However, the *Nogood* sender removes all *Nogoods* containing the assignment of the *Nogood* receiver and it does not pass any other agent contained in the *Nogood*. Thus, no information is actually lost by this change. Moreover, the number of times two agents can move in front of one another without a higher priority agent changing its assignment is bounded by their domain sizes.  $\square$

**theorem 2** *Retroactive ABT\_DO* is correct when  $n \geq k > 0$ .

In order to prove that *Retroactive ABT\_DO* is correct for the case that  $n \geq k > 0$  we need to show that infinite loops cannot occur. In the case of *Nogoods* which are smaller or equal to  $k$  the case is very simple. All agents involved in the *Nogood* continue to hold it, therefore the same assignment can never be produced again. The number of these *Nogoods* with a limited length is finite. In finite time the algorithm reaches a state in which no permanent *Nogoods* are added. In this state, agents do not move in front of the second last in the *Nogoods* generated and the previous proof holds.  $\square$

## 7 Experimental Evaluation

The common approach in evaluating the performance of distributed algorithms is to compare two independent measures of performance - time, in the form of steps of computation [Lyn97, Yok00], and communication load, in the form of the total number of messages sent [Lyn97].

Non concurrent steps of computation, are counted by a method similar to the clock synchronization algorithm of [Lam78]. Every agent holds a counter of computation steps. Every message carries the value of the sending agent's counter. When an agent receives a message it stores the data received together with the corresponding counter. When the agent first uses the received counter it updates its counter to the largest value between its own counter and the stored counter value which was carried by the message [ZM06b]. By reporting the cost of the search as the largest counter held by some agent at the end of the search, a measure of non-concurrent search effort that is close to Lamports logical time is achieved [Lam78]. If instead of steps of computation, the number of non concurrent constraint checks is counted (*NCCC*s), then the local computational effort of agents in each step is measured [MRKZ02, ZM06b].

The first set of experiments was conducted on random networks of constraints of  $n$  variables,  $k$  values in each domain, a constraints density of  $p_1$  and tightness  $p_2$  (which are commonly used in experimental evaluations of CSP algorithms [Smi96, Pro96]). The constraint networks were generated with 20 agents ( $n = 20$ ) each holding exactly

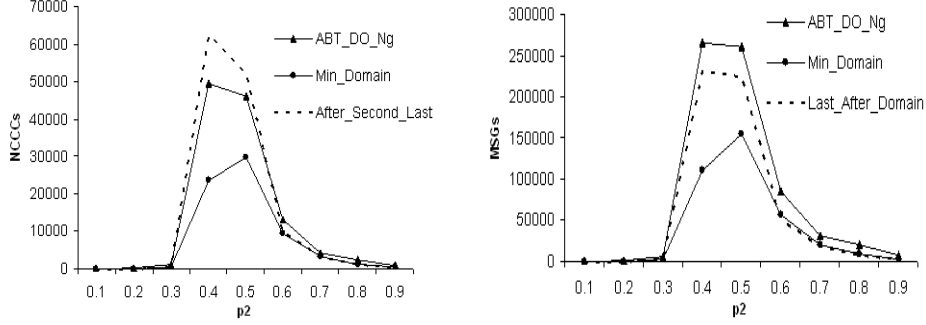


Figure 5: Non concurrent constraint checks performed and messages sent by Retroactive *ABT\_DO* and *ABT\_DO* on low density DisCSPs ( $p_1 = 0.4$ ).

one variable, 10 values for each variable ( $k = 10$ ) with two different constraints densities  $p_1 = 0.4$  and  $p_1 = 0.7$ . The tightness value  $p_2$ , is varied between 0.1 and 0.9, to cover all ranges of problem difficulty. For each pair of fixed density and tightness ( $p_1, p_2$ ) 50 different random problems were solved by each algorithm and the results presented are an average of these 50 runs.

In order to confirm the dependency of the performance on the size of the current domain of the moved agents, we compared *ABT\_DO* with *ABT\_DO* with a retroactive domain of the moved agents in which agents are not allocated any additional *Nogood* storage. Agents include in their messages the size of their current domains. This information is stored in the agent’s *Agent\_views* (as in [ZM06a]). A *Nogood* generator moves itself to be in a higher position than the culprit agent but it moves in front of an agent *only if its current domain is smaller* than the domain of that agent. Otherwise, it places itself right after the culprit agent as in standard *ABT\_DO*.

The Left Hand Side (LHS) of Figure 5 presents the results in *NCCCs* for *ABT\_DO* and Retroactive *ABT\_DO* with the above heuristic. The retroactive version of *ABT\_DO* (depicted in the figures as *min-domain*) improves the run-time performance of *ABT\_DO* (depicted as *ABT\_DO\_NG*). In order to emphasize the relation to the *Min-Domain* property, a third line in Figures 5, 6 and 8 represents retroactive *ABT\_DO* without checking the domain sizes (depicted in the figures as *After Second Last*). This version of retroactive *ABT\_DO* was the slowest among the three. Similar results for the number of messages sent are presented on the Right Hand Side (RHS) of Figure 5. In the case of network load, both versions of *Retroactive ABT\_DO* send less messages than standard *ABT\_DO*. For high density problems the difference between the algorithms is similar but smaller (Figure 6).

In order to further demonstrate the dependency of the domain size of agents on the success of the selected heuristic we performed an additional experiment on random problems in which the size limit for keeping *Nogoods* is varied<sup>1</sup>. A *Nogood* generator which created a *Nogood* of length larger than  $k$  places itself right after the *Nogood* receiver as in standard *ABT\_DO*. When the *Nogood* generator creates a *Nogood* smaller or equal to  $k$ , it places itself first in the priority order and sends the generated

<sup>1</sup>In this experiment the problems were smaller ( $n = 15$ ) since the algorithms run slower.

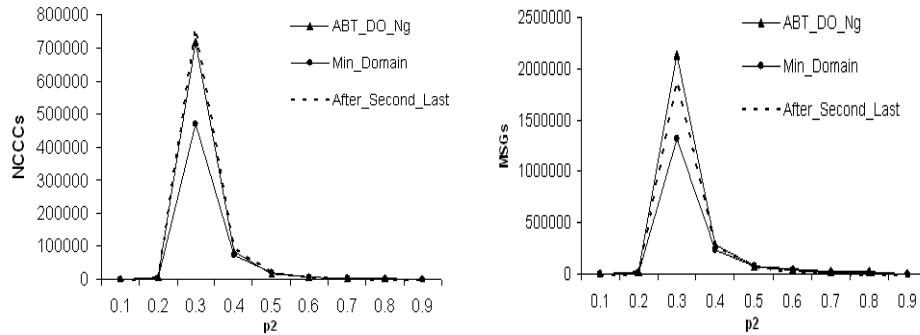


Figure 6: Non concurrent constraint checks performed and messages sent by Retroactive *ABT\_DO* and *ABT\_DO* on high density DisCSPs ( $p_1 = 0.7$ ).

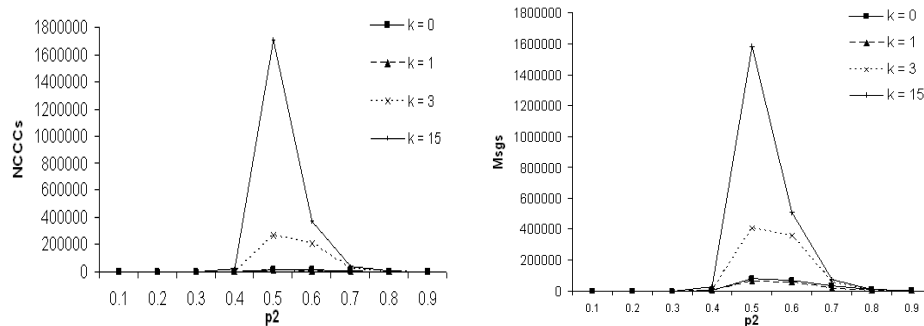


Figure 7: Non concurrent constraint checks performed and messages sent by Retroactive *ABT\_DO* with different limits on Nogood size ( $p_1 = 0.4$ ).

*Nogood* to all the participating agents. In the case of  $k = n$  the resulting algorithm is exactly *AWC*. In the case of  $k = 0$  the resulting algorithm is standard *ABT\_DO*. The LHS of Figure 7 presents the number of *NCCCs* performed by the algorithm with  $k$  equal to 0, 1, 3 and  $n$  ( $n = 15$ ). The results show similar performance when  $k$  is small. The performance of the algorithm deteriorates when  $k = 3$  and the slowest performance is when  $k = n$ . Similar results in the number of messages are presented on the RHS of Figure 7.

The fact that a larger storage which enables more flexibility of the heuristic actually causes a deterioration of the performance might come as a surprise. However, one must examine the effect of the specific heuristic used on the size of the domains of the agents which are moved up in the order of priorities. An agent creates a *Nogood* when its domain empties. After sending the *Nogood* it removes the assignment of the culprit agent from its *Agent.view* and returns to the domain only values whose eliminating *Nogood* included the removed assignment. When the agent is moved in front of other agents whose assignments were included in the generated *Nogood* it must return more values to its domain (the values whose explanation *Nogood* included the assignment of the agent which was passed). This of course does not happen for the case of a *Nogood* of size one and that is why for  $k = 1$  we get better results. Thus, moving an agent as




Figure 8: Non concurrent constraint checks performed and messages sent by Retroactive *ABT\_DO* and *ABT\_DO* on Random Meeting Scheduling Problems

high as possible in the priority order actually results in moving upwards an agent with a larger domain.

In the next set of experiments, the successful versions of ABT with retroactive dynamic ordering were compared on realistic structured problems. The generated problems were course scheduling in which each variable assigned to an agent represents a single course which is taken by a number of students. Two variables are constrained if there is a student attending both courses. The constraints are arrival constraints, i.e. if the length of a course is  $t_1$  and the time to get from one course to the other is  $t_2$ , then the beginnings of each two constrained courses must satisfy:  $ct_1 - ct_2 \geq t_1 + t_2$ . This problem is equivalent to the published *Meeting Scheduling Problems* in which each agent holds exactly one variable [GW99]. For a detailed description of how a large random benchmark of problems with these realistic properties can be produced the reader is referred to [GW99]. The results presented in Figure 8 show clearly that the advantage of the retroactive heuristic which takes into account the domain sizes is more pronounced for structured DisCSPs. Furthermore, on tight problems, the performance of the version of retroactive *ABT\_DO* which does not take the domain sizes deteriorates.

## 8 Discussion

The results in the previous section show clearly the relation between the examined heuristics and the *Min-Domain* property of the generated search tree. A well known fact from centralized *CSP* algorithms [HE80, Dec03] and from *DisCSP* algorithms with a sequential assignment protocol [BM04] is that the *Min-Domain* heuristic is very powerful and improves the run of the same algorithms using a static order. If we investigate the *Nogood-triggered* heuristic of [ZM05] we can see that in most cases this heuristic moves to higher priority, agents with smaller domains. This is because an agent whose domain was exhausted returns to its domain, after sending the *Nogood*, only the values in conflict with the assignment of the culprit agent. Thus, only a small number of values are returned to its domain. It is not surprising that this heuristic was found to be very successful in [ZM05, ZM06a]. On the other hand, when an agent is moved to a higher position than the agents in the *Nogood* it discovered, it must return additional values to its domain. This contradicts the properties of the *min-domain*

heuristic and was found to perform poorly in practice. The case of  $k = 1$  did show an improvement since the last assignment in a detected *Nogood* is removed from the *Agent\_view* of the agent which found the *Nogood* anyway.

In our best performing heuristic, agents are moved higher in the priority order as long as their domain size is smaller than the domains of the agents before them and as long as they do not pass the second last in the *Nogood* they have generated, which would result in returning more values to their domain. Since the agent moving to a higher position is not in conflict with the assignments of agents it has moved in front of, its move will not cause the loss of *Nogoods* and therefore the information it holds on the size of the current domains of these agents remains valid. The retroactive ordering version has improved the results of [ZM05, ZM06a] by a factor of 2. In the case of structured problems, this heuristic was found to improve the run of the standard *ABT.DO* by an even larger factor.

## 9 Conclusion

A general scheme for introducing retroactive heuristics into Asynchronous Backtracking with dynamic agent ordering was presented. The flexibility of the heuristic is dependent upon the amount of memory that agents are allowed to use. However, moving agents to the highest position possible was found to deteriorate the performance of the algorithm. The best heuristic, moves to a higher priority only agents whose variable's domains are small but avoid the return of values to domains as a result of reordering. This brings to an extreme the exploitation of the *Min-Domain* property, and improves the run of the best heuristic reported so far [ZM05, ZM06a] by a large factor.

## References

- [Bak94] Andrew B. Baker. The hazards of fancy backtracking. In *Proceedings of the 12th National Conference on Artificial Intelligence (AAAI '94), Volume 1*, pages 288–293, Seattle, WA, USA, July 31 - August 4 1994. AAAI Press.
- [BM04] I. Brito and P. Meseguer. Synchronous, asynchronous and hybrid algorithms for discsp. In *Workshop on Distributed Constraints Reasoning (DCR-04) CP-2004*, Toronto, September 2004.
- [BMBM05] 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, January 2005.
- [BMM01] C. Bessiere, A. Maestre, and P. Meseguer. Distributed dynamic backtracking. In *Proc. Workshop on Distributed Constraint of IJCAI01*, 2001.
- [BR96] C. Bessiere and J.C. Regin. Mac and combined heuristics: two reasons to forsake fc (and cbj?) on hard problems. In *Proc. CP 96*, pages 61–75, Cambridge MA, 1996.
- [Dec03] Rina Dechter. *Constraint Processing*. Morgan Kaufman, 2003.
- [Gin93] M. L. Ginsberg. Dynamic backtracking. *J. of Artificial Intelligence Research*, 1:25–46, 1993.
- [GW99] I.P. Gent and T. Walsh. Cspplib: a benchmark library for constraints. Technical report, Technical report APES-09-1999, 1999. Available from

- <http://csplib.cs.strath.ac.uk/>. A shorter version appears in the Proceedings of the 5th International Conference on Principles and Practices of Constraint Programming (CP-99).
- [Ham01] Y. Hamadi. Distributed interleaved parallel and cooperative search in constraint satisfaction networks. In *Proc. IAT-01*, Singapore, 2001.
  - [HE80] R. M. Haralick and G. L. Elliott. Increasing tree search efficiency for constraint satisfaction problems. *Artificial Intelligence*, 14:263–313, 1980.
  - [Lam78] L. Lamport. Time, clocks, and the ordering of events in distributed system. *Communication of the ACM*, 2:95–114, April 1978.
  - [Lyn97] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufmann Series, 1997.
  - [MRKZ02] A. Meisels, I. Razgon, E. Kaplansky, and R. Zivan. Comparing performance of distributed constraints processing algorithms. In *Proc. AAMAS-2002 Workshop on Distributed Constraint Reasoning DCR*, pages 86–93, Bologna, July 2002.
  - [MZ07] A. Meisels and R. Zivan. Asynchronous forward-checking for distributed csp. *Constraints*, 12(1), 2007.
  - [Pro96] P. Prosser. An empirical study of phase transitions in binary constraint satisfaction problems. *Artificial Intelligence*, 81:81–109, 1996.
  - [SF05] M. C. Silaghi and B. Faltings. Asynchronous aggregation and consistency in distributed constraint satisfaction. *Artificial Intelligence*, 161:1-2:25–54, January 2005.
  - [SG98] Barbara M. Smith and Stuart A. Grant. Trying harder to fail first. In *European Conference on Artificial Intelligence*, pages 249–253, 1998.
  - [SGM96] G. Solotorevsky, E. Gudes, and A. Meisels. Modeling and solving distributed constraint satisfaction problems (dcsp). In *Constraint Processing-96, (short paper)*, pages 561–2, Cambridge, Massachusetts, USA, October 1996.
  - [Sil06] M. C. Silaghi. Generalized dynamic ordering for asynchronous backtracking on discsp. In *AAMAS 2006, DCR workshop*, Hakodate, Japan, 2006.
  - [Smi96] B. M. Smith. Locating the phase transition in binary constraint satisfaction problems. *Artificial Intelligence*, 81:155 – 181, 1996.
  - [SSHF01] M. C. Silaghi, D. Sam-Haroud, and B. Faltings. Hybridizing abt and awc into a polynomial space, complete protocol with reordering. Technical Report 01/#364, EPFL, May 2001.
  - [Wal05] Richard J. Wallace. Factor analytic studies of csp heuristics. In *CP*, pages 712–726, Sigtes (Barcelona), Spain, 2005.
  - [YDIK98] M. Yokoo, E. H. Durfee, T. Ishida, and K. Kuwabara. Distributed constraint satisfaction problem: Formalization and algorithms. *IEEE Trans. on Data and Kn. Eng.*, 10:673–685, 1998.
  - [Yok95] M. Yokoo. Asynchronous weak-commitment search for solving distributed constraint satisfaction problems. In *Proc. 1st Internat. Conf. on Const. Progr.*, pages 88 – 102, Cassis, France, 1995.
  - [Yok00] M. Yokoo. Algorithms for distributed constraint satisfaction problems: A review. *Autonomous Agents & Multi-Agent Sys.*, 3:198–212, 2000.
  - [ZM05] R. Zivan and A. Meisels. Dynamic ordering for asynchronous backtracking on discsp. In *CP-2005*, pages 32–46, Sigtes (Barcelona), Spain, 2005.
  - [ZM06a] R. Zivan and A. Meisels. Dynamic ordering for asynchronous backtracking on discsp. *Constraints*, 11(2,3):179–197, 2006.
  - [ZM06b] R. Zivan and A. Meisels. Message delay and asynchronous discsp search. *Archives of Control Sciences*, 16(2):221–242, 2006.