Dynamic Ordering for Asynchronous Backtracking on DisCSPs

Roie Zivan · Amnon Meisels

Abstract An algorithm that performs asynchronous backtracking on distributed *CSPs*, with dynamic ordering of agents is proposed, *ABT_DO*. Agents propose reorderings of lower priority agents and send these proposals whenever they send assignment messages. Changes of ordering triggers a different computation of *Nogoods*. The dynamic ordered asynchronous backtracking algorithm uses polynomial space, similarly to standard *ABT*. The *ABT_DO* algorithm with three different ordering heuristics is compared to standard *ABT* on randomly generated *DisCSPs*. A *Nogood-triggered* heuristic, inspired by dynamic backtracking, is found to outperform static order *ABT* by a large factor in run-time and improve the network load.

Keywords Distributed Constraint satisfaction · Distibuted AI · Search

1. Introduction

Distributed constraint satisfaction problems (*DisCSP*s) are composed of agents, each holding its local constraint network. Agents 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. [23, 25]). 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 [2].

Distributed CSPs are an elegant model for many every day combinatorial problems that are distributed by nature. Take for example a large hospital that is composed of many wards. Each ward constructs a weekly timetable assigning its nurses to shifts. The construction of a weekly timetable involves solving a constraint satisfaction problem for each ward. Some of the nurses in every ward are qualified to work

R. Zivan · A. Meisels (✉)
Department of Computer Science, Ben-Gurion University of the Negev,
Beer-Sheva, 84–105, Israel
e-mail: am@cs.bgu.ac.il

R. Zivan
e-mail: zivanr@cs.bgu.ac.il

in the *Emergency Room*. Hospital regulations require a certain number of qualified nurses (e.g., for Emergency Room) in each shift. This imposes constraints among the timetables of different wards and generates a complex Distributed CSP [23].

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 the assignments to all variables of all agents 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 [25, 26] and was developed further and studied in [2, 3, 10, 20]. Agents in the *ABT* algorithm 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* specify the use of polynomial space memory in order to perform backjumping using the resolving of stored explanations, similarly to dynamic backtracking [2, 3]. The versions of asynchronous backtracking presented in all of the above studies use a static priority order among all agents.

In centralized *CSPs*, dynamic variable ordering is known to be an effective heuristic for gaining efficiency [6]. Recent studies have shown that the same is true for algorithms which perform sequential (synchronous) assignments in Distributed *CSPs* [5, 17]. These studies suggest heuristics of agent/variable ordering and empirically show large gains in efficiency over the same algorithms using static order. Thus we have a strong motivation for exploring the possibilities for dynamic re-ordering of asynchronous backtracking.

In [11] the authors present a distributed ordering algorithm, according to the properties of the constraints graph. Once the order is determined, the asynchronous backtracking algorithm uses this fixed order.

An asynchronous algorithm with dynamic ordering was proposed by [24], Asynchronous Weak Commitment (*AWC*). According to [25], *AWC* outperforms *ABT*. However, in order to be complete, *AWC* uses exponential space for storing *Nogoods* which makes it impractical for solving hard instances of even small *DisCSPs*.

An attempt to combine *ABT* with *AWC* was reported by [21]. In order to perform asynchronous finite reordering operations [21] suggest that the reordering operation will be performed by abstract agents. The results presented in [21] show minor improvements to static order *ABT*.

The present paper proposes a simple algorithm for dynamic ordering in asynchronous backtracking, *ABT_DO*. The proposed algorithm uses polynomial space, as standard *ABT*. In the proposed *ABT_DO* algorithm the agents of the *DisCSP* choose orders dynamically and asynchronously. Agents in *ABT_DO* perform according to the current, most updated order they hold. Each agent can change the order of all agents with lower priority. An agent can propose an order change each time it replaces its assignment. Each order is time-stamped according to agents assignment. The method of time-stamping for defining the most updated order is the same that was used in [17] for choosing the most updated partial assignment. A simple array of counters represents the priority of a proposed order, according to the global search tree.

Having established a correct algorithm for dynamic variable ordering in *ABT*, one needs to investigate ordering heuristics. Surprisingly, some of the heuristics which

are very effective for distributed algorithms using a sequential assignment protocol, do not improve the run-time of *ABT*. It turns out that an ordering heuristic, based on *Dynamic Backtracking* [9], is very successful on random *DisCSPs* (see Section 7).

Distributed *CSPs* are presented in Section 2. A description of the standard *ABT* algorithm is presented in Section 3. Section 4 presents a short reminder of variable ordering in standard *CSPs*. Asynchronous backtracking with dynamic ordering (*ABT_DO*) is presented in Section 5. Section 6 introduces a correctness and completeness proof for *ABT_DO*. An extensive experimental evaluation, which compares *ABT* to *ABT_DO* with several ordering heuristics is in Section 7. The experiments were conducted on randomly generated *DisCSPs*.

2. Distributed Constraint Satisfaction

A distributed constraint satisfaction problem—*DisCSP* is composed of a set of k agents $A_1, A_2, \ldots, A_k$. Each agent A_i contains a set of constrained variables X_{i_1}, $X_{i_2}, \ldots, X_{i_{n_i}}$. Constraints or **relations** R are subsets of the Cartesian product of the domains of the constrained variables. For a set of constrained variables $X_{i_k}, X_{j_l}, ..., X_{m_n}$, with domains of values for each variable $D_{i_k}, D_{j_l}, ..., D_{m_n}$, the constraint is defined as $R \subseteq D_{i_k} \times D_{j_l} \times ... \times D_{m_n}$. 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*, the agents variables are constrained with variables that belong to different agents [23, 26]. In addition, each agent has a set of constrained variables, i.e., a *local constraint network*.

An assignment (or a label) is a pair , 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* 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 might be in conflict [2]. 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" [2, 25].

The following assumptions are routinely made in studies of *DisCSPs* and are assumed to hold in the present study [2, 25].

1. All agents hold exactly one variable.
2. The amount of time that passes between the sending and the receiving of a message is finite.
3. Messages sent by agent A_i to agent A_j are received by A_j in the order they were sent.

3. Asynchronous Backtracking (*ABT*)

The *Asynchronous Backtracking* algorithm, was presented in several versions over the last decade and is described here in the form of the more recent papers [2, 25]. In the ABT algorithm, agents hold an assignment for their variables at all times, which is consistent with their view of the state of the system (i.e., the assignments of

their neighboring agents). When the agent cannot find an assignment which is consistent with the assignments of higher priorities neighboring agents, it changes its view by sending a *Nogood* to an agent with a conflicting assignment and eliminating this conflicting assignment from its current view. Then it makes another attempt to assign its variable [2, 25].

The code of the Asynchronous Backtracking algorithm (ABT) is presented in Fig. 1. ABT has a total order of priorities among agents. Agents hold a data structure called *Agent_view* which contains the most recent assignments received from agents with higher priority. The algorithm starts by each agent assigning its variable, and sending the assignment to neighboring agents with lower priority. When an agent receives a message containing an assignment (an **ok?** message [25]), it updates its *Agent_view* with the received assignment and removes inconsistent

when received (**ok?**, (x_j, d_j))
1.   add (x_j, d_j) to *agent_view*;
2.   remove non consistent *Nogoods*;
3.   **check_agent_view**;

when received (**nogood**, x_j, *nogood*)
1.   *old_value* ← *current_value*;
2.   **if**(*nogood* is consistent with *agent_view* ∪ *current_value*)
3.   store *nogood*;
4.   **when** (*nogood* contains an agent x_k that is not a neighbor)
5.   request x_k to add x_i as a neighbor,
6.   add (x_k, d_k) to *agent_view*;
7.   **check_agent_view**;
8.   **when** (*old_value* = *current_value*)
9.   send (**ok?**, $(x_i, current\_value)$) to x_j ;

procedure **check_agent_view**
1.   **when** (*agent_view* and *current_value* are not consistent)
2.   **if** (no value in D_i is consistent with *agent_view*)
3.   **backtrack**;
4.   **else**
5.   select $d \in D_i$ where *agent_view* and d are consistent;
6.   *current_value* ← d;
7.   send (**ok?**,(x_i, d)) to *low_priority_neighbors*;

procedure **backtrack**
1.   *nogood* ← **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.   send (**nogood**, x_i, *nogood*) to x_j;
7.   remove (x_j, d_j) from *agent_view*;
8.   remove all *Nogoods* containing (x_j, d_j);
9.   **check_agent_view**;

Fig. 1 Standard ABT algorithm

Nogoods. Then by calling procedure **check_agent_view**, it checks whether its assignment is still consistent or it needs to be replaced (first procedure in Fig. 1).

Agents that reassign their variable, inform their lower priority neighbors by sending them **ok?** messages (procedure **check_agent_view**, lines 5–7). Agents that cannot find a consistent assignment, send the inconsistent tuple in their *Agent_view* in a backtrack message (a *Nogood* message [25]) and remove from their *Agent_view* the assignment of the lowest priority agent in the inconsistent tuple (procedure **backtrack**). In the simplest form of the *ABT* algorithm, the complete *Agent_view* is sent as a *Nogood* [25]. The *Nogood* is sent to the lowest priority agent whose assignment is included in the *Nogood*. After the culprit assignment is removed from the *AgentView* the agent makes another attempt to assign its variable by calling procedure **check_agent_view** (procedure **backtrack** lines 5–9).

Agents that receive a *Nogood*, check its relevance against the content of their *Agent_view* [2]. If the *Nogood* is relevant the agent stores it, and tries to find a consistent assignment. If the agent receiving the *Nogood* keeps its assignment, it informs the *Nogood* sender by re-sending it an **ok?** message with its assignment. An agent A_i which receives a *Nogood* containing an assignment of agent A_j which is not included in its *Agent_view*, adds the assignment of A_j to its *Agent_view* and sends a message to A_j asking it to add a link between them, i.e., inform A_i about all assignment changes it performs in the future (lines 2–4 in the second procedure of Fig. 1).

The performance of *ABT* can be greatly improved by requiring agents to read all messages they receive before performing computation [2, 25]. This technique was found to improve the performance of *Asynchronous Backtracking* on the harder instances of randomly generated Distributed CSPs by a large factor [5, 27].

Another improvement to the performance of *ABT* can be achieved by using the method for resolving inconsistent subsets of the *Agent_view*, based on methods of dynamic backtrack. A version of *ABT* that uses this method was presented in [2] and was shown to use polynomial space for storing *Nogoods* (explanations). In all the experiments in this paper, a version of *ABT* which includes both of the above improvements is used. Agents read all incoming messages that were received before performing computation and *No goods* are resolved, using the dynamic back-tracking method.

4. Dynamic Ordering in centralized algorithms

Dynamic variable ordering was found to be a powerful tool for gaining efficiency in *CSP* algorithms. Many studies in the last two decades acknowledged this fact and searched for the best ordering heuristic [1, 4, 7, 8, 12]. In all of these studies, dynamic variable ordering is performed by using a heuristic in order to choose the next variable to be assigned. A simple observation on the way the dynamic ordering algorithm traverses the search tree is that the order of the variables in the current partial assignment of the algorithm was chosen one by one, each time the algorithm performed a successful assignment and decides on the next variable [1, 4, 7, 8, 12]. When the algorithm backtracks, it must do so according to the order of the current partial assignment. In a naive algorithm this would mean to backtrack to the last variable assigned. In the case of a conflict based backjumping algorithm [18], the algorithm backtracks to the last variable assigned in the conflicting set of variables,

according to the order the partial assignment was obtained [1, 4, 7, 8, 12]. Two simple rules can be derived from this observation:

1. Whenever the search algorithm moves forward, choose the next variable according to the desired heuristic.
2. Backtrack according to the order in which variables were assigned.

Distributed sequential assignments algorithms [5, 17] implement the same rules while performing dynamic agent ordering.

Although the above observation and rules are straightforward in centralized and sequential assignments algorithms they can form the basic idea that enables asynchronous backtracking to perform dynamic agent ordering. The proposed method of distributed asynchronous dynamic reordering that is described in the next section forces the agents to apply the above rules asynchronously.

5. 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, each time it replaces its assignment. This makes the sending of an ordering proposal message always coincide with an **ok?** message. 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.

Consider an example in which agent A_2 holds the following *Current_order*: $(1, 4)(2, 3)(3, 1)(4, 0)(5, 1)$. There are five agents $A_1 \ldots A_5$ and they are ordered according to their IDs from left to right. After replacing its assignment it changes the order to: $(1, 4)(2, 4)(4, 0)(5, 0)(3, 0)$. In the new order, agent A_1 which had higher priority than A_2 in the previous order keeps its place and the value of its counter does not change. A_2 also keeps its place and the value of its counter is incremented by one. The rest of the agents, which have lower priority than A_2 in the previous order, change places and are still located lower than A_2. The new order for these agents is A_4, A_5, A_3 and their counters 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*.

For simplicity of presentation we assume that agents send **order** messages to all lower priority agents. In the more realistic form of the algorithm, agents send **order** messages only to their lower priority *neighbors*. Both versions are proven correct in Section 6.

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 the agents IDs since the agent that performs the change does not change its own position and the positions of higher priority agents. In the above example the common prefix includes agents A_1 and A_2. Since the agent proposing the new order increases its own counter, when two different orders are compared, at least one of the time-stamp counters in the common prefix is different between the two orders. The more up-to-date order is the one for which the first different counter in the common prefix is larger. In the example above, any agent which will receive the new order will know it is more up-to-date than the previous order since the first pair is identical, but the counter of the second pair is larger.

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* 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 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*.

Figures 2 and 3 present the code of asynchronous backtracking with dynamic ordering (*ABT_DO*).

When an **ok?** message is received (first procedure in Fig. 2), the agent updates the *Agent_view* and removes inconsistent *Nogoods*. Then it calls **check_agent_view** to make sure its assignment is still consistent.

A new order received in an order message is accepted only if it is more up to date than the *Current_order* (second procedure of Fig. 2). If so, the received order is stored and **check_agent_view** is called to make sure the current assignment is consistent with the higher priority assignments in the *Agent_view*.

when received (ok?, (x_j, d_j)**):**
1.   add (x_j, d_j) to $agent\_view$;
2.   remove inconsistent $nogoods$;
3.   **check_agent_view;**

when received (order, $received\_order$**):**
1.   **if** ($received\_order$ is more updated than $Current\_order$)
2. $Current\_order \leftarrow received\_order$;
3.   remove inconsistent nogoods;
4.   **check_agent_view;**

when received (**nogood**, $x_j, nogood$)
1.   **if** ($nogood$ contains an agent x_k with lower priority than x_i)
2.   send (**nogood**, $(x_i, nogood)$) to x_k;
3.   send (**ok?,** $(x_i, current\_value)$ to x_j;
4.   **else**
5.   **if** ($nogood$ consistent with $\{Agent\_view \cup current\_assignment\}$)
6.   store $nogood$;
7.   **if** ($nogood$ contains an agent x_k that is not its neighbor)
8.   request x_k to add x_i as a neighbor;
9.   add (x_k, d_k) to $agent\_view$;
10.   **check_agent_view;**
11.   **else**
12.   send (**ok?,** $(x_i, current\_value)$) to x_j;

Fig. 2 The ABT_DO algorithm (first part)

When a *Nogood* is received (third procedure in Fig. 2) the agent first checks if it is the lowest priority agent in the received *Nogood*, according to the *Current_order*. If not, it sends the *Nogood* to the lowest priority agent and an **ok?** message to the *Nogood* sender (lines 1–3). If the receiving agent is the lowest priority agent it performs the same operations as in the standard *ABT* algorithm (lines 4–12).

Procedure **backtrack** (Fig. 3) is the same as in standard *ABT*. The *Nogood* is resolved and the result is sent to the lowest priority agent in the *Nogood*, according to the *Current_order*.

Procedure **check_agent_view** (Fig. 3) is very similar to the same procedure in standard *ABT* but the difference is important (lines 5–9). If the current assignment is not consistent and must be replaced and a new consistent assignment is found, the agent chooses a new order, according to the algorithms rules and the heuristic used, as its *Current_order* (line 7) and updates the corresponding time-stamp. Next, **ok?** messages are sent to all neighboring agents. The new order and its time-stamp counters are sent to all lower priority agents.

6. Correctness of *ABT_DO*

In order to prove the correctness of the *ABT_DO* algorithm we first establish two facts by proving the following lemmas:

Lemma 1 *The highest priority agent in the initial order remains the highest priority agent in all proposed orders.*

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.   *Current_order* $\leftarrow$ **choose_new_order**
8.   send (**ok?**,(x_i, d)) to *neighbors*;
9.   send (**order**,*Current_order*) to *lower priority agents*;

procedure **backtrack**
1.   *nogood* $\leftarrow$ **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.   send (**nogood**, x_i, *nogood*) to x_j;
7.   remove (x_j, d_j) from *agent_view*;
8.   remove all *Nogoods* containing (x_j, d_j);
9.   **check_agent_view**;

Fig. 3 The two procedures of the ABT_DO algorithm

The proof for Lemma 1 is immediate from the two rules of reordering. Since no agent can propose a new order which changes the priority of higher priority agents and its own priority, no agent including the first can move the highest priority agent to a lower position.

Lemma 2   *When the highest priority agent proposes a new order, it is more up to date than all previous orders.*

This proof is again immediate. In all previous orders the time-stamp counter of the first agent is smaller than the counter of the time-stamp counter of the first agent in the new proposed order.

To prove correctness of a search algorithm for *DisCSPs* one needs to prove that it is sound, complete and that it terminates. *ABT_DO*, like *ABT*, reports a solution when all agents are idle and no messages are sent. Its soundness follows from the soundness of *ABT* (see for example [2]). One point needs mentioning. Since no messages are traveling in the system in the idle state, all overriding messages have arrived at their destinations. This means that for every pair of constraining agents an agreement about their pairwise order has been achieved. One of each pair of constraining agents checks their constraint and no messages mean no violations, as in the proof for *ABT* [2].

Theorem 1   *ABT_DO is complete and it terminates.*

To prove the completeness and termination of *ABT_DO* we use induction on the number of agents (i.e., number of variables) in the *DisCSP*. For a single agent

DisCSP the order is static therefore the completeness and termination of *ABT* implies the same for *ABT_DO*. Assume *ABT_DO* is complete and terminates for every *DisCSP* with k agents where $k < n$. Consider a *DisCSP* with n agents. According to Lemma 1 the agent with the highest priority in the initial order will not change its place. The highest priority agent assigns its variable for the first time and sends it along with its order proposal to other agents. The remaining *DisCSP* has $n - 1$ agents and its initial order is that proposed by the first agent (all other orders are discarded according to Lemma 2). By the induction assumption the remaining *DisCSP* is complete and terminates. If a solution to the induced *DisCSP* is found, this means that the lower priority $n - 1$ agents are idle. So is the first (highest priority) agent since none of the others sends it any message. If a solution is not found, by the $n - 1$ lower priority agents, either an empty *Nogood* was found by one of the agents and the whole search is terminated, or a single assignment *Nogood* will be sent to the highest priority agent which will cause it to replace its assignment. The new assignment of the first agent and the new order proposed will induce a new *DisCSP* of size $n - 1$. The search on this new *DisCSP* of size $n - 1$ is also complete and terminates according to the induction assumption. The number of induced *DisCSPs*, created by the assignments of the highest priority agent is bound by the size of its domain. Therefore, the algorithm will terminate in a finite time. The *ABT_DO* algorithm is complete since a solution to the *DisCSP* must include one of the highest priority agent value assignments, which means that one of the induced *DisCSPs* includes a solution *iff* the original *DisCSP* includes a solution. This completes the correctness proof of *ABT_DO*.

If the network model, or privacy restrictions, enable agents to communicate only with their neighbors in the constraint network, some small changes are needed in order to maintain correctness. First, agents must be allowed to change only the order of lower priority *neighbors*. This means that the method **choose_new_order**, called in line 7 of procedure **check_agent_view**, changes the order by switching between the position of lower priority neighbors and leaving other lower priority agents at their current position. Second, whenever an updated order message is received, an agent informs its neighbors of its new *Current_order*.

In order to prove that the above two changes do not affect the correctness of the algorithm we first establish the correctness of Lemmas 1 and 2 under these changes. Lemma 1 is not affected by the change since the rules for changing agents positions have become more strict, and still do not allow to change the position of higher priority agents. Lemma 2 holds because the time-stamp mechanism which promises its correctness has not changed. These Lemmas are the basis for the correctness of the induction which proves the algorithm is complete and terminates. However, we still need to prove the algorithm is sound. One of the assumptions that our soundness proof depended on was that an idle state of the system would mean that every constrained pair of agents agrees on the order between them. This claim might not hold since the most up to date order is not sent to all agents. The following Lemma proves this claim is still true after the changes in the algorithm:

Lemma 3   *When the system reaches an idle state, every pair of constrained agents hold the same order.*

According to the changes described above, whenever one of the constrained agents receives an updated order message, it informs its neighbors. Therefore, all

Fig. 4 Non-concurrent constraints checks performed by *ABT* and *ABT_DO* using different order heuristics on low density DisCSPs ($p_1 = 0.4$)

agents which have constraints with it will be notified and hold the updated order. If two agents are not informed with the most updated order, this would mean that *both of them* are not lower priority neighbors of the reordering agent and as a result their current position in the order stays the same. Lemma 3 implies that the algorithm is sound for versions of *ABT_DO* that are restricted to send messages only between pairs of constraining agents.

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 [14, 25], and communication load, in the form of the total number of messages sent [14].

Non concurrent steps of computation, are counted by a method similar to the clock synchronization algorithm of [13]. 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 [28]. 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 [13]. If instead of steps of computation, the number of non-concurrent constraints check is counted (*NCCCs*), then the local computational effort of agents in each step is measured [15, 29].

Experiments were 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 [19, 22]). All three sets of experiments were conducted on networks with 20 agents ($n = 20$) each holding exactly one variable, 10 values for each variable ($k = 10$) and two values of constraints density $p_1 = 0.4$ and $p_1 = 0.7$. The tightness value p_2, is varied between

Fig. 5 Number of messages sent by *ABT* and *ABT_DO* on low density DisCSPs ($p_1 = 0.4$)

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.

ABT_DO is compared to the run of standard *ABT*. For ordering variables in *ABT_DO* three different heuristics were used.

1. Random: each time an agent changes its assignment it randomly orders all agents with lower priorities in its *Current_order*.
2. Domain-size: This heuristic is inspired by the heuristics used for sequential assigning algorithms in [5]. Domain sizes are calculated based on the fact that each agent that performs an assignment includes its current domain size in the sent **order** message to all other agents. Every agent that replaces an assignment, orders the lower priority agents according to their domain size from the smallest to the largest.
3. Nogood-triggered: Agents change the order of the lower priority agents only when they receive a *Nogood* which eliminates their current assignment. In this


Fig. 6 Non-concurrent constraints checks performed by *ABT* and *ABT_DO* using different order heuristics on high density DisCSPs ($p_1 = 0.7$)

**Fig. 7** Total number of messages sent by *ABT* and *ABT_DO* on high density DisCSPs ($p_1 = 0.7$)

case the agent moves the sender of the *Nogood* to be in front of all other lower priority agents. This heuristic was first used for dynamic backtracking in centralized *CSPs* [9].

Figure 4 presents the computational effort in number of non-concurrent constraints checks to find a solution, performed by *ABT* and *ABT_DO* using the above three heuristics. The algorithms solve low density *DisCSPs* with density of $p_1 = 0.4$. *ABT_DO* with random ordering slightly improves the results of standard *ABT*. *ABT_DO* that uses domain sizes to order the lower priority agents performs slightly better than the random version. The largest improvement is gained by using the *Nogood-triggered* heuristic. For the hardest *DisCSP* instances, *ABT_DO* with the *Nogood-triggered* heuristic improves the performance of standard *ABT* by a factor of 4.

Figure 5 presents the total number of messages sent by the algorithms for the same problems. While *ABT_DO* with random ordering heuristic shows a small improvement in the run time results over standard *ABT*, it sends more messages.


Fig. 8 Non-concurrent constraints checks performed by *AFC* using different order heuristics on low density DisCSPs ($p_1 = 0.4$)

Fig. 9 Non-concurrent constraints checks performed by *AFC* using different order heuristics on high density DisCSPs ($p_1 = 0.7$)

This can be expected since in *ABT_DO* agents send additional **order** messages and **ok?** messages to all their neighbors while in standard *ABT*, **ok?** messages are sent only to lower priority agents. *ABT_DO* with domain size ordering sends more messages than standard *ABT* but less than the random ordering version. The really interesting result is that *ABT_DO* with the *Nogood-triggered* heuristic *sends fewer messages than ABT*. Counting the additional **ok?** messages (sent to higher priority agents) and the **order** messages, it still sends fewer messages than standard *ABT* on the hardest *DisCSP* instances.

Figures 6 and 7 present the results of the *ABT_DO* algorithm using the same heuristics on high density *DisCSPs* with $p_1 = 0.7$. The factor of improvement in runtime of *ABT_DO* with the *nogood-triggered* heuristic over standard *ABT* is 2. Both algorithms send a similar number of messages. The *min-domain* and the *random* heuristic send more messages than standard *ABT* and perform a similar number of non-concurrent constraint checks.

**Fig. 10** Number of *Nogood* removed as a result of order changes by *ABT_DO* ($p_1 = 0.4$)

Fig. 11 Number of *Nogood* removed as a result of order changes by *ABT_DO* ($p_1 = 0.7$)

7.1. Heuristics Analysis

Dynamic ordering is a powerful heuristic used to improve the run-time of centralized *CSP* algorithms [6, 8] and of distributed *CSP* algorithms with sequential assignments [5, 17]. The results in the previous section, show that dynamic ordering must be combined with the right heuristic in order to improve the run-time and justify the overhead in message load of asynchronous backtracking. Surprisingly, ordering heuristics that are very effective for sequential assigning algorithms are not as effective when they are used by *Asynchronous Backtracking*.

Figures 8 and 9 present the results of the run of a sequential assignment algorithm, Asynchronous Forward Checking (*AFC*) [16]. In *AFC* agents assign their variables sequentially and perform consistency checks against the current partial assignment concurrently. Although the heuristics used by *AFC* are the same heuristics used by *ABT_DO* in the previous section, the results are very different. All dynamic ordering heuristics used by *AFC*, improve the run of static order *AFC*. The best heuristic is the *min-domain* heuristic.


Fig. 12 Non-concurrent constraints checks performed by different versions of the *ABT_DO* ordering heuristics ($p_1 = 0.4$)

Fig. 13 Non-concurrent constraints checks performed by different versions of the *ABT_DO* ordering heuristics ($p_1 = 0.7$)

It is interesting to try and understand the difference between the behavior of asynchronous backtracking with dynamic ordering and that of sequential assignment *DisCSP* algorithms like asynchronous forward checking To achieve some understanding one needs to remember that agents in asynchronous backtracking constantly and asynchronously perform assignments against their current view of the system. The state of the system viewed by an agent, includes its values, pruned by either *Nogood* or some current assignments of higher priority agents. In standard *ABT* a *Nogood* is discarded and its corresponding value is returned to the agent's current domain only when higher priority agents replace their assignments. In Dynamic Ordering *ABT*, *Nogoods* can be discarded due to a change of order even if the assignments included in the *Nogood* are not changed. For example, if agent A_i holds a Nogood *ng* which includes the assignment of a higher priority agent A_j, if agent A_j is moved to a lower priority than A_i, *ng* is no longer valid since values are discarded only when they conflict with assignments of higher priority agents.

Another look at the tested heuristics with the above insight in mind reveals that both the *random* heuristic and the *min-domain* heuristic do not take this property

Fig. 14 Total number of *Nogoods* which are removed as a result of order changes by *ABT_DO* using different versions of ordering heuristics ($p_1 = 0.4$)

Fig. 15 Total number of *Nogoods* which are removed as a result of order changes by *ABT_DO* using different versions of ordering heuristics ($p_1 = 0.7$)

into consideration and that reordering agents according to them may cause the loss of valid *Nogoods*. In contrast, *ABT_DO* using the *Nogood-triggered* heuristic rarely removes *Nogoods* due to changes of order. In *ABT*, an agent sends a *No good* to the lowest priority agent among the higher priority agents whose assignment is included in one of its *Nogoods* [2, 25]. This means that if an agent A_i is moved by agent A_j to a place immediately following A_j, all the assignments of agents that were previously ordered between A_i and A_j are removed from A_i's *Nogoods*. Since these assignments were not involved in the *No good*, all of A_i's previous *No goods* are still valid.

Figures 10 and 11 present the total number of *Nogoods* removed by *ABT_DO*, as a result of order changes. The *Nogood – triggered* heuristic loses a very small number of *Nogoods* as a result of order changes. The number of *Nogoods* removed by the random and the min domain heuristic is much larger.

In sequential assignment algorithms only the next variable to be assigned is selected by the heuristic. In *ABT_DO* all unassigned agents can be reordered. Figures 12 and 13 present a comparison between two versions of the *random* and the *min-domain* heuristics. Each heuristic was performed in two different versions. In one, after each assignment all lower priority agents are reordered according to the heuristic, in the other only the agent which will have the highest priority among the lower priority agents is chosen by the heuristic and the other agents keep their places from the previous heuristic (these heuristics are called *single* in the presented figures). The results are clearly in favor of the *single* version of the heuristics.

Figures 14 and 15 present a possible explanation for these results. It is clear that the *single* versions of the heuristics remove fewer *Nogoods* due to order changes than the heuristics that order all of the lower priority agents.

8. Conclusions

Most of the studies of *Asynchronous Backtracking* used a static order of agents and variables [2, 10, 20, 25]. An exponential space algorithm that used dynamic ordering

has shown improvement in run-time over *ABT* [25]. The only study that suggested dynamic ordering in *ABT* with polynomial space used a complex method including additional abstract agents [21]. The results presented in [21] show a minor improvement compared to standard, static order, *ABT*.

The present study proposes a simple way of performing dynamic ordering in *ABT* with polynomial space. The ordering is performed as in sequential assignment algorithms by each agent changing only the order of agents following it in the current order. A simple method of time-stamping [17] is used to determine the most updated proposed order.

When a heuristic order inspired by dynamic backtracking [9] is used to dynamically reorder agents, there is a significant improvement in run-time and network load over standard *ABT*.

References

1. Bacchus, F., & van Run, P. (1995). Dynamic variable ordering in csps. In *Proc. Principles and Practice of Constraint Programming (CP-95)* (pp. 258–275).
2. Bessiere, C., Maestre, A., Brito, I., & Meseguer, P. (2005). Asynchronous backtracking without adding links: A new member in the abt family. *Artificial Intelligence, 161*(1–2), 7–24 (January).
3. Bessiere, C., Maestre, A., & Messeguer, P. (2001). Distributed dynamic backtracking. In *Proc. Workshop on Distributed Constraint of IJCAI01*.
4. Bessiere, C., & Regin, J. C. (1995). Using bidirectionality to speed up arc-consistency processing. *Constraint Processing (LNCS 923)* (pp. 157–169).
5. Brito, I., & Meseguer, P. (2004). Synchronous,asnchronous and hybrid algorithms for discsp. In *Workshop on Distributed Constraints Reasoning (DCR-04) CP-2004*. Toronto (September).
6. Dechter, R. (2003). *Constraint Processing*. Morgan Kaufman.
7. Frost, D., & Dechter, R. (1994). In search of the best constraint satisfaction search. In *Proc. Twelfth National Conference of Artificial Intelligence (AAAI-94)* (pp. 301–306). Seattle, USA (August).
8. Gent, I. P., MacIntyre, E., Prosser, P., Smith, Barbara M., & Walsh, T. (1996). An empirical study of dynamic variable ordering heuristics for the constraint satisfaction problem. In *Principles and Practice of Constraint Programming* (pp. 179–193).
9. Ginsberg, M. L. (1993). Dynamic backtracking. *Journal of Artificial Intelligence Research, 1*, 25–46.
10. Hamadi, Y. (2001). Distributed interleaved parallel and cooperative search in constraint satisfaction networks. In *Proc. IAT-01*. Singapore.
11. Hamadi, Y., & Bessiere, C. (1998). Backtracking in distributed constraint networks. In *Proc. ECAI-98* (pp. 219–223). Brighton (August).
12. Haralick, R. M., & Elliott, G. L. (1980). Increasing tree search efficiency for constraint satisfaction problems. *Artificial Intelligence, 14*, 263–313.
13. Lamport, L. (1978). Time, clocks, and the ordering of events in distributed system. *Communications of the ACM, 2*, 95–114 (April).
14. Lynch, N. A. (1997). *Distributed Algorithms*. Morgan Kaufmann Series.
15. Meisels, A., Razgon, I., Kaplansky, E., & Zivan, R. (2002). Comparing performance of distributed constraints processing algorithms. In *Proc. AAMAS-2002 Workshop on Distributed Constraint Reasoning DCR* (pp. 86–93). Bologna (July).
16. Meisels, A., & Zivan, R. (2003). Asynchronous forward-checking for distributed csps. In W. Zhang (Ed.), *Frontiers in Artificial Intelligence and Applications*. IOS.
17. Nguyen, T., Sam-Hroud, D., & Faltings, B. (2004). Dynamic distributed backjumping. In *Proc. 5th Workshop on Distributed Constraints Reasoning DCR-04*. Toronto (September).
18. Prosser, P. (1993). Hybrid algorithms for the constraint satisfaction problem. *Computational Intelligence, 9*, 268–299.

19. Prosser, P. (1996). An empirical study of phase transitions in binary constraint satisfaction problems. *Artificial Intelligence, 81*, 81–109.
20. Silaghi, M. C., & Faltings, B. (2005). Asynchronous aggregation and consistency in distributed constraint satisfaction. *Artificial Intelligence, 161*(1–2), 25–54 (January).
21. Silaghi, M. C., Sam-Haroud, D., & Faltings, B. (2001). Hybridizing abt and awc into a polynomial space, complete protocol with reordering. Technical Report 01/#364, EPFL (May).
22. Smith, B. M. (1996). Locating the phase transition in binary constraint satisfaction problems. *Artificial Intelligence, 81*, 155–181.
23. Solotorevsky, G., Gudes, E., & Meisels, A. (1996). Modeling and solving distributed constraint satisfaction problems (dcsps). In *Constraint Processing-96, (short paper)* (pp. 561–562). Cambridge, Massachusetts, USA (October).
24. Yokoo, M. (1995). Asynchronous weak-commitment search for solving distributed constraint satisfaction problems. In *Proc. 1st Intrnat. Conf. on Const. Progr.* (pp. 88–102). Cassis, France.
25. Yokoo, M. (2000). Algorithms for distributed constraint satisfaction problems: A review. *Autonomous Agents & Multi-Agent Sys., 3*, 198–212.
26. Yokoo, M., Durfee, E. H., Ishida, T., & Kuwabara, K. (1998). Distributed constraint satisfaction problem: Formalization and algorithms. *IEEE Trans. Data Kn. Eng., 10*, 673–685.
27. Zivan, R., & Meisels, A. (2003). Synchronous vs asynchronous search on discsps. In *Proc. 1st European Workshop on Multi Agent System, EUMAS*. Oxford (December).
28. Zivan, R., & Meisels, A. (2006). Message delay and asynchronous discsp search. *Archives of Control*, (accepted for publication).
29. Zivan, R., & Meisels, A. (2006). Message delay and discsp search algorithms. *Annals of Mathematics and Artificial Intelligence (AMAI)*, (accepted for publication).