ADOPT-ing: Unifying Asynchronous Distributed Optimization with Asynchronous Backtracking

Marius C. Silaghi
Florida Institute of Technology
MSILAGHI@FIT.EDU

Makoto Yokoo
Kyushu University
YOKOO@IS.KYUSHU-U.AC.JP

Abstract

This article presents an asynchronous algorithm for solving Distributed Constraint Optimization problems (DCOPs). The proposed technique unifies asynchronous backtracking (ABT) and asynchronous distributed optimization (ADOPT) where valued nogoods enable more flexible reasoning and more opportunities for communication, leading to an important speed-up. While feedback can be sent in ADOPT by COST messages only to one predefined predecessor, our extension allows for sending such information to any relevant agent. The concept of valued nogood is an extension by Dago and Verfaille of the concept of classic nogood that associates the list of conflicting assignments with a cost and, optionally, with a set of references to culprit constraints.

DCOPs have been shown to have very elegant distributed solutions, such as ADOPT, distributed asynchronous overlay (DisAO), or DPOP. These algorithms are typically tuned to minimize the longest causal chain of messages as a measure of how the algorithms will scale for systems with remote agents (with large latency in communication). ADOPT has the property of maintaining the initial distribution of the problem. To be efficient, ADOPT needs a preprocessing step consisting of computing a Depth-First Search (DFS) tree on the constraint graph. Valued nogoods allow for automatically detecting and exploiting the best DFS tree compatible with the current ordering. To exploit such DFS trees it is now sufficient to ensure that they exist. Also, the inference rules available for valued nogoods help to exploit schemes of communication where more feedback is sent to higher priority agents. Together they result in an order of magnitude improvement.

1. Introduction

Distributed Constraint Optimization (DCOP) is a formalism that can model problems distributed due to their nature. These are problems where agents try to find assignments to a set of variables that are subject to constraints. The reason for the distribution of the solving process comes from the assumption that only a subset of the agents has knowledge of each given constraint. Nevertheless, in DCOPs it is assumed that agents try to maximize their cumulated satisfaction by the chosen solution. This is different from other related formalisms where agents try to maximize the satisfaction of the least satisfied among them (Yokoo, 1993). It is also different from formalisms involving self-interested agents (which wish to maximize their own utility individually).

The application of the distributed constraint optimization framework to modeling and solving multi-agent meeting scheduling problems is detailed in (Modi & Veloso, 2005; Franzin, Rossi, E.C., & Wallace, 2004; Maheswaran, Tambe, Bowring, Pearce, & Varakantham, 2004; Sultanik, Modi, & Regli, 2006). The application to Distributed Generator Maintenance is described in (Petcu & Faltings, 2006a). An application to oil pipelines is described in (Marcellino, Omar, & Moura, 2007), while an application to traffic light scheduling is described in (Walsh, 2007). These problems have in common the fact that some constraints are originally distributed among involved agents and are difficult to centralize due to privacy or due to other structural issues. Among the techniques for handling DCOPs we mention (Hirayama & Yokoo, 1997; Maheswaran et al.,

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

2004; Petcu & Faltings, 2005b; Ali, Koenig, & Tambe, 2005; Chechetka & Sycara, 2006; Greenstadt, Pearce, Bowring, & Tambe, 2006). ADOPT (Modi, Shen, Tambe, & Yokoo, 2005) is a basic DCOP solver.

ADOPT can be criticized for its strict message pattern that only provides reduced reasoning opportunities. ADOPT works with orderings on agents dictated by some Depth-First Search tree on the constraint graph, and allows cost communication from an agent only to its parent node. In this work we address the aforementioned critiques of ADOPT, showing that it is possible to define a message scheme based on a type of nogoods, called *valued nogoods* (Dago & Verfaillie, 1996; Dago, 1997), which besides automatically detecting and exploiting the DFS tree of the constraint graph coherent with the current order, helps to exploit additional communication leading to significant improvement in efficiency. The examples given of additional communication are based on allowing each agent to send feedback via valued nogoods to several higher priority agents in parallel. The usage of nogoods is a source of much flexibility in asynchronous algorithms. A nogood specifies a set of assignments that conflict with existing constraints (Stallman & Sussman, 1977). A basic version of the valued nogoods consists of associating each nogood with a cost, namely a cost limit violated due to the assignments of the nogood. Valued nogoods that are associated with a list of culprit constraints produce important efficiency improvements. Each of these incremental concepts is described in the following sections.

We start by defining the general DCOP problem, followed by introduction of the immediately related background knowledge consisting of the ADOPT algorithm and use of Depth-First Search trees in optimization. In Section 4.1 we also describe valued nogoods together with the simplified version of valued global nogoods. In Section 5 we present our new algorithm that unifies ADOPT with the older Asynchronous Backtracking (ABT). The algorithm is introduced by first describing the goals in terms of new communication schemes to be enabled. Then the data structures needed for such communication are explored together with the associated flow of data. Finally the pseudo-code and the proof of optimality are provided before discussing other existing and possible extensions. The different versions mentioned during the description are compared experimentally in the last section.

2. Background

Now we introduce in more detail the distributed constraint optimization problems, the ABT and ADOPT algorithms.

2.1 Distributed Constraint Optimization

A DCOP can be viewed as a distributed generalization of the common centralized Weighted Constraint Satisfaction Problems (WCSPs / Σ-VCSP) (Bistarelli, Fargier, Montanari, Rossi, Schiex, & Verfaillie, 1996; Bistarelli, Montanari, & Rossi, 1995; Schiex, Fargier, & Verfaillie, 1995; Bistarelli, Montanari, Rossi, Schiex, Verfaillie, & Fargier, 1999). We now give the definition of WCSPs, since the valued nogood concept we introduce next was initially defined for WCSPs.

Definition 1 (WCSP (Larrosa, 2002; Bistarelli et al., 1996)) *A Weighted CSP is defined by a triplet of sets (X, D, C) and a bound B. X specifies a set of variables $x_1, ..., x_n$ and D specifies their domains: $D_1, ..., D_n$. C=$\{c_1, ..., c_m\}$ is a set of functions, $c_i : D_{i_1} \times ... \times D_{i_{m_i}} \rightarrow \mathbb{N}^{\infty}$ where m_i is the arity of c_i.*

Its solution is $\epsilon* = \underset{\epsilon \in D_1 \times ... \times D_n}{\operatorname{argmin}} \sum_{i=1}^{m} c_i(\epsilon|_{X_i})$, if $\sum_{i=1}^{m} c_i(\epsilon * |_{X_i}) < B$, where $X_i = D_{i_1} \times ... \times D_{i_{m_i}}$.

Definition 2 (DCOP) *A distributed constraint optimization problem (DCOP), is defined by a set of agents $A_1, A_2, ..., A_n$, and a set X of variables, $x_1, x_2, ..., x_n$. Each agent A_i has a set of k_i functions $C_i = \{c_i^1, ..., c_i^{k_i}\}$, $c_i^j : X_{i,j} \rightarrow \mathbb{R}_+$, $X_{i,j} \subseteq X$, where only A_i knows C_i. We assume that x_i can only take values from a domain $D_i = \{1, ..., d\}$.*

Denoting with ϵ an assignment of values to all the variables in X, the problem is to find $\underset{\epsilon}{\operatorname{argmin}} \sum_{i=1}^{n} \sum_{j=1}^{k_i} c_i^j(\epsilon_{|X_{i,j}})$.

For simplification and without loss of generality, one typically assumes that $X_{i,j} \subseteq \{x_1, ..., x_i\}$.

By $\epsilon_{|X_{i,j}}$ we denote the projection of the set of assignments in ϵ on the set of variables in $X_{i,j}$.

2.2 DFS-trees

The primal graph of a DCOP is the graph having the variables in X as nodes and having an arc for each pair of variables linked by a constraint (Dechter, 2003). A Depth-First Search (DFS) tree associated with a DCOP is a spanning tree generated by the arcs used for first visiting each node during some Depth-First Traversal of its primal graph. DFS trees were first successfully used for distributed constraint satisfaction problems in (Collin, Dechter, & Katz, 2000). The property exploited there is that separate branches of the DFS-tree are completely independent once the assignments of common ancestors are decided. Nodes directly connected to a node in a primal graph are said to be its *neighbors*. The *ancestors* of a node are the nodes on the path between it and the root of the DFS tree, inclusively. If a variable x_i is an ancestor of a variable x_j, then x_j is a *descendant* of x_i.

2.3 ADOPT and ABT

ADOPT. ADOPT (Modi et al., 2005) is an asynchronous complete DCOP solver, which is guaranteed to find an optimal solution. Here, we only show a brief description of ADOPT. Please consult (Modi et al., 2005) for more details. First, ADOPT organizes agents into a Depth-First Search (DFS) tree, in which constraints are allowed between a variable and any of its ancestors or descendants, but not between variables in separate sub-trees.

ADOPT uses three kinds of messages: VALUE, COST, and THRESHOLD. A VALUE message communicates the assignment of a variable from ancestors to descendants that share constraints with the sender. When the algorithm starts, each agent takes a random value for its variable and sends appropriate VALUE messages. A COST message is sent from a child to its parent, which indicates the estimated lower bound of the cost of the sub-tree rooted at the child. Since communication is asynchronous, a cost message contains a context, i.e., a list of the value assignments of the ancestors. The THRESHOLD message is introduced to improve the search efficiency. An agent tries to assign its value so that the estimated cost is lower than the given threshold communicated by the THRESHOLD message from its parent. Initially, the threshold is 0. When the estimated cost is higher than the given threshold, the agent opportunistically switches its value assignment to another value that has the smallest estimated cost. Initially, the estimated cost is 0. Therefore, an unexplored assignment has an estimated cost of 0. A cost message also contains the information of the upper bound of the cost of the sub-tree, i.e., the actual cost of the sub-tree. When the upper bound and the lower bound meet at the root agent, then a globally optimal solution has been found and the algorithm is terminated.

ABT. Distributed constraint satisfaction problems are special cases of DCOPs where the constraints c_i^j can return only values in $\{0, \infty\}$. The basic asynchronous algorithm for solving distributed constraint satisfaction problems is asynchronous backtracking (ABT) (Yokoo, Durfee, Ishida, & Kuwabara, 1998). ABT uses a total priority order on agents where agents announce new assignments to lower priority agents using **ok?** messages, and announce conflicts to higher priority agents using **nogood** messages. New dependencies created by dynamically learned conflicts are announced using **add-link** messages. An important difference between ABT and ADOPT is that, in ABT, conflicts (the equivalents of cost) can be freely sent to any higher priority agent.

3. Comparison between the proposed technique (ADOPT-ing) and ADOPT/ABT

ADOPT-ing vs ADOPT The difference starts with adding justifications (SRCs) to ADOPT's messages. This explicitly bundles cost-related data into valued nogoods such that the destination of the nogood (cost) messages can include other agents besides the parent. Internal data management is also different:

1. The DFS tree can be dynamically detected (in the ADOPT-ing version called ADOPT-Y$_{--}$, shown later). It is based only on already used constraints.

2. ADOPT did not have **add-link** messages.

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

3. In ADOPT (as a result of not using SRCs and not having our rules on the order for combination of nogoods) messages could be sent only to the parent rather than to any ancestor.

4. ADOPT could not use explicit max-inference (because it did not maintain SRCs).

5. ADOPT did not maintain data structures (like lr and $lastSent$, presented later) to avoid resending the same message several times and easy the network load.

6. ADOPT did not provide guidelines for using any additional storage other than the minimal ones (ADOPT did not specify/have an equivalent of Lemma 6 with rules for using cost information).

7. New assignments arriving first via nogoods can be detected as such in ADOPT-ing (as in (Silaghi & Faltings, 2004)) while in ADOPT they had to be considered old.

8. The distributed termination detection via maintenance of upper bounds in ADOPT is presented as an option in ADOPT-ing, since the potential usage of ADOPT-ing simulators as WCSP solvers can detect termination by direct detection of quiescence. Also, the technique is presented in the version of ADOPT-ing in (Faltings, 2006) which unifies the original ADOPT termination detection with the solution detection based on spanning trees that we introduced in (Silaghi, Sam-Haroud, & Faltings, 2000, 2001b) for distributed CSPs.

9. The way to prepare the threshold in THRESHOLD messages is different in ADOPT-ing versus ADOPT. ADOPT-ing stores received costs in its ca structure and sends them back as thresholds when their context is reused.

The cost in nogoods are the same as the lower bounds transmitted with ADOPT. The upper bounds used in ADOPT to detect termination can be implemented similarly in ADOPT-ing, but we detected termination by detecting quiescence in the simulator (which theoretically, as confirmed by experiments, does not produce different results between the original implementation and our implementation of ADOPT). Another way to implement the upper bound in ADOPT-ing is described in (Faltings, 2006), namely based on a boolean value that simply specifies whether the upper bound equals the lower bound on that particular nogood.

The performance difference for the ADOPT-ing variants (other than ADOPT) comes from:

1. The sending of nogoods to earlier agent.

2. Improved inference by SRCs when nogoods are sent to earlier agents.

3. The lazy creation of the DFS tree (ensuring short trees).

4. Smaller and more targeted traffic by sending nogoods only where they are most needed.

ADOPT-ing vs ABT Unlike ABT:

- An ADOPT-ing agent may send possibly irrelevant messages to a given predecessor (its parent in the current DFS tree). It does this to guarantee optimality given the non-idempotent aggregation operation of DCOPs.

- The **nogood** messages have an associated cost and justification (SRCs). These are used to find the assignments with the least conflicts in case of an unsatisfiable problem.

- The solution detection based on spanning trees that we introduced for versions of ABT in (Silaghi et al., 2000, 2001b) also publishes the assignments in the found solution, and can lead to termination before quiescence. The optional version described for ADOPT-ing does not pass the local assignments and therefore does not gather in the root agent the assignments of the found solution. Unlike in (Silaghi et al., 2000), it also cannot reach early termination through solutions detected due to asynchronous changes, before quiescence (when partial solutions from agents happen to intersect). This would be possible in ADOPT-ing only if all non-infinite aggregated costs of each given agent would have the same value.

4. Preliminaries

4.1 Cost of nogoods

Previous flexible algorithms for solving distributed constraint satisfaction problems exploit the inference power of nogoods (e.g., ABT, AWC, ABTR (Yokoo, Durfee, Ishida, & Kuwabara, 1992; Yokoo et al., 1998; Silaghi et al., 2001b))[1]. A nogood $\neg N$ stands for a set N of assignments that was proven impossible, by inference, using constraints. If $N = (\langle x_1, v_1\rangle, ..., \langle x_t, v_t\rangle)$ where $v_i \in D_i$, then we denote by $\overline{N}$ the set of variables assigned in N, $\overline{N} = \{x_1, ..., x_t\}$.

4.1.1 Valued Global Nogoods

In order to apply nogood-based algorithms to DCOP, one redefines the notion of nogoods as follows. First, we attach a value to each nogood obtaining a *valued global nogood*. These are a simplified version of Dago&Verfaille's valued nogoods introduced next, and are basically equivalent to the content of COST messages in ADOPT.

Definition 3 (Valued Global Nogood) *A valued global nogood has the form* $[c, N]$*, and specifies that the (global) problem has cost at least* c*, given the set of assignments* N *for distinct variables.*

Given a valued global nogood $[c, (\langle x_1, v_1\rangle, ..., \langle x_t, v_t\rangle)]$, one can infer a *global cost assessment (GCA)* for the value v_t from the domain of x_t given the assignments $S = \langle x_1, v_1\rangle, ..., \langle x_{t-1}, v_{t-1}\rangle$. This GCA is denoted (v_t, c, S) and is semantically equivalent to an applied valued global nogood (i.e., the inference):

$$(\langle x_1, v_1\rangle, ..., \langle x_{t-1}, v_{t-1}\rangle) \rightarrow (\langle x_t, v_t\rangle \text{ has cost } c).$$

The following remark is used in our algorithms whenever an agent receives a valued global nogood.

Remark 1 *Given a valued global nogood* $[c, N]$ *known to some agent, that agent can infer the GCA* (v, c, N) *for any value* v *from the domain of any variable* x*, where* x *is not assigned in* N*, i.e.,* $x \notin \overline{N}$*.*

Proposition 1 (min-resolution) *Given a minimization WCSP, assume that we have a set of GCAs of the form* (v, c_v, N_v) *that has the property of containing exactly one GCA for each value* v *in the domain of variable* x_i *and that for all* k *and* j*, the assignments for variables* $\overline{N_k} \cap \overline{N_j}$ *are identical in both* N_k *and* N_j*. Then one can resolve a new valued global nogood:* $[\min_v c_v, \cup_v N_v]$.

4.1.2 Dago and Verfaille's Valued Nogoods

We would like to allow free sharing of nogoods between agents. The operator for aggregating the weights of constraints in DCOPs is $+$, which is not idempotent (i.e., in general $a+a \neq a$). Therefore a constraint cannot be duplicated and implied constraints cannot be added straightforwardly without modifying the semantic of the problem (which was possible with distributed CSPs (Schiex et al., 1995; Bistarelli et al., 1999))[2]. Two solutions are known. One solution is based on DFS trees (used by ADOPT), while the second is based on justifications. We will use both of them.

Remark 2 (DFS sub-trees) *Given two GCAs* (v, c'_v, S'_v) *and* (v, c''_v, S''_v) *for a value* v *in the domain of variable* x_i *of a minimization WCSP, if one knows that the two GCAs are inferred from different constraints, then one can infer a new GCA:* $(v, c'_v + c''_v, S'_v \cup S''_v)$*. This is similar to what ADOPT does to combine cost messages coming from disjoint problem sub-trees (Modi, Tambe, Shen, & Yokoo, 2002; Collin et al., 2000).*

1. Other algorithms, like AAS, exploit generalized nogoods (i.e., extensions of nogoods to sets of values for a variable), and the extension of the work here for that case is suggested in (Silaghi, 2002).
2. The aggregation method for fuzzy CSPs (a kind of VCSPs) (Schiex et al., 1995) is MIN, being idempotent. Therefore inferred global valued nogoods can be freely added in that framework.

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

The question is how to determine that the two GCAs are inferred from different constraints in a more general setting. This can be done by tagging cost assessments with the identifiers of the constraints used to infer them (the justifications of the cost assessments).

Definition 4 *A* set of references to constraints *(***SRC***) is a set of identifiers, each for a distinct constraint.*

Note that several constraints of a given problem description can be composed in one constraint (in a different description of the same problem).[3]

SRCs help to define a generalization of the concept of *valued global nogood* named *valued nogood* (Dago & Verfaillie, 1996; Dago, 1997).

Definition 5 (Valued Nogood) *A valued* nogood *has the form* $[R, c, N]$ *where* R *is a set of references to constraints having cost at least* c*, given a set of assignments,* N*, for distinct variables.*

Valued nogoods are generalizations of valued global nogoods. Valued global nogoods are valued nogoods whose SRCs contain the references of all the constraints.

Once we decide that a nogood $[R, c, (\langle x_1, v_1\rangle, ..., \langle x_i, v_i\rangle)]$ will be applied to a certain variable x_i, we obtain a cost assessment tagged with the set of references to constraints[4] R, denoted $(R, v_i, c, (\langle x_1, v_1\rangle, ..., \langle x_{i-1}, v_{i-1}\rangle))$.

Definition 6 (Cost Assessment (CA)) *A cost assessment of variable* x_i *has the form* (R, v, c, N) *where* R *is a set of references to constraints having cost with lower bound* c*, given a set of assignments* N *for distinct variables where the assignment of* x_i *is set to the value* v*.*

As for valued nogoods and valued global nogoods, cost assessments are generalizations of global cost assessments.

Remark 3 *Given a valued nogood* $[R, c, N]$ *known to some agent, that agent can infer the CA* (R, v, c, N) *for any value* v *from the domain of any variable* x*, where* x *is not assigned in* N*, i.e., where* $x \notin \overline{N}$*.*

We can now detect and perform the desired powerful reasoning on valued nogoods and/or CAs coming from disjoint sub-trees, mentioned in Remark 2.

Proposition 2 (sum-inference (Dago & Verfaillie, 1996; Dago, 1997)) *A set of cost assessments of type* (R_i, v, c_i, N_i) *for a value* v *of some variable, where* $\forall i, j : i \neq j \Rightarrow R_i \cap R_j = \emptyset$*, and the assignment of any variable* x_k *is identical in all* N_i *where* x_k *is present, can be combined into a new cost assessment. The obtained cost assessment is* (R, v, c, N) *such that* $R{=}\cup_i R_i$*,* $c{=}\sum_i(c_i)$*, and* $N{=}\cup_i N_i$*.*

The min-resolution proposed for GCAs translates straightforwardly for CAs as follows.

Proposition 3 (min-resolution (Dago & Verfaillie, 1996; Dago, 1997)) *Assume that we have a set of cost assessments for* x_i *of the form* (R_v, v, c_v, N_v) *that has the property of containing exactly one CA for each value* v *in the domain of variable* x_i *and that for all* k *and* j*, the assignments for variables* $\overline{N_k} \cap \overline{N_j}$ *are identical in both* N_k *and* N_j*. Then the CAs in this set can be combined into a new valued nogood. The obtained valued nogood is* $[R, c, N]$ *such that* $R{=}\cup_i R_i$*,* $c{=}\min_i(c_i)$ *and* $N{=}\cup_i N_i$*.*

3. For privacy, a constraint can be represented by several constraint references and several constraints of an agent can be represented by a single constraint reference.
4. This is called a *valued conflict list* in (Silaghi, 2002).

4.1.3 VALUED NOGOODS WITH UPPER BOUNDS

The cost associated with a valued nogood corresponds to a low bound on the cost of the constraints indicated by the attached SRC. Recent work suggests to also associate such nogoods with an additional information, namely whether this cost is an exact evaluation of the optimum (i.e., also an upper bound at minimization) (Modi et al., 2005). Note that the upper bounds in ADOPT either have the same value as the lower bound, or is infinite, so it can be replaced by a boolean variable (Faltings, 2006). ADOPT can use this information in its termination detection procedure: an agent terminates when it knows the exact optimum cost of its whole subtree. A valued nogood with exact upper bounds (VNE) takes the form $[R, exact, c, N]$, where $exact$ is a boolean value. A cost assessment for an assignment $x = v$ takes the form $(R, v, exact, c, N)$.

Remark 4 *Sum-inference on "valued nogoods with upper bounds" are similar to the ones for valued nogoods, with the addition that the value "$exact$" of the result is given by a logical AND of its value in the operands (in ADOPT it is an unknown operand that sets it sometimes to false). Min-resolution is also similar to the one on valued nogoods, and the value "$exact$" of the result is given by a logical OR on the its value in the operands with minimal cost.*

Valued nogoods with upper bounds can be used by ADOPT-ing algorithms for facilitating the immediate termination detection in a distributed system. Since simulators can detect termination by direct inspection of the communication channels, maintenance of "$exact$" values does not modify the measured logic-time performance in a simulator.

5. ADOPT with nogoods

We now present a distributed optimization algorithm whose efficiency is improved by exploiting the increased flexibility brought by the use of valued nogoods. The algorithm can be seen as an extension of both ADOPT and ABT, and will be denoted Asynchronous Distributed OPTimization with inferences based on valued nogoods (ADOPT-ing).

As in ABT, agents communicate with **ok?** messages proposing new assignments of the variable of the sender, **nogood** messages announcing a nogood, and **add-link** messages announcing interest in a variable. As in ADOPT, agents can also use **threshold** messages, but their content can be included in **ok?** messages.

For simplicity we assume in this algorithm that the communication channels are FIFO (as enforced by the Internet transport control protocol). Attachment of counters to proposed assignments and nogoods can also be used to ensure this requirement (i.e., older assignments and older nogoods for the currently proposed value are discarded).

5.1 Exploiting DFS trees for Feedback

In ADOPT-ing, agents are totally ordered as in ABT, A_1 having the *highest priority* and A_n the lowest priority. The *target* of a valued nogood is the position of the lowest priority agent among those that proposed an assignment referred by that nogood. Note that simple versions of ADOPT-ing do not need to maintain a DFS tree, but each agent can send messages with valued nogoods to any predecessor and the DFS tree is discovered dynamically. We also propose hybrid versions that can exploit an existing DFS tree. We have identified two ways of exploiting such an existing structure. The first is by having each agent send its valued nogood only to its parent in the tree. The obtained algorithm is equivalent to the original ADOPT. Another way is by sending valued nogoods only to ancestors. This later hybrid approach can be seen as a fulfillment of a direction of research suggested in (Modi et al., 2005), namely communication of costs to higher priority parents.

The versions of ADOPT-ing introduced in this article are differentiated using the notation **ADOPT-**$\mathcal{DON}$. $\mathcal{D}$ shows the destinations of the messages containing valued nogoods. $\mathcal{D}$ has one of the values $\{A, D, Y\}$ where A stands for *all predecessors*, while D and Y stand for *all ancestors in a DFS tree*. Y is as D but for a dynamically discovered DFS tree. $\mathcal{O}$ marks the optimization criteria used by the $sum_inference()$

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

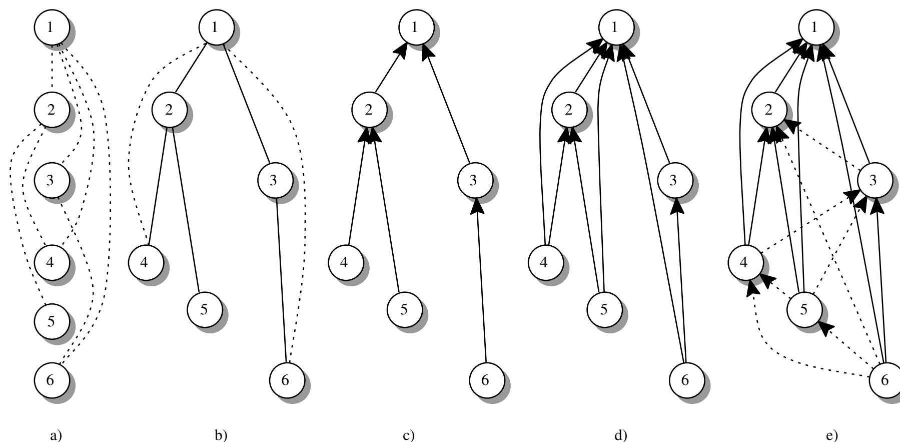

Figure 1: Feedback modes in ADOPT-ing. a) a constraint graph on a totally ordered set of agents; b) a DFS tree compatible with the given total order; c) ADOPT: sending valued nogoods only to parent (graph-based backjumping); d) ADOPT-D__, and ADOPT-Y__: sending valued nogoods to any ancestor in the tree; e) ADOPT-A__: sending valued nogoods to any predecessor agent.

function of Section 5.5 in selecting a nogood when the alternatives have the same cost (if the nogoods cannot be combined with the sum-inference of Proposition 2). For now we use a single criterion, denoted o, which consists of choosing the nogood whose target has the highest priority. $\mathcal{N}$ specifies the type of nogoods employed and has possible value $\{s\}$, where s specifies the use of valued nogoods.

The different schemes are described in Figure 1. The total order on agents is described in Figure 1.a where the constraint graph is also depicted with dotted lines representing the arcs. Each agent (representing its variable) is depicted with a circle. A DFS tree of the constraint graph which is compatible to this total order is depicted in Figure 1.b. ADOPT gets such a tree as input, and each agent sends COST messages (containing information roughly equivalent to a valued global nogood) only to its parent. The versions of ADOPT-ing that replicate this behavior of ADOPT when a DFS tree is provided will continue to be called simply ADOPT (SRCs do not produce any change of behavior in ADOPT since this scheme allows no opportunity to use the sum-inference that they enable). This method of announcing conflicts based on the constraint graph is depicted in Figure 1.c and is related to the classic Graph-based Backjumping algorithm (Dechter, 1990; Hamadi & Bessière, 1998).

In Figure 1.d we depict the nogoods exchange schemes used in ADOPT-D__ and ADOPT-Y__ where, for each new piece of information, valued nogoods are separately computed to be sent to each of the ancestors in the currently known DFS tree. These schemes are strongly boosted by valued nogoods and are shown by experiments to bring large improvements. Sometimes the underscores are dropped to improve readability. As for the initial version of ADOPT, the proof shows that the only mandatory nogood messages for guaranteeing optimality in this scheme are the ones to the parent agent. However, agents can infer from their constraints valued nogoods that are based solely on assignments made by shorter prefixes of the ordered list of ancestor agents. The agents try to infer and send valued nogoods separately for all such prefixes.

Figure 1.e depicts the simple version of ADOPT-ing, when a chain of agents is used instead of a DFS tree (ADOPT-A__), and where nogoods can be sent to all predecessor agents. The dotted lines show messages, which are sent between independent branches of the DFS tree, and which are expected to be redundant. Experiments show that valued nogoods help to remove the redundant dependencies whose introduction would otherwise be expected from such messages. The only mandatory nogood messages for guaranteeing optimality in this scheme are the ones to the immediately previous agent (parent in the chain). However, agents can

infer from their constraints valued nogoods that are based solely on assignments made by shorter prefixes of the ordered list of all agents. As in the other case, the agents try to infer and send valued nogoods separately for all such prefixes. Note that the original ADOPT can also run on any chain of the agents, but our experiments show that its efficiency decreases by 20% when it does not know the shortest DFS tree compatible with the current order, and is an order of magnitude less efficient than any of these two variants of ADOPT-ing. When no DFS tree is known in advance, ADOPT-Y$_{--}$ slightly improves on ADOPT-A$_{--}$ as it dynamically detects a tree with reduced depth.

5.2 Dynamic Discovery of Compatible DFS Tree in ADOPT-Y$_{--}$

Let us now assume that at the beginning, the agents only know the address of the agents involved in their constraints (their neighbors), as in ABT. Finding a DFS tree of a constraint graph is different from the minimal cycle cutset problem, whose distributed solutions have been studied in the past (Jagota & Dechter, 1997). We address the problem of computing a DFS tree *compatible* with a given total order on nodes, namely where the parent of a node precedes that node in the given total order. However, not any given total order on the variables is compatible with a DFS tree of the constraint graph. Given an agreed total order on agents that unknowingly happens to be compatible with a DFS tree, it is relatively simple (less than n rounds) to find the compatible DFS tree. When a compatible DFS tree does not exist, our technique adds a small set of arcs (total constraints) that keep the problem equivalent to the original one and then returns a DFS tree compatible with the new graph.

procedure initPreprocessing*()* do
1.1 $ancestors \leftarrow neighboring\ predecessors$;
    foreach A_j *in* $ancestors$ **do**
1.2   **send DFS**($ancestors$) to A_j;
1.3 parent $\leftarrow$ last agent in $ancestors$;

when *receive* **DFS***(induced) from* A_t **do**
1.4 **if** *(*predecessors in $induced$*)* $\not\subseteq$ $ancestors$ **then**
1.5 $ancestors \leftarrow ancestors \cup$ (predecessors in $induced$);
        foreach A_j *in* $ancestors$ **do**
1.6   **send DFS**($ancestors$) to A_j;
1.7   parent $\leftarrow$ last agent in $ancestors$;

Algorithm 1: Procedures of agent A_i during preprocessing for dynamic discovery of DFS tree.

Preprocessing for computing the DFS tree Algorithm 1 can be used for preprocessing the distributed problem. Each agent maintains a list with its $ancestors$ and starts executing the procedure **initPreprocessing**. The first step consists of initializing its $ancestors$ list with the neighboring predecessors (Line 1.1). The obtained list is broadcast to the known ancestors using a dedicated message named **DFS** (Line 1.2). On receiving a **DFS** message from A_t, an agent discards it when the parameter is a subset of its already known ancestors (Line 1.4). Otherwise the new ancestors induced because of A_t are inserted in the $ancestors$ list (Line 1.5). The new elements of the list are broadcast to all interested ancestors, namely ancestors that will have these new elements as their ancestors (Line 1.6). The parent of an agent is the last ancestor (Lines 1.3,1.7).

Lemma 4 *Algorithm 1 computes a DFS tree compatible with a problem equivalent to the initial DCOP.*

Proof. Let us insert in the initial constraint graph of the DCOP a new total constraint (constraint allowing everything) for each link between an agent and its parent computed by this algorithm, if no constraint existed already. A constraint allowing everything does not change the problem therefore the obtained problem is equivalent to the initial DCOP. Note that the arcs between each agent and its parent define a tree.

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

Now we can observe that there exists a DFS traversal of the graph of the new DCOP that yields the obtained DFS tree. Take three agents A_i, A_j, and A_k such that A_i is the obtained parent of both A_j and A_k. Our lemma is equivalent to the statement that no constraint exists between sub-trees rooted by A_j and A_k (given the arcs defining parent relations).

Let us assume (trying to refute) that an agent $A_{j'}$ in the sub-tree rooted by A_j has a constraint with an agent $A_{k'}$ in the sub-tree rooted by A_k. Symmetry allows us to assume without loss of generality that $A_{k'}$ precedes $A_{j'}$. Therefore $A_{j'}$ includes $A_{k'}$ in its $ancestors$ list and sends it to its parent, which propagates it further to its parent, and so on to all ancestors of $A_{j'}$. Let $A_{j''}$ be the highest priority ancestor of $A_{j'}$ having lower priority than $A_{k'}$. But then $A_{j''}$ will set $A_{k'}$ as its parent (Lines 1.3,1.7), making $A_{k'}$ an ancestor of $A_{j'}$. This contradicts the assumption that $A_{k'}$ and $A_{j''}$ are in different sub-trees of A_i. □

Note that for any given total order on agents, Algorithm 1 returns a single compatible DFS tree. This tree is built by construction, adding only arcs needed to fit the definition of a DFS tree. The removal of any of the added parent links leads to breaking the DFS-tree property, as described in the proof of the Lemma. Therefore, we infer that Algorithm 1 obtains the smallest DFS tree compatible with the initial order.

Remark 5 *The trivial approach to using the DFS construction algorithm as a preprocessing technique also requires the detection of the termination, to launch ADOPT-D__ when the preprocessing terminates. Some of our techniques efficiently avoid such detection.*

The preprocessing algorithm terminates, and the maximal casual chain of messages it involves has a length of at most n. That is due to the effort required to propagate ancestors from the last agent to the first agent. All messages travel only from low priority agents to high priority agents, and therefore the algorithm terminates after the messages caused by the agents in leaves reach the root of the tree[5].

Lemma 5 *If the total order on the agents is compatible with a known DFS tree of the initial DCOP, then all agent-parent arcs defined by the result of the above algorithm correspond to arcs in the original graph (rediscovering the DFS tree).*

Proof. Assume (trying to refute) that an obtained agent-parent relation, A_i–A_j, corresponds to an arc that does not exist in the original constraint graph (for the lowest priority agent A_i obtaining such a parent). The parent A_k of A_i in the known DFS tree must have a higher or equal priority than A_j; otherwise A_i (having A_k in his $ancestors$) would chose it as the parent in Algorithm 1 (Lines 1.3, 1.7). If A_k and A_j are not identical, it means that A_i has no constraint with A_j in the original graph (otherwise, the known DFS would not be correct). Therefore, A_j was received by A_i as an induced link from a descendant A_t which had constraints with A_j (all descendants being defined by original arcs due to the assumption). However, if such a link exists between a descendant A_t and A_j, then the known DFS tree would have been incorrect (since in a DFS pseudo-tree all predecessor neighbors of one's descendants must be ancestors of oneself). This contradicts the assumption and proves the Lemma. □

Remark 6 *If one knows that there exists a DFS tree of the initial constraint graph that is compatible with the order on agents, then the parent of each agent in that tree is its lowest priority predecessor neighbor. The agent can therefore compute its parent from the beginning without any message. This is at the basis of our implementation of the versions of ADOPT that exploit a previously known DFS tree, namely ADOPT-D__ and ADOPT, where we know that the input order is compatible with a DFS tree (being the same order as the one used by ADOPT) but we do not bother providing the tree to the solver. This assumption cannot be used, and is not used in the versions that dynamically discover the DFS tree, such as ADOPT-Y__.*

5. Or roots of the forest.

name	DFS tree	nogoods are sent to
ADOPT	received as input	only the parent in DFS tree
ADOPT-D$_{--}$	received as input	the parent in the DFS tree & the nogood target
ADOPT-A$_{--}$	not used	the predecessor agent & the nogood target
ADOPT-Y$_{--}$	dynamically detected	the current parent in the DFS tree & the nogood target

Table 1: Comparison of ADOPT-ing members.

Dynamic detection of DFS trees Intuitively, detecting a DFS tree in a preprocessing phase has three potential weaknesses which we can overcome. The first drawback is that it necessarily adds a preprocessing of up to n sequential messages. Second, it uses all constraints up-front while some of them may be irrelevant, at least for initial assignments of the agents (and shorter trees can be used to speed up search in the initial stages). Third, trivial DFS tree detection may also require an additional termination detection algorithm. Here we show how we address these issues in one of our next techniques.

Therefore, we propose to build a DFS tree only for the constraints used so far in the search. Therefore, agents in ADOPT-Y$_{--}$ do not start initializing their $ancestors$ with all neighboring predecessors, but with the empty set. Neighboring predecessors are added to the $ancestors$ list only when the constraint defining that neighborhood is actually used to increase the cost of a valued nogood[6]. On such an event, the new $ancestor$ is propagated further as on a receipt of new induced ancestors with a **DFS** message in Algorithm 1. The handling of **DFS** messages is also treated as before. The dynamic detection is run concurrently with the search and integrated with the search, thus circumventing the mentioned weaknesses of the previous version based on preprocessing. The payload of the **DFS** messages is attached to **nogood** messages.

Another problem consists of dynamically detecting the children nodes and how descendants are currently grouped in sub-trees by the dynamic DFS tree. In our solution, A_i groups agents A_k and A_t in the same sub-tree if it detects that its own descendants in the received lists of induced links from A_k and A_t do intersect. This is done as follows. A check is performed each time there is a new descendant agent A_u in the lists of induced links received from a descendant A_k. If A_u was not a previously known descendant of A_i, then A_u is inserted in the sub-tree of A_k. Otherwise, the previous sub-tree containing A_u is merged with the sub-tree containing A_k. Also, a new sub-tree is created for each agent from which we receive a nogood and that was not previously known as a descendant. The data structure employed by an agent A_i for this purpose consists of a vector of n integers, called $subtrees$. $subtrees[j]$ holds the ID of the sub-tree containing A_j, or 0 if A_j is not currently considered to be a descendant of A_i. Each agent generates a different unique ID (positive number) for each of its sub-trees (e.g., by incrementing a counter).

Remark 7 *If agents start ADOPT-Y$_{--}$ by inserting all their predecessor neighbors in their $ancestors$ list, the algorithm becomes equivalent to ADOPT-D$_{--}$ after less than n rounds.*

Experiments show that ADOPT-Y$_{--}$ performs very well according to several metrics, such as total number of messages and number of sequential messages. It is the best ADOPT-ing technique from the point of view of the network traffic, and competes tightly with ADOPT-A$_{--}$ and ADOPT-D$_{--}$ in terms of number of sequential messages.

ADOPT-ing is a framework defining a parametrized family of techniques based on inference with valued nogoods. The differences between the main variants are shown in Table 1. As mentioned above, each such major variant has subversions based on: the method used to combine nogoods, the methodology used to tag nogoods with SRCs.

5.3 Data Structures

Besides the $ancestors$ and $subtrees$ structures of ADOPT-Y$_{--}$, each agent A_i stores its *agent-view* (received assignments) and its *outgoing_links* (agents of lower priority than A_i and having constraints on x_i). The

6. More exactly, when a message is sent to that neighboring agent.

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

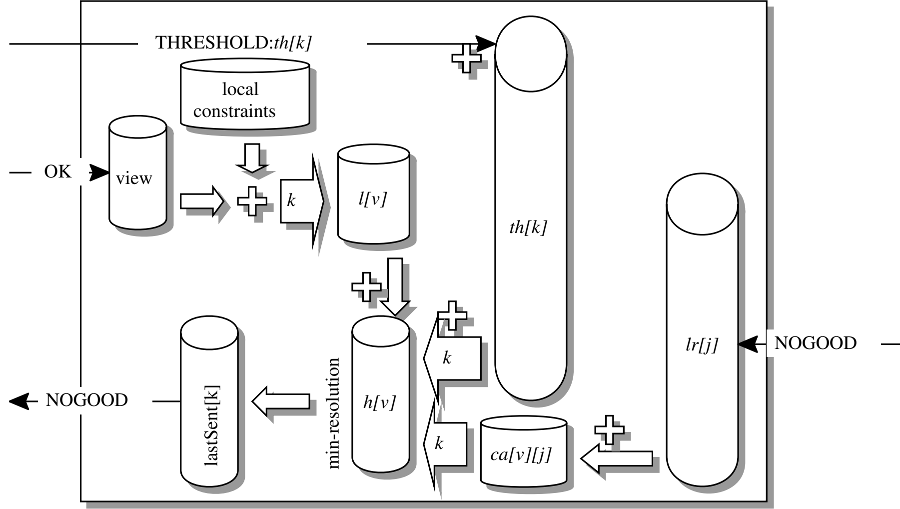

Figure 2: Schematic flow of data through the different data structures used by an agent A_i in ADOPT-ing.

instantiation of each variable is tagged with the value of a separate counter incremented each time the assignment changes. To manage nogoods and CAs, A_i uses matrices l[1..d], h[1..d], ca[1..d][i+1..n], th[1..i], lr[i+1..n] and $lastSent$[1..i-1] where d is the domain size for x_i. crt_val is the current value A_i proposes for x_i. These matrices have the following usage:

- l[k] stores a CA for $x_i = k$, which is inferred solely from the local constraints between x_i and prior variables.
- ca[k][j] stores a CA for $x_i = k$, which is obtained by sum-inference from valued nogoods received from A_j.
- th[k] stores nogoods coming via **threshold/ok?** messages from A_k.
- h[k] stores a CA for x_i=k, which is inferred from ca[k][j], l[k] and th[t] for all t and j.
- lr[k] stores the last valued nogood received from A_k.
- $lastSent$[k] stores the last valued nogood sent to A_k.

The names of the structures were chosen by following the relation of ADOPT with A* search (Silaghi, 2003a; Silaghi, Landwehr, & Larrosa, 2004). Thus, h stands for the “heuristic” estimation of the cost due to constraints maintained by future agents (equivalent to the $h()$ function in A*) and l stands for the part of the standard $g()$ function of A* that is “local” to the current agent. Here, as in ADOPT, the value for $h()$ is estimated by aggregating the costs received from lower priority agents. Since the costs due to constraints of higher priority agents are identical for each value, they are irrelevant for the decisions of the current agent. Thus, the function $f()$ of this version of A* is computed combining solely l and h. We currently store the result of combining h and l in h itself to avoid allocating a new structure for $f()$.

The structures lr and th store received valued nogoods, and ca stores intermediary valued nogoods used in computing h. The reason for storing lr, th and ca is that change of context may invalidate some of the nogoods in h while not invalidating each of the intermediary components from which h is computed. Storing these components (which is optional) saves some work and offers better initial heuristic estimations after a change of context. The cost assessments stored in $ca[v][j]$ of A_i also maintain the information needed for

threshold messages, namely the heuristic estimate for the value v of the variable x_i at successor A_j (to be transmitted to A_j if the value v is proposed again).

The array $lastSent$ is used to store at each index k the last valued nogood sent to the agent A_k. The array lr is used to store at each index k the last valued nogood received from the agent A_k. Storing them separately guarantees that in case of changes in context, they are discarded at the recipient only if they are also discarded at the sender. This property guarantees that an agent can safely avoid retransmitting to A_k messages duplicating the last sent nogood, since if it has not yet been discarded from $lastSent[k]$, then the recipients have not discarded it from $lr[k]$ either.

5.4 Data flow in ADOPT-ing

The flow of data through these data structures of an agent A_i is illustrated in Figure 2. Arrows $\Leftarrow$ are used to show a stream of valued nogoods being copied from a source data structure into a destination data structure. These valued nogoods are typically sorted according to some parameter such as the source agent, the target of the valued nogood, or the value v assigned to the variable x_i in that nogood (see Section 5.3). The $+$ sign at the meeting point of streams of valued nogoods or cost assessments shows that the streams are combined using sum-inference. The $\stackrel{+}{\Leftarrow}$ sign is used to show that the stream of valued nogoods is added to the destination using sum-inference, instead of replacing the destination. When computing a nogood to be sent to A_k, the arrows marked with $\boxed{<k}$ restrict the passage to allow only those valued nogoods containing solely assignments of the variables of agents $A_1, ..., A_k$. Our current implementation recomputes the elements of h and l separately for each target agent A_k by discarding the previous values.

5.5 ADOPT-ing pseudo-code and proof

The pseudo-code for the procedures in ADOPT-ing is given in Algorithms 2 and 3. To extract the cost of a CA, we introduce the function $cost()$, where $cost((R, v, c, N))$ returns c. The $min_resolution(j)$ function applies the min-resolution over the CAs associated with all the values of the variable of the current agent, but uses only CAs having no assignment from agents with lower priority than A_j. More exactly, it first recomputes the array h using only CAs in ca and l that contain only assignments from $A_1, ..., A_j$, and then applies min-resolution over the obtained elements of h. In the current implementation, we recompute l and h at each call to $min_resolution(j)$. An optimization is possible here, reusing the result[7] of computing $min_resolution(k-1)$ in the computation of $min_resolution(k)$ for $k < parent$ by adding only nogoods on x_k to it. Experiments show that this brings minor 4% improvements in simulator time (local computations) on hard problems.

The $sum_inference()$ function used in Algorithm 3 applies the sum-inference to its parameters whenever this is possible (it detects disjoint SRCs). Otherwise, it selects the nogood with the highest cost or the one whose lowest priority assignment has the highest priority (this has been previously used in (Bessiere, Brito, Maestre, & Meseguer, 2005; Silaghi et al., 2001b)). The function $vn2ca(vn, i)$ transforms a valued nogood vn in a cost assessment for x_i. Its inverse is function $ca2vn$. If vn has no assignment for x_i, then a cost assessment can be obtained according to Remark 3. The function $vn2ca(vn, i, v)$ translates vn into a cost assessment for the value v of x_i, using the technique in Remark 3 if needed. The function $target(N)$ gives the index of the lowest priority variable present in the assignment of nogood N. As with file expansion, when "*" is present in an index of a matrix, the notation is interpreted as the set obtained for all possible values of that index (e.g., ca[v][*] stands for $\{ca[v][t] \mid \forall t\}$). Given a valued nogood ng, the notation $ng_{|v}$ stands for vn2ca(ng) when ng's value for x_i is v, and $\emptyset$ otherwise.

5.5.1 Pseudo-code

This sub-section explains line by line the pseudocode in Algorithms 2 and 3. Each agent A_i starts by calling the init() procedure in Algorithm 3, which at Line 3.1 initializes l with valued nogoods inferred from

7. From applying Step 2 of Remark 8.

when *receive* **ok?**($\langle x_j, v_j \rangle$, tvn) **do**
2.1 integrate($\langle x_j, v_j \rangle$);
2.2 **if** *(tvn no-null and has no old assignment)* **then**
2.3 k:=target(tvn); // threshold tvn as common cost;
2.4 th[k]:=sum-inference(tvn,th[k]);
2.5 check-agent-view();

when *receive* **add-link**($\langle x_j, v_j \rangle$) *from* A_j **do**
2.6 add A_j to outgoing_links;
2.7 **if** ($\langle x_j, v_j \rangle$) is old, **send** new assignment to A_j;

when *receive* **nogood**(rvn, *t*, *inducedLinks*) *from* A_t **do**
2.8 insert new predecessors from $inducedLinks$ in $ancestors$, on change making sure interested predecessors will be (re-)sent **nogood** messages; //needed only in ADOPT-Y__;
2.9 **foreach** *new assignment a of a linked variable* x_j *in* rvn **do**
2.10 integrate(a); // counters show newer assignment;
2.11 lr[t]:=rvn;
2.12 **if** *(an assignment in rvn is outdated)* **then**
2.13 **if** *(some new assignment was integrated now)* **then**
2.14 check-agent-view();
2.15 return;
2.16 **foreach** *assignment a of a non-linked variable* x_j *in* rvn **do**
2.17 send **add-link**(a) to A_j;
2.18 **foreach** *value* v *of* x_i *such that* $rvn_{|v}$ *is not* $\emptyset$ **do**
2.19 $vn2ca(rvn, i, v) \rightarrow rca$ (a CA for the value v of x_i);
2.20 ca[v][t]:=sum-inference(rca,ca[v][t]);
2.21 update h[v] and retract changes to ca[v][t] if h[v]'s cost decreases;
2.22 check-agent-view();

Algorithm 2: Receiving messages of A_i in ADOPT-ing

local (unary) constraints. The agent assigns x_i to a value with minimal local cost, crt_val (Line 3.2), announcing the assignment to lower priority agents in *outgoing_links* (Line 3.3). The *outgoing_links* of an agent A_i initially holds the address of the agents enforcing constraints that involve the variable x_i. The agents answer to any received message with the corresponding procedure in Algorithm 2: "**when** receive **ok?**," "**when** receive **nogood**," and "**when** receive **add-link**."

When a new assignment of a variable x_j is learned from **ok?** or **nogood** messages, valued nogoods based on older assignments for the same variables are discarded (Lines 2.1,2.10) by calling the function *integrate()* in Algorithm 3. Within this function, all valued nogoods (cost assignments) stored by the agent are verified and those that contain an old assignment of x_j, which is no longer valid, are deleted (Line 3.17). Any discarded element of ca is recomputed from lr. Namely, if a cost assessment $ca[v][t]$ is deleted in this process while $lr[t]$ remains valid, the agent attempts to apply the nogood in $lr[t]$ to the value v and the obtained cost assessment is copied in $ca[v][t]$ (Line 3.18). This application of the nogood $lr[t]$ to v is possible either if it contains $x_i = v$ or if it contains no assignment for the variable x_i of the current agent (Remark 3). Eventually the new assignment is stored in the agent-view (Line 3.19).

Further, when an **ok?** message is received, it is checked for valid threshold nogoods (Line 2.2). The target k of any such nogood, i.e., the position of the owner of the lowest priority variable, is extracted at Line 2.3 with a procedure called *target*, to detect the place where the nogood should be stored. The newly received threshold nogood is stored at $th[k]$ by sum-inference with the current nogood found there (Lines 2.4,3.21). If no nogood is found in $th[k]$, the new nogood is simply copied there (Line 3.20). If a nogood is already stored in $th[k]$,

but its SRC intersects the one in the new nogood, then the behavior depends on the version of ADOPT-ing. Our pseudo-code illustrates the versions ADOPT-_o_, where the valued nogoods with the highest cost are retained (Line 3.22). In case of a tie, the one with the smallest target is maintained (Line 3.23) (Bessiere et al., 2005; Silaghi et al., 2001b).

After receiving a new value, like in ABT, the *check-agent-view* procedure is used to select a value or detect nogoods (Line 2.5). In this procedure, the agent first tries to compute a nogood for each of its predecessors (Line 3.4). For each such destination, a separate nogood is computed in l for each value v by considering only local constraints with that target agent and with its predecessors. Then, by considering these nogoods of l and all cost assessments in ca based only on assignments from the target agent and its predecessors, new elements of h are computed by sum-inference (Line 3.5). The order of the steps used in this computation is important for correctness and is described in detail later, in Remark 8. If all values of x_i have non-zero cost nogoods in h (Line 3.6), then all elements of h are combined via min-resolution and a nogood vn is obtained for the currently targeted destination (Line 3.7). If some value of x_i has a zero cost nogood in h, the nogood obtained by min-inference has a zero cost. When termination detection is based on $exact$ fields in VNEs, zero cost nogoods must still be sent to the parent/immediate predecessor (Line 3.6), generating some additional network traffic but without effect on the rest of the data structures and operations. However, the nogood vn is sent only if it is different from the last nogood sent to that same agent (Line 3.8). Repeating its sending would be redundant since the recipient holds it in its lr vector. In the versions described here, while not generally required for correctness in ADOPT-ing, the nogood is sent only if the lowest priority variable involved in it is the same as the one controlled by the destination (Line 3.9). The nogood is always sent to the parent in the DFS tree (with ADOPT-D__, and ADOPT-Y__) which is the immediate predecessor with ADOPT-A__. With ADOPT-Y__, when a nogood is sent for the first time to an agent A_k[8], A_k is added to the list $ancestors$ (Line 3.10). After the nogood is sent (Line 3.11), it is stored in $lastSent$ to help avoid immediate retransmission (Line 3.12). If some change was recently made to the $ancestors$ list, the change is propagated at Line 3.13 to all the ancestors that had not already been notified with **nogood** messages at Line 3.11.

The second part of the *check-agent-view* procedure deals with selecting opportunistically a value with the smallest estimated cost (Line 3.14), as common in ADOPT and ABT. We used the common mathematical notation $argmin_v(f(v))$ to denote a computation that returns the value v minimizing the function $f(v)$ passed as the parameter (here $cost(h[v])$). In case of a tie with the old value of x_i, our implementation of $argmin$ prefers to maintain the old value. If the value selected for x_i is different from the old value (Line 3.15), the new value is sent to all agents in $outgoing_links$ (Line 3.16).

When **nogood** messages are received, in the ADOPT-Y__ version we first insert new received induced links into $ancestors$ (Line 2.8). If the set of $ancestors$ was changed by this operation, we set a flag to make sure that *check-agent-view* is eventually called and will propagate the change to all current ancestors. The agent checks if the transported nogood has newer assignments than the ones it already knows. A new assignment can reach an agent as part of a nogood before the corresponding **ok?** message. This can be handled in two ways:

i The original solution of ADOPT and ABT (Yokoo et al., 1998; Modi et al., 2005) is to consider any assignment in a nogood that is different from the assignment known for that variable as being invalid. Assignments are re-announced after each received valid message. Therefore, later retransmission[9] of the nogood triggered by this scheme is guaranteed to correctly deliver each nogood eventually.

ii The other scheme identifies new assignments in **nogood** messages as such, and validates the nogoods on their first reception. The mechanism was used in several versions of ABT (Silaghi & Faltings, 2004). It works by letting each agent maintain a separate counter for each variable. The counter is incremented when the assignment is changed and tags each sent assignment. Each agent stores the last value of the counter it sees for each variable. An agent detects a new assignment by comparing its tag

8. Because the corresponding constraint increases for the first time the cost of the computed nogood.
9. Assuming no mechanism is used to block immediate retransmission of nogoods, such as our $lastSent$ structure.

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

with the previously seen value of that counter. Once detected (Line 2.9), new assignments in nogoods are integrated as on the arrival of their **ok?** message (Line 2.10).[10]

The last nogood received from some agent A_j is stored in $lr[j]$ (Line 2.11), such that it would not be lost as long as it is stored by A_j in its $lastSent$ (otherwise deadlocks could occur).[11] If some assignment in a nogood is considered old at Line 2.12 (with any mentioned scheme) the handling of the nogood is stopped and the nogood is discarded (Line 2.15). However, if some new assignment was integrated at Line 2.10, then the rest of the processing normally executed on **ok?** messages is performed by calling the *check-agent-view* procedure at Lines 2.13,2.14.

If a received nogood contains a variable not previously involved in constraints with the variable of the agent (Line 2.16), an **add-link** message is sent to the agent owning that variable (Line 2.17) to announce the creation of a new link between the two agents (Line 2.6) and to request updates on the values of that variable (Line 2.7). In ADOPT-ing, the assignment received in the nogood is attached to the **add-link** message. This allows the owner of that variable to spare a message by not sending this assignment to A_i if the assignment is still valid.

An agent can receive a nogood where its variable is not present and therefore where the nogood can be applied to all its values. Valid nogoods are projected on all values of A_i (Lines 2.18,2.19), and the result is added to the corresponding cost assessments in ca using the sum-inference procedure (Line 2.20). It is possible that by the quirks of the impact of disjoint SRCs on sum-inference, the addition of a new nogood leads to the decrease of the cost of the obtained cost assessment for the corresponding value. We prefer to enforce a monotonic behavior by withdrawing changes to ca in such situations (Line 2.21). For this purpose, the evaluation of the modification of the cost is done by computing h as when messages are prepared for the parent in the DFS tree (or immediate predecessor). After integrating the new nogood, *check-agent-view* is called at Line 2.22 to infer new nogoods and to select the best value of x_i.

5.5.2 Proof of Termination and Optimality

We prove the termination by induction on increasing sets of agents, starting from the lowest priority ones (suffixes of the list of agents). Then the solution optimality is proven based on a similar intermediary proof of knowledge at quiescence by each agent about the optimal cost for its successors.

Received nogoods are stored in matrices lr and th (Algorithm 2). A_i always sets its crt_val to the index with the lowest CA cost in vector h (preferring the previous assignment in case of ties). On each change that propagates to h, and for each ancestor A_j (or higher priority agent in versions not using DFS trees), the elements of h are recomputed separately by min-resolution(j) to generate new nogoods for A_j. The simultaneous generation and use of multiple nogoods is already known to be useful for the constraint satisfaction case (Yokoo & Hirayama, 1998).

The threshold valued nogood tvn delivered with **ok?** messages sets a common cost on all values of the receiver (see Remark 3), effectively setting a threshold on costs below which the receiver does not change its value. This achieves the effect of THRESHOLD messages in ADOPT.

The procedure described in the following remark is used in the proof of termination and optimality.

Remark 8 *The order of combining CAs to get h at Line 3.5 matters. To compute h[v]:*

1. *a) When maintaining DFS trees, for each value v, CAs are combined separately for each set s of agents defining a DFS sub-tree of the current node:*
   tmp[v][s]=sum-inference$_{t\in s}$(ca[v][t]).
   b) Otherwise, with ADOPT-A$_{--}$, we act as if we have a single sub-tree:
   tmp[v]=sum-inference$_{t\in[i+1,n]}$(ca[v][t]).

---

10. Assignments having the same value are considered identical, even if their tag differs (allowing for re-using old nogoods).
11. Note that with the first scheme (i), where assignments are not tagged with counters, ADOPT-ing should not delete old nogoods from lr (which is done with the second scheme), but checks them when **ok?** messages are received.

2. *CAs from step 1 (a or b) are combined:*
   In case (a) this means: $\forall v, s; h[v]$=sum-inference$_{\forall s}$(tmp[v][s]).
   Note that the SRCs in each term of this sum-inference are disjoint and therefore we obtain a valued nogood with cost given by the sum of the individual costs obtained for each DFS sub-tree.

   For case (b) we obtain h[v]=tmp[v].
   This makes sure that at quiescence the cost of $h[v]$ is at least equal to the total cost obtained at the next agent.

3. *Add l[v]: h[v]=sum-inference(h[v], l[v]).*

4. *Add threshold: h[v]=sum-inference(h[v], th[*]).*

Note that method (a) at Step 1 can be applied only to ADOPT-Y$_{--}$ and ADOPT-D$_{--}$ while method (b) can be applied to all versions. Experiments show that, when applicable, method (a) works only slightly (i.e. 1%) better than method (b).

Lemma 6 (Infinite Cycle) *At a given agent, assume that the agent-view no longer changes and that its array h (used for min-resolution and for deciding the next assignment) is computed only using cost assessments that are updated solely by sum-inference. In this case the costs of the elements of its h cannot be modified in an infinite cycle due to incoming valued nogoods.*

Proof. Valued nogoods that are updated solely by sum-inference have costs that can only increase (which can happen only a finite number of times). For a given cost, modifications can only consist of modifying assignments to obtain lower target agents, which again can happen only a finite number of times. Therefore, after a finite number of events, the cost assessments used to infer h will not be modified any longer and therefore h will no longer be modified. □

Corollary 6.1 *If ADOPT-ing uses the procedure in Remark 8, then for a given agent-view, the elements of the array h for that agent cannot be modified in an infinite cycle.*

Remark 9 *Since lr contains the last received valued nogoods via messages other than* **ok?** *messages, which change the agent-view, that array is updated by assignment with recently received nogoods without sum-inference. Therefore, it cannot be used directly to infer h.*

Note that with the described procedure, a newly arriving valued nogood can decrease the cost of certain elements of h (even if it does not decrease the cost of any of the elements from which h is computed). This is because, while increasing the cost of some element in ca, it can also modify its SRC and therefore forbid its composition by sum-inference with other cost assessments.

Remark 10 (Obtaining Monotonic Increase) *One can avoid the undesired aforementioned effect, where incoming nogoods decrease costs of elements in h. Namely, after a newly received valued nogood is added by sum-inference to the corresponding element of $ca[v]$ for some value v, if the cost of $h[v]$ decreases, then the old content of $ca[v]$ can be restored. Each new valued nogood is used for updating lr. On each change to some element in ca, one has to add to ca the elements found in lr and coming from children in the DFS tree (if they do not lead to a decrease in the cost of h). Experiments show that this technique can bring a small improvement of up to 2% in the number of cycles.*

Intuitively, the convergence of ADOPT-ing can be noticed from the fact that valued nogoods can only monotonically increase valuation for each subset of the search space, and this has to terminate since such valuations can be covered by a finite number of values. If agents A_j, $j<i$ no longer change their assignments, valued nogoods can only monotonically increase at A_i for each value in D_i: costs of the nogoods only increase since they only change by sum-inference.

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

```
procedure init do
3.1    h[v] := l[v]:=initialize CAs from unary constraints;
3.2    crt_val=argmin_v(cost(h[v]));
3.3    send ok?(⟨x_i, crt_val⟩,∅) to all agents in outgoing_links;
procedure check-agent-view() do
3.4    for every A_j with higher priority than A_i (respectively ancestor in the DFS tree, when one is main-
       tained) do
3.5        for every(v ∈ D_i) update l[v] and recompute h[v];
                   // with valued nogoods using only instantiations of {x_1, ..., x_j};
3.6        if ((j is parent/immediate predecessor (VNE)) or (h has non-null cost CA for all values of D_i)) then
3.7            vn:=min_resolution(j);
3.8            if (vn ≠ lastSent[j]) then
3.9                if ((target(vn) == j) or (j is parent/immediate predecessor)) then
3.10                   add j to ancestors (updating parent);// for ADOPT-Y__;
3.11                   send nogood(vn,i,ancestors) to A_j;
3.12                   lastSent[j] = vn;
3.13               on new ancestors, send nogood(∅,i,ancestors) to each ancestor A_h not yet announced;
3.14   crt_val=argmin_v(cost(h[v]));
3.15   if (crt_val changed) then
3.16       send ok?(⟨x_i, crt_val⟩, ca2vn(ca[crt_val][k]),i)
               to each A_k in outgoing_links;
procedure integrate(⟨x_j, v_j⟩) do
3.17   discard elements in ca, th, lastSent and lr based on other values for x_j;
3.18   use lr[t]_|v to replace each discarded ca[v][t];
3.19   store ⟨x_j, v_j⟩ in agent-view;
function sum-inference(vng1, vng2)
3.20   if either vng1 or vng2 has cost 0 then
           return the other one;
3.21   if vng1 and vng2 have disjoint SRCs then
           return the result of applying sum-inference on them;
3.22   if vng1 and vng2 have different costs then
           return the one with lower cost;
3.23   if vng1 and vng2 have different targets then
           return the one with smaller target;
       return vng1;
```

Algorithm 3: Procedures of A_i in ADOPT-ing

Lemma 7 *ADOPT-ing terminates in finite time.*

Proof. Given the list of agents $A_1, ..., A_n$, define the suffix of length m of this list as the last m agents. Then the result follows immediately by induction for an increasingly growing suffix (increasing m), assuming the other agents reach quiescence.

The basic case of the induction (for the last agent) follows from the fact that the last agent terminates in one step if the previous agents do not change their assignments.

Let us now assume that the induction assertion is true for a suffix of k agents. Based on this assumption, we now prove the induction step, namely that the property is also true for a suffix of $k+1$ agents: For each

assignment of the agent A_{n-k}, the remaining k agents will reach quiescence, according to the assumption of the induction step; otherwise, the assignment's CA cost increases. By construction, costs for CAs associated with the values of A_{n-k} can only grow (see Remark 10). Even without the technique in Remark 10, costs for CAs associated with the values of A_{n-k} will eventually stop being modified as a consequence of Lemma 6. After values are proposed in turn and the smallest cost reaches its highest estimate, agent A_{n-k} selects the best value and reaches quiescence. The other agents reach quiescence according to the assumption of the induction step. □

Lemma 8 *The last valued nogoods sent by each agent additively integrate the non-zero costs of the constraints of all of the agent's successors (or descendants in the DFS tree when a DFS tree is maintained).*

Proof. At quiescence, each agent A_k has received the valued nogoods describing the costs of each of its successors (or descendants in the DFS tree when a DFS tree is maintained).

The lemma results by induction for an increasingly growing suffix of the list of agents (in the order used by the algorithm): It is trivial for the last agent.

Assuming that it is true for agent A_k, it follows that it is also true for agent A_{k-1} since adding A_{k-1}'s local cost to the cost received from its children in the tree (A_k for ADOPT-A$_{--}$) will be higher (or equal when removing zero costs) than the result of adding A_{k-1}'s local cost to that of any descendants of those children. Respecting the order in Remark 8 guarantees that this value is obtained (according to the assumption of the induction step, costs from children will be higher than the ones from their descendants and prevail at Step 1, and therefore the result of Step 2 is the sum of the costs of the children). Therefore, the sum between the local cost and the last valued nogood coming from its children defines the last valued nogood sent by A_{k-1}.
□

Theorem 9 *ADOPT-ing returns an optimal solution.*

Proof. We prove by induction on an ever increasing suffix of the list of agents that this suffix converges to a solution that is optimal for the union of the sub-problems of the agents in that suffix.

The induction step is immediate for the suffix composed of the agent A_n alone. Assume now that it is true for the suffix starting with A_k. Following the previous two lemmas, one can conclude that at quiescence, A_{k-1} knows exactly the minimal cumulated cost of the problems of its successors for its chosen assignment, and therefore knows that this cumulated cost cannot be better for any of its other values.

Since A_{k-1} has selected the value leading to the best sum of costs (between its own local cost and the costs of all subsequent agents), it follows that the suffix of agents starting with A_{k-1} converged to an optimal solution for the union of their sub-problems. □

The space complexity is basically the same as for ADOPT. The SRCs do not change the space complexity of the nogoods. The largest space is required by the data structure used for storing potential payloads of future (equivalents of) THRESHOLD messages.

Theorem 10 *The space complexity of an agent in ADOPT-ing is O(dn^2).*

Proof. In an agent, A_i, the space for storing the $outgoing_links$, and the agent view (assignments) is linear in n, having at most one link and one assignment per agent. Six data structures in ADOPT-ing store valued nogoods (l[1..d], ca[1..d][i+1..n], th[1..i], h[1..d], lr[i+1..n], $lastSent$[1..i-1]). Therefore the space complexity is given by the complexity of the largest of them, ca, which stores O(dn) cost assessments that can be sent as threshold nogoods.

Each valued nogood contains a list of up to n assignments and a list of up to n SRCs, its space being linear in n. Therefore the total space requirement for an agent is O(dn^2). □

The space complexity for using the simulator of ADOPT-ing as a centralized WCSP solver is given by the sum of all the spaces of the n agents, which is O(dn^3). The simulator also maintains the queues of *traveling*

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

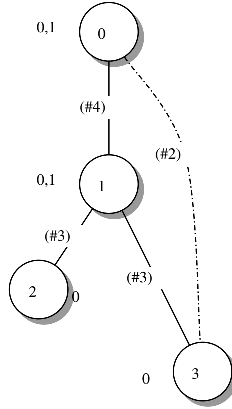

Figure 3: A DCOP with four agents and four inequality constraints. For example, the fact that the cost associated with not satisfying the constraint $x_0 \neq x_1$ is 4 is denoted by the notation ($\#4$).

messages, which can be compacted such that only the last sent message is stored for each channel (Silaghi, Sam-Haroud, & Faltings, 2000). There are O(n^2) bidirectional channels, each of them requiring at most a valued nogood (for an optimized simulator); therefore their total size is O(n^3), being smaller than the sum of the sizes of the agents.

We expect that one can further optimize the space of a centralized implementation by abandoning the message-passing paradigm of the simulator and by sharing the ca data structures of the agents, directly storing each inferred valued nogood at its final position in the structure ca. Additional improvements in space complexity are possible by simply discarding the ca storage in favor of more compact aggregations of its nogoods (where h and the structure for $f()$ mentioned in Section 5.3 are used alone without ca, integrating incoming nogoods directly in h), with a total space complexity of O(dn^2). However, some nogoods would be lost and may have to be recomputed, and threshold nogoods would no longer be available.

5.6 Optimizing valued nogoods

Both for the versions of ADOPT-ing using DFS trees, as well as for the version that does not use such DFS trees, if valued nogoods are used for managing cost inferences, then a lot of effort can be saved at context switching by keeping nogoods that remain valid (Ginsberg, 1993). The amount of effort saved is higher if the nogoods are carefully selected (to minimize their dependence on assignments for low priority variables, which change more often). We compute valued nogoods by minimizing the index of the least priority variable involved in the context. At sum-inference with intersecting SRCs, we keep the valued nogoods with lower priority target agents only if they have better costs. Nogoods optimized in a similar manner were used in several previous distributed CSP techniques (Bessiere et al., 2005; Silaghi et al., 2001b). A similar effect is achieved by computing min_resolution(j) with incrementally increasing j and keeping new nogoods only if they have higher cost than previous ones with lower targets. Between similar VNEs differing only in the $exact$ field, one has to keep the one where the value of $exact$ is $true$, to achieve the termination detection proposed in ADOPT.

5.7 Example

Next we detail and contrast the executions of ADOPT-Yos, and ADOPT-Aos illustrating the different types of inferences involved in them. The main description follows the run of ADOPT-Aos while describing differences with ADOPT-Yos when they occur. Take the problem in Figure 3, a trace of which is shown in Figure 4.

Sender	Message	Recipients
1. A_0	**ok?**$\langle x_0, 0\rangle$ $\longrightarrow$	A_1, A_3
2. A_1	**ok?**$\langle x_1, 0\rangle$ $\longrightarrow$	A_2, A_3
3. A_1	**ok?**$\langle x_1, 1\rangle$ $\longrightarrow$	A_2, A_3
4. A_2	**nogood**$[\vert F, F, T, F\vert, 3, \langle x_1, 0\rangle]$ $\longrightarrow$	A_1
5. A_3	**nogood**$[\vert F, F, F, T\vert, 2, \langle x_0, 0\rangle]$ $\longrightarrow$	A_0, A_2
6. A_3	**nogood**$[\vert F, F, F, T\vert, 5, \langle x_0, 0\rangle\langle x_1, 0\rangle]$ $\longrightarrow$	A_1, A_2
7. A_0	**ok?**$\langle x_0, 1\rangle$ $\longrightarrow$	A_1, A_3
8. A_2	**nogood**$[\vert F, F, F, T\vert, 2, \langle x_0, 0\rangle]$ $\longrightarrow$	A_0
9. A_2	**nogood**$[\vert F, F, T, T\vert, 5, \langle x_0, 0\rangle\langle x_1, 0\rangle]$ $\longrightarrow$	A_1
10. A_2	**add-link**$\langle x_0, 0\rangle$ $\longrightarrow$	A_0
11. A_2	**nogood**$[\vert F, F, T, T\vert, 8, \langle x_0, 0\rangle\langle x_1, 0\rangle]$ $\longrightarrow$	A_1
12. A_3	**nogood**$[\vert F, F, F, T\vert, 2, \langle x_0, 0\rangle]$ $\longrightarrow$	A_2
13. A_0	**ok?**$\langle x_0, 1\rangle$ $\longrightarrow$	A_2
14. A_1	**ok?**$\langle x_1, 0\rangle$ $\longrightarrow$	A_2, A_3
15. A_2	**nogood**$[\vert F, F, F, T\vert, 2, \langle x_0, 0\rangle]$ $\longrightarrow$	A_0, A_1
16. A_2	**nogood**$[\vert F, F, T, T\vert, 5, \langle x_0, 0\rangle\langle x_1, 0\rangle]$ $\longrightarrow$	A_1
17. A_2	**nogood**$[\vert F, F, T, F\vert, 3, \langle x_1, 0\rangle]$ $\longrightarrow$	A_1
18. A_3	**nogood**$[\vert F, F, F, T\vert, 3, \langle x_1, 0\rangle]$ $\longrightarrow$	A_1, A_2
19. A_1	**nogood**$[\vert F, T, F, T\vert, 3, \langle x_0, 1\rangle]$ $\longrightarrow$	A_0
20. A_1	**nogood**$[\vert F, T, T, T\vert, 4, \langle x_0, 1\rangle]$ $\longrightarrow$	A_0
21. A_1	**ok?**$\langle x_1, 1\rangle$ $\longrightarrow$	A_2, A_3
22. A_2	**nogood**$[\vert F, F, T, T\vert, 6, \langle x_1, 0\rangle]$ $\longrightarrow$	A_1
23. A_0	**ok?**$\langle x_0, 0\rangle$ $\longrightarrow$	A_1
24. A_0	**ok?**$\langle x_0, 0\rangle$**threshold** $[\vert F, F, F, T\vert, 2, \langle x_0, 0\rangle]$ $\longrightarrow$	A_2, A_3
25. A_3	**nogood**$[\vert F, F, F, T\vert, 2, \langle x_0, 0\rangle]$ $\longrightarrow$	A_0, A_2
26. A_2	**nogood**$[\vert F, F, F, T\vert, 2, \langle x_0, 0\rangle]$ $\longrightarrow$	A_0, A_1
27. A_1	**nogood**$[\vert F, T, F, T\vert, 2, \langle x_0, 0\rangle]$ $\longrightarrow$	A_0

Figure 4: Trace of ADOPT-Aos on the problem in Figure 3. Horizontal lines separate groups of messages with the same logic clock (i.e., messages that are part of the same round in a simulator based on rounds).

Identical messages sent simultaneously to several agents are grouped by displaying the list of recipients on the right hand side of the arrow. In our implementation, we decide to maintain a single reference for each agent's secret constraints. In our next description, the notation which refers to the constraints of the agent A_i in a SRC is J_i. In the messages of Figure 4, SRCs are represented as Boolean values in an array of size n. A value at index i in the array of SRCs set to T signifies that the constraints of A_i are used in the inference of that nogood (i.e., J_i is part of the justification of the valued nogood).

Initialization. The agents start selecting values for their variables and announce them to interested lower priority agents. There are no constraints between x_3 and x_2. Similarly, there is no constraint between x_0 and x_2; therefore, the first exchanged messages are **ok?** messages sent by A_0 to both successors A_1 and A_3 and which propose the assignment x_0=0. Concurrently, A_1 sends **ok?** messages to A_2 and A_3 proposing x_1=0. These are messages 1 and 2 in Figure 4. The messages in Figure 4 are grouped by their cycle in the simulator based on rounds (i.e., assuming constant communication latency and no cost for local computations). The simulator with variable message latencies can yield different traces function of the random latencies.

Handling data structures for ok? messages. On the receipt of the **ok?** messages, the agents update their agent-view with the new assignment. Each agent tries to generate valued nogoods for each prefix of its list of predecessor agents, such as: $\{A_0\}$, $\{A_0, A_1\}$, $\{A_0, A_1, A_2\}$. A_1 receives the assignment of x_0 and infers a valued nogood based on its constraint ($x_0 \neq x_1$). It is stored as cost assessment in its structure l, before being integrated in h. $h[1] = l[1] = [\{J_1\}, 4, \langle x_0, 0\rangle]$. l[1] (and h[1]) have cost 0 while l[0] and h[0] have cost 4. Therefore A_1 switches the value of x_1 to 1 and announces it to A_2 and A_3 via message 3. A_1 cannot compute any valued nogood to send to A_0.

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

After the agent A_2 gets message 2, it computes in $l[0]$ a valued nogood with cost 3 (conflict with $x_1 \neq x_2$). This valued nogood is copied in $h[0]$ and $lastSent[1]$ before being sent to A_1 via message 4. No nogood can be computed for A_0.

Remark 11 (ADOPT-Aos vs ADOPT-Yos) *In ADOPT-Yos this message would also include the current list of known $ancestors$ which here contains only A_1.*

When A_3 gets message 1, it tries to separately infer nogoods for the prefixes of the set of agents: $\{A_0\}$, $\{A_0, A_1\}$, and $\{A_0, A_1, A_2\}$. For the set $\{A_0\}$ it detects a conflict with its constraint $x_3 \neq x_0$ from which it infers a valued nogood stored as cost assessment in $l[0]$, copied to $h[0]$ and $lastSent[0]$ before being sent to A_0 via message 5. For the set $\{A_0, A_1\}$, the computed nogood is identical with the one for A_0 and its target does not coincide with A_1, the last agent of the corresponding set. Therefore ADOPT-Aos sends no message to A_1. Message 5 is also sent to A_2 according to the rule that an agent always attempts to send nogoods to its predecessor, to ensure optimality. Its nogood is stored by A_3 in $lastSent[2]$.

Remark 12 (ADOPT-Aos vs ADOPT-Yos) *With ADOPT-Yos, message 5 would not be sent to A_2, since the current parent of A_3 would be A_0.*

After receiving the assignment in message 2, A_3 detects a new conflict with its constraint $x_1 \neq x_3$. From its two constraints A_3 infers a new valued nogood, stored in its $l[0]$ and $h[0]$, and sent to A_1 and A_2 via message 6. Note that a nogood is not sent to A_0 as the nogood to be sent is identical to the last nogood sent to that destination (as recorded in $lastSent[0]$).

Remark 13 (ADOPT-Aos vs ADOPT-Yos) *With ADOPT-Yos, A_1 would become the parent of A_3 at this stage due to the non-zero cost of the constraint between x_3 and x_1. A_3's known ancestors would become A_0,A_1, and this list would be sent with all nogood messages.*

Handling data structures for nogoods. As a result of getting the nogood in message 5 from A_3, the agent A_0 stores that nogood in $lr[3]$, copies it to $ca[0][3]$ (which was empty), and copies it further in $h[0]$. Since now the cost of $h[0]$ is 2, A_0 decides to switch to its next value, 1. This assignment is announced via message 7.

After receiving message 5, A_2 registers that nogood in its $lr[3]$, $ca[0][3]$ and $h[0]$. Computing a nogood for A_0, the nogood of message 5 is stored in $lastSent[0]$ and sent to A_0 via message 8. Agent A_2 also computes a nogood for destination A_1, where it can also use the local constraint with x_1 which yields for $l[0]$ a nogood with cost 3. Combining $l[0]$ with $ca[0][3]$ by sum-inference, A_2 infers a nogood, which it stores in $h[0]$ and $lastSent[0]$ before sending it to A_1 via message 9. A_2 detects a new variable in the nogood in message 6, and sends an **add-link** message to A_0 asking to be notified of changes to the assignment $x_0 = 0$. The nogood in message 6 replaces the one stored in $lr[3]$. Since the new nogood cannot be combined by sum-inference with the old nogood in $ca[0][3]$ but has a higher cost, it also replaces that cost assessment and leads to the computation by sum-inference of message 11 to be sent to A_1.

In the following we skip the details of changes to data structures that are similar to steps that have already been presented. When the new assignment of x_1 in message 3 is received at agent A_3, the old nogoods based on x_1 are discarded from its $l[0]$. To send a nogood to A_0, a new $l[0]$ is computed based solely on the constraint $x_0 \neq x_3$. Nogoods computed for the other prefixes of agents do not differ from this one since the constraint with x_1 is satisfied. This nogood with cost 2 is sent via message 12 to the agent A_2. Note that the nogood does not need to be sent to A_0 because it is not different from the one just sent earlier (via message 5) and recorded in $lastSent[0]$. After getting message 7, A_1 deletes its nogoods in $l[0]$ and $ca[0][3]$, infers a new valued nogood in $l[1]$ with cost 4, and switches to the value 0 (announced via message 14).

Use of lr data structure. Let us assume that A_2 receives message 12 before message 3, which is possible and allows us to illustrate better the usage of the lr structure. On receiving message 12, agent A_2 stores it in $lr[0]$. However, A_2 does not propagate it further to $ca[0][3]$ since the current cost assessment had a higher cost and cannot be combined by sum-inference with the new one (sharing the reference to the constraints of

A_3). When A_2 receives message 3, it deletes its ca[0][3] and l[0], which are based on the older value of x_1, and uses lr[3]. After copying lr[3] through its ca[0][3] and h[0] data structures where all other nogoods were empty, it passes it further to A_0 and to A_1 via message 15 (storing it at $lastSent$[0] and $lastSent$[1]). Since A_0's value for x_0 is different from the one in the **add-link** message 10, A_0 answers to A_2 with the message 13.

Now A_2 receives message 14 and computes a new local nogood l[0] with cost 3 that is combined by sum-inference with the nogood received in message 12 to generate the nogood in message 16. No change appears in the nogood computed specially for the target A_0. However, after A_2 also receives message 13 it discards the nogood received via message 12 (which was based on an outdated assignment) and infers its h[0] solely based on l[0]. The result is sent to A_1 with message 17. After receiving the two assignments in messages 13 and 14 (in this order) the agent A_3 infers from its constraint $x_3 \neq x_1$ a valued nogood sent to A_1 and A_2 via message 18.

Min-resolution. Now our example encounters the first nontrivial min-resolution. When agent A_1 receives message 18, it stores that nogood in lr[3] and ca[0][3]. No other nogood is stored in ca at this point (the nogood received with message 15 in ca[0][1] has already been invalidated by the new assignment in message 7). The only other nogood held by A_1 at this moment is the one in $l[1] = [\{J_1\}, 4, \langle x_0, 1\rangle]$, which is due to its constraint with x_0. l[1] is copied in h[1] while ca[0][3] is copied in h[0]. The two are combined via min-resolution to generate the nogood in message 19 (also stored in $lastSent$[0]).

$$min_resolution([\{J_3\}, 3, \langle x_1, 0\rangle], [\{J_1\}, 4, \langle x_0, 1\rangle]) \rightarrow [\{J_1, J_3\}, 3, \langle x_0, 1\rangle\langle x_1, 0\rangle]$$

Message 16 is discarded at its destination because its assignment for x_0 is no longer valid. On the arrival of message 17 (which is concurrent with messages 16 and 18) its nogood is stored in lr[2] and ca[0][2]. Now, when computing the updated nogood to be sent to A_0, h[0] is computed by sum-inference on ca[0][2] and ca[0][3] obtaining $[\{J_2, J_3\}, 6, \langle x_1, 0\rangle]$.

$$sum_inference([\{J_2\}, 3, \langle x_1, 0\rangle], [\{J_3\}, 3, \langle x_1, 0\rangle]) \rightarrow [\{J_2, J_3\}, 6, \langle x_1, 0\rangle]$$

The obtained valued nogood has a higher cost than the one for h[1], causing the agent to switch the assignment of x_1 to 1 (announced via message 21). When min-resolution is applied on the two nogoods in h[0] and h[1], the obtained nogood is sent to A_0 via message 20.

$$\begin{aligned} min_resolution([\{J_2, J_3\}, 6, \langle x_1, 0\rangle], [\{J_1\}, 4, \langle x_0, 1\rangle]) \\ \rightarrow [\{J_1, J_2, J_3\}, 3, \langle x_0, 1\rangle\langle x_1, 0\rangle] \end{aligned}$$

Convergence. Agent A_2 also receives message 18, storing the nogood in lr[3] and in ca[0][3]. Its constraint $x_2 \neq x_1$ generates a nogood with cost 3 in l[0], which combined by sum-inference with the nogood in ca, leads to a nogood with total cost 6, visible in message 22.

Agent A_0 receives message 19 and registers the nogood in lr[1], ca[1][1], and h[1]. The cost assessment obtained in h[1] has a cost higher than the one in h[0], determining the switch of the assignment of x_0 to 0 (announced via messages 23 and 24). Message 24 also transports a threshold nogood obtained from ca[0][2] and ca[0][3] (received via messages 15 and 5). The agent A_3 evaluates its constraint $x_0 \neq x_3$ inferring a valued nogood in l[0], which propagates through its h[0], $lastSent$[2], $lastSent$[0] to messages 25. Similarly A_2 propagates this nogood to A_1, which propagates it further through its data structures and eventually delivers it to A_0 via message 27. Messages 25, 26 and 27 basically confirm the already known threshold nogoods. Further research may make it possible to avoid them[12].

We have modeled solved this example with our implementation for ADOPT-Aos with rounds.

12. E.g, by a mechanism for storing threshold nogoods in the $lastSent$ of the recipient and in the lr of the sender, resending the $lastSent$ when the threshold nogood does not apply.

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

5.8 Possible Extensions

We addressed ADOPT-ing as an asynchronous version of A*, more exactly a version of iterative deepening A*, where the heuristic is computed by recursively using ADOPT-ing itself, and where the composition of the results of recursive ADOPT-ing is based on backtracking.

A proposed extension to this work consists of composing the recursive asynchronous heuristic estimator by using consistency maintenance. This can be done with the introduction of *valued consistency nogoods*. Details and variations are described in (Silaghi, 2002, 2003b; Silaghi et al., 2004; Gershman, Meisels, & Zivan, 2006, 2007; Sultanik et al., 2006). The control of the space requirements for such extensions may be based on the use of consistency nogoods to simulate the distributed weighted arc consistency in (Silaghi et al., 2004), while the maintenance of this control of space in asynchronous search may be similar to the one for distributed CSPs described in (Silaghi & Faltings, 2004). Another possible extension is by further generalizing the nogoods such that each variable can be assigned a set of values. This type of aggregation was shown in (Silaghi & Faltings, 2004) to improve search, and the extension is detailed in (Silaghi, 2002).

In our implementation we concentrated on minimizing the logic time of the computation, evaluated as the number of rounds on a simulator. The optimization of local processing (which is polynomial in the number of variables) is not at the center of attention at this stage. Local computations can be optimized, for example, by reusing values of structures l and h computed at min-resolution for a given target agent in obtaining values of these structures at the min-resolution for messages sent to lower priority target agents. Further work can determine whether improvements could be made by storing separately the nogoods of h for each target k. The size of messages in ADOPT-Yos could be slightly reduced by appending a given content of the $ancestors$ list only once to each target. ADOPT-Yos is better than ADOPT-Aos in terms of simulated time. Agents in ADOPT-Yos could insert from the beginning all their neighboring predecessors in their ancestors list, obtaining from the first n rounds the DFS tree of ADOPT-Dos, thereby replicating the efficiency of ADOPT-Dos.

Other extensions seem possible by integrating additive branch and bound searches on DFS sub-trees, as proposed by (Chechetka & Sycara, 2006; Yeoh, Koenig, & Felner, 2007). This can be added to ADOPT-ing by maintaining solution-based nogoods as suggested in (Silaghi, 2002). It remains to be seen if the quality of solutions with a certain value can be predicted with the technique in (Petcu & Faltings, 2006b). Further improvements are possible by running ADOPT-ing in parallel for several orderings of the agents (Ringwelski & Hamadi, 2005; Benisch & Sadeh, 2006).

ADOPT-ing can be seen as an extension of ABT. The extension of ABT called ABTR (Silaghi, Sam-Haroud, & Faltings, 2001a; Silaghi, 2006) proposes a way to extend ABT-based algorithms to allow for dynamic ordering of the agents (Armstrong & Durfee, 1997). Work in the area consistent with this approach, but mainly favoring static ordering, appears in (Liu & Sycara, 1995; Chechetka & Sycara, 2005). Finding good heuristics was shown to be a difficult problem (Silaghi et al., 2001b; Zivan & Meisels, 2005) and here one will need to take into account the importance of the existence of a short DFS tree compatible to the current ordering.

6. Experiments

For experiments with random message latencies and for outputs not provided by the original implementation of ADOPT (e.g., ENCCCs), we had to provide the results of our implementation. While our implementation performs in general similarly to the original implementation of ADOPT, our technique solved in a few hours the instances for which the original ADOPT implementation was interrupted after some weeks, confirming that some differences in details may exist. Functional differences between our implementation and the original implementation of ADOPT may lie only in petty details not described in (Modi et al., 2005). In our experiments we detect the termination by detecting the quiescence in the simulator. Theoretically this detects the same termination moment as the detection based on upper-bounds described in ADOPT, and has no impact on the execution of the algorithm itself.

We also implemented a version of ADOPT (using our implementation of ADOPT with threshold nogoods) that uses a chain of agents rather than the DFS tree. This version is denoted *ADOPT.chain*.

We first verified experimentally the fact that ADOPT behaves identically if SRCs are added to contexts while sending the messages only to the same destinations. Then we send messages to additional ancestors (ADOPT-Dos), and we show an improvement. Then we detect the incremental improvement due to the dynamic detection of the tree, by depicting the three techniques on the same graphs. We also verify the incremental improvement of sending messages to all predecessors (ADOPT-Aos), and then the incremental improvement brought on this by detecting the DFS tree dynamically.

The algorithms are compared on the same problems that are used to report the performance of ADOPT in (Modi et al., 2005). To correctly compare our techniques with the original ADOPT, we have used the same order (or DFS trees) on agents for each problem. The impact of the existence of a good DFS tree compatible with the used order is tested separately by comparison with a random ordering. The set of problems distributed with ADOPT and used here contains 25 problems for each problem size. It contains problems with 8, 10, 12, 14, 16, 18, 20, 25, 30, and 40 agents, and for each of these numbers of agents, it contains test sets with density 20% and with density 30%. A smaller set of problems with density 40% is also available. The density of a (binary) constraint problem's graph with n variables is defined by the ratio between the number of binary constraints and $\frac{n(n-1)}{2}$. Results are averaged on the 25 problems with the same parameters.

We believe that the size of problems in this set is sufficiently large, given that the average simulated time (expected time of a real solver) for the instances with 40 agents at density 30% is between 3 hours and 27 hours, (and up to 10 days at 25 agents and density 40%), longer than what users are expected to wait for a solution.

Our simulator allows for defining the latency of each message. We performed experiments with random message latencies, and the efficiency is measured in sequential messages and in the number of total messages. The random latencies were generated in the range of common values for Internet communications via optical fiber between Israel and the United States which is between 150ms and 250ms (Neystadt & Har'El, 1997). To reproduce our results for the set of tests, one has to seed the standard C 'random()' function with the value 10000 and generate each latency as carried out in (Neystadt & Har'El, 1997):

$$latency = 150 + \frac{random() * 100.0}{LONG_MAX}(msec).$$

FIFO channels are ensured in the second set of tests by setting the delivery time of each message to the maximum between the value obtained using the latency yielded by the aforementioned computation and the delivery time of the last message sent on that particular communication channel. Messages with the same value for the delivery time are handed to the destination agent in a FIFO manner through a queue.

In graphs, an algorithm *ADOPT-*$\mathcal{DON}$ is typically shortened to $\mathcal{DON}$.

The length of the longest causal (sequential) chain of messages of each solver (the number of sequential messages), averaged on problems with density 30%, is given in Figure 5. Results for problems with density 20% are given in Figure 6. Results for density 40% are shown in Figure 7. We can note that version ADOPT-Yos of ADOPT-ing brought an improvement of approximately 10 times on problems with 40 agents and density 30%, and of approximately 12 times on problems with 25 agents and density 40%. The improvement at density 20% is 2 times when compared to ADOPT.[13] Therefore, sending nogoods only to the parent node is significantly worse (in number of message latencies), than sending nogoods to several ancestors. With respect to the number of message latencies, the use of SRCs with nogood contexts practically replaces the need to maintain the DFS tree since ADOPT-Aos is comparable in efficiency to ADOPT-Dos. New versions of ADOPT-ing are up to 14 times faster than ADOPT.chain, proving that ADOPT-ing is not a simple application of ADOPT to a chain of agents, but that justified valued nogoods literally succeed in dynamically discovering the DFS tree.

13. At density 20%, with synchronous rounds, the original implementation of ADOPT performs 3.5 times worse than ADOPT, i.e., 7 times worse than ADOPT-Yos. This may be explained by some inefficient detail in the original implementation of ADOPT, since the deviation from ADOPT does not appear at other densities.

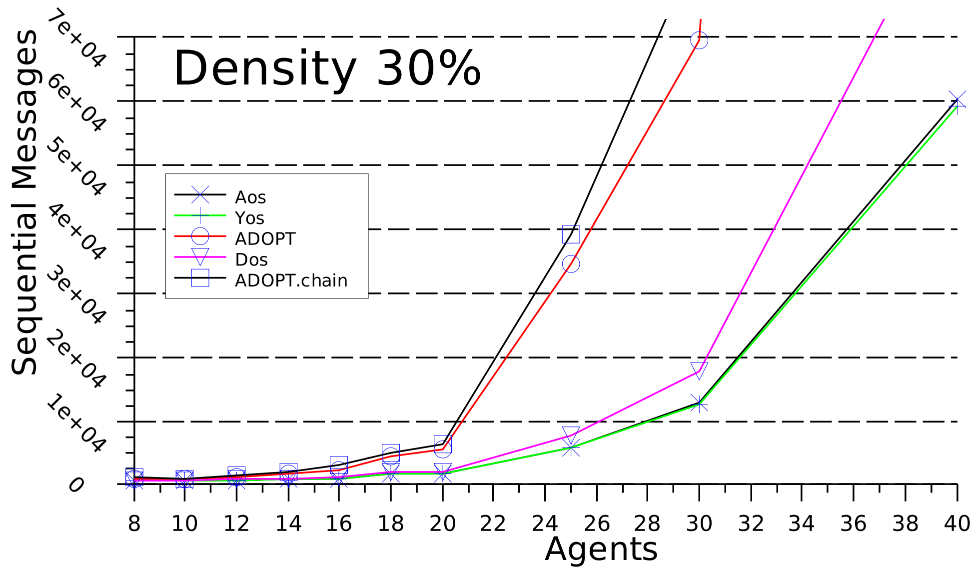

Figure 5: Sequential messages for problems with density 30%.

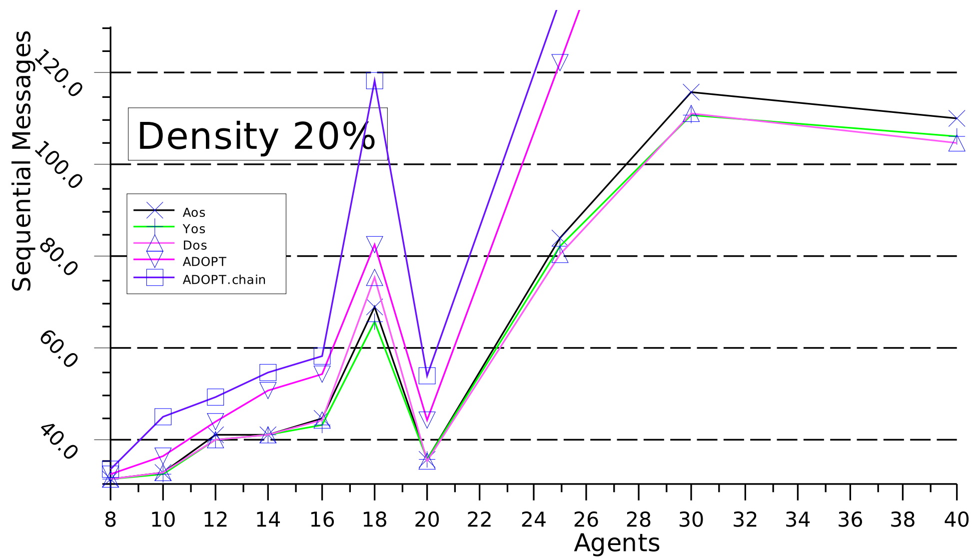

Figure 6: Sequential messages for problems with density 20%.

Remark 14 *Versions using DFS trees require fewer parallel/total messages, being more network friendly, as seen in Figure 8. Figure 8 shows that sending messages to other predecessors, as done in ADOPT-Aos, ADOPT-Yos and ADOPT-Dos, is 4 times better at density 30% than ADOPT in terms of total number of messages, while (as shown by previous graphs) also being more efficient in terms of message latencies. At density 40% ADOPT-Yos is 6 times better than ADOPT in terms of total number of messages. ADOPT-Yos is the most efficient algorithm in terms of total number of messages, being 30% better at density 30% than the second best algorithm, ADOPT-Aos. At density 40% it is 12% better than ADOPT-Aos.*

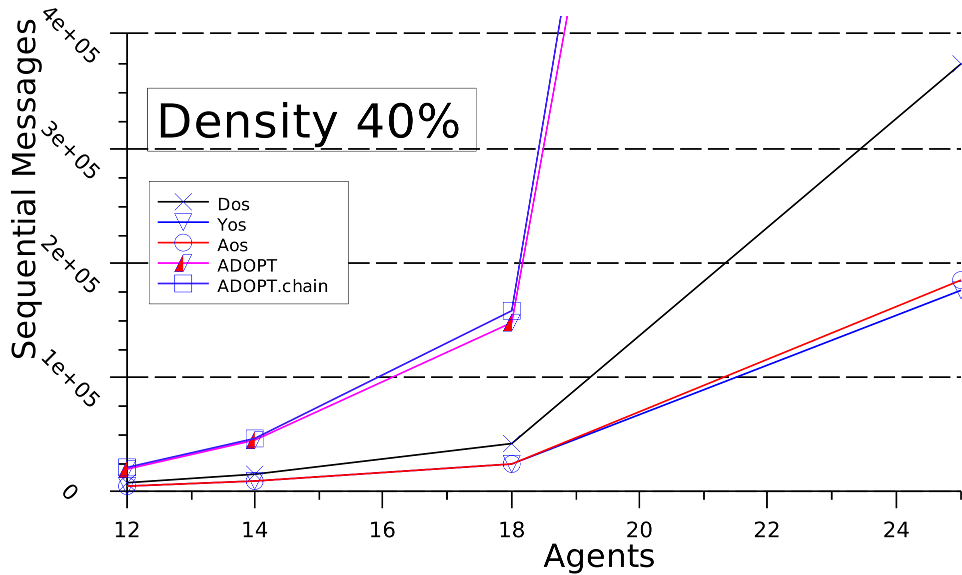

Figure 7: Sequential messages for problems with density 40%.

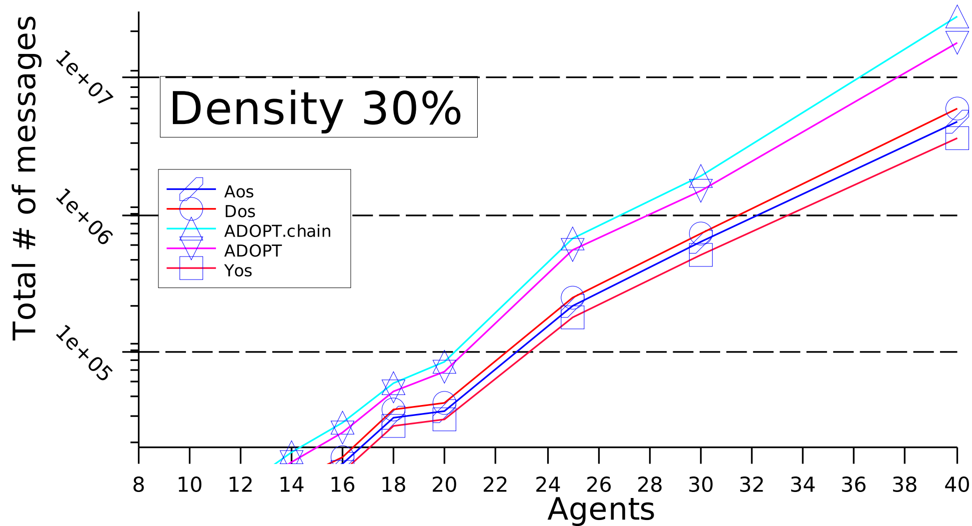

Figure 8: Total number of messages at density 30% (log scale).

While the importance of privacy is clear, evaluating the privacy loss is a controversial issue outside a concrete application. The versions of ADOPT-ing that send less total messages, are intuitively expected to perform better (since privacy is expected to be related to the total number of messages

We do not show run-time comparisons with the original implementation of ADOPT since our versions of ADOPT are implemented in C++, while the original ADOPT is in Java (which obviously leads to all our versions being an irrelevant order of magnitude faster). However, we provide run-time comparisons with our implementation of ADOPT. A comparison between the time required by versions of ADOPT-ing on a simulator is shown in Figure 9 for sequential messages. It reveals the computational load of the agents which, as expected, is related to the total number of exchanged messages.

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

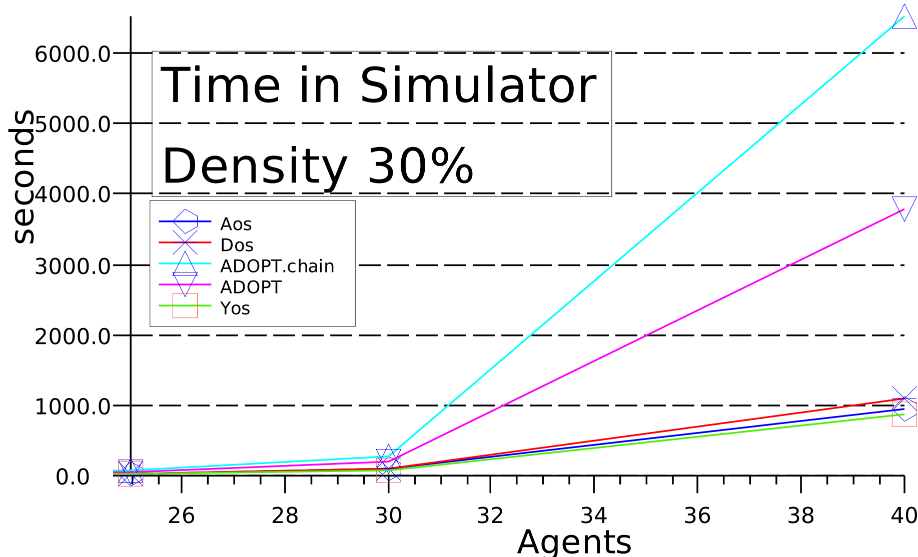

Figure 9: Actual time in seconds using our simulator as solver of centralized WCSPs.

Agents	16	18	20	25	30	40
ADOPT-Yos.b	701.72	1438	1345.56	5540.84	12394	59114.36
no threshold	872.76	1781.96	1708.72	7391.28	17531.36	92745.44

Table 2: Impact of threshold valued nogoods on the longest causal chain of messages (sequential messages) for versions of ADOPT-ing, averaged on problems with density 30%.

Agents	16	18	20	25	30	40
DFS compatible	708.8	1429.48	1357.07	5579.56	$12.4*10^3$	$60*10^3$
random order	4807.44	$15.6*10^3$	$33*10^3$	$219*10^3$	$708*10^3$	—

Table 3: Impact of choice of order according to a DFS tree on the longest causal chain of messages (sequential messages) for ADOPT-Yos, averaged on problems with density 30%.

A separate set of experiments was run for isolating and evaluating the contribution of threshold valued nogoods. Table 2 shows that the use of threshold nogoods almost halves the computation time. Another experiment, whose results are shown in Table 3, is meant to evaluate the impact of the guarantees that the ordering on agents is compatible with a short DFS tree. We evaluate this by comparing ADOPT-Yos with an ordering that is compatible with the DFS tree built by ADOPT, versus a random ordering. At 30 agents it was found to be 60 times more efficient to ensure that a DFS tree exists rather than to use a random ordering. The results show that random orderings are unlikely to be compatible with short DFS trees and that verifying the existence of a short DFS tree compatible to the ordering on agents to be used by ADOPT-ing is highly recommended.

The simulated time of the computations, where the random latencies of the messages are accumulated along the longest causal chain, is shown in Figure 10. The time taken for the local computation handling/generating each message (Figure 9) is hundreds of times smaller than the latency of the associated

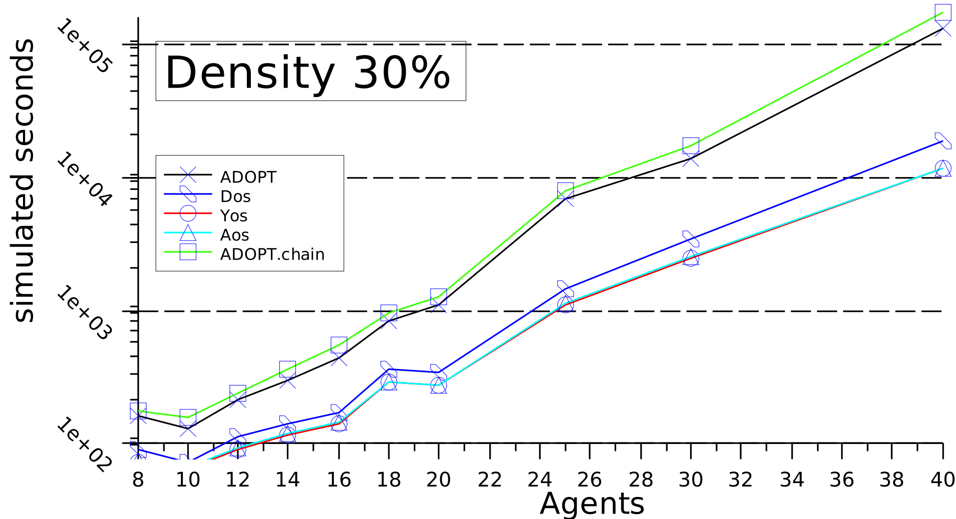

Figure 10: Simulated time in seconds, each latency being drawn randomly between 150ms and 250ms (for problems with density 30%).

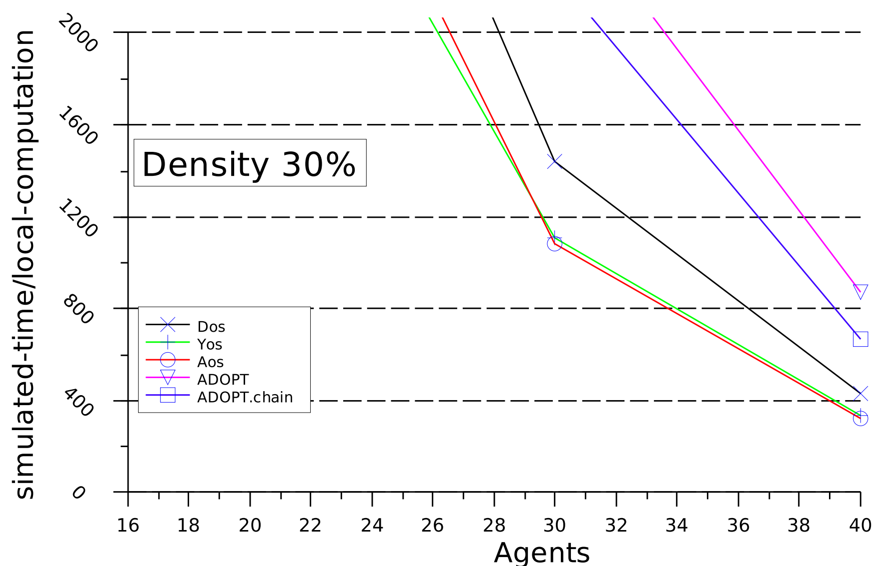

Figure 11: Ratio between total expected time where each latency is drawn randomly between 150ms and 250ms, and the local time of an agent (for problems with density 30%).

messages, falling close to the numerical precision of this accounting (Figure 11), confirming the relevance of the metrics we use here.

Figure 5 shows the behavior of our implementation of ADOPT. It took more than two weeks for the original ADOPT implementation to solve one of the problems for 20 agents and density 30%, and one of the problems for 25 agents and density 30% (at which moment the solver was interrupted). Therefore, it was evaluated using only the remaining 24 problems at those problem sizes. In (Silaghi & Yokoo, 2006) we

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

have also shown that SRCs bring improvements over versions with valued global nogoods, since SRCs allow detection of dynamically obtained independence.

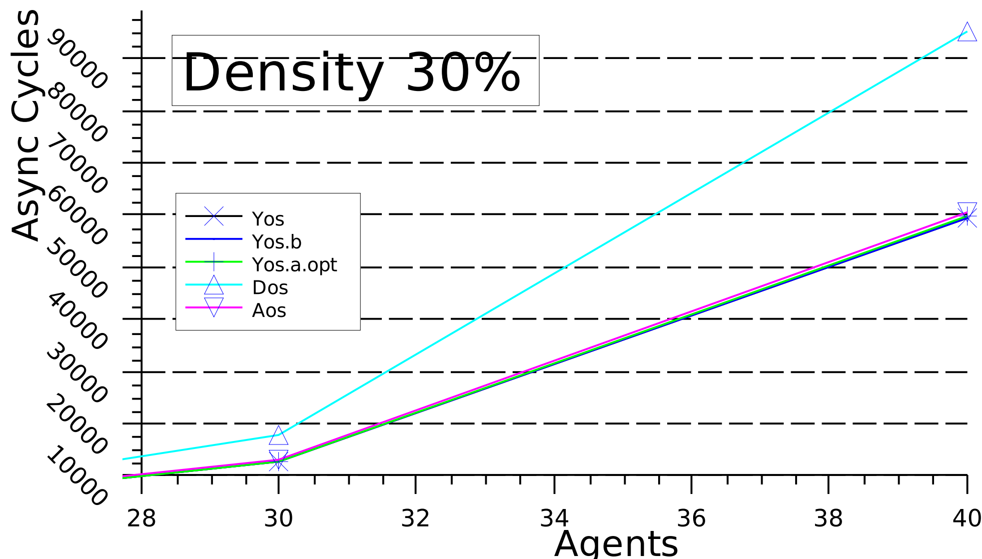

Figure 12: Local computations have little effect, but ADOPT-Yos is clearly better than ADOPT-Aos, competing with ADOPT-Dos. The optimized version of ADOPT-Yos is in average approximately 1% better than ADOPT-Dos (up to 15% better on some problem instances).

We tried to figure out the importance of using method (a) rather than method (b) in Step 1 of Remark 8 (comparing obtained versions ADOPT-Yos and ADOPT-Yos.b), and we found the two alternatives to be equally good (ADOPT-Yos being less than 1% better than ADOPT-Yos.b). We also evaluated the effects of optimizations in local computations, by computing the nogoods for an agent A_k based on the nogoods computed for higher priority agents rather than computing them from scratch (ADOPT-Yos.a.optim). The same figure shows the effect on sequential messages to be minor (approximately 1% worse than ADOPT-Yos). The effect on constraint checks is similarly minor (4%) and is not depicted here.

Figure 5 clearly shows that the highest improvement in number of sequential messages is brought by sending valued nogoods to other ancestors besides the parent. The next factor for improvement with difficult problems (density .3) is the use of SRCs. The use of the structures of the DFS tree makes slight improvements in number of message latencies (when nogoods reach all ancestors).

Experimental comparison with DPOP (Petcu & Faltings, 2005a, 2005b) is redundant since its performance can be easily predicted. DPOP is a good choice if the induced width γ of the graph of the problem is smaller than $\log_d T/n$ and smaller than $\log_d S$, where T is the available time, n the number of variables, d the domain size, and S the available computer memory. The usage of the DFS trees in DPOP is discussed in (Atlas & Decker, 2007).

7. Conclusions

With the ADOPT distributed constraint optimization algorithm, an agent can communicate feedback only to a predefined predecessor, its parent in the DFS tree. The extension proposed here enables agents to send feedback to any relevant agent (fulfilling a research direction suggested in the original publication of ADOPT), bringing significant speed-up, and embodying a version of ADOPT on which one can apply the results related to the main algorithm for distributed constraint satisfaction, ABT.

ADOPT-ing can dynamically discover a DFS tree based only on the constraints that had been proved relevant by the search up to that moment. It uses (Dago & Verfaillie, 1996)'s valued nogoods tagging contexts with costs and with sets of references to culprit constraints. The generalized algorithm is denoted ADOPT-ing. Tagging costs with sets of references to culprit constraints (SRCs) allows detection and exploitation of dynamically created independence between sub-problems. Such independence can be caused by assignments. Experimentation shows that it is important for an agent to infer and send in parallel several valued nogoods to different higher priority agents. Each inferred valued nogood is sent only to the highest priority agent that can handle it (its target). Precomputed DFS trees can still be used in conjunction with the valued nogood paradigm for optimization, thereby providing some additional improvements. ADOPT-ing versions detecting and/or exploiting DFS trees that we tested so far are also slightly better (in number of sequential messages and total messages) than the ones without DFS trees.

We isolated and evaluated the contribution of using threshold valued nogoods in ADOPT-ing. In addition, we determined the importance of precomputing and maintaining a short DFS tree of the constraint graph, or at least of guaranteeing that a DFS tree is compatible with the order on agents, which is almost an order of magnitude in our problems.

The use of SRCs to dynamically detect and exploit independence and the generalized communication of valued nogoods to several ancestors bring elegance and flexibility to the description and implementation of ADOPT in ADOPT-ing. They also produced experimental improvements of an order of magnitude.

Acknowledgments

We thank Judith Strother for her professional restyling of the paper. We also thank anonymous reviewers for suggesting particularly relevant references, clarifications, and experiments. This research was partially supported by the Ministry of Education, Science, Sports and Culture, Grant-in-Aid for Scientific Research (B), 17300049, 2005–2007.

References

Ali, S., Koenig, S., & Tambe, M. (2005). Preprocessing techniques for accelerating the DCOP algorithm ADOPT. In *AAMAS*.

Armstrong, A., & Durfee, E. F. (1997). Dynamic prioritization of complex agents in distributed constraint satisfaction problems. In *Proceedings of 15th IJCAI*.

Atlas, J., & Decker, K. (2007). A complete distributed constraint optimization method for non-traditional pseudotree arrangements. In *AAMAS*.

Benisch, M., & Sadeh, N. (2006). Examining dcsp coordination tradeoffs. In *AAMAS*.

Bessiere, C., Brito, I., Maestre, A., & Meseguer, P. (2005). Asynchronous backtracking without adding links: A new member in the abt family. *Artificial Intelligence*, *161*, 7–24.

Bistarelli, S., Montanari, U., & Rossi, F. (1995). Constraint solving over semirings. In *Proceedings IJCAI*, pp. 624–630, Montreal.

Bistarelli, S., Montanari, U., Rossi, F., Schiex, T., Verfaillie, G., & Fargier, H. (1999). Semiring-based CSPs and valued CSPs: Frameworks, properties, and comparison. *Constraints*, *4*(3), 199–240.

Bistarelli, S., Fargier, H., Montanari, U., Rossi, F., Schiex, T., & Verfaillie, G. (1996). Semiring-based csps and valued csps: Basic properties and comparison. In *Over-Constrained Systems*, pp. 111–150, London, UK. Springer-Verlag.

Chechetka, A., & Sycara, K. (2005). A decentralized variable ordering method for distributed constraint optimization. In *AAMAS*.

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

Chechetka, A., & Sycara, K. (2006). No-commitment branch and bound search for distributed constraint optimization. In *AAMAS*.

Collin, Z., Dechter, R., & Katz, S. (2000). Self-stabilizing distributed constraint satisfaction. *Chicago Journal of Theoretical Computer Science*, *3*(4).

Dago, P. (1997). Backtrack dynamique valué.. In *JFPLC*, pp. 133–148.

Dago, P., & Verfaillie, G. (1996). Nogood recording for valued constraint satisfaction problems.. In *ICTAI*.

Dechter, R. (1990). Enhancement schemes for constraint processing: Backjumping, learning, and cutset decomposition. *Artificial Intelligence*, *41*(3), 273 – 312.

Dechter, R. (2003). *Constraint Processing*. Morgan Kaufman.

Faltings, B. (2006). *Handbook of Constraint Programming (Foundations of Artificial Intelligence)*, chap. Distributed Constraint Programming. Elsevier, New York, NY, USA.

Franzin, M., Rossi, F., E.C., F., & Wallace, R. (2004). Multi-agent meeting scheduling with preferences: efficiency, privacy loss, and solution quality. *Computational Intelligence*, *20*(2).

Gershman, A., Meisels, A., & Zivan, R. (2006). Asynchronous forward-bounding for distributed constraints optimization. In *ECAI*.

Gershman, A., Meisels, A., & Zivan, R. (2007). Asynchronous forward-bounding with backjumping. In *IJCAI DCR Workshop*.

Ginsberg, M. L. (1993). Dynamic backtracking. *Journal of AI Research*, *1*.

Greenstadt, R., Pearce, J., Bowring, E., & Tambe, M. (2006). Experimental analysis of privacy loss in dcop algorithms. In *AAMAS*, pp. 1024–1027.

Hamadi, Y., & Bessière, C. (1998). Backtracking in distributed constraint networks. In *ECAI'98*, pp. 219–223.

Hirayama, K., & Yokoo, M. (1997). Distributed partial constraint satisfaction problem. In *Proceedings of the Conference on Constraint Processing (CP-97),LNCS 1330*, pp. 222–236.

Jagota, A., & Dechter, R. (1997). Simple distributed algorithms for the cycle cutset problem. In *SAC '97: Proceedings of the 1997 ACM symposium on Applied computing*, pp. 366–373, New York, NY, USA. ACM Press.

Larrosa, J. (2002). Node and arc consistency in weighted csp. In *AAAI-2002*, Edmonton.

Liu, J., & Sycara, K. P. (1995). Exploiting problem structure for distributed constraint optimization. In *ICMAS*.

Maheswaran, R., Tambe, M., Bowring, E., Pearce, J., & Varakantham, P. (2004). Taking DCOP to the real world: Efficient complete solutions for distributed event scheduling. In *AAMAS*.

Marcellino, F. M., Omar, N., & Moura, A. V. (2007). The planning of the oil derivatives transportation by pipelines as a distributed constraint optimization problem. In *IJCAI-DCR Workshop*, India.

Modi, P., & Veloso, M. (2005). Bumping strategies for the multiagent agreement problem. In *AAMAS*.

Modi, P. J., Shen, W.-M., Tambe, M., & Yokoo, M. (2005). ADOPT: Asynchronous distributed constraint optimization with quality guarantees. *AIJ*, *161*.

Modi, P. J., Tambe, M., Shen, W.-M., & Yokoo, M. (2002). A general-purpose asynchronous algorithm for distributed constraint optimization. In *Distributed Constraint Reasoning, Proc. of the AAMAS'02 Workshop*, Bologna. AAMAS.

Neystadt, J., & Har'El, N. (1997). Israeli internet guide (iguide). http://www.iguide.co.il/isp-sum.htm.

Petcu, A., & Faltings, B. (2005a). Approximations in distributed optimization. In *Principles and Practice of Constraint Programming CP 2005*.

Petcu, A., & Faltings, B. (2005b). A scalable method for multiagent constraint optimization. In *IJCAI*.

Petcu, A., & Faltings, B. (2006a). Distributed generator maintenance scheduling. In *Proceedings of the First International ICSC Symposium on Artificial Intelligence in Energy Systems and Power: AIESP'06*, Madeira, Portugal.

Petcu, A., & Faltings, B. (2006b). ODPOP: an algorithm for open/distributed constraint optimization. In *AAAI*.

Ringwelski, G., & Hamadi, Y. (2005). Multi-directional distributed search with aggregation. In *IJCAI-DCR*.

Schiex, T., Fargier, H., & Verfaillie, G. (1995). Valued constraint satisfaction problems: hard and easy problems.. In *Procs. IJCAI'95*, pp. 631–637.

Silaghi, M.-C. (2002). *Asynchronously Solving Distributed Problems with Privacy Requirements*. PhD Thesis 2601, (EPFL). `http://www.cs.fit.edu/~msilaghi/teza`.

Silaghi, M.-C. (2003a). Asynchronous PFC-MRDAC±Adopt —consistency-maintenance in ADOPT—. In *Distributed Constraint Reasoning Workshop at the International Joint Conference on Artificial Intelligence (IJCAI-DCR)*.

Silaghi, M.-C. (2003b). Howto: Asynchronous PFC-MRDAC –optimization in distributed constraint problems +/-ADOPT–. In *International Conference on Intelligent Agent Technology (IAT)*, Halifax.

Silaghi, M.-C. (2006). Framework for modeling reordering heuristics for asynchronous backtracking. In *International Conference on Intelligent Agent Technology (IAT)*.

Silaghi, M.-C., & Faltings, B. (2004). Asynchronous aggregation and consistency in distributed constraint satisfaction. *Artificial Intelligence*, *161*(1-2), 25–53.

Silaghi, M.-C., Landwehr, J., & Larrosa, J. B. (2004). *Selected Paper from the 2003 Distributed Constraint Reasoning Workshop*, Vol. 112 of *Frontiers in Artificial Intelligence and Applications*, chap. Asynchronous Branch & Bound and A* for DisWCSPs with heuristic function based on Consistency-Maintenance. IOS Press.

Silaghi, M.-C., Sam-Haroud, D., & Faltings, B. (2000). Asynchronous search with aggregations. In *Proc. of National Conference of the American Association of Artificial Intelligence (AAAI2000)*, pp. 917–922, Austin.

Silaghi, M.-C., Sam-Haroud, D., & Faltings, B. (2001a). ABT with asynchronous reordering. In *International Conference on Intelligent Agent Technology (IAT)*.

Silaghi, M.-C., Sam-Haroud, D., & Faltings, B. (2001b). Hybridizing ABT and AWC into a polynomial space, complete protocol with reordering. Tech. rep. #01/364, EPFL.

Silaghi, M.-C., Sam-Haroud, D., & Faltings, B. (2000). Maintaining hierarchical distributed consistency. In *Workshop on Distributed CSPs*, Singapore. 6th International Conference on CP 2000.

Silaghi, M.-C., & Yokoo, M. (2006). Nogood-based asynchronous distributed optimization. In *International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

Stallman, R. M., & Sussman, G. J. (1977). Forward reasoning and dependency-directed backtracking in a system for computer-aided circuit analysis. *Artificial Intelligence*, *9*, 135–193.

Sultanik, E., Modi, P. J., & Regli, W. (2006). Constraint propagation for domain bounding in distributed task scheduling. In *CP*.

Walsh, T. (2007). Traffic light scheduling: a challenging distributed constraint optimization problem. In *DCR*, India.

Yeoh, W., Koenig, S., & Felner, A. (2007). Idb-adopt : A depth first search dcop algorithm. In *IJCAI DCR Workshop*.

Yokoo, M. (1993). Constraint relaxation in distributed constraint satisfaction problem. In *ICDCS'93*, pp. 56–63.

Submitted to JAAMAS on May 9, 2007. First notification received on Jan 28, 2008.
Resubmitted on Feb 20, 2008.

Yokoo, M., Durfee, E. H., Ishida, T., & Kuwabara, K. (1992). Distributed constraint satisfaction for formalizing distributed problem solving. In *ICDCS*, pp. 614–621.

Yokoo, M., Durfee, E. H., Ishida, T., & Kuwabara, K. (1998). The distributed constraint satisfaction problem: Formalization and algorithms. *IEEE TKDE*, *10*(5), 673–685.

Yokoo, M., & Hirayama, K. (1998). Distributed constraint satisfaction algorithm for complex local problems. In *Proceedings of 3rd ICMAS'98*, pp. 372–379.

Zivan, R., & Meisels, A. (2005). Dynamic ordering for asynchronous backtracking on discsps. In *CP*, pp. 161–172.