

Storing learnt (no)goods in ROBDDs for solving structured CSPs

Karim Boutaleb and Philippe Jégou and Cyril Terrioux

LSIS - UMR CNRS 6168

Université Paul Cézanne (Aix-Marseille 3)

Avenue Escadrille Normandie-Niemen

13397 Marseille Cedex 20 (France)

{ karim.boutaleb, philippe.jegou, cyril.terrioux }@univ-cezanne.fr

Abstract

It was shown that constraint satisfaction problems (CSPs) with a low width can be solved effectively by structural methods. In particular, the BTM method which exploits the concepts of goods and nogoods makes it possible to solve efficiently difficult instances. However, the memory space required for the storage of these (no)goods may make difficult or impossible the resolution of certain problems. We propose here to represent goods and nogoods with Binary Decision Diagrams (BDD). BDDs are data structures which efficiently represent informations in a compact and canonical form. Then, the practical interest of this trade-off which allows to save space memory to the detriment of time is assessed.

Introduction

The CSP formalism (Constraint Satisfaction Problem) offers a powerful framework for representing and solving efficiently many problems. Modeling a problem as a CSP consists in defining a set X of variables x_1, x_2, \dots, x_n , which must be assigned in their respective finite domain, by satisfying a set C of constraints which express restrictions between the different possible assignments. A solution is an assignment of every variable which satisfies all the constraints. Determining if a solution exists is a NP-complete problem.

The usual method for solving CSPs is based on backtracking search. This approach, often efficient in practice, has an exponential theoretical time complexity in $O(e.d^n)$ for an instance having n variables and e constraints and whose largest domain has d values. Several works have been developed, in order to provide better theoretical complexity bounds according to particular features of the instance. The best known complexity bounds are given by the "tree-width" of a CSP (often denoted w). This parameter is related to some topological properties of the constraint graph which represents the interactions between variables via the constraints. It leads to a time complexity in $O(n.d^{w+1})$. Different methods have been proposed to reach this bound like *Tree-Clustering* (Dechter & Pearl 1989) (see (Gottlob, Leone, & Scarcello 2000) for a survey and a theoretical comparison of these methods). They rely on the notion of tree-decomposition of the constraint graph. They aim to cluster

variables such that the cluster arrangement is a tree. Depending on the instances, we can expect a significant gain w.r.t. enumerative approaches. Yet, the space complexity, often linear for enumerative methods, may make such an approach unusable in practice. It can be reduced to $O(n.s.d^s)$ with s the size of the largest minimal separators of the graph (Dechter & Fattah 2001). Several works based on this approach have been performed. Most of them only present theoretical results. Except (Jégou & Terrioux 2003; Gottlob, Hutle, & Wotawa 2002), no practical results have been provided. Clearly, this lack of practical results is mostly explained by the greediness of these methods in terms of memory space. Even for implemented methods, the question about the amount of required memory is raised since these methods are not able to solve any instance. For instance, the BTM method (Jégou & Terrioux 2003), whose recorded informations (namely structural goods and nogoods) are memorized in hash tables, requires sometimes more memory than available for solving some instances.

A solution to this problem may consist in using Binary Decision Diagram (BDD) (Bryant 1986). BDDs are data structures which efficiently represent informations in a compact and canonical form. They are used in many areas, like circuit design, combinatorial logic, ... This approach was already used besides in solving Valued CSP with an extension of BTM (Sachenbacher & Williams 2005). However, the presented results did not make it possible to precisely evaluate the interest of this approach, in particular for the case of structured problems. In this article, we empirically study the use of BDDs for a method like BTM. We thus try to better understand their contribution. We note in particular that the profit in space is very significant. It makes it possible to solve problems by avoiding the saturation of the memory. However, the time required for the management of BDDs results in less interesting performances. That leads us to propose orientations of research to optimize the use of BDDs within this framework.

This paper is organized as follows. The next section provides the basic notions about CSPs and methods based on the tree-decomposition notion. In the third section, we remind of the BDD framework. The fourth section explains how BDDs are exploited in the BTM method. Before the conclusion, we give some experimental results to assess the practical interest of our propositions. In the last section, we

conclude and discuss about relevant works.

Preliminaries

A *constraint satisfaction problem* (CSP) is defined by a tuple (X, D, C) . X is a set $\{x_1, \dots, x_n\}$ of n variables. Each variable x_i takes its values in a finite domain from D (d denotes the size of the largest domain). The variables are subject to the constraints from C . Given an instance (X, D, C) , the CSP problem consists in determining if there is an assignment of each variable which satisfies each constraint. This problem is NP-complete. In this paper, without loss of generality, we only consider binary constraints (i.e. constraints which involve two variables). So, the structure of a CSP can be represented by the graph (X, C) , called the *constraint graph*. The vertices of this graph are the variables of X and an edge joins two vertices if the corresponding variables share a constraint.

Tree-Clustering (Dechter & Pearl 1989) is the reference method for solving CSPs thanks to the structure of its constraint graph. It is based on the notion of tree-decomposition of graphs (Robertson & Seymour 1986). Let $G = (X, C)$ be a graph, a *tree-decomposition* of G is a pair (E, \mathcal{T}) where $\mathcal{T} = (I, F)$ is a tree with nodes I and edges F and $E = \{E_i : i \in I\}$ a family of subsets of X , such that each subset (called cluster) E_i is a node of \mathcal{T} and verifies:

- $\cup_{i \in I} E_i = X$,
- for each edge $\{x, y\} \in E$, there exists $i \in I$ with $\{x, y\} \subseteq E_i$,
- for all $i, j, k \in I$, if k is in a path from i to j in \mathcal{T} , then $E_i \cap E_j \subseteq E_k$.

The width of a tree-decomposition (E, \mathcal{T}) is equal to $\max_{i \in I} |E_i| - 1$. The *tree-width* w of G is the minimal width over all the tree-decompositions of G .

The time complexity of Tree-Clustering is $O(n.d^{w+1})$ while its space complexity can be reduced to $O(n.s.d^s)$ with s the size of the largest minimal separators of the graph (Dechter & Fattah 2001). Note that Tree-Clustering does not provide interesting results in practical cases. So, an alternative approach, also based on tree-decomposition of graphs was proposed in (Jégou & Terrioux 2003). This method is called BTM and seems to provide empirical results among the best ones obtained by structural methods.

The BTM method (for Backtracking with Tree-Decomposition) proceeds by an enumerative search guided by a pre-established partial order induced by a tree-decomposition of the constraint-network. So, the first step of BTM consists in computing a tree-decomposition. The computed tree-decomposition induces a partial variable ordering which allows BTM to exploit some structural properties of the graph and so to prune some parts of the search tree, what distinguishes BTM from other enumerative methods. More precisely, such a variable ordering is produced thanks to a depth-first traversal of the cluster tree. So, BTM begins with the variables of the root cluster E_1 . Inside a cluster E_i , it proceeds classically like any backtracking algorithm by assigning a value to a variable, checking constraints and backtracking if a failure occurs.

When all the variables of the cluster E_i are assigned, BTM keeps on the search with the first son of E_i (if there is one). More generally, let us consider a son E_j of E_i . Given the current assignment \mathcal{A} on $E_i \cap E_j$, BTM checks whether the assignment \mathcal{A} corresponds to a structural good or nogood. A *structural good* (respectively *nogood*) of E_i with respect to E_j is a consistent assignment \mathcal{A} on $E_i \cap E_j$ such that there exists (respectively does not exist) a consistent extension of \mathcal{A} on $Desc(E_j)$. $Desc(E_j)$ denotes the variables which belong to the descent of the cluster E_i rooted in E_j . If \mathcal{A} corresponds to a good, we already know that the assignment \mathcal{A} can be consistently extended on $Desc(E_j)$ and so BTM does not solve again the subproblem corresponding to $Desc(E_j)$. It keeps on the search with the next cluster according to the considered depth-first traversal of the root cluster. In case \mathcal{A} corresponds to a nogood, we already know that there exists no consistent extension of \mathcal{A} on $Desc(E_j)$. Then BTM does not solve again the subproblem corresponding to $Desc(E_j)$ and a backtrack occurs. Finally, if \mathcal{A} corresponds neither to a good nor to a nogood, BTM solves the subproblem rooted in E_j . If BTM succeeds in extending consistently \mathcal{A} on $Desc(E_j)$, \mathcal{A} is recorded as a new structural good on $E_i \cap E_j$. Otherwise, \mathcal{A} is memorized as a new structural nogood. Note that a structural nogood is a particular kind of nogood justified by structural properties of the constraint network. Thanks to the recording and the exploitations of both goods and nogoods which allow it to prune some redundant parts of the search space, BTM offers an interesting theoretical time complexity bound in $O(n.d^{w+1})$ while classical enumerative algorithms have a time complexity in $O(e.d^n)$ ($w+1 \leq n$). Unfortunately, the space complexity, generally linear for classical enumerative algorithms, is in $O(n.s.d^s)$, what is the main drawback of structural methods like BTM. Due to the amount of required memory, few structural methods have been implemented and used successfully. The experimental results about BTM given in (Jégou & Terrioux 2003) have been obtained by using a hash table for each separator. However, this solution does not allow to solve any problem. In some cases, the amount of available memory is not sufficient for solving some problems. Hence, the use of BDDs for recording goods and nogoods may allow us to reduce the amount of required memory.

ROBDDs for partial assignments

In the framework of BDDs, Reduced Ordered Binary Decision Diagrams (ROBDD (Bryant 1986; 1992)) are commonly exploited. ROBDDs aim to represent boolean functions under the shape of oriented graphs without circuit. The OBDDs offer a powerful setting for solving boolean equation systems or for the treatment of various operations on boolean functions. More generally, they make it possible to represent sets in a concise way, such as for example of the sets of assignments. Their principles and mechanisms are described in details in (Akers 1978; Bryant 1986; Brace, Rudell, & Bryant 1990; Bryant 1992; Madre & Billon 1988) which are giving some building optimization. We recall in this sections, their principles and mechanisms of construction.

Given a boolean formula F and X its set of variables, we consider a total order (x_1, \dots, x_n) on X . The decision tree associated to F is a labeled path to nodes representing all interpretations of F . Internal nodes are labeled by elements of X , while leaves or *terminal nodes* are labeled by 0 or 1. These labels are noted $var(s)$ for each node s compatible with the order on X : a node of the i^{th} level in the graph is labeled $var(s) = x_i$, the root is labeled x_1 . The internal nodes s possess two children corresponding to the interpretations of $var(s)$: the *left child* $lc(s)$ ($var(s)$ is interpreted to 0) and the *right child* $rc(s)$ ($var(s)$ is interpreted to 1). One calls vertices $(s, lc(s))$ and $(s, rc(s))$ respectively the left vertex and the right vertex. Thus, every *maximal path* joining the root to a leaf is equivalent to an interpretation; it is a model if the label of the leaf is 1 (*positive maximal path*) and an counter-model if the label is 0.

The OBDD representing a boolean function F corresponds to one concise expression of the decision tree of F . It is a directed graph without circuit but can possess cycles.

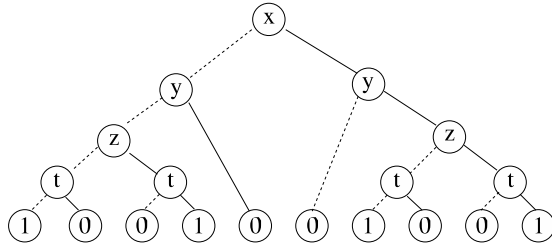


Figure 1: Function and decision tree for the formula $(x \leftrightarrow y) \wedge (z \leftrightarrow t)$ according to the order (x, y, z, t) (Bryant 1986). The left edges are in dotted lines, the right edges in solid lines.

The OBDD is the smallest graph which satisfies the following properties:

- it contains at most two terminal nodes: one labeled 1 and the other 0 ; if the represented function is a tautology (or a function with no model), the graph is reduced to a unique node labeled 1 (or 0).
- for any internal node s , $var(s) < var(lc(s))$ and $var(s) < var(rc(s))$, but if $var(s) = x_i$, we do not have necessarily neither $var(lc(s)) = x_{i+1}$, nor $var(rc(s)) = x_{i+1}$, nor $var(lc(s)) = var(rc(s))$.

Figure 1 gives an example of an OBDD. Every maximal path of the OBDD corresponds to a partial instantiation, restricted to variable labels of nodes of the path. If the label of the last node is 1 (or 0), all the extensions of this interpretation are models (or counter-models) of the represented function. Conversely, to any interpretation corresponds an unique maximal path in the OBDD. We will note I_c the interpretation associated to the maximal path c . The consistency check of a function F is achieved by verifying if the OBDD is reduced to the terminal node 0. To verify if an interpretation is a model, it is sufficient to browse the OBDD from the root while achieving the branchings corresponding to the interpretation. The time complexity is linear in the

number of variables. A model can be obtained by searching a positive maximal path. Its complexity is $O(|B_F|)$ where $|B_F|$ is the size of the OBDD associated to F .

The size of a OBDD can be significantly reduced using other reductions in order to obtain a ROBDD, that is a Reduced OBDD. The reduction of the graph associated to a formula F relies on an elimination of redundant nodes. This reduction does not modify the satisfiability of the formula coded, but allows to reduce considerably the OBDD size compared to the decision tree. The elimination of redundancy in the graph representing the coded function is defined by this three transformations:

1. duplicated external node elimination: all external nodes labeled 0 (or 1) are merged in only one node labeled 0 (or 1).
2. duplicated internal node elimination: if two internal nodes u and v are such that $var(u) = var(v)$, $lc(u) = lc(v)$ and $rc(u) = rc(v)$, then these nodes are merged.
3. redundant internal node elimination: an internal node u verifying $lc(u) = rc(u)$ is eliminated, the retractable incident edge of u being directed towards $lc(u)$.

The graph representing a boolean function is *reduced* if it contains no internal node u such that $lc(u) = rc(u)$, and if it does not contain two distinct internal nodes u and v such that the sub-graphs rooted by u and v are isomorphic (i.e. they represent a same function). A ROBDD is a reduced graph representing a boolean function. The reduced graphs possess some properties (Bryant 1986):

- For every reduced graph, for every node u of this graph, the sub-graph rooted by u is a reduced graph.
- Given a boolean function F and an order on the variables of F , there is an unique (up to isomorphism) reduced graph representing this function ; it is the ROBDD representing F . Any other graph representing F contains more nodes.

The ROBDD reduction depends on the variable ordering. The order impact on the size of the ROBDD can be significant (Bryant 1992). For example, for boolean functions representing the addition of integer numbers, the size of the ROBDD can grow linear to exponential. Furthermore, there are some pathological cases, as boolean functions representing the multiplication of integer numbers, for some order, the size of the ROBDD is exponential. Figure 2 provides two examples of reduction according to two different orders for the formula considered in Figure 1.

In order to build the ROBDD associated to a function F writing itself by $f < op > g$ where $< op >$ is an boolean operator, one has to compose the sub-graphs B_f and B_g associated to f and g . The time complexity is in $O(|B_f| \cdot |B_g|)$ where $|B_f|$ and $|B_g|$ denote respectively the number of ROBDDs nodes for B_f and B_g . Especially, if F is a function having a variable x in its scope, the computation of the ROBDD coding the restriction of F to x , $F < and > x$ (case where $x = 1$) will be linear in the size of the ROBDD.

There exist several extension of BDD. Each extension depends on its application area. We can cite, for example,

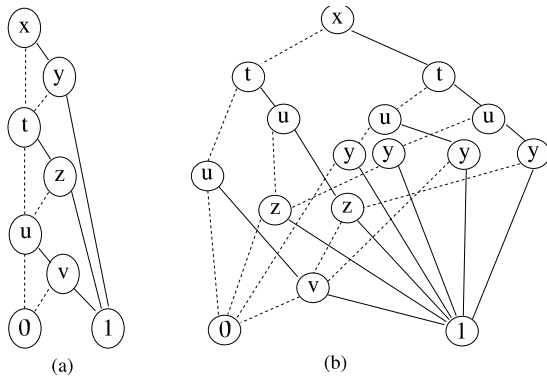


Figure 2: Influence of the order in BDD size.

FDD, ADD, BED, MTBDD, BMD, KMDD, BGD,... In our approach, we use the MDD extension (Srinivasan *et al.* 1990). It represents a discrete function whose input variables are multi-valued. MDD is a rooted, directed, ordered acyclic graph. Each internal node corresponds to a multi-valued variable and each leaf node represents one value of the function. Each internal node has d edges such that each edge corresponds to one of the d possible values for a variable.

Good and nogood represented by BDD

Solving a CSP instance thanks to the BTD method often requires to record a large amount of informations (namely goods and nogoods). The goods and the nogoods allow to save significantly time but consume a great quantity of memory. In fact, currently, for the empirical results presented in (Jégou & Terrioux 2003; Jégou, Ndiaye, & Terrioux 2005), goods and nogoods are vectors of values memorized in hash tables. When we use the hash tables, we often memorized redundant informations. Indeed, in most cases, as shown in the next paragraphs, partial instantiations can appear several times in the assignments. So, we clearly see the interest to use a compact and effective structure which makes it possible to reduce the size of the recorded data. Our choice is related to an extension of BDD to finite domains.

The adequate version of BDD to our problem is Multi-valued Decision Diagrams (MDD) (Srinivasan *et al.* 1990). This extension makes it possible to represent canonically a set of finite domains. We exploit the package extracted from VIS¹. This package has been developed at the Colorado University. It also uses the CUDD package². We note that, in most of the applications, MDDs are built with sets of ROBDDs in the internal structures. Each multi-valued variable is then decomposed in a set of binary variables. For example, in figure 3, we represent $x \in \{0, 1, 2\}$ by two binary variables. More generally, we decompose x in $\log_2(\max_{a \in D_x}(a))$ binary variables (D_x denotes the

value domain of x). In this manner, the set of the values taken by a multi-valued variable is built on a micro-structure ROBDD. Of course, there exist packages that implement directly MDDs without passing by ROBDD structure (Miller & Drechsler 1998). Unfortunately they suffer in general from the problem of optimization (Schmiedle, Günther, & Drechsler 2000).

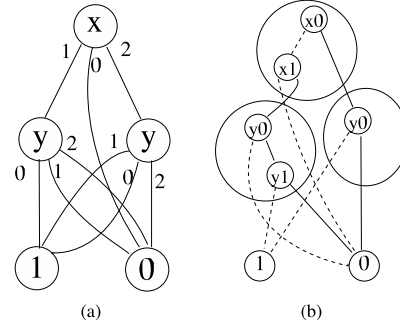


Figure 3: Mapping from a MDD to a ROBDD (Kam *et al.* 1998).

In order to obtain good results, the variables are ordered according to the variable ordering induced by the tree-decomposition. This order is static. Indeed, we have observed that the results with dynamic orders for all binary variables which optimize the required memory space do not show a great usefulness: the saved amount of memory space with a dynamic order does not exceed 15% with respect to a static order, while we may spend 40% of additional time.

Experimental results

Before presenting the empirical results, we first describe the experimental protocol.

Experimental protocol Applying a structural method on an instance generally assumes that this instance presents some particular topological features. So, our study is performed on instances having a close to ideal structure. In practice, the two current ways of recording goods and nogoods are compared here on random partial structured CSPs in order to point up the best one w.r.t. the CSP solving runtime and the required amount of memory space. For building a random partial structured instance of a class (n, d, w, t, s, ns, p) , the first step consists in producing randomly a structured CSP according to the model described in (Jégou & Terrioux 2003). This structured instance consists of n variables having d values in their domain. Its constraint graph is a clique tree with ns cliques whose size is at most w and whose separator size does not exceed s . Each constraint forbids t tuples. Then, the second step removes randomly $p\%$ edges from the structured instance. The experimentations are performed on a Linux-based PC with a Pentium IV 2.8GHz and 512MB of memory. For each considered class, the presented results are the average on 50 instances. We limit the runtime to 30 minutes. Above, the solver is stopped and the involved instance is considered as unsolved.

¹Verification Interacting with Synthesis.
<http://vlsi.colorado.edu/~vis>

²Colorado University Decision Diagrams:
<http://vlsi.colorado.edu/~fabio>

In the following tables, the symbol $>$ denotes that at least one instance cannot be solved within 30 minutes and so the mean runtime is greater than the provided value. The letter M means that at least one instance cannot be solved because it requires more than 512MB of memory.

In (Jégou, Ndiaye, & Terrioux 2005), a study was performed on triangulation algorithms to find out the best way to compute a good tree-decomposition w.r.t. CSP solving. As MCS (Tarjan & Yannakakis 1984) obtains the best results, we use it to compute tree-decompositions in this study.

Given a tree-decomposition, we choose as root cluster the cluster which minimizes the ratio of the expected number of partial solutions of the cluster over its size. Likewise, for each cluster, its sons are ordered according this decreasing ratio. Inside a cluster, the unassigned variables are ordered thanks to the *dom/deg* heuristic. This heuristic chooses as next variable the variable x which minimizes the ratio number of the remaining values for x over the degree of x in the constraint graph.

Experimental results In this part, we compare two versions of the BTM method. These two versions differ in the way they store the goods and the nogoods. On the one hand, the goods and nogoods are stored in several hash tables (one per separator). It is the initial version of BTM (Jégou & Terrioux 2003). On the other hand, in the version proposed in this paper, these informations are recorded in several MDDs (one per separator). This comparison only focuses on the runtime and the required amount of memory space. In particular, we do not need to consider other datas like the number of visited nodes or the number of performed constraint checks. Indeed, the two versions exactly obtain the same results if they exploit the same heuristics for choosing the root cluster, the next son cluster or the next variable to visit. Regarding the required amount of memory space, for the version based on hash tables, we assess it by counting the total number of recorded values. For instance, we count 3 for a good which involves 3 variables. For the version based on MDDs, we count the total number of binary nodes. For instances, the MDD of figure 3 is represented by 5 binary nodes.

Tables 1 and 2 provide the obtained results for a limited separator size and an unlimited one. One of the main interests of the restriction of the separator size consists in limiting the amount of required memory space. Indeed, with smaller separators, the size of goods and nogoods and their potential number decrease. Without such a limitation, the BTM version based on hash tables turns sometimes to be unable to solve some instances by lack of memory space.

Table 1 highlights the great performances of the version based on MDDs in terms of memory space. Such a result is not surprising because the recorded goods and nogoods often share values. MDDs consume at least 15 times less memory space as hash tables. For the considered classes of instances, this rate is often greater than 50. The comparison between the number of recorded values in the hash tables and the number of binary nodes in the MDDs clearly shows how much the informations related to (no)goods are compressed in MDDs. From a practical viewpoint, several bi-

nary nodes are required for representing a value. However, a binary node can be used in the representation of several values which appear in different (no)goods. In particular, in some cases, a binary node can be used several times in the representation of a single value which belongs to different (no)goods. That explains the small number of binary nodes with respect to the number of recorded values and so the significant gain in memory we have observed.

Regarding the runtime presented in Table 1, we observe an inverse behaviour but the rate is less important. The version based on MDDs is at most twice as slow than one based on hash tables due to the cost of the main operations. For hash tables, the memorization of a new good or nogood can be achieved in linear time (w.r.t. the size of the considered good or nogood) while checking if a good or a nogood is present in the hash table requires a time close to linear as soon as the goods and nogoods are fairly distributed in the hash table. For MDDs, the addition or the check can be performed in $O(a * \log_2(a))$ where a is the size of the considered good or nogood. The $\log_2(a)$ factor comes from the decomposition of the multi-valued variables in binary variables. Hence, a direct representation as a MDD (i.e. without a mapping to BDD) would be more interesting here. However, if, by so doing, we save a $\log_2(a)$ factor for the runtime, we consume more memory with the same factor. We note that the two versions solve all the instances, except one instance of the class (250,20,20,99,10,25,0.1) for the version based on MDDs. This instance cannot be solved within the time limit.

Given the promising results obtained thanks to MDDs in terms of required memory space, we assess the behaviour of the two versions for an unlimited separator size. By so doing, the size and the number of goods and nogoods increase and so we can expect a greater benefit from MDDs. Table 2 presents the observed results. Like previously, the version based on MDDs outperforms one based on hash tables w.r.t. the required memory space while it spends more time for solving the instances. This additional time corresponds again to the cost of managing goods and nogoods in the MDDs. Nonetheless, unlike for a limited separator size, the version based on hash tables does not succeed in solving all the instances. The amount of required memory space prevents from solving several instances. Note that the compactness of the MDD representation allows to solve these same instances.

The compactness of recorded informations allows to reduce the amount of required memory space but it requires some additional runtime. Hence, unlike the results about the memory space, the runtime obtained by using MDDs is not competitive enough with respect to one of the initial version of BTM based on hash tables. As explained above, this additional cost results from the construction and the management of MDDs mapped to BDDs. However, in spite of the non-competitive runtime, the BTM version based on MDDs remains interesting. Indeed, it is often possible to spend more time for solving an instance whereas we cannot consume more memory than available and, unfortunately, we cannot foresee the amount of needed memory space.

Instances (n, d, w, t, s, ns, pr)	Memory Space				Time	
	Size				Hash	MDD
	Hash		MDD			
	# values	MB	# nodes	MB		
(150,25,15,215,5,15,0.1)	16,168	6.75	6,795	0.11	2.30	2.69
(150,25,15,237,5,15,0.2)	22,799	7.64	7,652	0.12	1.79	2.38
(150,25,15,257,5,15,0.3)	29,448	9.46	7,412	0.18	1.01	1.80
(150,25,15,285,5,15,0.4)	5,418	13.12	3,764	0.06	0.40	0.52
(250,20,20,107,5,20,0.1)	47,836	9.11	8,558	0.14	10.39	11.70
(250,20,20,117,5,20,0.2)	59,392	10.33	9,516	0.15	8.52	10.46
(250,20,20,129,5,20,0.3)	48,135	11.63	5,408	0.09	5.82	7.91
(250,20,20,146,5,20,0.4)	90,180	14.83	8,250	0.13	3.81	6.03
(250,20,20,99,10,25,0.1)	1,554,308	25.22	100,696	1.61	58.21	>82.36
(250,25,15,211,5,25,0.1)	70,326	11.37	18,968	0.31	7.12	9.49
(250,25,15,230,5,25,0.2)	72,645	12.78	19,472	0.32	4.13	6.30
(250,25,15,253,5,25,0.3)	85,627	15.79	13,713	0.22	4.00	6.30
(250,25,15,280,5,25,0.4)	60,960	21.42	17,041	0.27	1.61	3.27

Table 1: Number of recorded value in hash tables, number of binary nodes in the MDDs, required memory space in MB for hash tables and MDDs, and runtime in seconds for a separator size limited to 5. For the class (250,20,20,99,10,25,0.1), one instance cannot be solved within the time limit by the version based on MDDs. For this class, the reported MDD size and the ratio correspond to the mean over the 49 solved instances.

Conclusion and discussion

In this article, we have studied the resolution of structured CSP. In particular, we have been interested in the BTM method (Jégou & Terrioux 2003) whose efficiency results from the exploitation of structural goods and nogoods learnt and recorded during the search. Whereas, in its initial version, BTM represented goods and nogoods in extension with hash tables, we have studied here from a practical viewpoint the interest which can present a memorization of these informations in a compact structure like BDDs (MDDs precisely).

A similar work has already been performed by (Sachenbacher & Williams 2005) with an extension of BTM for solving Valued CSPs. However, this work does not make it possible to determine the real interest of the use of the BDDs (ADDs in this work), in particular for the case of structured CSPs. Here, we present a study which aims to better assess this interest with respect to the saved amount of memory, but also the runtime.

Concerning the structured CSP, we have observed a very significant profit in terms of required memory space. Indeed, several problems which could not be solved by BTM with the hash tables are now manageable. More generally, one observes a systematic profit for space on all the problems. We have also noted that this profit is still better when the problems are inconsistent since, in this case, the search space must be completely traversed. So, one can hope that the potential profits will be increased if we consider optimization problems for which the whole search space must be treated.

Concerning the runtime, we have observed a degradation of the efficiency. Indeed, the time devoted to the management of BDDs, in particular for the addition of goods and nogoods, slows down significantly the effectiveness of the

approach. This report leads us to continue this work while trying to better manage space. In particular, we should propose an approach which would improve significantly the runtime.

On the level of the other prospects to this work, we will evaluate this approach on real problems. However, it still seems more interesting to us to focus our study on optimization problems like Valued CSP rather than decision ones. That is possible by exploiting ADDs (Bahar *et al.* 1993) like proposed in (Sachenbacher & Williams 2005). However, such an extension could also pass by the design of a new kind of BDDs better adapted to the resolution of optimization problems.

Acknowledgments

This work is supported by a "programme blanc" ANR grant (STAL-DEC-OPT project).

References

- Akers, S. B. 1978. Binary decision diagrams. *IEEE Trans. Computers* 27(6):509–516.
- Bahar, R. I.; Frohm, E. A.; Gaona, C. M.; Hachtel, G. D.; Macii, E.; Pardo, A.; and Somenzi, F. 1993. Algebraic decision diagrams and their applications. In *Proceedings of the 1993 IEEE/ACM international conference on Computer-aided design*, 188–191.
- Brace, K. S.; Rudell, R. L.; and Bryant, R. E. 1990. Efficient Implementation of a BDD Package. In *Proceedings of the ACM/IEEE Design Automation Conference (DAC)*, 40–45.
- Bryant, R. E. 1986. Graph-based algorithms for boolean function manipulation. *IEEE Transactions on Computers* 35(8):677–691.

Instances (n, d, w, t, s, ns, pr)	Memory Space				Time	
	Size				Hash	MDD
	Hash		MDD			
	# values	MB	# nodes	MB		
(150,25,15,215,5,15,0.1)	188,510	16.04	90,042	1.44	2.57	8.25
(150,25,15,237,5,15,0.2)	340,683	20.58	123,594	1.98	2.70	12.65
(150,25,15,257,5,15,0.3)	252,311	23.60	86,961	1.39	1.55	8.71
(150,25,15,285,5,15,0.4)	M	-	55,573	0.89	M	3.44
(250,20,20,107,5,20,0.1)	1,898,500	31.82	317,018	5.07	18.17	43.30
(250,20,20,117,5,20,0.2)	2,614,225	41.70	274,648	4.39	13.67	37.91
(250,20,20,129,5,20,0.3)	2,731,434	46.86	363,931	5.82	12.54	61.26
(250,20,20,146,5,20,0.4)	463,786	39.54	150,649	2.41	2.37	15.64
(250,20,20,99,10,25,0.1)	M	-	960,291	15.36	M	139.00
(250,25,15,211,5,25,0.1)	317,138	26.26	202,155	3.23	6.04	18.07
(250,25,15,230,5,25,0.2)	2,235,933	45.53	356,295	5.70	24.90	33.04
(250,25,15,253,5,25,0.3)	M	-	236,044	3.77	M	30.84
(250,25,15,280,5,25,0.4)	3,173,339	64.89	452,737	7.24	12.81	81.27

Table 2: Number of recorded value in hash tables, number of binary nodes in the MDDs, required memory space in MB for hash tables and MDDs, and runtime in seconds for an unlimited separator size.

Bryant, R. E. 1992. Symbolic Boolean manipulation with ordered binary-decision diagrams. *ACM Computing Surveys* 24(3):293–318.

Dechter, R., and Fattah, Y. E. 2001. Topological Parameters for Time-Space Tradeoff. *Artificial Intelligence* 125:93–118.

Dechter, R., and Pearl, J. 1989. Tree-Clustering for Constraint Networks. *Artificial Intelligence* 38:353–366.

Gottlob, G.; Hutle, M.; and Wotawa, F. 2002. Combining hypertree, bicomp and hinge decomposition. In *Proc. European Conference on Artificial Intelligence*, 161–165.

Gottlob, G.; Leone, N.; and Scarcello, F. 2000. A Comparison of Structural CSP Decomposition Methods. *Artificial Intelligence* 124:343–282.

Jégou, P., and Terrioux, C. 2003. Hybrid backtracking bounded by tree-decomposition of constraint networks. *Artificial Intelligence* 146:43–75.

Jégou, P.; Ndiaye, S. N.; and Terrioux, C. 2005. Computing and exploiting tree-decompositions for solving constraint networks. In *Proc. of the 11th International Conference on Principles and Practice of Constraint Programming*, 777–781.

Kam, T.; Villa, T.; Brayton, R. K.; and Sangiovanni-Vincentelli, A. L. 1998. Multi-valued decision diagrams: Theory and applications. *International Journal on Multiple-Valued Logic* 4(1-2):9–62.

Madre, J.-C., and Billon, J.-P. 1988. Proving circuit correctness using formal comparison between expected and extracted behaviour. In *Proceedings of the ACM/IEEE Design Automation Conference (DAC)*, 205–210. IEEE Computer Society Press.

Miller, D., and Drechsler, R. 1998. Implementing a multiple-valued decision diagram package. In *ISMVL '98: Proceedings of the The 28th International Symposium on*

Multiple-Valued Logic, 52. Washington, DC, USA: IEEE Computer Society.

Robertson, N., and Seymour, P. 1986. Graph minors II: Algorithmic aspects of tree-width. *Algorithms* 7:309–322.

Sachenbacher, M., and Williams, B. C. 2005. Bounded Search and Symbolic Inference for Constraint Optimization. In *Proceedings of IJCAI*, 286–291.

Schmiedle, F.; Günther, W.; and Drechsler, R. 2000. Dynamic re-encoding during mdd minimization. In *ISMVL*, 239–244.

Srinivasan, A.; Kam, T.; Malik, S.; and Brayton, R. K. 1990. Algorithms for discrete function manipulation. In *ICCAD*, 92–95.

Tarjan, R., and Yannakakis, M. 1984. Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM Journal on Computing* 13 (3):566–579.