# Improving Decision Diagrams for Decision Theoretic Planning 

Jean-Christophe Magnan and Pierre-Henri Wuillemin<br>Laboratoire d'Informatique de Paris VI<br>firstname.lastname@lip6.fr


#### Abstract

In the domain of decision theoretic planning, the factored framework (FMDP) has produced optimized algorithms using Decision Trees (SVI, SPI) and Algebraic Decision Diagrams (SPUDD). However, the state-of-the-art SPUDD algorithm requires i) the problem to be specified with binary variables and ii) the data structures to share a common order on variables. In this article, we propose a new algorithm within the factored framework that eliminates both these requirements. We compare our approach to the SPUDD algorithm. Experimental results show that our algorithm allows significant gains in time, illustrating a better trade-off between theoretical complexity of algorithms and size of representation.


## Introduction

In Decision Theoretic Planning, Markov Decision Process (MDP) is an important modeling tools. This framework models a decision problem as sets of states and actions, probabilistic transitions from state to state under actions and rewards triggered by specific states on certain transitions. Two main algorithms, value iteration (VI) and policy iteration (PI), exploit this representation and enable to efficiently solve planning problems. Indeed, these algorithms have a linear complexity in regard of the state space size (Puterman 1994).

Unfortunately, state spaces for realistic problems are often too large to find out optimal solutions in reasonable time. To deal with that issue, some promising solutions have emerged around the concept of abstraction : considering and treating similar states together leads to a reduction in the complexity of the state space enumeration. Several variants of MDPs use that approach, for instance hierarchical MDPs (Guestrin 2002) or Factored MDPs (FMDPs: Boutilier, Dean, and Hanks 1999).

FMDPS rely on the description of the system by a set of variables. Each state of the system is represented by a unique instantiation of those variables. The different inputs of the model are then factored: probability distributions with dynamic Bayesian networks (dBNs: Dean and Kanazawa 1989) and rewards with additive decomposition. However, conditional probability tables used in a dBN can still grow expo-

[^0]nentially. Furthermore, the whole enumeration of the state space during VI and PI remains an issue.

Boutilier, Dearden, and Goldszmidt proposed a solution to those problems using decision trees to represent conditional probability tables, and reward functions. Such compact representation of each function allows to design efficient algorithms that avoid the exhaustive enumeration of all states: based on VI and PI, Structured Value Iteration (SVI) and Structured Policy Iteration (SPI) significantly improve the experimental time and space complexity. To further improve these complexities, Hoey et al. proposed another and more compact graphical representation : Algebraic Decision Diagrams (ADD). ADDs are a reduced version of decision trees where isomorphic subgraphs are merged together (Bryant 1986). From there, research in this domain has mainly focused on gains obtained by approximations : approximated ADDs (St-aubin, Hoey, and Boutilier 2000), basis functions (Guestrin, Parr, and Venkataraman 2003), etc.

However ADDs hold some drawbacks. First, all the variables in the model have to be binary. Variables with three or more values have to be decomposed in binary variables to be inserted in ADDs. This recasting artificially increases the state space. The second drawback is that ADDs compactness strongly depends on their inherent order on variables. Finding an order on variables that would produce an optimally compacted ADD is NP-hard (Friedman and Supowit 1990). Moreover, the algorithm for operations on two ADDs (for instance addition, multiplication or maximization) requires that orders on those ADDs are the same. As a consequence, a global order on variables is needed for the algorithm and can be largely sub-optimal for certain ADDs, leading to an artificial increase in the complexity for these operations.
The purpose of this paper is to propose new ways to adress those issues. Its main contributions are i) to propose a more efficient data structure to represent functions of multi-valued variables, and ii) to present a new algorithm for operations on two decision diagrams without imposing a common order on both diagrams. This article is organized as follows: Section 2 covers MDPs and the use of compact representation to find optimal policy. ADDs and the limitations they brought will be tackle there, as well as the new model we propose to use instead. Section 3 presents the new algorithm for the operation on two decision diagrams. Finally, the ex-
perimentations in Section 4 illustrate the efficiency of these improvements.

## Factored Markov Decision Processes

We assume that our planning problem can be modeled by a fully-observable MDP. Let $S$ and $A$ be respectively the set of states and the set of actions. The conditional probability $\mathrm{P}\left(\mathrm{s} \mid \mathrm{a}, \mathrm{s}^{\prime}\right)$ represents the probability of reaching state $\mathrm{s}^{\prime}$ when action $a$ is performed in state $s$. In order to simplify the presentation, the reward is formalized as a real-value function $R(s)$ depending only on the current state $s$.

A stationary policy $\pi: S \rightarrow A$ defines the action a to execute when system reaches the state s. Assuming that the objective is to find an optimal policy over an infinite time horizon, we can compare two policies upon the expected total discounted reward defined on every state:

$$
\begin{equation*}
\mathrm{V}_{\pi}(\mathrm{s})=\mathrm{R}(\mathrm{~s})+\gamma \sum_{s^{\prime} \in S} \mathrm{P}\left(\mathrm{~s}, \pi(\mathrm{~s}), \mathrm{s}^{\prime}\right) \cdot \mathrm{V}_{\pi}\left(\mathrm{s}^{\prime}\right) \tag{1}
\end{equation*}
$$

where $\gamma$ is used to discount future rewards. $V_{\pi}$ is called the value function for policy $\pi$. An optimal policy $\pi^{*}$ is a policy that verifies: $\forall \pi, \forall s \in S, V_{\pi^{*}}(s) \geq V_{\pi}(s)$.

To find the optimal policy, VI algorithm consists in converging toward optimal value function by iterative updating:

$$
\begin{align*}
\mathrm{V}^{0}(\mathrm{~s}) & =\mathrm{R}(\mathrm{~s}) \\
\mathrm{V}^{\mathrm{n}+1}(\mathrm{~s}) & =\mathrm{R}(\mathrm{~s})+\gamma \cdot \max _{\mathrm{a} \in \mathrm{~A}} \sum_{s^{\prime} \in S} \mathrm{P}\left(s^{\prime} \mid a, s\right) \cdot \mathrm{V}^{\mathrm{n}}\left(s^{\prime}\right) \tag{2}
\end{align*}
$$

At each iteration of the algorithm, $\mathrm{V}^{\mathrm{n}}$ has to be updated for every state $s \in S$. The linear complexity of each iteration is an issue because $S$ happens to be large even for simple problems. A solution is to enhance the compactness of functions $P(. \mid s, a), R$ and $V$ with graphical and factored representations using discrete variables.

Let $X_{i}$ be a multi-valued variable taking its values over a finite discrete domain $\mathrm{D}_{X_{i}}$. Whenever $X_{i}$ is instantiated (i.e. set to a given value in $\mathrm{D}_{\mathrm{X}_{i}}$ ), it will be noted $x_{i}$.

## Definition (Decomposability of state space).

A state space S is decomposable if and only if there exists a set of discrete and finite variables $X=\left\{\mathrm{X}_{1}, \ldots, \mathrm{X}_{\mathrm{n}}\right\}$ that unequivocally characterizes $S$.

Each state $s \in S$ is then an instantiation of those variables $\left(s=\left\{x_{1}, \ldots, x_{n}\right\}\right)$.

When $S$ is decomposable, transition probabilities, value function and rewards are expressed as functions of $\left\{X_{1}, \ldots, X_{n}\right\}$ and, as a consequence, can be compactly represented by function graphs.

Let $f$ be a function over the variables $X_{1}, \ldots, X_{n}$. We denote $\left.f\right|_{X_{i}=b}$ the restriction of $f$ on $X_{i}=b$. The support of $f$ is the set of variables that $f$ really depends on, i.e.,

$$
\operatorname{support}(f)=\left\{X_{i} \mid \exists u, v \in D_{X_{i}} \text { s.t. }\left.f\right|_{X_{i}=u} \neq\left. f\right|_{X_{i}=v}\right\}
$$

Note that $\forall X_{i}$, support $\left(\left.f\right|_{X_{i}=b}\right) \subseteq \operatorname{support}(f) \backslash\left\{X_{i}\right\}$. Indeed the support of the restriction discards all non relevant variables and not only the variable $X_{i}$.

## Definition (Function Graph).

Let f be a function over $\left\{\mathrm{X}_{1}, \ldots, \mathrm{X}_{\mathrm{n}}\right\}$. A directed acyclic graph (DAG) $\mathrm{G}_{\mathrm{f}}(\mathrm{N}, \mathrm{A})$ is a function graph of f

- if f constant then $\mathrm{G}_{\mathrm{f}}$ has a unique node $\mathrm{r} \in \mathrm{N}, \mathrm{r}$. val $=\mathrm{f}$
- if f non constant then $\mathrm{G}_{\mathrm{f}}$ has a unique node $\mathrm{r} \in \mathrm{N}$ without parent. Moreover,
- r.var $\in \operatorname{support}(f)$
- $\forall u \in D_{r . v a r}, \exists!n_{u} \in N$ such that:
- $\left(r, n_{u}\right) \in A$
- $\operatorname{subgraph}\left(n_{u}\right)=G_{\left.\right|_{\text {r.var }=u}}$

Note that in a function graph, if a node $n$ is terminal (i.e. without children) then $n$ is bound to a constant value ( $n . v a l$ ) else $n$ is associated to a variable (n.var).

For any node $n$ of $G_{f}, \operatorname{subgraph}(n)$ is a function graph for the restriction of $f$ defined by the instantiation of every variables crossed on the path from the root to $n$. We note $\left.\mathrm{f}\right|_{n}$ this restriction. Note that several path may lead to the same subgraph, characterizing the fact that several restrictions may be equal. For instance, in Figure 1(c),
$\left.f\right|_{\text {terminal node } 2}=\left.f\right|_{X=0, Y=2}=\left.f\right|_{X=1, Y=1}=2$

(a)

(b)

(c)

Figure 1: Two function graphs for the same function.
As shown in Figure 1, there may be several function graphs for one function. But they all share a same property of compactness: no irrelevant variable can appear in a function graph due to the use of support(.). Thus, two states that share the same value for represented function but differ on an irrelevant variable will be viewed as a single abstract state by that representation.

Functions graphs give a compact and efficient way to represent the whole state space. Furthermore, dedicated algorithms for operations (addition, multiplication, maximization) on function graphs enable to create composed functions without an explicit enumeration of all states.

## Function Graphs in MDPs

The optimal policy search using VI algorithm can be reformulated using function graphs. Indeed, let $\mathrm{V}, \mathrm{R}$ and $\mathrm{P}_{\mathrm{a}}$ be respectively the function graphs for the value function, the reward, and the probabilistic transitions for each action $a \in A$. The update will be done exactly like in equation 2. The only differences are the data structures and the dedicated operations used at each iteration :

$$
\mathrm{V}^{0}=\mathrm{R} \quad \text { and } \quad \mathrm{V}^{\mathrm{n}+1}=\mathrm{R}+\gamma \cdot \max _{\mathrm{a} \in \mathcal{A}} \sum_{s^{\prime} \in S} \mathrm{P}_{\mathrm{a}} \cdot \mathrm{~V}^{\mathrm{n}}
$$

In a factored MDP (FMDP: Boutilier, Dearden, and Goldszmidt 1995), the state space is decomposable. Besides, the assumption of some conditional independences allow to factorize the probability distribution: $P_{a}\left(s^{\prime} \mid s\right)=\Pi_{i} P_{a}\left(X_{i}^{\prime} \mid s\right)$.

The first data structure used as function graph for optimal policy search was the decision tree (as in Figure 1(b)). SVI and SPI rely on such data structure and on the specific algorithm for operations on it (Boutilier, Dean, and Hanks 1999).

As shown in Figures 1(b), decision trees may contain several isomorphic subgraphs. Indeed due to the tree structure, those subgraphs can not be merged. That duplication unnecessarily increases the graph size.

Rather than tree, Hoey et al. proposed then to use Algebraic Decision Diagrams (Bahar et al. 1993) as function graphs (as in Figure 1(c)). ADDs are a generalization of Bi nary Decision Diagrams (BDDs: Bryant 1986) used to represent real functions of boolean variables $\left(f: \mathbb{B}^{n} \rightarrow \mathbb{R}\right)$. Their particularities are that they are reduced and ordered.

## Definition (Reduced Function Graph).

$\mathrm{G}_{\mathrm{f}}$ is reduced $\Longleftrightarrow \forall$ nodes $\mathrm{n} \neq \mathrm{n}^{\prime},\left.\mathrm{f}\right|_{\mathrm{n}} \neq\left.\mathrm{f}\right|_{n^{\prime}}$
When a function graph is reduced, two isomorphic subgraphs are necessarily merged together.

## Definition (Ordered Function Graph).

A function graph $\mathrm{G}_{\mathrm{f}}$ is ordered $\Leftrightarrow \exists \succ_{\mathrm{G}_{\mathrm{f}}}$ complete order on support( f ), s.t. $\forall$ nodes $\mathrm{n}_{1}, \mathrm{n}_{2}$ non terminal of $\mathrm{G}_{\mathrm{f}}$,

$$
\mathrm{n}_{2} \in \operatorname{desc}\left(\mathrm{n}_{1}\right) \Rightarrow \mathrm{n}_{1} \cdot \operatorname{var} \succ_{\mathrm{G}_{\mathrm{f}}} \mathrm{n}_{2} \cdot v \operatorname{ar}
$$

When a function graph is ordered, the algorithm for reducing it is polynomial.

Bryant describes the algorithm for the operations on BDDs. The ADDs algorithm relies on the same principle. In SPUDD, Hoey et al. proposes a version of VI algorithm using ADDs as data structure.

However, ADDS represent only functions of boolean variables. As a consequence, all multi-valued variables have to be encoded with binary variables. A first issue is raised by such an encoding ; the state space size is artificially increased. Indeed to code a variable with $n$ values, one needs $\left\lceil\log _{2} n\right\rceil$ binary variables. Those $\left\lceil\log _{2} n\right\rceil$ variables will generates $\left\lceil 2^{\log _{2} n}\right\rceil$ possibles states. Which means that $\left\lceil 2^{\log _{2} n}\right\rceil-n$ states are artificially created and have no real existence. Yet the algorithms (VI, PI, etc.) will search an optimal policy for those states.

In this article, we investigate the use of Multi-valued Decision Diagrams (MDDs: Srinivasan et al. 1990). MDDs simply generalize the concept of ADDs to multi-valued variables (see for instance, Figure 1(c) where Y is ternary). The artificial increase in variables is then avoided, reducing therefore both graph compactness and computation time.

We also propose to add a default arc mechanism : wherever a majority of arcs are pointing to the same node, those arcs are replaced by a default arc. This had the advantage of simplifying the graph and easing the computations. For instance, in the MDD of Figure 4(a), four default arcs (dashed lines) replace fourteen arcs.

As already evoked, the second drawback of ADDs (and of any reduced and ordered function graphs, i.e. MDDs) is that their compactness strongly depends on the inherent variable order. However, the current algorithm for operations on ADDs impose that they share the same order. As a consequence, that common order may not be the optimal one for
each of them or for the result of the operation. Taken that an operation between two ADDs $D_{1}$ and $D_{2}$ is in $\mathcal{O}\left(\left|D_{1}\right| \cdot\left|D_{2}\right|\right)$, this common order tends to arbitrarly increase the computation complexity.

In the next section, we describe a new algorithm of combination on MDDs that will not impose such a common order.

## Operations on Decision Diagrams

Let $G_{1}, G_{2}$ and $G$ be three reduced and ordered function graphs (BDDs, ADDs or MDDs) such that $G=G_{1} \odot G_{2}$ ( $\odot$ being either addition, or multiplication or maximization). We're looking for an algorithm that build $G$ from $G_{1}$ and $\mathrm{G}_{2}$, without imposing that orders $\succ_{1}$ from $\mathrm{G}_{1}$ and $\succ_{2}$ from $\mathrm{G}_{2}$ are the same. First, we will shortly present the state-of-the-art algorithm with common order (for further details, please refer to (Bryant 1986)).

## Operations with Common Order Constraint

Let $\succ$ be the common order imposed on $\mathrm{G}_{1}, \mathrm{G}_{2}$ and G for the operation. The algorithm of combination relies on simultaneous depth-first explorations of both $\mathrm{G}_{1}$ and $\mathrm{G}_{2}$. Each step of the algorithm is a recursive call to a same function. This function takes two node $n_{1} \in G_{1}$ and $n_{2} \in G_{2}$, starting from the root of both graphs. First, it determines how they should be explored :

1. if $n_{1}$ and $n_{2}$ are both terminal, then $n_{1} . v a l \odot n_{2} . v a l$ is computed,
2. if only one node is non-terminal, then exploration is only done on this node,
3. if both nodes are non-terminal then
a. if $n_{1} \cdot v a r=n_{2} \cdot v a r$, the exploration is done simultaneously on both nodes,
b. if $n_{1} \cdot v a r \succ n_{2} \cdot v a r$ (resp. $n_{2} \cdot v a r \succ n_{1} . v a r$ ), then the exploration is done only on $n_{1}$ (resp. $n_{2}$ ).
Exploration consists in calling the function again on each child of the visited node $n$ (one for each value of $n$.var, $n$ being either $n_{1}$, or $n_{2}$ ). The called node from the other diagram remains unchanged, unless it is simultaneously explored. In which case, it is its child selected by the current value of n.var that is called upon.

When exploration on every children is over, or a value has been computed, a node $n_{G}$ is inserted in $G$. If a value was computed, it is bound to that terminal node. Otherwise, variable associated to $n$ is bound to $\mathfrak{n}_{G}$. The children of $n_{G}$ are the resulting nodes from the explorations on $\mathfrak{n}$ children. This procedure ensures that every variables will be in the correct order in G.

Due to the assumption of a common order, the algorithm is quite simple. Its complexity is clearly in $\mathcal{O}\left(\left|\mathrm{G}_{1}\right| \cdot\left|\mathrm{G}_{2}\right|\right)$ when both $G_{1}$ and $G_{2}$ are trees. For BDDs, ADDs and MDDs, a pruning mechanism is needed in order to keep this complexity (see below for more details).

## Building an Order for Resulting Decision Diagram

When removing the common order constraint, we have to deal with different orders: $\succ_{1}$ for $G_{1}, \succ_{2}$ for $G_{2}$ and $\succ_{G}$
for $G=G_{1} \odot G_{2}$. The question about how to build the order $\succ_{\mathrm{G}}$ immediately raises.

With the objective of still performing a deep-first recursive and simultaneous exploration on $\mathrm{G}_{1}$ and $\mathrm{G}_{2}$, one has to analyze a new case at each step: let $n_{1} \in G_{1}$ and $n_{2} \in G_{2}$ be the considered nodes at any step. It may now happen that $n_{1} . v a r \succ_{1} n_{2} . v a r$ and $n_{2} . v a r \succ_{2} n_{1} . v a r$. In that case, the questions are which variable will precedes the other in $\succ_{\mathrm{G}}$ and in which order do we perform exploration on $\mathrm{n}_{1}$ and $n_{2}$.

## Definition (Retrograde Variable).

Let $\succ_{1}$ and $\succ_{2}$ be two orders on a set of variables X. A variable $X_{r} \in X$ is said to be retrograde in $\succ_{2}$ w.r.t. $\succ_{1}$ if $\exists$ $X_{p} \in X$ s.t. $X_{r} \succ_{1} X_{p}$ and $X_{p} \succ_{2} X_{r}$.
Corollary. $X_{p}$ is retrograde in $\succ_{1}$ w.r.t $\succ_{2}$ because of $X_{r}$ (at least).

We note the set of retrograde variables:
$\Re_{1,2}=\left\{X_{i} \in X, X_{i}\right.$ retrograde in $\succ_{2}$ w.r.t. $\left.\succ_{1}\right\}$.
$\mathrm{D}_{\mathfrak{R}_{1,2}}$ is the domain of that set. The size of that domain is $\left|D_{\mathfrak{R}_{1,2}}\right|=\prod_{X_{r} \in \mathfrak{R}_{1,2}}\left|X_{r}\right|$.

Note that in general $\Re_{1,2} \neq \mathfrak{R}_{2,1}$ and $\left|D_{\Re_{1,2}}\right| \neq\left|D_{\Re_{2,1}}\right|$. Whenever $\Re_{1,2} \neq \emptyset$, our algorithm will have to select for $\succ_{G}$ to be compatible with $\succ_{1}$ or $\succ_{2}$. We propose to arbitrarily privilege $\succ_{1}$ (see below for a discussion about that choice).
$\succ_{\mathrm{G}}$ will then have the following properties: (i) $\succ_{\mathrm{G}}$ extends $\succ_{1}$ and (ii) $\succ_{\mathrm{G}}$ extends $\succ_{2} \mathfrak{R}_{1,2}$. These two properties are sufficient to build $\succ_{\mathrm{G}}$ from $\succ_{1}$ and $\succ_{2}$.

## Exploration and construction

Now that G is ordered by $\succ_{\mathrm{G}}$, its construction upon exploration of $G_{1}$ and $G_{2}$ can be performed. As for the algorithm with common order constraint, exploration will be recursive and simultaneous on both decision diagrams.

Since $\succ_{\mathrm{G}}$ extends $\succ_{1}$, the set of retrograde variables that may be encountered is $\mathfrak{R}_{1,2}$.

By definition, for all $X_{r} \in \Re_{1,2}$, there exists at least one variable $X_{p} \in X$ which verifies that $X_{r} \succ_{G} X_{p}$ but $X_{p} \succ_{2}$ $X_{r}$. Note that occurrences of $X_{p}$ preceeding $X_{r}$ only happen in $G_{2}$.
Lemma. During any recursive algorithm to build G, any exploration on variable $X_{r}$ must have begun before any exploration on variable $\mathrm{X}_{\mathrm{p}}$.

Proof. Exploration is performed by recursive calls and a resulting node is created at the end of each call. That node has precedence over all the nodes recursively created during that call.

Then, in order to verify that $X_{r} \succ_{G} X_{p}$, any needed recursive call on $X_{r}$ must have begun before any recursive call on nodes bound to $X_{p}$ starts.

This lemma implies that our algorithm will have to possibly anticipate an exploration on $X_{r}$ whenever an exploration on $X_{p}$ is required in $G_{2}$. This has a consequence in term of complexity that will be analyzed below.

An anticipated exploration of $X_{r}$ on a node $n_{2}$ consists in performing a normal exploration on an artificially inserted node $n_{r}$ such that $n_{r} . v a r=X_{r}$ and all its children are $n_{2}$.

The end of this section technically characterizes the specific situation where this anticipation is required. Table 1 presents the core function for the exploration and the different cases it has to deal with.

Let $\left(n_{1}, n_{2}\right)$ be the visited nodes at current step of our algorithm. If subgraph $\left(n_{2}\right)$ contains $X_{r}{ }^{1}$, the algorithm possibly have to anticipate an exploration on $X_{r}$. On the contrary, if subgraph $\left(n_{2}\right)$ does not contain $X_{r}$, the algorithm can normally perform the exploration on $n_{2}$.

An exploration on $X_{p}$ in $G_{2}$ is required if and only if $n_{2} \cdot v a r=X_{p}$ and $n_{2} \cdot v a r \succ_{G} n_{1} . v a r$. Before that, exploration goes on normally. In particular, $X_{r}$ can be crossed in $G_{1}$. It is then normally explored regardless its presence in $\mathrm{G}_{2}$.

Assume that exploration on $X_{p}$ is required and that $X_{r}$ possibly have to be anticipated. Two cases can occur :
Case (a). If $\mathrm{X}_{\mathrm{r}} \in$ current explored path on $\mathrm{G}_{1}$ then $\mathrm{X}_{\mathrm{r}}$ has already been instantiated. Hence no anticipated exploration on $\mathrm{X}_{\mathrm{r}}$ is needed.
Case (b). If $\mathrm{X}_{\mathrm{r}} \notin$ current explored path on $\mathrm{G}_{1}$ then $\mathrm{X}_{\mathrm{r}}$ has not yet been explored. Hence the anticipated exploration of $X_{r}$ is needed before exploring $X_{p}$ on $G_{2}$.


Figure 2: In case (a), no anticipated exploration is needed: $X_{r}$ is already being explored on $r_{1}$. In case (b), anticipated exploration is needed: $X_{r}$ has not been crossed. Current visited nodes are $\left(n_{1}, n_{2}\right) . r_{1} . v a r=r_{2} . v a r=X_{r}$.

Note that both cases imply that during the explorations in $G_{2}$, a node bound to $X_{r}$ will be crossed. The algorithm will then immediately skips onto the child selected by the current value of $X_{r}$.

## Pruning and complexity

In the algorithm with common order constraint, several explorations of a same pair of subgraphs always lead to the same resulting structure. Paths that brought to their root do not alter the result. As a consequence, once a pair of nodes has been visited, the resulting node is stored in a table along with this pair as a key. Pruning consists then in looking for the pair of nodes in the table, and taking the result. This guarantees a complexity in $\mathcal{O}\left(\left|\mathrm{G}_{1}\right| \cdot\left|\mathrm{G}_{2}\right|\right)$, since every pair of nodes are visited only once.

In the new algorithm, pruning can't be done so easily. Suppose that for the current nodes $\left(n_{1}, n_{2}\right), n_{2} \cdot v a r=X_{p}$ and $X_{r} \in \operatorname{subgraph}\left(n_{2}\right)$. As evoked just above, when a node bound to $X_{r}$ is encountered in $G_{2}$, exploration will automatically jump onto the child selected by the current value

[^1]```
procedure Explore ( \(\mathfrak{n}_{1}, \mathfrak{n}_{2}\), FixedVars) :
    CASE: \(\mathfrak{n}_{1}\).isTerminal and \(n_{2}\).isTerminal
        return Terminal \(\left(\mathrm{n}_{1} . v a l \odot \mathrm{n}_{2} . \mathrm{val}\right)\)
    CASE: \(\mathfrak{n}_{2}\) non terminal and \(\exists \mathfrak{n}_{r}, \mathfrak{n}_{r} \in \mathfrak{n}_{2}\). Descendants
                and \(\mathfrak{n}_{r}\).var \(\in \mathfrak{R}_{1,2}\)
        \(\forall\) modality \(\mathrm{m} \in \mathrm{D}_{\mathrm{n}_{\mathrm{r}}, v a r}\) :
        \(n_{m}=\operatorname{Explore}\left(n_{1}, n_{2}\right.\), FixedVars \(\left.\cup\left\{n_{r} . v a r=m\right\}\right)\);
        RETURN NonTerminal ( \(\mathfrak{n}_{r} . v a r\), children \(=\left\{n_{m}\right\}_{m \in D_{n_{r} . v a r}}\) );
    CASE: \(n_{1} . v a r \succ_{G} n_{2} . v a r\) or \(n_{2}\).isTerminal
        \(\forall\) modality \(\mathrm{m} \in \mathrm{D}_{\mathrm{n}_{1}, v a r}\) :
        \(n_{m}=\operatorname{Explore}\left(n_{1} . \operatorname{child}(m), n_{2}\right.\), FixedVars \(\cup\left\{n_{r} \cdot v a r=m\right\}\) );
        return NonTerminal \(\left(n_{1} . v a r\right.\), children \(\left.=\left\{n_{m}\right\}_{m \in D_{n_{1} . v a r}}\right)\);
    CASE: \(n_{2} . v a r \succ_{1} \odot_{2} n_{1} . v a r\) or \(n_{1}\).isTerminal
        If \(\exists \mathrm{m}, \mathfrak{n}_{2} . v\) ar \(=\mathrm{m} \in\) FixedVars
        return Explore( \(\mathrm{n}_{1}, \mathrm{n}_{2}\).child( m ), FixedVars);
        Else
        \(\forall\) modality \(\mathrm{m} \in \mathrm{D}_{\mathrm{n}_{2}, v a r}\) :
            \(\mathfrak{n}_{\mathfrak{m}}=\operatorname{Explore}\left(\mathfrak{n}_{1}, \mathfrak{n}_{2} . \operatorname{child}(\mathfrak{m})\right.\), FixedVars \() ;\)
        return NonTerminal( \(n_{2} . v a r\), children \(=\left\{n_{m}\right\}_{\mathrm{m} \in \mathrm{D}_{\mathrm{n}_{2} . v a r}}\) );
    CASE: \(\mathfrak{n}_{1} \cdot v a r=n_{2} \cdot v a r\)
        \(\forall\) modality \(\mathrm{m} \in \mathrm{D}_{\mathrm{n}_{1}, \text { var }}\) :
            \(n_{m}=\operatorname{Explore}\left(n_{1} . \operatorname{child}(m), n_{2} . c h i l d(m)\right.\), FixedVars \() ;\)
        return \(\left.\operatorname{NonTerminal(n_{1}.var,~children~}=\left\{n_{m}\right\}_{\mathrm{m}} \in \mathrm{D}_{\mathrm{n}_{1} . v a r}\right)\);
end procedure
```

Table 1: Operation between MDD without common order.
of $X_{r}$ : only a part of subgraph $\left(n_{2}\right)$ is then explored. The consequence is that subgraph $\left(n_{2}\right)$ has to be re-explored for each value of $X_{r}$.

Unnecessary explorations can still be pruned : once an exploration is done for a value of the retrograde variable, there's no need to repeat this exploration. So the key used to know if a subgraph has already been visited simply has to be extended. It needs to indicate which value the retrograde variable had when the exploration was performed.

The multiple explorations due to retrograde variables affect the complexity of the algorithm. The increase in complexity is by the size of the domain of the retrograde variables $D_{\Re_{1,2}}$. Then complexity is now in $\mathcal{O}\left(\left|\mathrm{G}_{1}\right| \cdot\left|\mathrm{G}_{2}\right|\right.$. $\left.\left|D_{\Re_{1,2}}\right|\right)$.

But this worst case complexity has to be pondered. Firstly our algorithm only re-explore subgraphs of $\mathrm{G}_{2}$ when needed. Yet the given complexity is determined as if all the re-explorations concerned the whole graph. Unfortunately, a more accurate upper bound is difficult to obtain because it would demand a topological analysis of the graphs. Secondly, $G_{1}$ and $G_{2}$ have now their own order. Then the size of $G_{1}$ and $G_{2}$ can be smaller in this complexity than in the complexity of the algorithm with common order. Thirdly, $\succ_{\mathrm{G}}$ is compatible either with $\succ_{1}$ or with $\succ_{2}$ (we arbitrarily chose $\succ_{1}$ for presentation purposes). As a consequence, we have to deal either with $D_{\mathfrak{R}_{1,2}}$ or with $D_{\mathfrak{R}_{2,1}}$. But since $\left|D_{\Re_{1,2}}\right| \neq\left|D_{\Re_{2,1}}\right|$, another tradeoff can be found here.

This discussion is confirmed in the experiments described in the next section.

## Experiments and Observation

We have implemented MDD and SPUmDD (the implementation of our algorithm) using the aGrUM C++ library devel-
oped at LIP6 ${ }^{2}$. Since the standard implementation of SPUDD uses an highly optimized library for ADD (but not for MDD), we have coded our own version of SPUDD in order to compare the algorithms on a time basis.

Meanwhile, size of computed diagrams is the most interesting measure. Indeed, as seen before, the complexity of operations depends on the size of the diagrams. The size of the computed value function is particularly relevant: on each iteration, value function is used to compute various other decision diagrams that are themselves aggregated using Equation 2. Therefore, the size of value function is a good indicator of the efficiency of representations.

No reordering is performed during execution in both algorithms. For SPUDD, the common order on variables is fixed and specified at the beginning. For SPUmDD, each MDD has its own order. That order is chosen so that the overall complexity to obtained the MDD is minimal (see previous section).

Table 2 shows result of value iteration using SPUDD and SPUMDD on various MDPs. State space size gives the total number of states (including the ones induced by binarization of multi-valued variables). Internal nodes gives the number of non-terminal nodes inside the computed value functions at last iteration (the number of terminal nodes is the same for both methods). And time (in second) is the average time to reach stopping criterion ( $\varepsilon$ was set to $10^{-5}$ ) over 30 runs.

We examine the efficiency of our algorithm on two standard problems: coffee robot and factory. The interest of coffee robot planning problem is that it contains only binary variables. It allows to see if SPUmDD remains efficient on such cases. Results show that SPUmDD got same behavior than SPUDD. Yet it is slightly slower, showing that on purely and small binary problems ADDs are sufficient.

In factory, the interest resides in the mixed of binary and ternary variables. The conversion of ternary variables in binary variables generates an increase in the number of variables as much as an increase in state space size. Results shows clearly that advantage can be taken of by SPUmDD.

Notice that factoryl and factory 2 only differ on one variable that is not relevant for value function (it has no incidence on other variable). Both factories got eventually the same structure, showing clearly that MDD can eliminate non relevant variables as ADDs does.


Figure 3: Maze examples: blocked cases are in dark gray. Impossible states generated by binarization in light gray.

To examine the behavior on problems with multi-valued variables, two mazes have been created. First maze (Figure 3 ) has 30 cases, 8 of them being blocked. It only requires

[^2]|  | SPUDD |  |  | SPUmDD |  |  | SPUmDD in comparison with SPUDD |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | States Space | Internal <br> Nodes | $\begin{aligned} & \text { Time (in } \\ & \text { s) } \end{aligned}$ | States Space | Internal <br> Nodes | $\begin{aligned} & \text { Time (in } \\ & \text { s) } \end{aligned}$ | States <br> Space | Internal <br> Nodes | Time |
| Coffee Robot | 64 | 21 | 2.06 | 64 | 21 | 2.19 | 100.0\% | 100.0\% | 106.3\% |
| Factory | 131072 | 1269 | 1751.45 | 55296 | 473 | 661.93 | 42.2\% | 37.3\% | 37.8\% |
| Factory 0 | 524288 | 2131 | 2447.92 | 221184 | 733 | 1596.51 | 42.2\% | 34.4\% | 65.2\% |
| Factory 1 | 2097132 | 2889 | 8796.10 | 884736 | 1283 | 5181.56 | 42.2\% | 44.4\% | 58.9\% |
| Factory 2 | 4194304 | 2889 | 8916.83 | 1769472 | 1283 | 5986.89 | 42.2\% | 44.4\% | 67.1\% |
| Factory 3 | 33554432 | 3371 | 27240.00 | 10616832 | 2001 | 17439.90 | 31.6\% | 59.3\% | 64.0\% |
| Maze 5x6 | 64 | 34 | 2.51 | 30 | 6 | 0.86 | 46.9\% | 17.7\% | 34.3\% |
| Maze 8x8 | 64 | 58 | 3.32 | 64 | 9 | 1.05 | 100.0\% | 15.5\% | 31.6\% |

Table 2: Results using SPUDD and SPUmDD. The last columns illustrate the improvements in space and time using SPUmDD.


Figure 4: Maze Optimal Policy (a) with SPUmDD (dashed line stands for default arc) and (b) with SPUDD (dashed line from node $X$ stands for $\bar{x}$ ).
two multi-valued variables ( X and Y ) of 5 and 6 modalities to represent its 30 possible states. However, its translation in binary variables demands 3 variables on each axes. Those variables generate a grid of 64 states where 34 are impossible. Second maze is a 8 by 8 , and thus generates no impossible states on translation into a binary problem.

Here again, SPUmDD shows itself better than SPUDD, gaining both on time and size representation. Figure 4 shows the simplification of the policy obtained for the first maze with SPUmDD compared to SPUDD.

## Conclusion

In this paper, we propose new improvements in the factored framework for decision theoretic planning.

First, we investigate the use of multi-valued decision diagrams which avoids the transformations of multi-valued variables into sets of binary variables. Such transformations increase the diagram size and constrain the algorithm to deal with states that are impossible for the modeled system. Our extension to multi-valued decision diagrams allows computations on simpler structures.

We then propose a new algorithm for operations on such decision diagrams that eliminates the second main drawback of the state-of-the-art. Our algorithm does not force the diagrams to share a common order of variables. In spite of an increase in worst case complexity, SPUmDD shows a significant gain in time and in size when compared to SPUDD.

With these improvements, the optimal policy search operates directly on the inputs of the problem (no binarization, no reordering) and produces more readable solutions. This is an opening for future works on incremental algorithms. Particularly, we will look further in dynamical reordering for the policies and incremental learning of the model.

## References

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.
Boutilier, C.; Dean, T.; and Hanks, S. 1999. Decision-theoretic planning: Structural assumptions and computational leverage. JAIR 11:1-94.
Boutilier, C.; Dearden, R.; and Goldszmidt, M. 1995. Exploiting structure in policy construction. In IJCAI-95, pp. 11041111.
Bryant, R. E. 1986. Graph-based algorithms for boolean function manipulation. IEEE Transactions on Computers 35:677-691.
Dean, T., and Kanazawa, K. 1989. A model for reasoning about persistence and causation. Comput. Intell. 5(3):142-150.
Friedman, S. J., and Supowit, K. J. 1990. Finding the optimal variable ordering for binary decision diagrams. IEEE Trans. Comput. 39(5):710-713.
Guestrin, C.; Parr, R.; and Venkataraman, S. 2003. Efficient solution algorithms for factored mdps. Journal of Artificial Intelligence Research 19:399-468.
Guestrin, C. 2002. Distributed planning in hierarchical factored mdps. In In Proceedings of the Eighteenth Conference on Uncertainty in Artificial Intelligence, 197-206.
Hoey, J.; St-aubin, R.; Hu, A.; and Boutilier, C. 1999. Spudd: Stochastic planning using decision diagrams. In In Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence, 279-288. Morgan Kaufmann.
Puterman, M. L. 1994. Markov decision processes: Discrete stochastic dynamic programming.
Srinivasan, A.; Kam, T.; Malik, S.; and Brayton, R. K. 1990. Algorithms for discrete function manipulation. In ICCAD'90, 92-95.
St-aubin, R.; Hoey, J.; and Boutilier, C. 2000. Apricodd: Approximate policy construction using decision diagrams. In In Proceedings of Conference on Neural Information Processing Systems, 1089-1095.


[^0]:    Copyright © $C^{20} 13$, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[^1]:    ${ }^{1}$ More precisely, $X_{r} \in \operatorname{support}\left(\operatorname{subgraph}\left(n_{2}\right)\right)$

[^2]:    ${ }^{2}$ Laboratoire d'Informatique de Paris VI

