

Semiring Labelled Decision Diagrams, Revisited: Canonicity and Spatial Efficiency Issues*

Hélène Fargier¹, Pierre Marquis², Nicolas Schmidt^{1,2}

¹IRIT-CNRS, Université de Toulouse, France

²CRIL-CNRS, Université d'Artois, Lens, France

fargier@irit.fr {marquis,schmidt}@cril.fr

Abstract

Existing languages in the valued decision diagrams (VDDs) family, including ADD, AADD, and those of the SLDD family, prove to be valuable target languages for compiling multivariate functions. However, their efficiency is directly related to the size of the compiled formulae. In practice, the existence of canonical forms may have a major impact on the size of the compiled VDDs. While efficient normalization procedures have been pointed out for ADD and AADD the canonicity issue for SLDD formulae has not been addressed so far. In this paper, the SLDD family is revisited. We modify the algebraic requirements imposed on the valuation structure so as to ensure tractable conditioning, optimization and normalization for some languages of the revisited SLDD family. We show that AADD is captured by this family. Finally, we compare the spatial efficiency of some languages of this family, from both the theoretical side and the practical side.

1 Introduction

In configuration problems of combinatorial objects (like cars), there are two key tasks for which short, guaranteed response times are expected: conditioning (propagating the end-user's choices: version, engine, various options ...) and optimization (maintaining the minimum cost of a feasible car satisfying the user's requirements). When the set of feasible objects and the corresponding cost functions are represented as valued CSPs (VCSPs for short see [Schiex *et al.*, 1995]), the optimization task is NP-hard in the general case, so short response times cannot be ensured.

Valued decision diagrams (VDDs) from the families ADD [Bahar *et al.*, 1993], EVBDD [Lai and Sastry, 1992; Lai *et al.*, 1996; Amilhastre *et al.*, 2002] and their generalization SLDD [Wilson, 2005], and AADD [Tafertshofer and Pedram, 1997; Sanner and McAllester, 2005] do not have such a drawback and appear as interesting representation languages for compiling mappings associating valuations with assignments of discrete variables (including utility functions and probability

distributions). Indeed, those languages offer tractable conditioning and tractable optimization (under some conditions in the SLDD case). However, the efficiency of these operations is directly related to the size of the compiled formulae. Following [Darwiche and Marquis, 2002], the choice of the target representation language for the compiled forms must be guided by its succinctness. From the practical side, normalization (and all the more canonicity) are also important: subformulae in normalized form can be more efficiently recognized and the canonicity of the compiled formulae facilitates the search for compiled forms of optimal size (see the discussion about it in [Darwiche, 2011]). Indeed, the ability to ensure a unique form for subformulae prevents them from being represented twice or more.

In this paper, the SLDD family [Wilson, 2005] is revisited, focusing on the canonicity and the spatial efficiency issues. We extend the SLDD setting by relaxing some algebraic requirements on the valuation structure. This extension allows us to capture the AADD language as an element of e-SLDD, the revisited SLDD family. We point out a normalization procedure which extends the AADD's one to some representation languages of e-SLDD. We also provide a number of succinctness results relating some elements of e-SLDD with ADD and AADD. We finally report some experimental results where we compiled some instances of an industrial configuration problem into each of those languages, thus comparing their spatial efficiency from the practical side.

The rest of the paper is organized as follows. Section 2 gives some formal preliminaries on valued decision diagrams. Section 3 presents the e-SLDD family and describes our normalization procedure. In Section 4, succinctness results concerning ADD, some elements of the e-SLDD family, and AADD are pointed out. Section 5 gives and discusses our empirical results about the spatial efficiency of those languages. Finally, Section 6 concludes the paper.¹

2 Valued Decision Diagrams

Given a finite set $X = \{x_1, \dots, x_n\}$ of variables where each variable $x \in X$ ranges over a finite domain D_x , we are interested in representing mappings associating an element from a valuation set E with assignments $\vec{x} = \{(x_i, d_i) \mid d_i \in$

*This work is partially supported by the project BR4CP ANR-11-BS02-008 of the French National Agency for Research.

¹A full-proof version of the paper is available at <ftp://ftp.irit.fr/IRIT/ADRIA/ijcai13FMS.pdf>

$D_{x_i}, i = 1, \dots, n\}$ ($\vec{X}$ will denote the set of all assignments over X). E is the carrier of a valuation structure $\mathcal{E}$, which can be more or less sophisticated from an algebraic point of view. A representation language given X w.r.t. a valuation structure $\mathcal{E}$ is mainly a set of data structures. The targeted mapping is called the *semantics* of the data structure and the data structure is a *representation* of the mapping:

Definition 1 (representation language) (inspired from [Gogic et al., 1995]) Given a valuation structure $\mathcal{E}$, a representation language $\mathcal{L}$ over X w.r.t. $\mathcal{E}$ is a 4-tuple $\langle C_{\mathcal{L}}, Var_{\mathcal{L}}, I_{\mathcal{L}}, s_{\mathcal{L}} \rangle$ where $C_{\mathcal{L}}$ is a set of data structures α (also referred to as $C_{\mathcal{L}}$ formulae), $Var_{\mathcal{L}} : C_{\mathcal{L}} \rightarrow 2^X$ is a scope function associating with each $C_{\mathcal{L}}$ formula the subset of X it depends on, $I_{\mathcal{L}}$ is an interpretation function associating with each $C_{\mathcal{L}}$ formula α a mapping $I_{\mathcal{L}}(\alpha)$ from the set of all assignments of $Var_{\mathcal{L}}(\alpha)$ to E , and $s_{\mathcal{L}}$ is a size function from $C_{\mathcal{L}}$ to $\mathbb{N}$ providing the size of any $C_{\mathcal{L}}$ formula.

Different formulae can share the same semantics:

Definition 2 (equivalent formulae) Let $\mathcal{L}_1$ (resp. $\mathcal{L}_2$) be a representation language over X w.r.t. $\mathcal{E}_1$ (resp. $\mathcal{E}_2$) where $E_1 = E_2$. $\alpha \in \mathcal{L}_1$ is equivalent to $\beta \in \mathcal{L}_2$ iff $Var_{\mathcal{L}_1}(\alpha) = Var_{\mathcal{L}_2}(\beta)$ and $I_{\mathcal{L}_1}(\alpha) = I_{\mathcal{L}_2}(\beta)$.

In this paper, we are specifically interested in data structures of the form of *valued decision diagrams*:

Definition 3 (valued decision diagram) A valued decision diagram (VDD) over X w.r.t. $\mathcal{E}$ is a finite DAG α with a single root, s.t. every internal node N is labelled with a variable $x \in X$ and if $D_x = \{d_1, \dots, d_k\}$, then N has k outgoing arcs $a_1, \dots, a_k$, so that the arc a_i of α is valued by $v(a_i) = d_i$. We note $out(N)$ (resp. $in(N)$) the arcs outgoing from (resp. incoming to N). Nodes and arcs can also be labelled by elements of E : if N (resp. a_i) is node (resp. an arc) of α , then $\phi(N)$ (resp. $\phi(a_i)$) denotes the label of N (resp. a_i). Finally, each VDD α is a read-once formula, i.e., for each path from the root of α to a sink, every variable $x \in X$ occurs at most once as a node label.

When ordered VDDs are considered, a total ordering over X is chosen and for each path from the root of α to a sink, the associated sequence of internal node labels is required to be compatible w.r.t. this variable ordering.

The key problems we focus on are the *conditioning problem* (given a $C_{\mathcal{L}}$ formula α over X w.r.t. $\mathcal{E}$ and an assignment $\vec{y} \in \vec{Y}$ where $Y \subseteq X$, compute a $C_{\mathcal{L}}$ formula representing the restriction of $I_{\mathcal{L}}(\alpha)$ by $\vec{y}$) and the *optimization problem* (given a $C_{\mathcal{L}}$ formula α over X w.r.t. $\mathcal{E}$, find an assignment $\vec{x}^* \in \vec{X}$ such that $I_{\mathcal{L}}(\alpha)(\vec{x}^*)$ is not dominated w.r.t. some relation $\succeq$ over E – typically, $\succeq$ is a total order). Conditioning is an easy operation on a VDD α . Mainly, for each $(y, d_i) \in \vec{y}$, just by-pass in α every node N labeled by y by linking directly each of its parents to the child N_i of N such that $v((N, N_i)) = d_i$ (N and all its outgoing arcs are thus removed). However, optimization is often more demanding, depending on the family of VDDs under consideration.

ADD, SLDD, and AADD are representation languages composed of valued decision diagrams. The scope functions Var_{ADD} , Var_{SLDD} , and Var_{AADD} are the same ones and

they return the set of variables $Var(\alpha)$ from X where each $x \in Var(\alpha)$ labels at least one node in α . The size functions s_{ADD} , s_{SLDD} , and s_{AADD} are closely related: the size of a (labelled) decision graph α is the size of the graph (number of nodes plus number of arcs) plus the sizes of the labels in it. The main difference between ADD, SLDD, and AADD lies in the way the decision diagrams are labelled and interpreted.

For ADD, no specific assumption has to be made on the valuation structure $\mathcal{E}$, even if $E = \mathbb{R}$ is often considered:

Definition 4 (ADD) ADD is the 4-tuple $\langle C_{ADD}, Var_{ADD}, I_{ADD}, s_{ADD} \rangle$ where C_{ADD} is the set of ordered VDDs α over X such that sinks S are labelled by elements of E , and the arcs are not labelled; I_{ADD} is defined inductively by: for every assignment $\vec{x}$ over X ,

- if α is a sink node S , labelled by $\phi(S) = e$, then $I_{ADD}(\alpha)(\vec{x}) = e$,
- else the root N of α is labelled by $x \in X$; let $d \in D_x$ such that $(x, d) \in \vec{x}$, $a = (N, M)$ the arc such that $v(a) = d$, and β the ADD formula rooted at node M in α ; we have $I_{ADD}(\alpha)(\vec{x}) = I_{ADD}(\beta)(\vec{x})$.

Optimization is easy on an ADD formula: every path from the root of α to a sink labelled by a non-dominated valuation among those labeling the sinks of α can be read as a (usually partial) variable assignment which can be extended to a (full) optimal assignment.

In the SLDD framework [Wilson, 2005], the valuation structure $\mathcal{E}$ must take the form of a commutative semiring $\langle E, \oplus, \otimes, 0_s, 1_s \rangle$: $\oplus$ and $\otimes$ are associative and commutative mappings from $E \times E$ to E , with identity elements (respectively) 0_s and 1_s , $\otimes$ left and right distributes over $\oplus$, and 0_s is an annihilator for $\otimes$ ($\forall a \in E, a \otimes 0_s = 0_s \otimes a = 0_s$).

Definition 5 (SLDD) Let $\mathcal{E} = \langle E, \oplus, \otimes, 0_s, 1_s \rangle$ be a commutative semiring. SLDD is the 4-tuple $\langle C_{SLDD}, Var_{SLDD}, I_{SLDD}, s_{SLDD} \rangle$ where C_{SLDD} is the set of VDDs α over X with a unique sink S , satisfying $\phi(S) = 1_s$, and such that the arcs are labelled by elements of E , and I_{SLDD} is defined inductively by: for every assignment $\vec{x}$ over X ,

- if α is the sink node S , then $I_{SLDD}(\alpha)(\vec{x}) = 1_s$,
- else the root N of α is labelled by $x \in X$; let $d \in D_x$ such that $(x, d) \in \vec{x}$, $a = (N, M)$ the arc such that $v(a) = d$, and β the SLDD formula rooted at node M in α ; we have $I_{SLDD}(\alpha)(\vec{x}) = \phi(a) \otimes I_{SLDD}(\beta)(\vec{x})$.

SLDD languages are not specifically suited to optimization w.r.t. any relation $\succeq$. Specifically, [Wilson, 2005] considers the following *addition-is-max-or-min* assumption about $\oplus$:

$$\forall a, b \in E, a \oplus b \in \{a, b\}.$$

Under this assumption, $\oplus$ is idempotent and the relation $\trianglelefteq$ defined by $a \trianglelefteq b$ iff $a \oplus b = a$ is total. [Wilson, 2005] shows that, when $\succeq$ coincides with $\trianglelefteq$, computing the valuation of $I_{SLDD}(\alpha)$ maximal w.r.t. $\succeq$ amounts to performing $\oplus$ -variable elimination; this can be achieved in polynomial time under the linear-time computability assumption for $\otimes$ and $\oplus$.

Sanner and Mc Allester's AADD framework [2005] focuses on the valuation set $E = \mathbb{R}^+$ but enables decision graphs into which the arcs are labelled with *pairs* of values from $\mathbb{R}^+$ and considers *two* operators, namely $+$ and $\times$:

Definition 6 (AADD) AADD is the 4-tuple $\langle C_{\text{AADD}}, \text{Var}_{\text{AADD}}, I_{\text{AADD}}, s_{\text{AADD}} \rangle$ where C_{AADD} is the set of ordered VDDs α over X with a unique sink S , satisfying $\phi(S) = 1$, and such that the arcs are labelled by pairs $\langle q, f \rangle$ in $\mathbb{R}^+ \times \mathbb{R}^+$; I_{AADD} is defined inductively by: for every assignment $\vec{x}$ over X ,

- if α is the sink node S , then $I_{\text{AADD}}(\alpha)(\vec{x}) = 1$,
- else the root N of α is labelled by $x \in X$; let $d \in D_x$ such that $(x, d) \in \vec{x}$, $a = (N, M)$ the arc such that $v(a) = d$ and $\phi(a) = \langle q, f \rangle$, and β the AADD formula rooted at node M in α ; we have

$$I_{\text{AADD}}(\alpha)(\vec{x}) = q + (f \times I_{\text{AADD}}(\beta)(\vec{x})).$$

For the normalization purpose, each α is equipped with a pair $\langle q_0, f_0 \rangle$ from $\mathbb{R}^+ \times \mathbb{R}^+$ (the "offset", labeling the root of α); the interpretation function of the resulting "augmented" AADD is given by, for every assignment $\vec{x}$ over X , $I_{\text{AADD}}^{(q_0, f_0)}(\alpha)(\vec{x}) = q_0 + (f_0 \times I_{\text{AADD}}(\alpha)(\vec{x}))$.

Conditioning and optimization are also tractable on AADD formulae (see [Sanner and McAllester, 2005]).

3 Revisiting the SLDD Framework

In the following, we extend the SLDD framework in two directions: we relax the algebraic requirements imposed on the valuation structure and we point out a normalization procedure which extends the AADD's one to some representation languages of e-SLDD, the extended SLDD family.

A first useful observation is that, in the SLDD framework, $\oplus$ is not used for defining the SLDD language. Actually, different $\oplus$ may be considered over the same formula (e.g., when SLDD is used to compile a Bayesian net, $\oplus = +$ can be used for marginalization purposes and $\oplus = \max$ can be considered when a most probable explanation is looked for). This explains why the requirements imposed on $\oplus$ in the SLDD setting can be relaxed. Let us recall that a *monoid* is a triple $\langle E, \otimes, 1_s \rangle$ where E is a set endowed with an associative binary operator $\otimes$ with identity element 1_s :

Definition 7 (e-SLDD) For any monoid $\mathcal{E} = \langle E, \otimes, 1_s \rangle$, e-SLDD is the 4-tuple $\langle C_{\text{e-SLDD}}, \text{Var}_{\text{e-SLDD}}, I_{\text{e-SLDD}}, s_{\text{e-SLDD}} \rangle$, defined as the SLDD one, except that, for the normalization purpose, each e-SLDD formula α is associated with a value $q_0 \in E$ (the "offset" of the data structure, labeling its root); the interpretation function $I_{\text{e-SLDD}}^{q_0}$ of the extended SLDD setting is given by, for every assignment $\vec{x}$ over X ,

$$I_{\text{e-SLDD}}^{q_0}(\alpha)(\vec{x}) = q_0 \otimes I_{\text{e-SLDD}}(\alpha)(\vec{x}).$$

Several choices for $\otimes$ remain usually possible when E is fixed; we sometimes make the notation of the language more precise (but not too heavy) and write e-SLDD $_{\otimes}$ instead of e-SLDD.

Obviously, the e-SLDD framework captures the SLDD one: when $\langle E, \oplus, \otimes, 0_s, 1_s \rangle$ is a commutative semiring, then $\langle E, \otimes, 1_s \rangle$ is a monoid, and every SLDD formula can be interpreted as an e-SLDD one (choose $q_0 = 1_s$). Interestingly, the e-SLDD framework also captures the AADD language:

Proposition 1 Let $E = \mathbb{R}^+ \times \mathbb{R}^+$, $1_s = \langle 0, 1 \rangle$ and $\otimes = \star$ be defined by $\forall b, b', c, c' \in E, \langle b, c \rangle \star \langle b', c' \rangle = \langle b + c \times b', c \times c' \rangle$. $\mathcal{E} = \langle E, \otimes, 1_s \rangle$ is a monoid.

The correspondence between AADD and e-SLDD $_{\star}$ is made precise by the following proposition:

Proposition 2 Let α be an AADD formula, also viewed as an e-SLDD $_{\star}$ formula. We have: $\forall \vec{x} \in \vec{X}$, if $I_{\text{AADD}}(\alpha)(\vec{x}) = a$ and $I_{\text{e-SLDD}_{\star}}(\alpha)(\vec{x}) = \langle b, c \rangle$, then $a = b + c$.

Observe that $\star$ is *not* commutative: the relaxation of the commutativity assumption is necessary to capture the AADD framework within the e-SLDD family.

Let us now switch to the normalization/canonicity issues for e-SLDD. When compiling a formula, normalization (and all the more canonicity) are important for computational reasons: in practice, subformulae in reduced, normalized form which have been already encountered and cached can be more efficiently recognized. Besides, when the canonicity property is ensured, the recognition issue boils down to a simple equality test. Thus, canonicity is more demanding and is achieved for ordered VDDs, only: reduced ADD formulae and normalized and reduced AADD formulae (which are ordered VDDs) offer the canonicity property. Contrastingly, though some simplification rules have been considered in [Wilson, 2005], no normalization procedure and canonicity conditions for SLDD have been pointed out so far.

The idea at work for normalizing AADD formulae is to propagate from the sink to the root of the diagram the *minimum* valuations of the outgoing arcs. In our more general framework, minimality is characterized by an idempotent, commutative and associative operator $\oplus$, which induces the binary relation $\succeq$ over E given by:

$$\forall a, b \in E, a \succeq b \text{ iff } a \oplus b = b.$$

The fact that $\oplus$ is associative (resp. commutative, idempotent) implies that the induced relation $\succeq$ is transitive (resp. antisymmetric, reflexive), hence an order over E .

Definition 8 ($\oplus$ -normalisation, $\oplus$ -reduction) An e-SLDD formula α is $\oplus$ -normalized iff for any node N of α , $\bigoplus_{a \in \text{out}(N)} \phi(a) = 1_s$ (by convention, we define $\bigoplus_{a \in \emptyset} \phi(a) = 1_s$). An e-SLDD formula α is $\oplus$ -reduced iff it is $\oplus$ -normalized, and reduced, i.e., it does not contain any (distinct) isomorphic nodes² and any redundant nodes.³

To allow to propagate valuations in VCSPs, where $\succeq$ is a total order, $\otimes$ is commutative and $\otimes$ is monotonic w.r.t. $\succeq$ (i.e., $\otimes$ is distributive over $\oplus$), [Cooper and Schiex, 2004] assume a "fairness" property of $\otimes$ w.r.t. $\oplus$: for any valuations $a, b \in E$ such that $a \oplus b = b$, there exists a unique valuation which is the maximal element w.r.t. $\succeq$ among the $c \in E$ satisfying $b \otimes c = c \otimes b = a$.

Here, we relax these conditions so as to be able to encompass the case of the (possibly partial) relation $\succeq$ induced

² N and M are isomorphic when they are labelled by the same variable and there exists a bijection f from $\text{out}(N)$ to $\text{out}(M)$ such that $\forall a \in \text{out}(N)$, a and $f(a)$ have the same end node and $\phi(a) = \phi(f(a))$.

³ N is redundant when all outgoing arcs a are labelled by the same value $\phi(a)$ and reach the same end node.

by $\oplus$. Let us state that $\otimes$ is *left-distributive* over $\oplus$ iff $\forall a, b, c \in E, c \otimes (a \oplus b) = (c \otimes a) \oplus (c \otimes b)$, and $\otimes$ is *left-fair* w.r.t. $\oplus$ iff $\forall a, b \in E$, if $a \oplus b = b$, then there exists a unique valuation of E , noted $a \otimes^{-1} b$, which is the maximal element w.r.t. $\succeq$ among the $c \in E$ satisfying $b \otimes c = a$.

Definition 9 (extended SLDD condition) A valuation structure $\mathcal{E} = \langle E, \oplus, \otimes, 1_s \rangle$ satisfies the extended SLDD condition iff $\langle E, \otimes, 1_s \rangle$ is a monoid, $\oplus$ is a mapping from $E \times E$ to E , which is associative, commutative, and idempotent, $\otimes$ is left-distributive over $\oplus$ and left-fair w.r.t. $\oplus$.

The extended SLDD condition is close to the commutative semiring assumption for SLDD. However, it requires neither the commutativity of $\otimes$, nor an annihilator for $\otimes$, and left-distributivity of $\otimes$ over $\oplus$ is less demanding than (full) distributivity; on the other hand, the left-fairness condition of $\otimes$ w.r.t. $\oplus$ is imposed. The idempotence of $\oplus$ is also less demanding than the "addition-is-max-or-min" condition.

The valuation considered in the AADD framework satisfies the extended SLDD condition:

Proposition 3 The valuation structure $\mathcal{E} = \langle \mathbb{R}^+ \times \mathbb{R}^+, \oplus, \star, \langle 0, 1 \rangle \rangle$ where $\oplus = \min_\star$ is defined by $\forall b, b', c, c' \in E: \langle b, b' \rangle \min_\star \langle c, c' \rangle = \langle \min(b, c), \max(b + b', c + c') - \min(b, c) \rangle$, satisfies the extended SLDD condition.

In the AADD case, $E = \mathbb{R}^+ \times \mathbb{R}^+$ is not totally ordered by $\succeq$ (for instance, none of $\langle 0, 2 \rangle \succeq \langle 1, 2 \rangle$ and $\langle 1, 2 \rangle \succeq \langle 0, 2 \rangle$ hold since $\langle 0, 2 \rangle \min_\star \langle 1, 2 \rangle = \langle 1, 2 \rangle \min_\star \langle 0, 2 \rangle = \langle 0, 3 \rangle$). When $\langle a, a' \rangle \succeq \langle b, b' \rangle$ holds, we have:

- $\langle a, a' \rangle \star^{-1} \langle b, b' \rangle = \langle 1, 0 \rangle$ if $b' = 0$,
- $\langle a, a' \rangle \star^{-1} \langle b, b' \rangle = \langle \frac{a-b}{b'}, \frac{a'}{b'} \rangle$ if $b' > 0$.

e-SLDD $_\star$ denotes the corresponding e-SLDD language.

Weighted finite automata and edge-valued binary decision diagrams are captured by using $\mathcal{E} = \langle \mathbb{R}^+, \min, +, 0 \rangle$. The following pairs, consisting of a valuation structure – a representation language, can actually be considered:

- $\mathcal{E} = \langle \mathbb{R}^+, \min, +, 0 \rangle$ – e-SLDD $_+$.
- $\mathcal{E} = \langle \mathbb{R}^+, \max, \times, 1 \rangle$ – e-SLDD $_\times$.
- $\mathcal{E} = \langle \mathbb{R}^+ \cup \{+\infty\}, \max, \min, +\infty \rangle$ – e-SLDD $_{\min}$.
- $\mathcal{E} = \langle \mathbb{R}^+, \min, \max, 0 \rangle$ – e-SLDD $_{\max}$.

Proposition 4 The valuation structures $\mathcal{E} = \langle \mathbb{R}^+, \min, +, 0 \rangle$, $\mathcal{E} = \langle \mathbb{R}^+, \max, \times, 1 \rangle$, $\mathcal{E} = \langle \mathbb{R}^+ \cup \{+\infty\}, \max, \min, +\infty \rangle$ and $\mathcal{E} = \langle \mathbb{R}^+, \min, \max, 0 \rangle$ satisfy the extended SLDD condition.

We are now ready to extend the AADD normalization procedure to the e-SLDD language, under the extended SLDD condition. Algorithm 1 is the normalization procedure. This procedure proceeds backwards (i.e., from the sink to the root). Figure 1 gives an e-SLDD $_\star$ formula and the corresponding $\min_\star$ -reduced formula.

Proposition 5 Assume that $\mathcal{E} = \langle E, \oplus, \otimes, 1_s \rangle$ satisfies the extended SLDD condition. If $\oplus$ satisfies the addition-is-max-or-min property then, for any e-SLDD formula α , a $\oplus$ -reduced e-SLDD formula equivalent to α can be computed in polynomial time provided that $\otimes, \otimes^{-1}$ and $\oplus$ can be computed in linear time.


Algorithm 1: normalize(α)

input : an e-SLDD $_\otimes$ formula α , with offset q_0
output: an e-SLDD $_\otimes$ formula which is $\oplus$ -normalized and equivalent to α

```

1 for each node  $N$  of  $\alpha$  in inverse topological ordering do
2    $q_{\min} := \oplus_{a \in \text{out}(N)} \phi(a)$ 
3   for each  $a \in \text{out}(N)$  do
4     if  $\phi(a) == q_{\min}$  then
7        $\phi(a) := 1_s$ 
     else
7        $\phi(a) := \phi(a) \otimes^{-1} q_{\min}$ 
5   for each  $a \in \text{in}(N)$  do
7      $\phi(a) := \phi(a) \otimes q_{\min}$ 
6  $q_0 := q_0 \otimes q_{\min}$ 
7 return  $\alpha$ 

```



(0, 1) (0, 1)

Figure 1: An e-SLDD $_\star$ formula (left) and the corresponding $\min_\star$ -reduced e-SLDD $_\star$ formula (right). x, y and z are Boolean variables. A (resp. plain) edge corresponds to the assignment of the variable labeling its source to 0 (resp. 1).

Clearly, the linear-time computability assumptions are satisfied by the operators $\otimes, \otimes^{-1}$, and $\oplus$ associated with e-SLDD $_+$, e-SLDD $_\times$, e-SLDD $_{\min}$, e-SLDD $_{\max}$. Thus, the formulae from all these languages can be $\oplus$ -reduced in polynomial time.

Interestingly, addition-is-max-or-min is not a necessary condition for ensuring a normalized form; *left-cancellativity* of $\otimes$ ($\forall a, b, c \in E$, if $c \otimes a = c \otimes b$ and c is not an annihilator for $\otimes$, then $a = b$) is also enough:

Proposition 6 Assume that $\mathcal{E} = \langle E, \oplus, \otimes, 1_s \rangle$ satisfies the extended SLDD condition. If $\otimes$ is left-cancellative, then for any e-SLDD formula α , a $\oplus$ -reduced e-SLDD formula equivalent to α it can be computed in polynomial time provided that $\otimes, \otimes^{-1}$ and $\oplus$ can be computed in linear time.

Furthermore, when $\otimes$ is left-cancellative, the canonicity property is ensured for ordered e-SLDD formulae (even if $\succeq$ is not total):

Proposition 7 Assume that $\mathcal{E} = \langle E, \oplus, \otimes, 1_s \rangle$ satisfies the extended SLDD condition. If $\otimes$ is left-cancellative, then two ordered e-SLDD formulae are equivalent iff they have the

same $\oplus$ -reduced form.

Especially, since $+$, $\times$ and $\star$ are left-cancellative, the ordered $e\text{-SLDD}_+$ (resp. $e\text{-SLDD}_\times$, $e\text{-SLDD}_\star$) formulae offer the canonicity property.

Let us finally switch to conditioning and optimization. First, conditioning does not preserve the $\oplus$ -reduction of a formula in the general case, but this is computationally harmless since the $\oplus$ -reduction of a conditioned formula can be done in polynomial time. As to optimization, when $\succeq$ is total, any $\oplus$ -reduced $e\text{-SLDD}$ formula α contains a path the arcs of which are labelled by 1_s . The (usually partial) variable assignment along this path can be extended to a full minimal solution $\vec{x}^*$ w.r.t. $\succeq$, and the offset of α is equal to $I_{e\text{-SLDD}}(\alpha)(\vec{x}^*)$. However, in the general case, the ordering $\succeq$ is not equal to $\succeq$, so the normalization procedure does not help for determining a minimal solution $\vec{x}^*$ w.r.t. $\succeq$ (or equivalently, a maximal solution w.r.t. the inverse ordering $\preceq$). Nevertheless, a simple left-monotonicity condition over the valuation structure is enough for ensuring that a minimal solution $\vec{x}^*$ w.r.t. $\succeq$ can be computed in time polynomial in the size of the $e\text{-SLDD}$ formula, using dynamic programming. The result of [Wilson, 2005] indeed can be extended as follows:

Proposition 8 *For any monoid $\mathcal{E} = \langle E, \otimes, 1_s \rangle$ such that E is totally pre-ordered by $\succeq$, if $\otimes$ is left-monotonic w.r.t. $\succeq$ (for any $a, b, c \in E$, if $a \succeq b$ then $c \otimes a \succeq c \otimes b$), then for any $e\text{-SLDD}$ formula α , a solution $\vec{x}^*$ minimal w.r.t. $\succeq$ can be computed in time polynomial in the size of α .*

4 Succinctness of VDDs: Theoretical Results

Let $\mathcal{L}_1$ (resp. $\mathcal{L}_2$) be a representation language over X w.r.t. $\mathcal{E}_1$ (resp. $\mathcal{E}_2$). The notion of succinctness and of translations usually considered over propositional languages (see [Darwiche and Marquis, 2002]) can be extended as follows:

Definition 10 (succinctness) $\mathcal{L}_1$ is at least as succinct as $\mathcal{L}_2$, denoted $\mathcal{L}_1 \leq_s \mathcal{L}_2$, iff there exists a polynomial p such that for every $\alpha \in C_{\mathcal{L}_2}$, there exists $\beta \in C_{\mathcal{L}_1}$ which is equivalent to α and such that $s_{\mathcal{L}_1}(\beta) \leq p(s_{\mathcal{L}_2}(\alpha))$.

Definition 11 (linear / polynomial translation) $\mathcal{L}_2$ is linearly (resp. polynomially) translatable into $\mathcal{L}_1$, denoted $\mathcal{L}_1 \leq_l \mathcal{L}_2$ (resp. $\mathcal{L}_1 \leq_p \mathcal{L}_2$), iff there exists a linear-time (resp. polynomial-time) algorithm f from $C_{\mathcal{L}_2}$ to $C_{\mathcal{L}_1}$ such that for every $\alpha \in C_{\mathcal{L}_2}$, α is equivalent to $f(\alpha)$.

$<_s$ (resp. $<_p$, $<_l$) denotes the asymmetric part of $\leq_s$ (resp. $\leq_p$, $\leq_l$), and $\sim_s$ (resp. $\sim_p$, $\sim_l$) denotes the symmetric part of $\leq_s$ (resp. $\leq_p$, $\leq_l$). By construction, $\sim_s$, $\sim_p$, $\sim_l$ are equivalence relations.

We have obtained the following result showing that every ADD is linearly translatable into any $e\text{-SLDD}$ (sharing the same valuation set E):

Proposition 9 $e\text{-SLDD} \leq_l \text{ADD}$.

As to the valuation set $E = \mathbb{R}^+$, we get:

Proposition 10

- $\text{ADD} \sim_p e\text{-SLDD}_{\max}$.
- $e\text{-SLDD}_\times \not\leq_s e\text{-SLDD}_+$ and $e\text{-SLDD}_+ \not\leq_s e\text{-SLDD}_\times$.

- $\text{AADD} <_s e\text{-SLDD}_+ <_s \text{ADD}$.
- $\text{AADD} <_s e\text{-SLDD}_\times <_s \text{ADD}$.

Similarly, for $E = \mathbb{R}^+ \cup \{+\infty\}$, $\text{ADD} \sim_p e\text{-SLDD}_{\min}$ holds.

5 Succinctness of VDDs: Empirical Results

While succinctness is a way to compare representation languages w.r.t. the concept of spatial efficiency, it does not capture all aspects of this concept, for two reasons (at least). On the one hand, succinctness focuses on the worst case, only. On the other hand, it is of qualitative (ordinal) nature: succinctness indicates when an exponential separation can be achieved between two languages but does not enable to draw any quantitative conclusion on the sizes of the compiled forms. This is why it is also important to complete succinctness results with some size measurements.

To this aim, we made some experiments. We designed a bottom-up ordered $e\text{-SLDD}$ compiler. This compiler takes as input VCSP instances in the XML format described in [Roussel and Lecoutre, 2009] or Bayesian networks conforming to the XML format given in [Cozman, 2002]. When VCSP instances are considered, the compiler generates a data structure equivalent to each valued constraint of the instance, under the form of a reduced $e\text{-SLDD}_+$ formula, and incrementally combines them w.r.t. $+$ using a simplified version of the *apply*($+$) procedure described in [Santer and McAllester, 2005]. Similarly, when Bayesian network instances are considered, the conditional probability tables are first compiled into reduced $e\text{-SLDD}_\times$ formulae, which are then combined using $\times$. At each combination step, the current $e\text{-SLDD}$ formula is reduced. We developed a toolbox which also contains procedures for transforming any $e\text{-SLDD}_+$ (resp. $e\text{-SLDD}_\times$) formula into an equivalent ADD formula, and any ADD formula into an equivalent $e\text{-SLDD}_+$ (resp. $e\text{-SLDD}_\times$, AADD) formula; the transformation procedure from $e\text{-SLDD}_+$ (resp. $e\text{-SLDD}_\times$) formulae to ADD formulae roughly consists in pushing the labels from the root to the last arcs of the diagram. The transformation procedures from ADD formulae to $e\text{-SLDD}_+$, $e\text{-SLDD}_\times$ and AADD formulae are basically normalization procedures.

We considered two families of benchmarks. The VCSP instances we used concern car configurations problems;⁴ these instances contain hard constraints and soft constraints, with valuations representing prices, to be aggregated additively. They have the following characteristic features:

- Small: #variables=139; max. domain size=16; #constraints=176 (including 29 soft constraints)
- Medium: #variables=148; max. domain size=20; #constraints=268 (including 94 soft constraints)
- Big: #variables=268; max. domain size=324; #constraints=2157 (including 1825 soft constraints)

We also compiled only the soft constraints of the benchmarks, leading to three other instances, referred to as {Small, Medium, Big} Price only. As to

⁴These instances have been built in collaboration with the french car manufacturer Renault; they are described in more depth in [Aste-sana et al., 2013].

Table 1: Compilation of VCSPs into $e\text{-SLDD}_+$, and transformations into ADD, $e\text{-SLDD}_\times$ and AADD.

	$e\text{-SLDD}_+$		ADD	$e\text{-SLDD}_\times$	AADD
Instance	nodes (edges)	time (s)	nodes (edges)	nodes (edges)	nodes (edges)
Small Price only	36 (108)	< 1	4364 (7439)	3291 (7439)	36 (108)
Medium Price only	169 (499)	< 1	37807 (99280)	33595 (99280)	168 (495)
Big Price only	3317 (9687)	18	m-o	-	3317 (9687)
Small	2344 (5584)	1	299960 (637319)	14686 (33639)	2344 (5584)
Medium	6234 (17062)	6	752466 (2071474)	129803 (314648)	6234 (17062)
Big	198001 (925472)	79043	m-o	-	198001 (925472)

Table 2: Compilation of Bayesian networks into $e\text{-SLDD}_\times$, and transformations into ADD, $e\text{-SLDD}_+$ and AADD.

	$e\text{-SLDD}_\times$		ADD	$e\text{-SLDD}_+$	AADD
Instance	nodes (edges)	time (s)	nodes (edges)	nodes (edges)	nodes (edges)
Cancer	13 (25)	< 1	38 (45)	23 (45)	11 (21)
Asia	23 (45)	< 1	415 (431)	216 (431)	23 (45)
Car-starts	41 (83)	< 1	42741 (64029)	19632 (39265)	38 (77)
Alarm	1301 (3993)	< 1	m-o	-	1301 (3993)
Hailfinder25	32718 (108083)	8	m-o	-	32713 (108063)

Bayesian networks, which are of multiplicative nature (joint probabilities are products of conditional probabilities), we used some standard benchmarks [Cozman, 2002].

Each configuration (resp. Bayesian net) instance has been compiled into an $e\text{-SLDD}_+$ formula (resp. an $e\text{-SLDD}_\times$ formula), and then transformed into an ADD formula, an $e\text{-SLDD}_\times$ formula (resp. an $e\text{-SLDD}_+$ formula), and an AADD formula – the time needed for the compilation and the sizes on the compiled formulae are reported in Table 1 (resp. Table 2). In order to determine a variable ordering, we used the *Maximum Cardinality Search* heuristic [Tarjan and Yannakakis, 1984] in reverse order, as proposed in [Amilhastre, 1999] for the compilation of (classical) CSPs. This heuristic is easy to compute and efficient; experiments reported in [Amilhastre, 1999] show that it typically outperforms several standard CSP variable ordering heuristics.

We ran all our experiments on a computer running at 800MHz with 256Mb of memory. "m-o" means that the available memory has been exhausted, and that the program aborted for this reason.

Our experiments confirm some of the theory-oriented succinctness results, especially the fact that the succinctness of $e\text{-SLDD}_+$ and of $e\text{-SLDD}_\times$ are incomparable but each of them is strictly more succinct than ADD. Unsurprisingly, when the values of the soft constraints are to be aggregated additively as this is the case for configuration instances (resp. multiplicatively, as this is the case for Bayesian nets), $e\text{-SLDD}_+$ (resp. $e\text{-SLDD}_\times$) performs better than $e\text{-SLDD}_\times$ (resp. $e\text{-SLDD}_+$). AADD does not prove to be better than $e\text{-SLDD}_+$ in the additive case, or better than $e\text{-SLDD}_\times$ in the multiplicative case.⁵ Thus, targeting the AADD language

does not lead to much better compiled formulae from the spatial efficiency point of view, when the mapping to be represented is additive or multiplicative in essence, but not both.

6 Conclusion

In this paper, we have extended the SLDD family to the $e\text{-SLDD}$ family, thanks to a relaxation of some requirements on the valuation structure, which is harmless for the conditioning and optimization purposes. The $e\text{-SLDD}$ family is general enough to capture AADD as a specific element. We have pointed out a normalization procedure and a canonicity condition for formulae from some $e\text{-SLDD}$ languages, including $e\text{-SLDD}_+$ and $e\text{-SLDD}_\times$. We have also compared the spatial efficiency of some elements of the $e\text{-SLDD}$ family, i.e., $e\text{-SLDD}_+$ and $e\text{-SLDD}_\times$, with ADD and AADD from both the theoretical side and the practical side. Though $e\text{-SLDD}_+$ (resp. $e\text{-SLDD}_\times$) is less succinct than AADD from a theoretical point of view, it proves space-efficient enough for enabling the compilation of cost-based configuration problems (resp. Bayesian networks).

Interestingly, one of the conditions pointed out in the $e\text{-SLDD}$ setting for tractable normalization (and reduction) does not impose the valuation set E to be *totally* ordered. Clearly, this paves the way for the compilation of multi-criteria objective functions as $e\text{-SLDD}$ representations. Investigating this issue is a major perspective for future works. Another important issue for further research is to draw the full knowledge compilation map for VDD languages, which will require to identify the tractable queries and transforma-

⁵On the Bayesian net instances, the resulting ADD and AADD formulae are larger than the ones obtained by [Sanner and McAllester, 2005]. This is due to the way numeric labels are merged (remember that reals are approximated by finite-precision floating-point num-

bers on a computer). Indeed, in our implementation, e_1 and e_2 are considered identical whenever $e_1 - e_2 < 10^{-9} \cdot e_1$ (where $e_1 \geq e_2$). Since $e_1, e_2 \leq 1$ (they represent probabilities) the standard merging condition $e_1 - e_2 < 10^{-9}$ considered in [Sanner and McAllester, 2005] is subsumed by ours. This explains the size discrepancy.

tions of interest, depending on the algebraic properties of the valuation structure.

References

- [Amilhastre *et al.*, 2002] Jérôme Amilhastre, Hélène Fargier, and Pierre Marquis. Consistency restoration and explanations in dynamic CSPs application to configuration. *Artif. Intell.*, 135(1-2):199–234, 2002.
- [Amilhastre, 1999] Jérôme Amilhastre. *Représentation par automate d'ensemble de solutions de problèmes de satisfaction de contraintes*. PhD thesis, Université de Montpellier II, 1999.
- [Astesana *et al.*, 2013] Jean-Marc Astesana, Laurent Cosserat, and Hélène Fargier. Business recommendation for configurable products (BR4CP) project: the case study. <http://www.irit.fr/~Helene.Fargier/BR4CPBenchs.html>, March 2013.
- [Bahar *et al.*, 1993] R. Iris Bahar, Erica A. Frohm, Charles M. Gaona, Gary D. Hachtel, Enrico Macii, Abelardo Pardo, and Fabio Somenzi. Algebraic decision diagrams and their applications. In *Proc. of ICCAD'93*, pages 188–191, 1993.
- [Cooper and Schiex, 2004] Martin C. Cooper and Thomas Schiex. Arc consistency for soft constraints. *Artif. Intell.*, 154(1-2):199–227, 2004.
- [Cozman, 2002] Fabio Gagliardi Cozman. JavaBayes Version 0.347, Bayesian Networks in Java, User Manual. Technical report, dec 2002. Benchmarks at <http://sites.poli.usp.br/pmr/ltd/Software/javabayes/Home/node3.html>.
- [Darwiche and Marquis, 2002] Adnan Darwiche and Pierre Marquis. A knowledge compilation map. *J. Artif. Intell. Res. (JAIR)*, 17:229–264, 2002.
- [Darwiche, 2011] A. Darwiche. SDD: A new canonical representation of propositional knowledge bases. In *Proc. of IJCAI'11*, pages 819–826, 2011.
- [Gogic *et al.*, 1995] G. Gogic, H.A. Kautz, Ch.H. Papadimitriou, and B. Selman. The comparative linguistics of knowledge representation. In *Proc. of IJCAI'95*, pages 862–869, 1995.
- [Lai and Sastry, 1992] Yung-Te Lai and Sarma Sastry. Edge-valued binary decision diagrams for multi-level hierarchical verification. In *Proc. of DAC'92*, pages 608–613, 1992.
- [Lai *et al.*, 1996] Yung-Te Lai, Massoud Pedram, and Sarma B. K. Vrudhula. Formal verification using edge-valued binary decision diagrams. *IEEE Trans. on Computers*, 45(2):247–255, 1996.
- [Roussel and Lecoutre, 2009] Olivier Roussel and Christophe Lecoutre. XML Representation of Constraint Networks: Format XCSP 2.1. Technical report, CoRR abs/0902.2362, feb 2009.
- [Sanner and McAllester, 2005] Scott Sanner and David A. McAllester. Affine algebraic decision diagrams (AADDs) and their application to structured probabilistic inference. In *Proc. of IJCAI'05*, pages 1384–1390, 2005.
- [Schiex *et al.*, 1995] Thomas Schiex, Hélène Fargier, and Gérard Verfaillie. Valued constraint satisfaction problems: Hard and easy problems. In *Proc. of IJCAI'95*, pages 631–639, 1995.
- [Tafertshofer and Pedram, 1997] Paul Tafertshofer and Massoud Pedram. Factored edge-valued binary decision diagrams. *Formal Methods in System Design*, 10(2/3):243–270, 1997.
- [Tarjan and Yannakakis, 1984] Robert E. Tarjan and Mihalis Yannakakis. Simple linear-time algorithms to test chordality of graphs, test acyclicity of hypergraphs, and selectively reduce acyclic hypergraphs. *SIAM J. Comput.*, 13(3):566–579, July 1984.
- [Wilson, 2005] Nic Wilson. Decision diagrams for the computation of semiring valuations. In *Proc. of IJCAI'05*, pages 331–336, 2005.