

Data Structures for Symbolic Multi-Valued Model-Checking

Marsha Chechik, Arie Gurfinkel, Benet Devereux, Albert Lai and Steve Easterbrook
Department of Computer Science, University of Toronto,
Toronto, ON M5S 3G4, Canada.
Email: {chechik, arie, benet, trebla, sme}@cs.toronto.edu

Draft of June 7, 2006

Abstract

Multi-valued logics provide an interesting alternative to classical boolean logic for modeling and reasoning about systems. Such logics can be used for reasoning about partially-specified systems, effectively encode vacuity detection and query-checking problems, help in detecting inconsistencies, and many others.

In our earlier work, we identified a useful family of multi-valued logics: those specified over finite distributive lattices where negation preserves involution, i.e., $\neg\neg a = a$ for every element a of the logic. Such structures are called *quasi-boolean algebras*, and model-checking over these not only extends the domain of applicability of automated reasoning to new problems, but can also speed up solutions to some classical verification problems.

Symbolic model-checking over quasi-boolean algebras can be cast in terms of operations over multi-valued sets: sets whose membership functions are multi-valued. In this paper, we propose and empirically evaluate several choices for implementing multi-valued sets with decision diagrams. In particular, we describe two major approaches: (1) representing the multi-valued membership function canonically, using MDDs or ADDs; (2) representing multi-valued sets as a collection of classical sets, using a vector of either MBTDDs or BDDs. The naive implementation of (2) includes having a classical set for each value of the algebra. We exploit a result of lattice theory to reduce the number of such sets that need to be represented.

The major contribution of this paper is the evaluation of the different implementations of multi-valued sets, done via a series of experiments and using several case studies.

1 Introduction

Multi-valued logics provide an interesting alternative to classical boolean logic for modeling and reasoning about systems. By allowing truth values other than just the classical true and false, these logics support the explicit modeling of uncertainty and disagreement. For example, Bruns and Godefroid [12] use three-valued Kleene logic [42] for model-checking of partially specified systems: the in-between value represents a transition which may become either enabled or disabled in further refinements of the specification. This logic has also been used for building abstractions [52] and performing counterexample-based

abstraction-refinement [34, 56] and is starting to be used in software model-checking. Hazelhurst and Seger [40] use four-valued Belnap logic [8] to model abstraction and inconsistency in Symbolic Trajectory Evaluation (STE). Pazos-Arias et al. [48] use a six-valued logic to formalize a more detailed view of partial specification: the additional values are used to distinguish between possible refinements, and also to indicate where an inconsistency has arisen. In addition, multi-valued modeling and reasoning has been recently applied to a variety of other verification problems. For example, *query-checking* [18] – the process of discovering properties of models – can be cast into a multi-valued model-checking problem using upset lattices of propositional formulas [13, 38]. It is also possible to determine whether a given formula is *vacuous* in a subformula of interest [44] by reasoning with values “Vacuously True”, “Vacuously False”, “Non-Vacuously True” and “Non-Vacuously False” [37]. The above logics fall into a family of *quasi-boolean* [11] (or *De Morgan* [31]) algebras. These are the algebras defined over finite quasi-boolean lattices with an additional negation operator \neg , so that $\neg\neg a = a$ for every element a of the algebra.

Clearly, multi-valued algebras with a finite number of values do not add expressive power to the modeling. Yet, modeling with such algebras is often more intuitive than staying within the confines of classical logic. In all of the above cases, casting the problem into many values not only simplified the existing approaches but gave insights that led to considerable improvements of the solutions.

In our previous work [20], we formulated the general multi-valued model-checking problem: given a particular quasi-boolean algebra, a state machine model, and a correctness property expressed in a temporal logic enriched with multi-valued semantics, a multi-valued model-checker returns the truth value that the property has in the model, and a counterexample, if applicable. In particular, we defined CTL model-checking (subsequently extended to the propositional μ -calculus [36]), proved correctness of our algorithms and studied its complexity. We also explored a few additional applications of multi-valued modeling.

How should multi-valued model-checking be implemented? In our recent work, we showed that it can be reduced to several questions to the classical model-checker [36]. However, in many cases the direct approach results in significantly smaller state spaces than the alternative of introducing additional boolean predicates (see [20] for a few examples). Further, doing the verification directly allows us to compute a counterexample [35]. That is, if a property φ evaluates to some value ℓ of the algebra, we explain why the answer is no more true than ℓ , and why it is no less true than ℓ . Such counterexamples can be used in abstraction-refinement frameworks [5], as witnesses to non-vacuity [7], as explanations of query-checking results [38], etc. For these reasons, we have implemented symbolic multi-valued model-checking directly, in a tool called χ Chek [22].

There are a number of alternatives for symbolic encoding of multi-valued operations. Which of these should be used? In this paper, we define such operations using one generic data structure called *mv-sets*, describe four implementations of mv-sets and evaluate these implementations experimentally. Thus, the goal of this paper is to apply and evaluate the well-tried approach of performing symbolic reasoning using decision diagrams to the new problem of multi-valued model-checking. This paper makes the following major contributions:

1. Success of symbolic model-checking in a given domain (probabilistic, multi-valued, timed, etc.) depends on efficient algorithms for manipulating the sets of states in which a property holds. We need to calculate union, intersection, complement, and quantify over these sets. For model checking in a given multi-valued algebra, we can treat these as operations over *multi-valued sets (mv-sets)*: sets whose membership functions are multi-valued.




Figure 1: Some lattices. Join-irreducible elements (see Definition 7) are indicated by filled circles.

2. Mv-sets have multi-valued membership functions, and thus can be naturally encoded using MDDs, the multi-valued extension of binary decision diagrams [15]. MDDs are multi-terminal decision diagrams where each node has a multi-valued branching factor. Mv-sets can also be represented and manipulated as a collection of classical sets. The easiest such encoding is the one that associates one classical set with each element of the algebra. Our second contribution is in using the join-irreducible decomposition of lattice elements to encode mv-set manipulations. This decomposition is always better than the one above. The encoding can be represented symbolically using boolean-terminal decision diagrams, and in particular, using Multi-Valued Boolean-Terminal diagrams (MBTDDs). MBTDDs is a special case of MDDs where the input variables are multi-valued but the terminal nodes have only the classical truth values 0 or 1. Preliminary version of this work appeared in [21].
3. In order to compare the MDD- and the MBTDD-based implementations of mv-sets, we devised a new method for generating multi-valued benchmarks from classical ones.
4. Using these benchmarks, we have performed experiments to explore the tradeoffs in time and space between the two proposed representations, in the specific domain of model-checking.

The rest of this paper is organized as follows: Sections 2-4 give background material for this paper, defining quasi-boolean algebras, multi-valued sets and relations, and multi-valued model-checking, respectively. We discuss the high-level implementation of the multi-valued model-checking algorithm in Section 5. The MDD-based and the MBTDD-based implementations of mv-sets are described in Sections 6 and 7, respectively. Section 8 compares the performance of mv-set operations on three case studies. The experiments and the theoretical assessment of the complexity of mv-set operations under different decision diagram representations are given in Section 9. We conclude in Section 10 with the summary of the paper and an outline for future work.

2 Quasi-Boolean Algebras

In this section we review multi-valued logics, where values are elements of distributive lattices, conjunction and disjunction is defined as the lattice operations meet and join, and negation preserves De Morgan's laws. The resulting algebras are called *quasi-boolean*, or *De Morgan* [11, 31]. For details, see our previous work [20].

Definition 1 A partially ordered set, $(\mathcal{L}, \sqsubseteq)$, is a lattice if a unique greatest lower bound and least upper bound exist for every finite subset of \mathcal{L} .

Given lattice elements a and b , their greatest lower bound is referred to as *meet* and denoted $a \sqcap b$, and their least upper bound is referred to as *join* and denoted $a \sqcup b$. Every (finite) lattice has a top (\top) and a bottom (\perp) element, and in this paper we assume that all lattices are finite. Additional operations are defined as follows:

$$\begin{aligned} a \rightarrow b &\triangleq \neg a \sqcup b && \text{material implication} \\ a \leftrightarrow b &\triangleq (a \rightarrow b) \sqcap (b \rightarrow a) && \text{equivalence} \\ a = b &\triangleq (a \sqsubseteq b) \wedge (b \sqsubseteq a) && \text{equality} \end{aligned}$$

A lattice is *distributive* iff \sqcap and \sqcup distribute over each other, and in this paper we use only distributive lattices. In defining functions over lattice elements, we distinguish the elements \top and \perp , since they are present in every lattice:

Definition 2 Let $(\mathcal{L}, \sqsubseteq)$ be a lattice. An element $x \in \mathcal{L}$ is called *crisp* iff $x \in \{\top, \perp\}$. A function $f : \mathcal{L}^n \rightarrow \mathcal{L}$ is called *crisp-output* iff its only outputs are \top and \perp : that is, $f(\mathcal{L}^n) \subseteq \{\perp, \top\}$. It is called *crisp-input* iff only inputs in $\{\perp, \top\}$ can lead to a non- \perp output: $f(\mathcal{L}^n \setminus \{\perp, \top\}^n) = \{\perp\}$. A function f is called *crisp* if it is both *crisp-input* and *crisp-output*.

Definition 3 [51, 20] A quasi-boolean algebra is a tuple $(\mathcal{L}, \sqcap, \sqcup, \neg)$, where:

- $(\mathcal{L}, \sqsubseteq)$ is a finite distributive lattice, with $a \sqsubseteq b$ iff $a \sqcap b = a$;
- Conjunction (\sqcap) and disjunction (\sqcup) are meet and join operators of $(\mathcal{L}, \sqsubseteq)$, respectively;
- Negation \neg is a function $\mathcal{L} \rightarrow \mathcal{L}$ such that every element $a \in \mathcal{L}$ corresponds to a unique element $\neg a \in \mathcal{L}$ satisfying the following conditions:

$$\begin{aligned} \neg(a \sqcap b) &= \neg a \sqcup \neg b && \text{De Morgan} && \neg\neg a &= a && \neg \text{involution} \\ \neg(a \sqcup b) &= \neg a \sqcap \neg b && && a \sqsubseteq b &\Leftrightarrow \neg a \sqsupseteq \neg b && \neg \text{antimonotonic} \end{aligned}$$

where $b \in \mathcal{L}$. $\neg a$ is called a quasi-complement of a .

The negation operator satisfying the above properties is a *lattice dual isomorphism* with period 2 [10].

Definition 4 [20] A product of two algebras, $L_1 = (\mathcal{L}_1, \sqcap_1, \sqcup_1, \neg_1)$ and $L_2 = (\mathcal{L}_2, \sqcap_2, \sqcup_2, \neg_2)$, is an algebra, $L_1 \times L_2 = (\mathcal{L}_1 \times \mathcal{L}_2, \sqcap, \sqcup, \neg)$, where the operations, including \top , \perp , and the order operator, are the component-wise extensions of their individual counterparts, e.g.

$$(a, b) \sqsubseteq (a', b') \Leftrightarrow a \sqsubseteq_1 a' \wedge b \sqsubseteq_2 b' \quad \sqsubseteq \text{ of pairs}$$

A product of two quasi-boolean algebras is quasi-boolean [20].

We now give some example quasi-boolean algebras using the lattices in Figure 1.

1. The lattice in Figure 1(a), with $\neg\top = \text{F}$ and $\neg\text{F} = \top$, gives us classical logic, which we refer to as **2**. Note that in this case, \sqcup and \sqcap are conventionally written \vee and \wedge , respectively.
2. The three-valued logic **3** is defined on the lattice in Figure 1(b), where $\neg\top = \text{F}$, $\neg\text{F} = \top$, $\neg\text{M} = \text{M}$.

3. The lattice in Figure 1(c) shows the product algebra $\mathbf{2x2}$, where $\neg TF = FT$ and $\neg FT = TF$. The underlying lattice is isomorphic to the one in Figure 1(d), but the resulting quasi-boolean algebras are not isomorphic, because of the choice of negations.
4. The lattice in Figure 1(e) shows a nine-valued logic constructed as the product algebra $\mathbf{3x3}$.

We now define the embedding from $\mathbf{2}$ into an arbitrary quasi-boolean algebra.

Definition 5 Let $L_1 = (\mathcal{L}_1, \sqcap_1, \sqcup_1, \neg_1)$ and $L_2 = (\mathcal{L}_2, \sqcap_2, \sqcup_2, \neg_2)$ be two quasi-boolean algebras. A function $h : \mathcal{L}_1 \rightarrow \mathcal{L}_2$ is a quasi-boolean algebra homomorphism iff h is a lattice homomorphism between $(\mathcal{L}_1, \sqsubseteq_1)$ and $(\mathcal{L}_2, \sqsubseteq_2)$, and $\forall x \in \mathcal{L}_1 \cdot h(\neg_1 x) = \neg_2 h(x)$.

Since any quasi-boolean algebra homomorphism is also a lattice homomorphism, it preserves top and bottom elements; therefore, there is a unique homomorphism from $\mathbf{2}$ to any quasi-boolean algebra taking T to \top and F to \perp .

Definition 6 Let L be a quasi-boolean algebra. An embedding of classical logic into L is the unique homomorphism $\alpha_L : \mathbf{2} \rightarrow L$.

Intuitively, α_L corresponds to an embedding of classical truth-values into an arbitrary quasi-boolean algebra L . For instance, $\alpha_{\mathbf{2x2}} : \mathbf{2} \rightarrow \mathbf{2x2}$ maps T to TT and F to FF.

We extend this notion of embedding to arbitrary boolean functions: if $f : \mathbf{2}^n \rightarrow \mathbf{2}$ is a boolean function, then $\alpha_L(f) : L^n \rightarrow L$ is the embedding of f into a quasi-boolean algebra L .

We now review the concept of join-irreducibility for distributive lattices.

Definition 7 [29] An element j in a lattice $(\mathcal{L}, \sqsubseteq)$ is join-irreducible iff $j \neq \perp$ and for any x and y in \mathcal{L} , $j = x \sqcup y$ implies $j = x$ or $j = y$. Dually, an element m is meet-irreducible iff $m \neq \top$ and for any x and y , $m = x \sqcap y$ implies $m = x$ or $m = y$.

In other words, a join-irreducible element j cannot be further decomposed into the join of other elements in the lattice, and a meet-irreducible element m cannot be further decomposed into the meet of other elements in the lattice, just as prime numbers cannot be further factored into the product of other natural numbers. For example, the join-irreducible elements of the lattice $\mathbf{3x3}$ in Figure 1(e), indicated by unfilled circles, are $\{MF, TF, FM, FT\}$. We denote the sets of all join-irreducible and meet-irreducible elements of a lattice $(\mathcal{L}, \sqsubseteq)$ by $\mathcal{J}(\mathcal{L}, \sqsubseteq)$ and $\mathcal{M}(\mathcal{L}, \sqsubseteq)$, respectively.

Every element of a finite lattice is equal to a join of all join-irreducible elements below it:

Theorem 1 [29] Let ℓ be any element in a lattice $(\mathcal{L}, \sqsubseteq)$. Then $\ell = \bigsqcup\{j \in \mathcal{J}(\mathcal{L}, \sqsubseteq) \mid j \sqsubseteq \ell\}$.

For example, in the lattice shown in Figure 1(e), $TT = \bigsqcup\{MF, TF, FM, FT\}$, $TM = \bigsqcup\{MF, TF, FM\}$, $FF = \bigsqcup\emptyset$, and so on. Thus, each element of a lattice can be represented uniquely by a subset of join-irreducibles. For example, $TT = \{MF, TF, FM\}$.

Given an algebra, $L = (\mathcal{L}, \sqcap, \sqcup, \neg)$, we use notation $\mathcal{J}(L)$ and $\mathcal{M}(L)$ to refer to $\mathcal{J}(\mathcal{L}, \sqsubseteq)$ and $\mathcal{M}(\mathcal{L}, \sqsubseteq)$, respectively. Further, we often use L as the elements comprising the algebra, so $|L|$ means the size of the algebra and $\forall \ell \in L$ quantifies over all elements of the algebra.




Figure 2: **Sample**: a simple χ Kripke structure.

3 Multi-Valued Sets and Relations

Multi-valued model-checking algorithms can be naturally described using a data structure that encapsulates reasoning about operations on sets of states in which a property holds. Such operations include union, intersection, complement, and backward image for computing predecessors. Given an algebra, we can treat these as operations over multi-valued sets: sets whose membership functions are multi-valued.

In this section, we define mv-sets – data types for storing and manipulating sets with varying membership degrees. A familiar instance of such mv-sets is a probability distribution on a set of independent events. In the case of multi-valued model-checking, the entities are states and transitions.

3.1 Multi-Valued Sets

In classical set theory, a set is defined by a boolean predicate, also called a *membership function*. Typically, it is written using the *set comprehension notation*: a predicate P defines the set $S = \{x \mid P(x)\}$. For instance, if $P = \lambda x \in \text{nat} \cdot 0 \leq x \leq 10$, then S is the set of all integers between 0 and 10 inclusive. If, instead of using a boolean membership function, we allow the membership function to range over an algebra with more than two values, we obtain a *multi-valued* set theory in which it is possible to make statements like “element x is more in a set \mathbb{S} than element y ”. We call the resulting sets *mv-sets* and refer to them using a special font, e.g., \mathbb{S} .

Definition 8 Given an algebra, $L = (\mathcal{L}, \sqcap, \sqcup, \neg)$, and a classical set, S , an L -valued set on S , referred to as \mathbb{S} , is a total function $S \rightarrow \mathcal{L}$.

When the underlying algebra L is clear from context, we refer to an L -valued set just as an mv-set. For an mv-set, \mathbb{S} , and a candidate element, x , we use $\mathbb{S}(x)$ to denote the membership degree of x in \mathbb{S} . In the classical case, this amounts to representing a set by its characteristic function.

We illustrate mv-sets and other concepts in this paper using a simple state machine **Sample**, shown in Figure 2. This machine uses the quasi-boolean algebra $\mathbf{2x2}$ where the logical values form the lattice in Figure 1(c), and exemplifies a χ Kripke structure – a multi-valued generalization of Kripke structures, defined formally in Section 4.2. In classical symbolic model-checking, each (boolean-valued) expression, x , partitions the state space into states where x is true and states where it is false. Likewise, we use multi-valued expressions to partition the state space of the system. For example, the variable, c , partitions the states of the χ Kripke structure **Sample**: for each value, ℓ of $\mathbf{2x2}$, we get the set of states where c has value




Figure 3: (a) The 4-valued set of states reflecting values of c in **Sample**; (b) The 4-valued set of states reflecting values of a ; (c) A union of mv-sets in (a) and (b); (d) A multi-valued complement of the mv-set in (a).

ℓ : it has value TT in $\{s_3\}$, TF in $\{s_2\}$, FF in $\{s_0, s_1\}$, and FT nowhere. The resulting mv-set, referred to as $\llbracket c \rrbracket$, can be graphically represented as shown in Figure 3(a), where the structure corresponds to that of the underlying lattice.

We extend some standard set operations to the multi-valued case by lifting the lattice meet and join operations, as follows¹:

$$\begin{aligned} (\mathbb{S} \cap_L \mathbb{S}')(x) &\triangleq (\mathbb{S}(x) \sqcap \mathbb{S}'(x)) && \text{multi-valued intersection} \\ (\mathbb{S} \cup_L \mathbb{S}')(x) &\triangleq (\mathbb{S}(x) \sqcup \mathbb{S}'(x)) && \text{multi-valued union} \end{aligned}$$

For example, in computing the union of mv-sets $\llbracket c \rrbracket$ and $\llbracket a \rrbracket$ given in Figure 3(a) and (b), respectively, we note that in state s_1 , c is FF and a is TT. Thus, $(\llbracket c \rrbracket \cup_L \llbracket a \rrbracket)(s_1) = \text{FF} \sqcup \text{TT} = \text{TT}$. The set representing this union is given in Figure 3(c).

We can also extend the notion of *set complement* to the multi-valued case, by defining it in terms of the quasi-complement of L , and denoting it with a bar:

$$\overline{\mathbb{S}}(x) \triangleq \neg(\mathbb{S}(x)) \quad \text{multi-valued complement}$$

Mv-set $\overline{\llbracket c \rrbracket}$ is given in Figure 3(d). We then obtain the expected properties:

$$\begin{aligned} \overline{\mathbb{S} \cup_L \mathbb{S}'} &= \overline{\mathbb{S}} \cap_L \overline{\mathbb{S}'} && \text{De Morgan 1} && \mathbb{S} \subseteq_L \mathbb{S}' &= \overline{\mathbb{S}'} \subseteq_L \overline{\mathbb{S}} && \text{antimonotonicity} \\ \overline{\mathbb{S} \cap_L \mathbb{S}'} &= \overline{\mathbb{S}} \cup_L \overline{\mathbb{S}'} && \text{De Morgan 2} \end{aligned}$$

We now define some additional concepts for mv-sets that are not needed in the classical set theory but are used later in this paper.

Definition 9 *Support, Core, ℓ -cut, and ℓ -clip of a multi-valued set \mathbb{S} :*

$$\begin{aligned} \sigma(\mathbb{S}) &\triangleq \{x \mid \mathbb{S}(x) \neq \perp\} && \text{Support} && \mathcal{C}(\mathbb{S}) &\triangleq \{x \mid \mathbb{S}(x) = \top\} && \text{Core} \\ \uparrow_\ell(\mathbb{S}) &\triangleq \{x \mid \ell \sqsubseteq \mathbb{S}(x)\} && \ell\text{-cut} && \downarrow_\ell(\mathbb{S}) &\triangleq \{x \mid \mathbb{S}(x) \sqsubseteq \ell\} && \ell\text{-clip} \end{aligned}$$

¹The subscript on mv-set operations \cup_L, \cap_L , and so forth, refers to a given algebra, $L = (\mathcal{L}, \sqcap, \sqcup, \neg)$.

All of these operations create classical sets. Support, cut, and core are standard concepts from fuzzy set theory [61]; clip is a new operation, which we need to prove a later result. For example,

$$\begin{aligned}\sigma(\llbracket c \rrbracket) &= \{s_3, s_2\} & \mathcal{C}(\llbracket c \rrbracket) &= \{s_3\} \\ \uparrow_{\text{TF}}(\llbracket c \rrbracket) &= \{s_3, s_2\} & \uparrow_{\text{FT}}(\llbracket c \rrbracket) &= \{s_3\} \\ \downarrow_{\text{TF}}(\llbracket c \rrbracket) &= \{s_0, s_2, s_1\} & \downarrow_{\text{FT}}(\llbracket c \rrbracket) &= \{s_0, s_1\}\end{aligned}$$

Following the conventions of fuzzy set theory, we identify an explicit universe of discourse U , rather than use an undefinable set of all possible entities. Mv-sets can be thought of as functions from elements of U to the underlying algebra; therefore, $U = \uparrow_{\perp}(\mathbb{S})$.

3.2 Multi-Valued Relations

Now we extend the concept of degrees of membership in an mv-set to degrees of relatedness of two entities. This concept, formalized by *multi-valued relations*, allows us to define multi-valued transitions in state machine models.

Definition 10 For a given algebra, L , an L -valued relation \mathbb{R} on two sets S and T is an L -valued set on $S \times T$.

Let S be the set of states of the χ Kripke structure **Sample**, shown in Figure 2. The multi-valued relation over $S \times S$ represents values of transitions between pairs of states of **Sample** and is shown in Figure 4(a). We refer to this transition as \mathbb{A} . For example, the value of the transition (s_3, s_2) is TF, so $\mathbb{A}((s_3, s_2)) = \text{TF}$.

Definition 11 Given an algebra, L , an L -valued relation, \mathbb{R} , on sets S and T , and an L -valued set, \mathbb{S} , on S , t forward image of \mathbb{S} under \mathbb{R} , denoted $\vec{\mathbb{R}}(\mathbb{S})$, is an L -valued set on T , defined as:

$$\vec{\mathbb{R}}(\mathbb{Q}) \triangleq \lambda t \in T \cdot \bigsqcup_{s \in S} (\mathbb{S}(s) \sqcap \mathbb{R}(s, t))$$

and for an L -valued set, \mathbb{T} , on T , backward image of \mathbb{T} under \mathbb{R} is

$$\overleftarrow{\mathbb{R}}(\mathbb{T}) \triangleq \lambda s \in S \cdot \bigsqcup_{t \in T} (\mathbb{T}(t) \sqcap \mathbb{R}(s, t))$$

Intuitively, forward image of an mv-set \mathbb{S} under the relation represents all elements reachable from \mathbb{S} by \mathbb{R} , where multi-valued memberships of \mathbb{R} and \mathbb{S} are taken into consideration. Similarly, backward image of an mv-set \mathbb{T} under \mathbb{R} represents all elements that can reach \mathbb{T} by \mathbb{R} .

We now compute forward and backward images of $\llbracket c \rrbracket$ (see Figure 3(a)) under the multi-valued relation \mathbb{A} between the pairs of states of the χ Kripke structure **Sample**. These are shown in Figures 4(b) and (c), respectively. For example, when we compute backward image of s_0 , we get

$$\bigsqcup_{t \in S} \llbracket c \rrbracket(t) \sqcap \mathbb{T}(s_0, t) = (\text{FF} \sqcap \text{TT}) \sqcup (\text{FF} \sqcap \text{FF}) \sqcup (\text{TT} \sqcap \text{FF}) \sqcup (\text{TF} \sqcap \text{TT}) = \text{TF}$$

which indicates that there exists a TT transition from state s_0 to another state (actually, s_2), where c is TF.

In this paper we do not use forward image. It is defined here for symmetry.




Figure 4: (a) \mathbb{A} : The multi-valued relation between pairs of states of **Sample**; (b) forward image of $\llbracket c \rrbracket$ over \mathbb{A} ; (c) backward image of $\llbracket c \rrbracket$ over \mathbb{A} .

4 CTL Model-Checking: from Classical to Multi-Valued

In this section we review classical CTL model-checking and then briefly describe its multi-valued extensions.

4.1 CTL Model-Checking

CTL model-checking is an automatic technique for verifying properties expressed in a propositional branching-time temporal logic called *Computation Tree Logic* (CTL) [27]. A model is a Kripke structure, and properties are evaluated on a tree of infinite computations produced by the model. The standard notation $M, s \models \varphi$ indicates that a formula φ holds in a state s of a model M . If a formula holds in the initial state, it is considered to hold in the model.

A Kripke structure consists of a set of states, S , a transition relation, $R \subseteq S \times S$, an initial state, $s_0 \in S$, a set of atomic propositions, A , and a labeling function, $I : S \rightarrow 2^A$. R must be total, i.e., $\forall s \in S, \exists t \in S$, such that $(s, t) \in R$. Finite computations are modeled by adding a self-loop to the final state of the computation. For each $s \in S$, the labeling function provides a set of atomic propositions which hold in the state S .

The syntax of CTL is as follows:

1. Every atomic proposition $a \in A$ is a CTL formula.
2. If φ and ψ are CTL formulas, then so are $\neg\varphi, \varphi \wedge \psi, \varphi \vee \psi, EX\varphi, AX\varphi, EF\varphi, AF\varphi, E[\varphi U \psi], A[\varphi U \psi], AG\varphi, EG\varphi$.

The logic connectives \neg, \wedge and \vee have their usual meanings. The existential and universal quantifiers E and A are used to quantify over paths. The operator X means “in the next state”, F represents “sometime in the future”, U is “until”, and G is “globally”. For example, $EX\varphi$ is true in state s if φ holds in some immediate successor of s , while $AX\varphi$ is true if φ holds in every immediate successor of s . $EF\varphi$ is true in s if φ holds in the future along some path from s ; $E[\varphi U \psi]$ is true in s if along some path from s , φ continuously holds until ψ becomes true. $EG\varphi$ hold in s if φ holds in every state along some path from

$AG\varphi$	$= \nu Z.(\varphi \wedge AXZ)$	AG fixpoint
$EG\varphi$	$= \nu Z.(\varphi \wedge EXZ)$	EG fixpoint
$AF\varphi$	$= \mu Z.(\varphi \vee AXZ)$	AF fixpoint
$EF\varphi$	$= \mu Z.(\varphi \vee EXZ)$	EF fixpoint
$A[\varphi U \psi]$	$= \mu Z.(\psi \vee (\varphi \wedge AXZ))$	AU fixpoint
$E[\varphi U \psi]$	$= \mu Z.(\psi \vee (\varphi \wedge EXZ))$	EU fixpoint

Figure 5: Fixpoint formulations of CTL operators. Note: $\mu Z.f(Z)$ and $\nu Z.f(Z)$ indicate the least and the greatest fixpoints of f , respectively.

s . $AF\varphi$, $A[\varphi U \psi]$ and $AG\varphi$ are defined similarly, replacing the quantification over some paths by the one over all paths. Formally,

$$\begin{aligned}
M, s \models a & \text{ iff } a \in I(s) \\
M, s \models \neg\varphi & \text{ iff } M, s \not\models \varphi \\
M, s \models \varphi \wedge \psi & \text{ iff } M, s \models \varphi \wedge M, s \models \psi \\
M, s \models \varphi \vee \psi & \text{ iff } M, s \models \varphi \vee M, s \models \psi \\
M, s \models EX\varphi & \text{ iff } \exists t \in S, (s, t) \in R \wedge M, t \models \varphi \\
M, s_i \models EG\varphi & \text{ iff there exists some path } s_i, s_{i+1}, \dots \text{ s.t. } \forall j \geq i \cdot M, s_j \models \varphi \\
M, s_i \models E[\varphi U \psi] & \text{ iff there exists some path } s_i, s_{i+1}, \dots, \text{ s.t.} \\
& \exists j \geq i \cdot M, s_j \models \psi \wedge \forall k \cdot i \leq k < j \Rightarrow M, s_k \models \varphi
\end{aligned}$$

Note that these definitions give us a “strong until”, that is, $E[\varphi U \psi]$ is true only if ψ eventually occurs. Further, note that we have used EG , EX and EU as an adequate set of temporal operators, following [41, 25]. The remaining temporal operators are defined in terms of these:

$$\begin{aligned}
A[\varphi U \psi] & \triangleq \neg E[\neg\psi U \neg\varphi \wedge \neg\psi] \wedge \neg EG\neg\psi & \text{def. of } AU \\
AX\varphi & \triangleq \neg EX\neg\varphi & \text{def. of } AX \\
AF\varphi & \triangleq A[\top U \varphi] & \text{def. of } AF \\
EF\varphi & \triangleq E[\top U \varphi] & \text{def. of } EF \\
AG\varphi & \triangleq \neg EF\neg\varphi & \text{def. of } AG
\end{aligned}$$

Alternatively, CTL operators can be described using their fixpoint formulations, as shown in Figure 5. This description is most useful for symbolic model-checking [46].

4.2 Multi-Valued Model-Checking

We proceed by defining multi-valued Kripke structures, which we call χ Kripke structures, and multi-valued CTL (χ CTL).

Definition 12 M is a χ Kripke structure if $M = (S, s_0, R, I, A, L)$, where:

- $L = (\mathcal{L}, \sqcap, \sqcup, \neg)$ is a quasi-boolean algebra, used for all mv-sets in the model;

- A is a (finite) set of atomic propositions that evaluate to elements of the algebra, L ;
- S is a (finite) set of states;
- $s_0 \in S$ is the initial state;
- $R : S \times S \rightarrow \mathcal{L}$ is the multi-valued transition relation;
- $I : S \rightarrow (A \rightarrow \mathcal{L})$ is a (total) labeling function that maps states in S into L -valued sets on A .

Intuitively, for any atomic proposition, $a \in A$, $(I(s))(a) = \ell$ means that the variable a has value ℓ in state s . Given an atomic proposition, $a \in A$, $I'_a : S \rightarrow \mathcal{L}$ is a (total) multi-valued characteristic function for an mv-set on S . I'_a is defined as follows:

$$I'_a \triangleq \lambda s \in S. (I(s))(a)$$

Note that a χ Kripke structure is a completely connected graph. As with classical model-checking, we ensure that all traces have infinite length by requiring that there is at least one non- \perp transition out of each state (if necessary, by adding a non- \perp self-loop to terminal states). To avoid clutter, when we present finite-state machines graphically, we follow the convention of not showing \perp transitions. An example χ Kripke structure **Sample**, shown in Figure 2, was introduced in Section 3. The multi-valued transition relation for this model is shown in Figure 4(a).

Multi-valued CTL, or χ CTL, is a temporal logic where CTL operators are extended to operate on a given χ Kripke structure over a given quasi-boolean algebra. We start defining χ CTL by giving the semantics of propositional operators. We use the double-brace notation, adopted from denotational semantics, and write $\llbracket \varphi \rrbracket$ to denote the mv-set of states representing the degree to which φ holds in a given state. Note that we have already used this notation when illustrating mv-sets in Section 3.

The semantics is as follows:

$$\begin{aligned} \llbracket a \rrbracket &\triangleq I'_a & \llbracket \neg \varphi \rrbracket &\triangleq \overline{\llbracket \varphi \rrbracket} \\ \llbracket \varphi \wedge \psi \rrbracket &\triangleq \llbracket \varphi \rrbracket \cap_L \llbracket \psi \rrbracket & \llbracket \varphi \vee \psi \rrbracket &\triangleq \llbracket \varphi \rrbracket \cup_L \llbracket \psi \rrbracket \end{aligned}$$

We proceed by defining the EX operator. In classical CTL, this operator is defined using existential quantification over next states. We extend the notion of existential quantification for multi-valued reasoning through the use of disjunction. This treatment of quantification is standard [8, 51]. The semantics of EX is:

$$\llbracket EX \varphi \rrbracket \triangleq \overset{\leftarrow}{\mathbb{R}} (\llbracket \varphi \rrbracket) \quad \text{def. of } EX$$

For example, $\llbracket EX c \rrbracket$ for the model **Sample** is shown in Figure 4(c), i.e., this property has value FF in state s_1 , TT in s_2 and s_3 , and TF in state s_0 .

We further define EG and EU using the EG and EU fixpoint properties of classical CTL:

$$\begin{aligned} \llbracket EG \varphi \rrbracket &\triangleq \nu Z. \llbracket \varphi \rrbracket \cap_L \llbracket EX Z \rrbracket & \text{def. of } EG \\ \llbracket E[\varphi U \psi] \rrbracket &\triangleq \mu Z. \llbracket \psi \rrbracket \cup_L (\llbracket \varphi \rrbracket \cap_L \llbracket EX Z \rrbracket) & \text{def. of } EU \end{aligned}$$

Then, $A[\varphi U \psi]$ becomes

$$\llbracket A[\varphi U \psi] \rrbracket \triangleq \overline{\llbracket E[\neg \psi U \neg \varphi \wedge \neg \psi] \rrbracket} \cap_L \overline{\llbracket EG \neg \psi \rrbracket} \quad \text{def. of } AU$$

Name	Description	Example
MvSet	an multi-valued set	see Figure 3
MvRelation	same as MvSet	see Figure 4(a)
Var	an atomic proposition	$c, 1, a$
Val	a logic value	TF, TT, <i>null</i>
Vector	a vector of type Val	(TF, FT, FF), (<i>null</i> , TT, <i>null</i>)
Oper	a logic operator	\wedge, \vee, \neg
CTL	a CTL expression	$AG\ c$

Table 1: Types used by the mv-set library.

and the remaining χ CTL operators are defined as their classical counterparts:

$$\begin{aligned}
\llbracket AX\varphi \rrbracket &\triangleq \overline{\llbracket EX\neg\varphi \rrbracket} && \text{def. of } AX \\
\llbracket AF\varphi \rrbracket &\triangleq \llbracket A[\top\ U\ \varphi] \rrbracket && \text{def. of } AF \\
\llbracket EF\varphi \rrbracket &\triangleq \llbracket E[\top\ U\ \varphi] \rrbracket && \text{def. of } EF \\
\llbracket AG\varphi \rrbracket &\triangleq \overline{\llbracket EF\neg\varphi \rrbracket} && \text{def. of } AG
\end{aligned}$$

Under the above definitions, χ CTL operators reduce to their corresponding CTL counterparts when defined over the classical logic **2**, AX and EX are monotone, and the multi-valued model-checking problem is decidable [20]. Its complexity is $O(|M| \times |S| \times |\mathcal{J}(L)| \times |\varphi|)$, where $|M| = |S| + |R|$ is the size of the χ Kripke structure M , $\mathcal{J}(L)$ is the set of join-irreducible elements of the algebra L , and φ is the χ CTL property to be checked [20]. Fair χ CTL and model-checking over it are also defined in [20].

Theorem 2 [20] *Fixpoint properties of (derived) χ CTL operators are the same as for CTL operators. That is,*

- (1) $\llbracket AG\varphi \rrbracket = \nu Z. \llbracket \varphi \rrbracket \cap_L \llbracket AXZ \rrbracket$ *AG* fixpoint
- (2) $\llbracket AF\varphi \rrbracket = \mu Z. \llbracket \varphi \rrbracket \cup_L \llbracket AXZ \rrbracket$ *AF* fixpoint
- (3) $\llbracket EF\varphi \rrbracket = \mu Z. \llbracket \varphi \rrbracket \cup_L \llbracket EXZ \rrbracket$ *EF* fixpoint
- (4) $\llbracket A[\varphi\ U\ \psi] \rrbracket = \mu Z. \llbracket \psi \rrbracket \cup_L (\llbracket \varphi \rrbracket \cap_L \llbracket AXZ \rrbracket)$ *AU* fixpoint

Thus, χ CTL can be checked symbolically, using mv-set operations. The remainder of the paper gives the high-level overview of the mv-set library and evaluates choices for decision-diagram implementation of mv-set operators.

5 Implementing a Symbolic χ CTL Model-Checker

Symbolic model checkers for boolean logic (e.g. [25, 16]) are naturally extended to the multi-valued case. Figure 6 shows the high-level implementation of our multi-valued model-checker χ Chck.

Table 1 lists types used in the algorithm and its underlying mv-set library. Type **Oper** represents binary or unary logic operator, and **CTL** represents a CTL expression. **Var** represents an atomic proposition (a variable), where we adopt the following conventions: (a) for every variable v in the source state there is a

```

function checkCTL(CTL  $p$ ) : MvSet
  case  $p \in A$ :      return projection( $p$ )
  case  $p = \neg\varphi$ :   return ptwiseApply( $\neg$ , checkCTL( $\varphi$ ))
  case  $p = \varphi \wedge \psi$ : return ptwiseApply( $\wedge$ , checkCTL( $\varphi$ ), checkCTL( $\psi$ ))
  case  $p = \varphi \vee \psi$ : return ptwiseApply( $\vee$ , checkCTL( $\varphi$ ), checkCTL( $\psi$ ))
  case  $p = EX \varphi$ :   return backwardImg(checkCTL( $\varphi$ ),  $Rel$ )
  case  $p = E[\varphi U \psi]$ : return checkEU(checkCTL( $\varphi$ ), checkCTL( $\psi$ ))
  case  $p = EG \varphi$ :  return checkEG(checkCTL( $\varphi$ ))

function checkEG(MvSet  $\varphi$ ) : MvSet
  result0 = constant( $\top$ )
  do
    result $i+1$  = ptwiseApply( $\wedge$ ,  $\varphi$ , backwardImg(result $i$ ,  $Rel$ ))
  while result $i$   $\neq$  result $i+1$ 
  return result $i$ 

function checkUntil(MvSet  $\varphi$ , MvSet  $\psi$ ) : MvSet
  result0 = constant( $\perp$ )
  do
    result $i+1$  = ptwiseApply( $\vee$ ,  $\psi$ , ptwiseApply( $\wedge$ ,  $\varphi$ , backwardImg(result $i$ ,  $Rel$ )))
  while result $i$   $\neq$  result $i+1$ 
  return result $i$ 

function buildTransitionRelation(Relation  $\mathbb{R}$ ) : MvRelation
  MvSet trans = constant( $\perp$ )
  foreach edge  $e = (s, t, \ell) \in \mathbb{R}$ 
    newEdge = point( $s, t, \ell$ )
    trans = ptwiseApply( $\vee$ , trans, newEdge)
  return trans

function ModelCheck(CTL  $p$ , Vector  $s_0$ , Relation  $\mathbb{R}$ ) : Val
  Global Var[] SourceVars
  Global Var[] DestVars = variableMap(SourceVars)
  Global MvRelation  $Rel$  = buildTransitionRelation( $\mathbb{R}$ )
  return cofactor(checkCTL( $p$ ),  $s_0$ )

```

Figure 6: The model-checking algorithm.

corresponding variable v' in the destination state, and (b) there exists a variable ordering, such that each variable has a unique index in it. Function `variableMap(Var[] source)` receives names of variables in the source state and returns names of variables in the destination state. For the purposes of this presentation, we do not distinguish between a variable name and its index in the variable ordering. Type `Val` represents logic values, with a special value `null` used to denote a “don’t care” value. `Vector` is a vector of elements

of type `Val`, such that for any `Vector` \vec{s} and any atomic proposition v , both $\vec{s}[v]$ and $\vec{s}[v']$ are defined. In this and all other examples in this section, we assume a variable ordering $a < b < c < a' < b' < c'$. For example, state s_1 in **Sample** (see Figure 2) under our variable ordering is represented as `(TT, FF, FF, null, null, null)`. That is, $s_1[b] = \text{FF}$, whereas $s_1[b'] = \text{null}$. To avoid clutter, we often remove trailing `nulls`, so s_1 is denoted as `(TT, FF, FF)`. We also use the name of a state interchangeably with a vector of values of its variables.

`MvSet` is a type representing an mv-set, i.e., a mapping between vectors and their values. Type `MvRelation` represents an mv-relation between two mv-sets. Its representation is the same as that of `MvSet`, but we use a different name for clarity of the presentation.

We further assume that the transition relation \mathbb{R} is represented as a set of *edges* (\vec{s}, \vec{t}, ℓ) , indicating that an edge between states \vec{s} and \vec{t} has value ℓ . For example, a TT transition between s_1 and s_0 under our variable ordering is represented as `((TT,FF,FF), (FF,FF,FF), TT)`, or $\mathbb{R}(s_1, s_0) = \text{TT}$.

The algorithm `ModelCheck` receives a χ CTL formula p to verify, the transition relation \mathbb{R} of the χ Kripke structure, and its initial state s_0 . `ModelCheck` assumes that the list of variables `SourceVars` is available and computes, for each variable $v \in \text{SourceVars}$, its associated $v' \in \text{DestVars}$ using a function `variableMap`. It proceeds by computing *Rel* – an `MvRelation` representation of the transition relation. Then it calls the function `checkCTL` which computes the mv-set representing the partition of the state-space with respect to p . `checkCTL` uses *Rel* for the computation of `backwardImage`. Finally, `ModelCheck` constrains the result to return the value that p has in the initial state s_0 of the system.

Function `checkCTL` recursively goes through the structure of the χ CTL property p , associating each subformula of p with an mv-set representing the value that p has in each state of the χ Kripke structure. `checkEG` and `checkEU` carry out the fix-point computations of the corresponding χ CTL operators, as defined in Section 4.1.

Functions of the mv-set library are discussed below.

- Function `constant(Val ℓ)` returns an mv-set \mathbb{S} where membership of each element is ℓ : $\forall \vec{x} \cdot \mathbb{S}(\vec{x}) = \ell$. For example, `constant(TF)` returns an mv-set where $\forall \vec{x} \cdot \mathbb{S}(x) = \text{TF}$.
- Function `projection(Var v)` creates an mv-set \mathbb{S} where membership of each element is determined by the value of variable v . Suppose the universe of discourse for mv-sets is the set of assignments to two variables, q and r , of type `2x2`; elements are shown as ordered pairs (x, y) , with x the assignment to q and y the assignment to r . Then `projection(r)` returns an mv-set where the pair (x, y) has value y .
- Function `point(Vector \vec{x} , Val ℓ)` returns an mv-set \mathbb{S} in which \vec{x} has membership ℓ , and all other elements have membership \perp : $\forall \vec{y} \cdot \text{if } \mathbb{S}(\vec{y}) = \vec{x} \text{ then } \ell \text{ else } \perp$. For example, `point((TT,TF,TF), (FT), restricted to the universe of variables a, b and c)` is an mv-set where `(TT,TF,TF)` has value `TF`, while the remaining $4^3 - 1$ states have value `FF`.
- Function `point(Vector \vec{x} , Vector \vec{y} , Val ℓ)` computes a representation for one edge. For example, to compute a representation of a TT transition between s_1 and s_0 , we call `point((TT,FF,FF), (FF,FF,FF), TT)` which produces a vector $\vec{x} = (\text{TT}, \text{FF}, \text{FF}, \text{FF}, \text{FF}, \text{FF})$ representing values of source variables in s_1 together with values of source variables in s_0 moved to the destination position. `point(\vec{x} , TT)` then computes the desired mv-set.

- Function `ptwiseApply` (`Oper op`, `MvSet S`) returns an mv-set \mathbb{T} s.t. $\forall \vec{x} \cdot \mathbb{T}(\vec{x}) = op\ S(\vec{x})$. For example, result `ptwiseApply`(\neg , $\llbracket c \rrbracket$) is shown in Figure 3(d). Note that we abuse notation somewhat, using $\llbracket \varphi \rrbracket$ as an mv-set and as its `MvSet` representation. Membership of each element in the mv-set returned by `ptwiseApply` (`Oper op`, `MvSet S`, `MvSet T`) is determined by a pairwise application of `op` to each element of \mathbb{S} and \mathbb{T} . See Figure 3(c) for a result of `ptwiseApply`(\wedge , $\llbracket c \rrbracket$, $\llbracket a \rrbracket$).
- Function `cofactor`(`MvSet S`, `Vector x`) evaluates \mathbb{S} on either a total or partial assignment, specified by \vec{x} , to its variables. For every variable v , $\vec{x}[v]$ is either *null*, indicating a “don’t-care”, or an algebra value. The `cofactor` function iterates over each index i from 0 to $|\vec{x}|$, and either does nothing (in the case of a *null*) or reduces \mathbb{S} by restricting variable i to $\vec{x}[i]$ by calling the auxiliary function `cofactor_var`(`MvSet S`, `Var v`, `Val l`). This function returns an mv-set which is the same as \mathbb{S} except that variable v ’s value is fixed to be ℓ . For instance, if \mathbb{S} is `projection`(r) (with variables q and r as in the previous example), `cofactor_var`(\mathbb{S} , q , `TT`) is \mathbb{S} again, since \mathbb{S} did not depend on q ; whereas `cofactor_var`(\mathbb{S} , r , `TF`) is the constant function `TF`, since fixing the value of r completely determines the value of \mathbb{S} at any point. Thus, for the model **Sample**, `cofactor`($\llbracket c \rrbracket$, s_2) = `TF`. Recall that s_2 is represented by the vector (`FF`, `TT`, `TF`) (for assignments to a , b , c respectively), and $\llbracket c \rrbracket$ is the mv-set produced by `projection`(c).
- Function `backwardImg` (`MvSet S`, `MvRelation R`) returns an mv-set corresponding to $\overleftarrow{\mathbb{R}}(\mathbb{S})$ in Definition 11.

Function `buildTransitionRelation` of the model-checker uses the mv-set library to construct the transition relation Rel in an `MvRelation` form. Rel is a disjunction of edges, where the mv-set for each edge is a conjunction of mv-sets representing each state and the value of the edge. For example, computation of `buildTransitionRelation` on **Sample** shown in Figure 4(a). Suppose we are checking a property $EX\ c$ on **Sample**. `checkCTL` returns $\overleftarrow{\mathbb{R}}(\llbracket c \rrbracket)$, shown in Figure 4(c). `cofactor`, called on this result and s_0 , the initial state of the **Sample**, yields `TF`, which is returned to the user.

	Boolean-terminal	Multi-terminal
Boolean branching factor	BDD-vector	ADD
Multi-valued branching factor	MBTDD-vector	MDD

Table 2: Choices of Decision Diagram Packages for Implementing mv-sets.

In the remainder of the paper, we discuss alternatives for implementing the mv-set library. The library is built on top of several decision diagram packages. The choices we face with these is whether each node has two successors (boolean) or many successors (multi-valued), and whether the mv-set membership function is represented by a single multi-valued functions or by a collection of boolean functions. The first choice is referred to as *the branching factor*, and the second – deciding on the number of terminal nodes: *boolean-terminal* or *multi-terminal*. These choices are summarized in Table 2. The entries of the tables are the names of decision diagram packages supporting this implementation. For example, MDDs are multi-valued multi-terminal decision diagrams. When *boolean-terminal* diagrams are used, the mv-set membership function must be encoded by a collection of diagrams. We refer to this collection as a *decision diagram vector*. For example, BDD-vector refers to a representation where a collection of *binary decision diagrams* is used to encode a single mv-set membership function.

All four varieties of decision diagrams listed in Table 2 have been proposed in the literature: MDDs were first described by Srinivasan et al [58]. They included MBTDDs as a special case, but these are discussed in more detail by Sasao and Butler [54]. ADDs were proposed by Bahar et al [3] (these are also known under the name MTBDDs [33]), and BDDs were introduced by Akers [1] and later by Bryant [14], who suggested the added properties of reducedness and orderedness to guarantee canonicity. Technology for the diagrams with a branching factor of 2, i.e., ADDs and BDDs, is very mature and is supported by standard libraries such as CUDD [57]. It has also been shown in [47] that the difference between a decision diagram vector and a multi-terminal decision diagram is reducible to the variable ordering problem for them. There is theoretical analysis [30, 43] of the trade-offs between different decision diagram variants for functional representation, largely for boolean and numeric functions. This analysis depends on classifying functions based on the effectiveness of the use of a given representation; without first generalizing such classification to multiple-valued functions, and then determining *a priori* which classes of functions are going to be used most, it is impossible for us to apply this analysis directly. We therefore set out to define and evaluate performance of multi-valued model-checking over all four data structures to identify trade-offs between them in the model-checking domain. The rest of the paper discusses multi-terminal decision diagram representation in Section 6 and boolean-terminal representation in Section 7.

6 Mv-set Implementation Using MDDs

Multi-Valued Decision Diagrams (MDDs) provide a viable option for implementing mv-sets and thus multi-valued model-checking. In this section we review definitions and properties of MDDs and discuss the implementation of mv-sets using these.

6.1 Multi-Valued Decision Diagrams

There is an extensive literature dealing with MDDs [58], mostly in the field of circuit design. To our knowledge, the logics used in that literature are given by total orders (such as the integers modulo n) and not by arbitrary quasi-boolean lattices, but we concede that this is a minor difference. Also, as far as we know, they have not been used in formal verification before, so we briefly describe them below. We assume a basic knowledge of binary decision diagrams (BDDs) [15].

The central notion in the construction of BDDs is the Shannon expansion. A boolean function f of n variables can be expressed relative to a variable a_0 , by computing f on $n-1$ variables with a_0 set to \top , and the same function with a_0 set to \perp . These functions are defined as $f|_{a_0=\top}$ and $f|_{a_0=\perp}$, respectively. We write this expansion as $f(a_0, \dots, a_{n-1}) \rightarrow f|_{a_0=\top}(a_1, \dots, a_{n-1}), f|_{a_0=\perp}(a_1, \dots, a_{n-1})$, and generalize it as follows:

Definition 13 [58] *Given a finite domain D , the generalized Shannon expansion of a function $f : D^n \rightarrow D$, with respect to the first variable in the ordering, is:*

$$f(a_0, a_1, \dots, a_{n-1}) \rightarrow f_0, \dots, f_{|D|-1}$$

where $\forall d_i \in D \cdot f_i = f|_{a_0=d_i}$ is the function obtained by substituting a literal d_i for a_0 in f . These functions are called cofactors.




Figure 7: Representation of function $f = x \wedge y$ in algebra **3**: (a) using MDD; (b) using ADD.

Definition 14 Assuming a finite set D and an ordered set of variables A , a multi-valued decision diagram (MDD) is a tuple $(V, E, \text{var}, \text{child}, \text{value})$ where:

- $V = V_t \cup V_n$ is a set of nodes, where V_t and V_n indicate a set of terminal and non-terminal nodes, respectively;
- $E \subseteq V_n \times V$ is a set of directed edges;
- $\text{var} : V_n \rightarrow A$ is a variable labeling function.
- $\text{child} : V_n \rightarrow D \rightarrow V$ is an indexed successor function for nonterminal nodes;
- $\text{value} : V_t \rightarrow D$ is a total function that maps each terminal node to a value of the algebra.

We describe constraints on the elements of an MDD below. Although D may be any finite set, we are interested only in lattices; so instead of D , we refer to the elements of a quasi-boolean algebra L . Note that in general we do not distinguish between a single node in an MDD and the subgraph rooted from it, referring to both indiscriminately as u .

Consider the function $f = x \wedge y$, with $\ell_0 = F, \ell_1 = M, \ell_2 = T$. The MDD for this expression is shown in Figure 7(a). The diagram is constructed by Shannon expansion, first with respect to x , and then (for each cofactor of f) with respect to y . The dashed arrows point to initial states of MDDs representing f and its cofactors. For example, the MDD rooted at a node with label x represents f . When $x = F$, f evaluates to F . Thus, $f_0 = F$. f_1 corresponds to setting $x = M$, and is indicated in Figure 7(a) by a dashed arrow that points to the left subtree rooted at y . This means that the value of f is further determined by y : when $y = F$, corresponding to $f_{1,0}$ cofactor, f has value F ; otherwise, i.e., for cofactors $f_{1,1}$ and $f_{1,2}$, f has value M .

Definition 15 *The function computed by an MDD u is denoted $f^u : L^n \rightarrow L$, and is defined recursively as follows:*

$$\begin{aligned} u \in V_t &\Rightarrow f^u(s_0, \dots, s_{n-1}) = \text{value}(u) && \text{terminal constants} \\ u \in V_n &\Rightarrow f^u(s_0, \dots, s_{n-1}) = f^{\text{child}_{s_i}(u)}(s_0, \dots, s_{n-1}), \\ &\quad \text{where } a_i = \text{var}(u) \text{ and } \vec{s} \in L^n && \text{cofactor expansion} \end{aligned}$$

Consider the MDD in Figure 7(a). To compute $f = x \wedge y$ with $x = T$ and $y = M$ using this diagram, we want to find $f(\vec{s})$ where $\vec{s} = (T, M)$. We begin at the root node. Its var is x , so we choose the first argument, which is T , and descend to the node $\text{child}_T(f)$, indicated by the arrow to f_2 (which represents the function $T \wedge y$). Now we compute $f_2(M)$ by choosing $\text{child}_M(f_2)$, which is a node in V_t , so we stop and return M . Thus, we conclude that $f(T, M) = M$. We further define size of an MDD:

$$\text{size}(u) \triangleq |\{v \mid v \text{ is reachable from } u\}|$$

For the function f in Figure 7(a), $\text{size}(f) = 6$.

The following properties hold for all MDDs:

$$\begin{aligned} \forall u_0 \in V_n \cdot \text{out}(u_0) = |L| \quad \wedge \quad \forall u_1 \in V_t \cdot \text{out}(u_1) = 0 && \text{semantics of nodes} \\ \forall u_0, u_1 \in V \cdot \exists \ell \in L \cdot (u_0, u_1) \in E \Rightarrow \text{child}_\ell(u_0) = u_1 && \text{semantics of edges} \end{aligned}$$

where $\text{out}(u)$ stands for the number of children of u . Several further properties are required for the data structure to be usable:

$$\begin{aligned} \forall u_0, u_1 \in V_n \cdot (u_0, u_1) \in E \wedge \text{var}(u_0) = a_i \wedge \text{var}(u_1) = a_j \Rightarrow i < j && \text{orderedness} \\ \forall u_0, u_1 \in V \cdot f^{u_0} = f^{u_1} \Rightarrow u_0 = u_1 && \text{reducedness} \\ \forall u_0, u_1 \in V_n \cdot ((\text{var}(u_0) = \text{var}(u_1)) \wedge \\ (\forall \ell \in L \cdot \text{child}_\ell(u_0) = \text{child}_\ell(u_1))) \Rightarrow u_0 = u_1 && \text{uniqueness 1} \\ \forall u_0, u_1 \in V_t \cdot (\text{value}(u_0) = \text{value}(u_1)) \Rightarrow u_0 = u_1 && \text{uniqueness 2} \end{aligned}$$

In general, the efficiency of decision diagrams, binary or multi-valued, comes from the properties of reducedness and orderedness (defined above). Orderedness is also required for termination of many algorithms on the diagrams. Uniqueness implies reducedness [58] – MDDs are unique by construction, and thus reduced.

MDDs have the same useful property as BDDs: given a variable ordering, there is precisely one MDD representation of a function. This allows for constant-time checking of function equality.

Theorem 3 Canonicity [58] *For any quasi-boolean algebra L , any nonnegative integer n , and any function $f : L^n \rightarrow L$, there is exactly one reduced ordered MDD u such that $f^u = f(a_0, \dots, a_{n-1})$.*

We now formalize a notion of *partial assignment* which is used to express partial evaluation of a function or a decision diagram.

Definition 16 *Let A be a set of variables, and L be an algebra. Then a partial assignment $\rho \in 2^{A \times L}$ is a set of pairs $\rho_i = (a_i, \ell_i)$ where $i \neq j \Rightarrow a_i \neq a_j$. A function f may be partially evaluated, given a partial assignment:*

$$f|_\rho = f|_{a_1=\ell_1, \dots, a_{|\rho|}=\ell_{|\rho|}}$$

This is a cofactor of f , with more than a single variable assigned.

```

function constant(Val  $\ell$ ) : MvSet
    return createMvSetFromMDD(MDDconstant( $\ell$ ))

function ptwiseApply (Oper  $op$ , MvSet  $\mathbb{S}$ , MvSet  $\mathbb{T}$ ) : MvSet
    return createMvSetFromMDD(MDDptwiseApply( $op$ ,  $\mathbb{S}$ .dd,  $\mathbb{T}$ .dd))

function point (Vector  $\vec{s}$ , Vector  $\vec{t}$ , Val  $\ell$ ) : MvSet
    return ptwiseApply( $\wedge$ , point( $\vec{s}$ ,  $\top$ ), exchange(point( $\vec{t}$ ,  $\ell$ ), SourceVars, DestVars))

function backwardImg (MvSet  $\mathbb{S}$ , MvRelation  $\mathbb{R}$ ) : MvSet
    return existAbstract (ptwiseApply ( $\wedge$ ,  $\mathbb{R}$ , exchange ( $\mathbb{S}$ , SourceVars, DestVars), DestVars))

Global Var[] SourceVars, DestVars

```

Figure 8: Partial implementation of the mv-set library using MDD operations.

For example, given a function $f(x, y) = x \wedge y$ in Figure 7(a) and a partial assignment $\rho = \{(x, \top)\}$, the partial evaluation of f is $f|_{\rho} = f|_{x=\top} = f(\top, y) = \top \wedge y = y$.

A partial assignment ρ' is *more complete than* ρ if $\rho \subseteq \rho'$; the most complete such assignment is called *total*. The total assignment for a function provides a value for every variable in this function. Thus, an assignment $\rho = \{(x, \perp), (y, \top)\}$ is more complete than $\rho' = \{(x, \perp)\}$, and is a total assignment for any function on two variables. A partial assignment ρ is said to be *unambiguous* if $f|_{\rho}$ is constant, and is said to be *minimally unambiguous* if, for any $\rho' \subset \rho$, $f|_{\rho'}$ is non-constant. For example, for a function $f(x, y) = x \wedge y$ in Figure 7(a), a partial assignment $\rho = \{(x, \perp), (y, \top)\}$ is unambiguous, since $f|_{\rho} = f|_{x=\perp, y=\top} = \perp$; and a partial assignment $\rho' = \{(x, \perp)\}$ is minimally unambiguous, since $f|_{\rho'} = f|_{x=\perp} = \perp$. An unambiguous assignment for a function does not need to provide values for each variable of this function; thus, every total assignment is an unambiguous assignment, whereas the converse is not true. Note that in a decision diagram, a minimally unambiguous partial assignment corresponds to a path from the root of the diagram to a terminal node.

6.2 Implementation of mv-sets

Here we discuss the implementation of the mv-set library on top of MDDs. We begin by defining the mv-set layer, and then describe the direct implementation of the MDD package in Section 6.2.2 and the implementation of MDDs on top of the standard ADD library [3] in Section 6.2.3.

6.2.1 The mv-set library on top of MDDs

Some of the functions of the mv-set library implementation on top of MDDs are given in Figure 8. We assume that each MvSet \mathbb{S} has a field \mathbb{S} .dd, for the decision diagram representation of this mv-set, and a function createMvSetFromMDD that allows to convert between the two representations. We use function names prefixed with “MDD” to indicate MDD library functions. As seen in Figure 8, there is a one-to-one correspondence between MDD and mv-set library functions constant and ptwiseApply. The same correspondence also holds for other mv-set functions, such as projection and cofactor. Two

functions of the mv-set library, a 3-parameter `point` and `backwardImg`, are defined using low-level routines `exchange` and `existAbstract`:

- Function `exchange(MvSet S, Var[] oldVars, Var[] newVars)` takes a decision diagram representation of an mv-set S and replaces labels of variables from `oldVars` into those of `newVars`.
- Function `existAbstract(MvSet S, Var[] vars)` computes $\exists \text{vars}[0], \text{vars}[1], \dots \cdot S$ – the existential quantification used in the computation of backward and forward image.

`exchange` and `existAbstract` correspond one-to-one to their MDD library counter-parts.

The function `cofactor` performs partial evaluation of an mv-set S . A partial assignment ρ is represented by a `Vector` \vec{s} : for any variable v , if $(v, \ell) \in \rho$ then $\vec{s}[v] = \ell$; if v is not assigned any value by ρ , then $\vec{s}[v] = \text{null}$.

6.2.2 MDD library

In general, algorithms for manipulating BDDs are easily extensible to the multi-valued case, provided they do not use any optimizations that depend on a two-valued boolean logic (e.g. complemented edges [57]). The public methods required for implementation of the mv-set library are: `MDDconstant`, `MDDprojection` and `MDDpoint`, to construct an MDD from an mv-set; `MDDptwiseApply`, to compute \wedge, \vee and \neg of MDDs; `MDDexistAbstract`, to existentially quantify over variables; and `MDDcofactor` to compute a cofactor. These functions use a private method `makeUnique`, defined in Figure 9, to guarantee uniqueness, and thus reducedness [58], of MDD diagrams.

`MDDconstant`, `MDDprojection` and `MDDpoint` ensure orderedness of MDDs while they are being constructed, and the other functions preserve it. `MDDptwiseApply` is shown in Figure 9. Note that it interfaces with the lattice library by calling the method `Lattice.doOp` to compute \wedge or \vee of two terminal nodes.

6.2.3 MDDs on Top of ADDs

Algebraic Decision Diagrams (ADDs) [3] are multi-terminal decision diagrams where each node has a boolean branching factor. In model-checking, ADDs are typically used for representing enumerated types, and are part of the CUDD decision diagram library [57].

Encoding of an MDD into an ADD follows the same algorithm as encoding of enumerated types and integer subranges into BDDs [46]. An ADD for a multi-valued algebra L has up to $|L|$ terminal nodes, one for each element of the algebra. To turn a multi-valued variable into a boolean variable, we can use the concept of join-irreducibility. Recall from Section 2 that every element of an algebra can be represented as a set of join-irreducibles of the corresponding lattice. For example, the join-irreducibles of the algebra $\mathbf{3}$ are M and T , so each element of the algebra can be uniquely represented as follows: $F = \emptyset$, $M = \{M\}$ and $T = \{M, T\}$. The characteristic function for this representation allows to encode each value $\ell \in L$ using $|\mathcal{J}(L)|$ boolean variables, one for each join-irreducible. For example, each 3-valued variable x can be encoded using a pair of variables x_0 and x_1 : $x = F$ is represented as (\perp, \perp) and (\perp, \top) , $x = M$ as (\top, \perp) , and $x = T$ as (\top, \top) . With this encoding, the ADD corresponding to the MDD in Figure 7(a) is shown in Figure 7(b).

```

function makeUnique(Var name, MDD[] children) : MDD
  find (create if not found) a node  $u$  s.t.
     $\text{var}(u) = \text{name} \wedge \forall \ell \cdot \text{child}_\ell(u) = \text{children}(\ell)$ 
  return  $u$ 

function MDDptwiseApply(Oper op, MDD  $u_1$ , MDD  $u_2$ ) : MDD
// applies the logic operation op to the MDDs  $u_1$  and  $u_2$ 
  Global int[ $u_1$ ][ $u_2$ ]  $G$ 
  return apply'(op,  $u_1$ ,  $u_2$ )

function apply'(Oper op, MDD  $u_1$ , MDD  $u_2$ ) : MDD
// helper function for MDDptwiseApply which actually does the work
  if  $G[u_1][u_2]$  non-empty
    then return  $G[u_1][u_2]$ 
  else
    if  $u_1 \in L \wedge u_2 \in L$ 
      then  $u = \text{Lattice.doOp}(u_1, u_2, \text{op})$ 
    else if  $\text{var}(u_1) = \text{var}(u_2)$ 
      then foreach  $\ell \in L$ 
         $\text{children}(\ell) = \text{apply}'(\text{op}, \text{child}_\ell(u_1), \text{child}_\ell(u_2))$ 
         $u = \text{makeUnique}(\text{var}(u_1), \text{children})$ 
    else if  $\text{var}(u_1) < \text{var}(u_2)$ 
      then foreach  $\ell \in L$ 
         $\text{children}(\ell) = \text{apply}'(\text{op}, \text{child}_\ell(u_2), u_2)$ 
         $u = \text{makeUnique}(\text{var}(u_1), \text{children})$ 
    else
      foreach  $\ell \in L$ 
         $\text{children}(\ell) = \text{apply}'(\text{op}, u_1, \text{child}_\ell(u_2))$ 
         $u = \text{makeUnique}(\text{var}(u_1), \text{children})$ 
     $G[u_1][u_2] = u$ 
  return  $u$ 

```

Figure 9: The MDD algorithm for MDDptwiseApply for binary operators, together with its helper functions. MDDptwiseApply for unary operators is defined similarly.

The number of variables in an ADD is $\mathcal{J}(L)$ times larger than that of the corresponding MDD. Clearly, other encodings of MDDs into ADDs are also possible, with the ratio between the number of variables in an ADD and the one in the corresponding MDD ranging from $\lceil \log_2(|L|) \rceil$ to $|L|$.

7 mv-set Implementation Using MBTDDs

In this section, we discuss alternatives for encoding operations on multi-valued sets using multi-valued branching boolean-terminal diagrams (MBTDDs). The major advantages of this approach over the one

described in Section 6 is that for a fixed algebra L , all operations where one operand is a constant take $O(1)$ time, and that the model-checking algorithm can potentially be distributed over $|\mathcal{J}(L)|$ machines.

The naive approach is to represent a multi-valued set as a collection of classical sets, one for each value of the algebra. For example, $\llbracket c \rrbracket$ (see Figure 3(a), Section 3) can be encoded using four classical sets as follows:

$$\begin{aligned} \{s_3\} & \quad // \text{ for value TT} \\ \{s_2\} & \quad // \text{ for value TF} \\ \emptyset & \quad // \text{ for value FT} \\ \{s_0, s_1\} & \quad // \text{ for value FF} \end{aligned}$$

Suppose we are interested in computing a union of this set and $\llbracket a \rrbracket$, shown in Figure 3(b). From the algebra we know that $a \vee b = \text{TT}$ iff $a = \text{TT}$, or $b = \text{TT}$, or both $a = \text{TF}$ and $b = \text{FT}$; thus the following is the computation of the TT set in $\llbracket c \rrbracket \vee \llbracket a \rrbracket$:

$$\begin{aligned} \{s_3\} \cup & \quad // \text{ TT set of } \llbracket c \rrbracket \\ \{s_1, s_3\} \cup & \quad // \text{ TT set of } \llbracket a \rrbracket \\ (\{s_2\} \cap \emptyset) \cup & \quad // \text{ TF set of } \llbracket c \rrbracket \text{ and FT set of } \llbracket a \rrbracket \\ (\emptyset \cap \emptyset) & \quad // \text{ FT set of } \llbracket c \rrbracket \text{ and TF set of } \llbracket a \rrbracket \end{aligned}$$

which results in $\{s_1, s_3\}$. Computations of the remaining sets are done similarly, driven by *inverse tables*, i.e., sets of values of operands that yield the desired value in the result. Our first implementation of the model-checker [23] was based on this encoding.

In Section 6.2.3, we used the fact that each element of the lattice can be uniquely represented as a subset of join-irreducibles to efficiently implement MDDs on top of ADDs. The same idea is used in the rest of this section: we first show how to use this encoding to perform the required mv-set functions and recover the complete mv-set, and then proceed with the implementation of the mv-set library.

7.1 Encoding mv-sets Using j -cuts

As indicated in Section 2, each element of a lattice can be encoded using the join-irreducible elements below it. We aim to use join-irreducibles to provide a “prime factorization” for mv-sets.

Lemma 1 [29] *Let j , x , and y be elements of a distributive lattice $(\mathcal{L}, \sqsubseteq)$, with j being a join-irreducible. Then $j \sqsubseteq x \sqcup y$ iff $j \sqsubseteq x$ or $j \sqsubseteq y$.*

Lemma 2 *Let ℓ be any element of a quasi-boolean algebra $(\mathcal{L}, \sqcap, sqcup, \neg)$. Then*

$$\{\neg x \mid x \in \mathcal{L}, \ell \sqsubseteq x\} = \{x \mid x \in \mathcal{L}, x \sqsubseteq \neg\ell\}$$

Proof:

$$\begin{aligned}
& y \in \{\neg x \mid x \in \mathcal{L}, \ell \sqsubseteq x\} \\
\Leftrightarrow & \text{double negation} \\
& \neg(\neg y) \in \{\neg x \mid x \in \mathcal{L}, \ell \sqsubseteq x\} \\
\Leftrightarrow & \text{set specification} \\
& \ell \sqsubseteq \neg y \\
\Leftrightarrow & \text{negation is order-reversing; double negation} \\
& y \sqsubseteq \neg \ell \\
\Leftrightarrow & \text{set specification} \\
& y \in \{x \mid x \in \mathcal{L}, x \sqsubseteq \neg \ell\}
\end{aligned}$$

□

The above two lemmas describe properties of join-irreducible elements that we exploit in our representation of mv-sets.

We then proceed to define up-sets and down-sets of lattice elements.

Definition 17 *Let ℓ be any element of a quasi-boolean lattice $(\mathcal{L}, \sqsubseteq)$. Then,*

$$\begin{aligned}
\uparrow \ell & \triangleq \{x \mid x \in \mathcal{L} \cdot \ell \sqsubseteq x\} && \text{up-set of } \ell \\
\downarrow \ell & \triangleq \{x \mid x \in \mathcal{L} \cdot x \sqsubseteq \ell\} && \text{down-set of } \ell
\end{aligned}$$

For example, for a lattice $\mathbf{2x2}$, $\uparrow \text{TF}$ is $\{\text{TF}, \text{TT}\}$, and $\downarrow \text{TF}$ is $\{\text{TF}, \text{FF}\}$.

Given an algebra L , we aim to encode mv-sets defined over L using a collection of j -cuts for $j \in \mathcal{J}(L)$. We then show that the j -cuts of an mv-set contain sufficient information to rebuild the original mv-set, and prove that all the necessary operations for symbolic model-checking can be easily carried out in this encoding.

Our first result, which follows trivially from Theorem 1, states that the membership degree of an element in an mv-set is the join of all join-irreducibles whose cuts (see Definition 9) contain the element.

Theorem 4 *For any mv-set \mathbb{S} over a quasi-boolean algebra L ,*

$$\mathbb{S}(x) = \bigsqcup \{j \mid j \in \mathcal{J}(L), x \in \uparrow_j(\mathbb{S})\}$$

For example, consider encoding $\llbracket c \rrbracket$ using the join-irreducible elements TF and FT of the algebra $\mathbf{2x2}$. $\uparrow_{\text{TF}}(\llbracket c \rrbracket) = \{s_2, s_3\}$, $\uparrow_{\text{FT}}(\llbracket c \rrbracket) = \{s_3\}$. Thus, $\llbracket c \rrbracket(s_3) = \text{TF} \sqcup \text{FT} = \text{TT}$, which corresponds to the value of s_3 in Figure 3(a).

Our goal is to define all operations over mv-sets in terms of operations over their j -cuts. We begin with multi-valued unions and intersections and show that the ℓ -cut of a multi-valued intersection is the intersection of the ℓ -cuts of the individual mv-sets. For $j \in \mathcal{J}(L)$, the j -cut of a multi-valued union is the union of the j -cuts.

Theorem 5 *Let L be a quasi-boolean algebra, $\ell \in L$ and $j \in \mathcal{J}(L)$. Then,*

$$\begin{aligned}
\uparrow_\ell(\mathbb{S} \cap_L \mathbb{T}) &= \uparrow_\ell(\mathbb{S}) \cap \uparrow_\ell(\mathbb{T}) && \text{cut-intersection} \\
\uparrow_j(\mathbb{S} \cup_L \mathbb{T}) &= \uparrow_j(\mathbb{S}) \cup \uparrow_j(\mathbb{T}) && \text{cut-union}
\end{aligned}$$

Proof:

We show both equations by set extensionality:

$x \in \uparrow_\ell(\mathbb{S} \cap_L \mathbb{T})$ \Leftrightarrow Definition 9 $\ell \sqsubseteq (\mathbb{S} \cap_L \mathbb{T})(x)$ \Leftrightarrow multi-valued intersection $\ell \sqsubseteq \mathbb{S}(x) \cap \mathbb{T}(x)$ \Leftrightarrow lattice properties $\ell \sqsubseteq \mathbb{S}(x)$ and $\ell \sqsubseteq \mathbb{T}(x)$ \Leftrightarrow Definition 9 $x \in \uparrow_\ell(\mathbb{S})$ and $x \in \uparrow_\ell(\mathbb{T})$ \Leftrightarrow set theory $x \in \uparrow_\ell(\mathbb{S}) \cap \uparrow_\ell(\mathbb{T})$	$x \in \uparrow_j(\mathbb{S} \cup_L \mathbb{T})$ \Leftrightarrow Definition 9 $j \sqsubseteq (\mathbb{S} \cup_L \mathbb{T})(x)$ \Leftrightarrow multi-valued union $j \sqsubseteq \mathbb{S}(x) \sqcup \mathbb{T}(x)$ \Leftrightarrow Lemma 1 $j \sqsubseteq \mathbb{S}(x)$ or $j \sqsubseteq \mathbb{T}(x)$ \Leftrightarrow Definition 9 $x \in \uparrow_j(\mathbb{S})$ or $x \in \uparrow_j(\mathbb{T})$ \Leftrightarrow set theory $x \in \uparrow_j(\mathbb{S}) \cup \uparrow_j(\mathbb{T})$
---	--

□

For the example in Figure 3(a)-(c),

$$\uparrow_{\text{TF}}(\llbracket c \rrbracket \sqcup \llbracket a \rrbracket) = \{s_1, s_2, s_3\} = \uparrow_{\text{TF}}(\llbracket c \rrbracket) \cup \uparrow_{\text{TF}}(\llbracket a \rrbracket) = \{s_2, s_3\} \cup \{s_1, s_3\}$$

Formulation of complementation uses two observations. First is a trivial corollary of Lemma 2, indicating that multi-valued negation turns each ℓ -cut into a $\neg\ell$ -clip of the complement of the original set:

Lemma 3 *Let L be a quasi-boolean algebra. For every $\ell \in L$,*

$$\uparrow_\ell(\mathbb{S}) = \downarrow_{\neg\ell}(\overline{\mathbb{S}})$$

For the example in Figure 3(a),(d), $\{s_2, s_3\} = \uparrow_{\text{TF}}(\llbracket c \rrbracket) = \downarrow_{\text{FT}}(\overline{\llbracket c \rrbracket})$.

Second, every up-set of a join-irreducible j is the complement (in the elements of the algebra) of the down-set of some meet-irreducible m :

Lemma 4 *Let L be a quasi-boolean algebra. Then, there exists a bijection $f : \mathcal{J}(L) \rightarrow \mathcal{M}(L)$ such that $\forall j \in \mathcal{J}(L) \cdot \uparrow j = L \setminus \downarrow f(j)$ and $\forall m \in \mathcal{M}(L) \cdot \downarrow m = L \setminus \uparrow f^{-1}(m)$.*

Proof:

Since $L = (\mathcal{L}, \sqcap, \sqcup, \neg)$ is a quasi-boolean algebra, $(\mathcal{L}, \sqsubseteq)$ is a finite distributive lattice and thus, by Proposition 9.4 of [29], there exists a bijection $g : \mathcal{J}(L) \rightarrow \{\downarrow m \mid m \in \mathcal{M}(L)\}$ given by $g(j) = \mathcal{L} \setminus \uparrow j$. Further, since $\sqcup : \{\downarrow m \mid m \in \mathcal{M}(L)\} \rightarrow \mathcal{M}(L)$ is a bijection, we obtain a bijection $f : \mathcal{J}(L) \rightarrow \mathcal{M}(L)$ where $f(j) = \sqcup(\mathcal{L} \setminus \uparrow j)$. It is easy to show that its inverse f^{-1} is given by $f^{-1}(m) = \sqcap(\mathcal{L} \setminus \downarrow m)$, where $m \in \mathcal{M}(L)$. Converting this back into logic terms, we get $f^{-1}(m) = \sqcap(L \setminus \downarrow m)$, where $m \in \mathcal{M}(L)$.

Now we show that $\uparrow j = L \setminus \downarrow f(j)$ and $\downarrow m = L \setminus \uparrow f^{-1}(m)$ by set extensionality:

$x \in L \downarrow f(j)$ \Leftrightarrow Definition 17 $x \not\subseteq f(j)$ \Leftrightarrow definition of f $x \not\subseteq \sqcup(L \uparrow j)$ \Leftrightarrow lattice properties and contraposition $x \notin L \uparrow j$ \Leftrightarrow double negation $x \in \uparrow j$	$x \in L \uparrow f^{-1}(m)$ \Leftrightarrow Definition 17 $f^{-1}(m) \not\subseteq x$ \Leftrightarrow definition of f^{-1} $\cap(L \downarrow m) \not\subseteq x$ \Leftrightarrow lattice properties and contraposition $x \notin L \downarrow m$ \Leftrightarrow double negation $x \in \downarrow m$
--	---

□

We can generalize the above result to mv-sets:

Corollary of Lemma 4 *Let \mathbb{S} be an mv-set over some universe U and quasi-boolean algebra L . Let $j \in \mathcal{J}(L)$ and $m \in \mathcal{M}(L)$. Then,*

$$\begin{aligned} \uparrow_j(\mathbb{S}) &= U \setminus \downarrow_{f(j)}(\mathbb{S}) && \text{cut of join-irreducible} \\ \downarrow_m(\mathbb{S}) &= U \setminus \uparrow_{f^{-1}(-m)}(\mathbb{S}) && \text{clip of meet-irreducible} \end{aligned}$$

We use the above Corollary to construct the j -cut of a multi-valued complement of an mv-set. The j -cut of the multi-valued complement of \mathbb{S} is the classical complement (in U) of the $f^{-1}(-j)$ cut of \mathbb{S} :

Theorem 6 *Let \mathbb{S} be an mv-set over some universe U and quasi-boolean algebra L , let $j \in \mathcal{J}(L)$, and $\text{neg}(j) \triangleq f^{-1}(-j)$. Then,*

$$\uparrow_j(\overline{\mathbb{S}}) = U \setminus \uparrow_{\text{neg}(j)}(\mathbb{S})$$

Proof:

$$\begin{aligned} &\uparrow_j(\overline{\mathbb{S}}) \\ = &\text{Lemma 3} \\ &\downarrow_{\neg j}(\overline{\mathbb{S}}) \\ = &\text{Involution} \\ &\downarrow_{\neg j}(\mathbb{S}) \\ = &\text{By Corollary of Lemma 4, since } \neg j \in \mathcal{M}(L), \text{ and by clip of meet-irreducible} \\ &U \setminus \uparrow_{f^{-1}(-j)}(\mathbb{S}) \\ = &\text{Definition of neg} \\ &U \setminus \uparrow_{\text{neg}(j)}(\mathbb{S}) \end{aligned}$$

□

For example, in logic **2x2**, neg is an identity; therefore,

$$\{s_0, s_1, s_2\} = \uparrow_{\text{FT}}(\overline{[c]}) = U \setminus \uparrow_{\text{FT}}([c]) = \{s_0, s_1, s_2, s_3\} \setminus \{s_3\}$$

These results allow us to represent each mv-set compactly as a family of its j -cuts. Theorem 4 guarantees that this representation of mv-sets preserves all information; it also tells us how to recover mv-set membership from the j -cuts.

The advantage of this representation is in the computation of multi-valued intersections and unions. If mv-sets were represented by families of pieces, i.e., indexed by all values of the algebra, then intersection and union would both take as many as $O(|L|^2)$ classical set operations. In contrast, the cut-intersection and the cut-union theorems imply that our representation requires just $O(|\mathcal{J}(L)|)$ classical set operations.

```

function constant(Val  $\ell$ ) : MvSet
  foreach  $j \in \mathcal{J}(L)$ 
    if ( $j \sqsubseteq \ell$ ) then result[ $j$ ] = MBTconstant( $\top$ )
    else result[ $j$ ] = MBTconstant( $\perp$ )
  return createMvSetFromMBTDD(result)

function projection(Var  $v$ ) : MvSet
  foreach  $j \in \mathcal{J}(L)$ 
    result[ $j$ ] = projectionElement( $v$ , UpSet[ $j$ ])
  return createMvSetFromMBTDD(result)

function projectionElement(Var  $v$ , Val[] UpSet) : MBTDD
  result0 =  $\perp$ 
  foreach  $u_j \in \text{UpSet}$ 
    result $i+1$  = MTBptwiseApply( $\vee$ , MBTpoint( $v := u_j, \top$ ), result $i$ )
  return result $i+1$ 

function point(Vector  $\vec{x}$ , Val  $\ell$ ) : MvSet
  foreach  $j \in \mathcal{J}(L)$ 
    if ( $j \sqsubseteq \ell$ ) then result[ $j$ ] = MBTpoint( $x, \top$ )
    else result[ $j$ ] = MBTconstant( $\perp$ )
  return createMvSetFromMBTDD(result)

function ptwiseApply (Oper  $op$ , MvSet  $\mathbb{S}$ , MvSet  $\mathbb{T}$ ) : MvSet
  foreach  $j \in \mathcal{J}(L)$ 
    result[ $j$ ] = MBTpointwiseApply( $op$ ,  $\mathbb{S}$ .dd[ $j$ ],  $\mathbb{T}$ .dd[ $j$ ])
  return createMvSetFromMBTDD(result)

function ptwiseApply (Oper  $\neg$ , MvSet  $\mathbb{S}$ ) : MvSet
  foreach  $j \in \mathcal{J}(L)$ 
    result[ $j$ ] = MBTpointwiseApply( $\neg$ ,  $\mathbb{S}$ .dd[neg[ $j$ ]])
  return createMvSetFromMBTDD(result)

Global Val[] UpSet = computeUp( $\mathcal{J}(L)$ )
Global Val[] neg

```

Figure 10: Implementation of the mv-set library using MBTDD operations.

7.2 Implementation of mv-sets

A single MBTDD computes a boolean-output function. By creating an MBTDD for each j -cut of a multi-valued function, we can represent that function as a *vector* of MBTDDs. These vectors and their semantics are defined formally as follows, using Theorem 4:

Definition 18 *Let \hat{u} be an MBTDD-vector indexed by the join-irreducibles of an algebra L . Each $\hat{u}[j]$*

defines a boolean-output function by the semantics of Definition 15. The function defined by the vector is computed as

$$f^{\hat{u}} = \bigsqcup_{j \in \mathcal{J}(L)} f^{\hat{u}[j]} \sqcap j$$

and conversely, for all j ,

$$f^{\hat{u}[j]} = \uparrow_j (f^{\hat{u}})$$

The size of an MBTDD-vector is defined as follows:

$$\text{size}(\hat{u}) = |\{v \mid v \text{ is reachable from } \hat{u}[j] \text{ for some } j \in \mathcal{J}(L)\}|$$

The mv-set library is implemented on top of MBTDDs as described in Figure 10. As in the MDD implementation, we assume that each `MvSet` \mathbb{S} has a field `\mathbb{S}.dd`, for a decision diagram representation of this mv-set. However, here `dd` is an *array* of decision diagrams, one for each j -cut of \mathbb{S} : it corresponds to \hat{u} from Definition 18. Function `createMvSetFromMBTDD` allows the construction of an mv-set from its decision diagram representation. In this figure, all function names prefixed with “MBT” are part of the MBTDD package. Functions `backwardImg` and `point` on three parameters are the same as shown in Figure 8 and thus are omitted here. In Figure 10, $\mathcal{J}(L)$ is an ordered list of join-irreducible elements of the algebra, and `UpSet` is the list of up-sets of join-irreducibles, indexed by $\mathcal{J}(L)$:

$$\text{UpSet}[j] \triangleq \uparrow j = \{\ell \in L \mid j \sqsubseteq \ell\}$$

`UpSet` is computed by a function `computeUp`. We use the notation $v \vec{:=} \ell$ to denote an element of type `Vector` corresponding to a partial assignment ρ consisting of a single tuple (v, ℓ) .

Additional functions, such as `cofactor`, `exchange` and `existAbstract` follow the same pattern as that of `ptwiseApply` in Figure 10. Note that all operations are performed w.r.t. j -cuts of mv-sets. Handling of \wedge and \vee follows from Theorem 5. The same theorem also serves as proof of correctness for `existAbstract`, since it is a disjunction over values of existentially-quantified variables. Handling of the \neg operator uses a precomputed `neg` table for a quick lookup of $f^{-1}(\neg j)$, for any join-irreducible element j , as required by Theorem 6. It is precomputed from the constructive proof of Lemma 4 where f is specified: first find f using $f(j) = \sqcup(L \setminus \uparrow j)$, then fill in the `neg` table using $\text{neg}[\neg f(j)] = j$.

Figure 11(a) shows a T-cut and an M-cut for the representation of a function $x \wedge y$ over an algebra $\mathbf{3}$ using an MBTDD-vector. The direct implementation of the MBTDD package is very similar to that of the MDD package described in Section 6.2.2, and is omitted here for brevity. MBTDD-vectors can also be easily implemented on top of a standard BDD library, such as CUDD [57]. Each MBTDD in the vector is represented using a BDD following the same process as described in the MDD to ADD conversion (see Section 6.2.3). For example, Figure 11(b) shows a BDD-vector representation for the function $f = x \wedge y$. Note that if had we used a different variable ordering, i.e., $x_0 > y_0 > x_1 > y_1$, then we would have been able to achieve an even more compact representation via the structure sharing between the two j -cuts of this function.

8 Case Studies

In the first part of this paper we defined mv-sets and discussed how to implement them using four decision diagram encodings: as MDDs and ADDs (Section 6) and as MBTDD-vectors and BDD-vectors

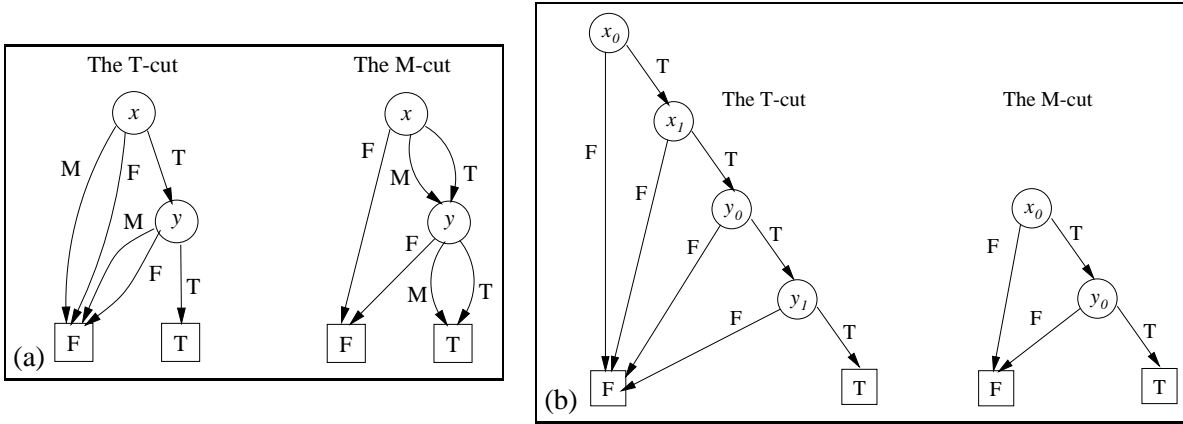


Figure 11: Representation of function $f = x \wedge y$ in $\mathbf{3}$: (a) using an MBTDD-vector; (b) using a BDD-vector.

(Section 7). In this section we study the effects of the different representations of mv-sets on some realistic problems using our multi-valued model-checker χ Chek [22]. The problems come with a series of CTL questions, thus providing us with a workload and letting us explore the relationship between the size of decision diagrams, discussed in detail in the previous section, and the running times of the verification.

This section includes three case studies of the use of multi-valued model-checking: checking abstractions of systems [20], checking compositions of features to determine which of these interact, and checking systems containing disagreements [32]. Results of the case studies are given in Table 3. For each analyzed model (“3-floor”, “5-floor”, “feature interaction”, “phone-system”), Table 3 shows verification times for several properties and the size of the transition relation. Since the latter does not change from property to property, it is shown only once for each model. All experiments were done on a Pentium III with 850 MHz processor and 256 MB RAM, running Sun JDK 1.3 under Linux 2.2.19.

8.1 Checking Abstractions

This case study uses the SMV elevator model of Plath & Ryan [50]. This model consists of a single elevator which accepts requests made by users pressing buttons on the floor landings and from inside the elevator. The elevator moves up and down between floors and opens and closes its doors in response to these requests according to Single Button Collective Control (SBCC) [9]. The model is implemented as several SMV modules. The `Main` module declares several instances of the module `Button` (one per floor, called `landingButi`), parameterized by the condition under which the request is considered fulfilled (`reset`), and one instance of the module `Lift`, called `lift`. The `Lift` module declares the variables `floor`, `door` and `direction` as well as further instances of `Button` to indicate requests from within the elevator (also one per floor, called `liftButi`).

In this case study we abstracted the original module `Button`, shown in Figure 12(a). In the original module, once a button is pressed, it latches and remains pressed until the elevator fulfills the request. First, we specified the latching explicitly: each variable `pressed` is decomposed into two variables, with `button` representing the actual button that users can press, and `pressed` representing the latching. The modified button is shown in Figure 12(b), and its state machine in Figure 13(a). The model has eight possible states, corresponding to the evaluation of the tuple $(\text{button}, \text{pressed}, \text{reset})$. To simplify the

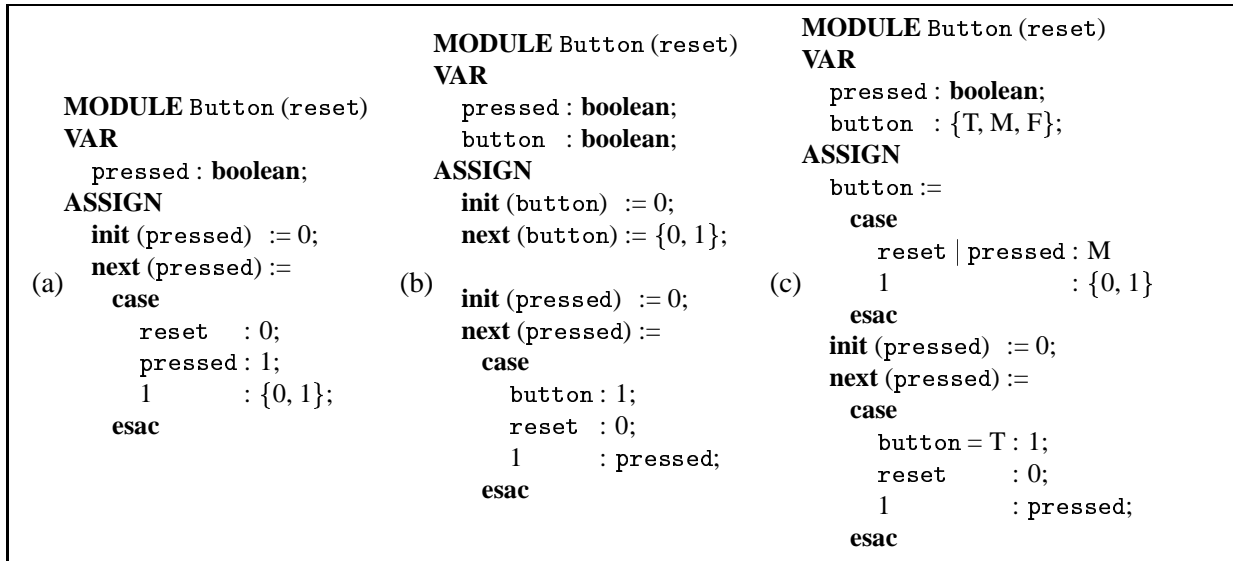


Figure 12: Three models of the elevator button in SMV: (a) the original module Button of Plath & Ryan; (b) a modified module Button; (c) an abstracted module Button.

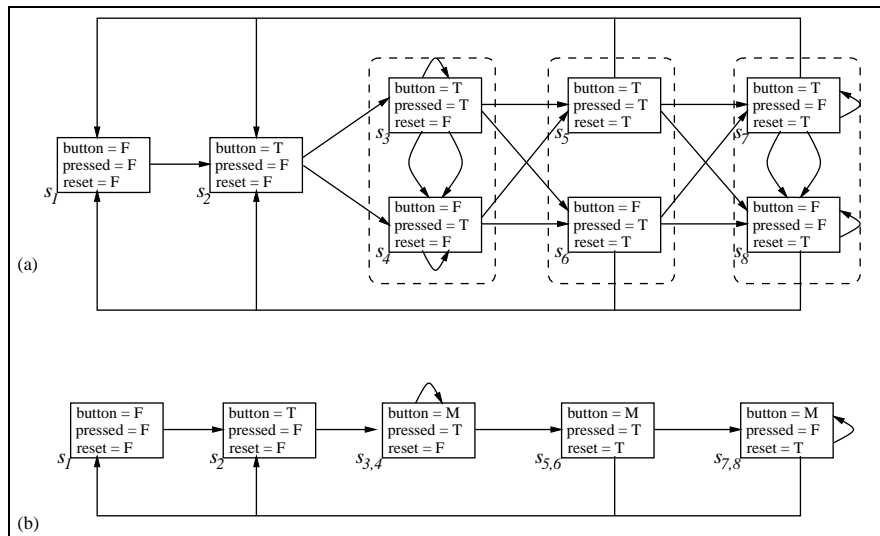


Figure 13: State machines: (a) of the modified module Button; (b) of the abstracted module Button.

presentation, in the state-machine description we assume that a button cannot be reset until it has been latched, i.e., **reset** cannot become \top if **pressed** is \perp .

Each of the pairs of states, $\{s_3, s_4\}$, $\{s_5, s_6\}$ and $\{s_7, s_8\}$, (indicated by dashed lines in Figure 13(a)) has a symmetric transition relation and thus can be abstracted. This corresponds to the value of **button** being irrelevant when **pressed** or **reset** are \top . Thus, we can model **button** by a 3-valued variable, as shown in Figure 12(c). The state machine model of the abstract system is shown in Figure 13(b). When this module is composed with the rest of the elevator model, we get a 3-valued model which cannot be directly verified using a classical model-checker.

1. “If the door closes, it will eventually open”:
 $AG(\text{lift.door} = \text{closed} \rightarrow AF \text{ lift.door} = \text{open})$
2. “If a button inside the lift is pressed, the lift will eventually arrive at the corresponding floor”:
 $AG(\text{lift.liftBut2.button} \rightarrow AF(\text{lift.floor} = 2 \wedge \text{lift.door} = \text{open}))$
3. “Pressing a landing button guarantees that the lift will arrive at that landing and open its doors”:
 $AG(\text{landingBut2.button} \rightarrow AF(\text{lift.floor} = 2 \wedge \text{lift.door} = \text{open}))$
4. “The lift may stop at floor 2 for landing calls when traveling downwards”:
 $\neg AG((\neg \text{lift.floor} = 2 \wedge \neg \text{lift.liftBut2.button} \wedge \text{lift.direction} = \text{down}) \rightarrow \text{lift.door} = \text{closed})$
5. “Whenever a button indicator is on, a button is being pressed”:
 $AG(\text{lift.liftBut2.pressed} \rightarrow \text{lift.liftBut2.button})$

Figure 14: Properties of the elevator system.

The properties of the elevator system that we verified appear in Figure 14. These are taken directly from [50], with the following exceptions: (a) we replaced `pressed` by `button` in all terms involving landing or elevator buttons because of our change in the module `Button`; (b) we selected only a subset of properties from [50]. Similar properties can be formulated for all other floors and all other landing and elevator buttons. We further parameterized the model by the number of floors and ran our experiments with 3 and 5 floors. The results are summarized in Table 3 under “3-floor” and “5-floor” models. For example, verification of property 1 of the 5-floor elevator can range from 19.009 s when mv-sets are implemented using MDDs, to 66.713 s, when they are implemented via a BDD-vector. Note that the MDD and the MBTDD-vector implementations yield the same transition relation sizes in the 3- and 5-floor elevator models. MBTDDs perform better in the 3-floor model, whereas MDDs perform better in the 5-floor case.

8.2 Finding Feature Interactions

This case study is based on the work of Plath and Ryan [50, 49] in the feature interaction domain.

Let P be a system, and δ_1 and δ_2 be two features. Then $P + \delta_1$ is a system composed with the first feature, and $P + \delta_1 + \delta_2$ is the system composed with two features, where δ_1 is applied first. We assume that the feature descriptions as well as the actual procedure to combine a feature and the system are the ones outlined in the work of Plath and Ryan. We are interested in different interactions between features. For example, for a property of interest φ , we want to know if the features commute relative to φ . That is, we want to know whether $([P + \delta_1 + \delta_2] \models \varphi) \wedge ([P + \delta_2 + \delta_1] \models \varphi)$. A naive approach is to break this problem into two independent model-checking problems. However, with χChek , a different solution is possible. First, we notice that the state space of $[P + \delta_1 + \delta_2]$ and $[P + \delta_2 + \delta_1]$, referred to as S , is exactly the same. This allows us to combine the two systems into one 4-valued (over algebra $\mathbf{2x2}$) model $M = (\hat{S}, \hat{R})$, where \hat{S} is the result of embedding S into the algebra $\mathbf{2x2}$ ($\hat{S} = \alpha_{\mathbf{2x2}}(S)$), and \hat{R} is constructed by merging transition relations of $[P + \delta_1 + \delta_2]$ and $[P + \delta_2 + \delta_1]$ using the process described in detail in [20].

Model	CTL Property Number	Result	χ Chek			
			MDD	MBTDD	ADD	BDD
3-floor	1.	F	0.505 s	0.423 s	1.222 s	1.007 s
	2.	T	0.194 s	0.125 s	1.306 s	0.23 s
	3.	T	0.197 s	0.122 s	1.171 s	0.233 s
	4.	T	0.497 s	0.591 s	1.406 s	1.196 s
	5.	M	0.202 s	0.125 s	0.596 s	0.23 s
Size of trans. relation			954	954	2467	1991
5-floor	1.	F	19.009 s	21.251 s	30.89 s	66.713 s
	2.	T	2.254 s	2.407 s	6.978 s	5.156 s
	3.	T	2.331 s	2.406 s	5.21 s	4.929 s
	4.	T	14.853 s	19.394 s	79.753 s	63.978 s
	5.	M	1.017 s	0.862 s	4.263 s	1.725 s
Size of trans. relation			4870	4870	12615	10114
Feature-interaction (3 floors)	1.	TT	1.081 s	1.001 s	2.854 s	0.716 s
	2.	FF	1.318 s	1.484 s	3.415 s	1.119 s
	3.	TF	2.449 s	2.445 s	6.169 s	1.95 s
	4.	FF	0.95s s	0.957 s	2.463 s	0.695 s
	5.	TT	2.257 s	2.413 s	5.776 s	2.362 s
Size of trans. relation			6596	7099	14868	9401
Phone system (9-valued)	1.	MM	0.048 s	0.066 s	2.3 s	0.048 s
	2.	MF	0.044 s	0.074 s	2.135 s	0.052 s
	3.	MF	0.046 s	0.068 s	2.613 s	0.037 s
Size of trans. relation			108	142	453	599
Phone system (4-valued)	1.	TT	0.031 s	0.023 s	0.024s	0.029 s
	2.	FF	0.031 s	0.027 s	0.042 s	0.036 s
	3.	FF	0.031 s	0.027 s	0.041 s	0.031 s
Size of trans. relation			84	99	178	207

Table 3: Case studies: sizes of transition relation and running times of the four implementations of χ Chek.

In the case where the two systems share the same state space, it can be easily shown that our model-checker χ Chek, run on an arbitrary property φ , returns TT if and only if the result of model-checking φ on each individual model is T; χ Chek returns TF if and only if the result of model-checking φ on the first model is T and on the second model is F, etc. Thus, instead of solving two boolean model-checking problems, we can solve one multi-valued problem.

1. “If the door closes, it may stay closed forever”:

$$\neg AG(\text{lift.door} = \text{closed} \rightarrow AF \text{lift.door} = \text{open})$$
2. “If a button inside the lift is pressed, the lift will eventually arrive at the corresponding floor”:

$$AG(\text{lift.liftBut2.button} \rightarrow AF(\text{lift.floor} = 2 \wedge \text{lift.door} = \text{open}))$$
3. “Lift calls have precedence when the lift is 2/3 full (indicated by `tt_full`)”:

$$AG((\text{lift.tt_full} \wedge \text{lift.liftBut2.button} \wedge \neg \text{lift.liftBut3.button})$$

$$\rightarrow A[(\text{lift.floor} \neq 2 \vee \text{lift.door} = \text{closed}) U$$

$$(\text{lift.floor} = 2 \wedge \text{lift.door} = \text{open})$$

$$\vee \neg \text{lift.tt_full} \vee \text{lift.liftBut3.button}])$$
4. “Pressing a landing button guarantees that the lift will arrive at that landing and open its doors”:

$$AG(\text{landingBut2.button} \rightarrow AF(\text{lift.floor} = 2 \wedge \text{lift.door} = \text{open}))$$
5. “The lift may stop at floor 2 for landing calls when traveling downwards”:

$$\neg AG((\neg \text{lift.floor} = 2 \wedge \neg \text{lift.liftBut2.button} \wedge \text{lift.direction} = \text{down})$$

$$\rightarrow \text{lift.door} = \text{closed})$$

Figure 15: Properties of the combined elevator system.

In this case study, we composed the base elevator system (SBCC) with two features:

Lift-2/3-full. When the elevator detects that it is more than two-thirds full, it does not stop in response to external calls, since it is unlikely to be able to accept more passengers. Instead, it gives priority to passengers already inside the elevator, as serving them will help reduce its load.

Executive Floor. The elevator gives priority to calls from the executive floor.

and ran the model-checker over the properties that appear in Figure 15. Note that most of these properties are the same as in Figure 14. We expect these properties to hold in the composed system. However, property 3 is likely to reveal feature interaction when the third floor is designated as the executive floor: when the elevator is 2/3 full and there is a call from the executive floor and from within the elevator, which call should get priority?

Results of the verification of the feature interaction of the three-floor elevator are shown in Table 3 under “feature interaction”. Note that each property except 3 yields TT or FF, indicating that the two features commute w.r.t. this property. Value TF for property 3 indicates that it holds if Lift-2/3-full is the last feature added, and does not if Executive Floor is the last feature added.

In this example, the implementation based on a BDD-vector performs the best. The implementations that use MDDs and MBTDD-vectors exhibit similar performance.

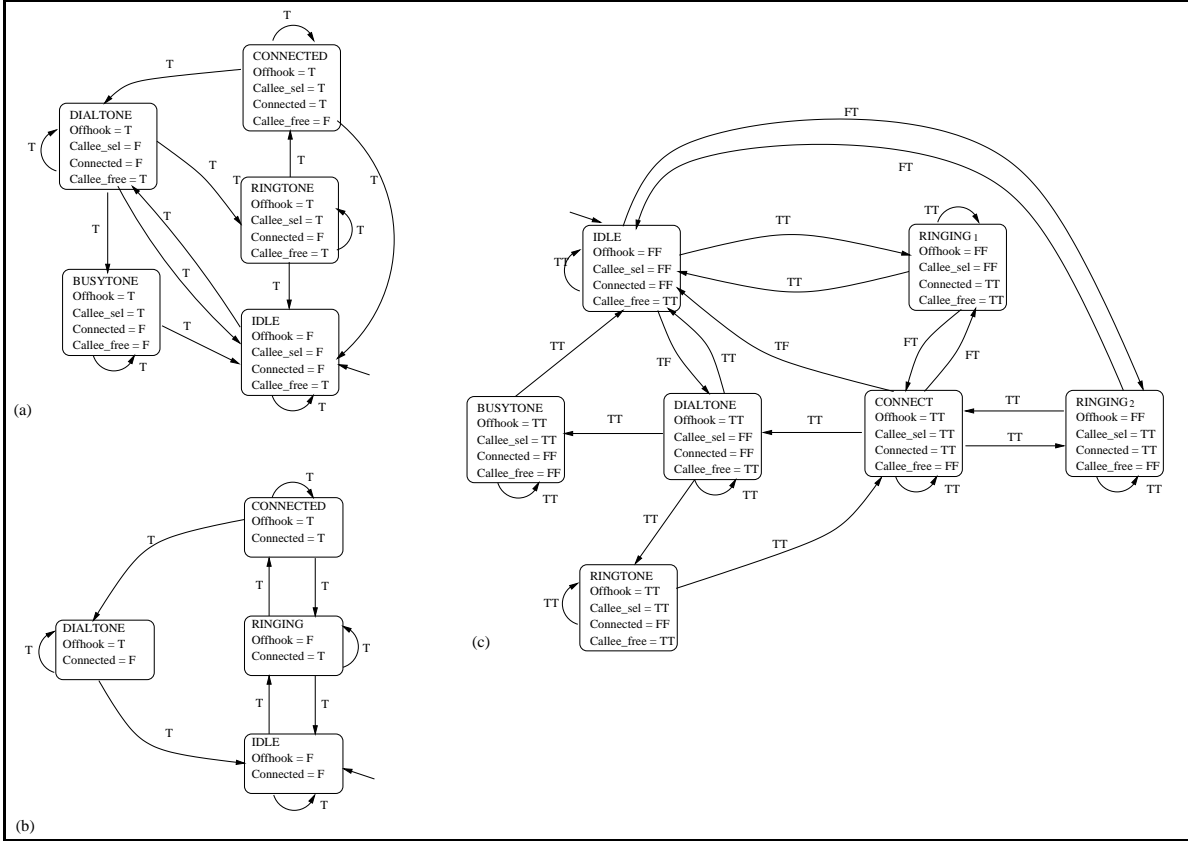


Figure 16: Partial descriptions of a phone system: (a) Caller's perspective; (b) Callee's perspective; (c) combined system.

8.3 Reasoning about Disagreement

In this case study, we use multi-valued algebras to model disagreement between descriptions of two aspects of a system. Model-checking is then used to determine whether these disagreements matter [32].

Figure 16(a)-(b) shows the individual (partial) descriptions of the caller's and the callee's perspectives on the phone system. The combined model, constructed using the technique described in [32, 53], is given in Figure 16(c). The combined model can be described either using a 9-valued algebra 3×3 or the four-valued algebra 2×2 , with values other than TT and FF representing disagreements between the two viewpoints, either on the values of variables in the states, or the values of the transitions. We can determine whether these disagreements matter by checking the combined system against the correctness criteria of the phone system which appear in Figure 17.

We give verification results of the phone system in the 9-valued and the 4-valued algebra in Table 3 under "Phone system". Note that the transition relation has a fairly small size in all four representations when the four-valued algebra is used. Further, the implementation that uses ADDs performs better than the others. However, in the 9-valued algebra case, the implementations based on ADD and a BDD-vector exhibit a roughly 250% size increase, as opposed to a 50% increase for the MDDs and the MBTDD-vectors. Further, in this case the ADD representation performs the worst, whereas the BDD representation

1. “If you are connected, you can hang up”:
 $AG(\text{connected} \rightarrow EX\neg\text{offhook})$
2. “A connection doesn’t start until you pick up the phone”:
 $AG(\neg\text{connected} \rightarrow \neg E[\neg\text{offhook} U (\text{connected} \wedge \neg\text{offhook})])$
3. “If you are offhook and connected, hanging up always disconnects you”:
 $AG((\text{offhook} \wedge \text{connected}) \rightarrow AX(\neg\text{offhook} \rightarrow \neg\text{connected}))$

Figure 17: Properties of the phone system.

Name	Description
abp4	alternating bit protocol
dartes	communication protocol of an Ada program
dme2-16	distributed mutual exclusion protocol
dpd75	dining philosophers
ftp3	file transfer protocol
furnace17	remote furnace program
key10	keyboard/screen interaction in a window manager
mmgt20	distributed memory manager
over12	automated highway system overtake protocol

Table 4: List of NuSMV models used in the experiments.

gives the best performance.

9 Experimental Comparisons

In Section 8, we gave experimental results of the four implementations of mv-sets on three simple case studies, and noticed that different implementations performed better on different examples. Interested in studying this relationship further but lacking a library of benchmark multi-valued problems, we conducted a set of experiments using randomly generated models with varying parameters. We describe them below. In Section 9.1, we outline the types of experiments we perform. We propose to study the impact of two factors on the size of underlying decision diagram representation: (a) encoding a characteristic function of an mv-set as a multi-valued function or as a vector of boolean functions (we refer to this as “the number of terminal nodes”, meaning that if the underlying decision diagram has two terminal nodes, \top and \perp , then the characteristic function is boolean; otherwise, it is multi-valued); and (b) the branching factor of the decision diagrams. Each factor is analyzed in detail in Sections 9.2 and 9.3, respectively. In Section 9.4, we relate the decision diagram size and the running time of mv-set operations.

9.1 Methodology

Theoretical estimates of space and time complexity for decision diagrams are very loose. Further, it has been shown [59] that the performance of decision diagrams for most *randomly chosen* functions is very poor. Yet, it is generally known that functions used in formal verification have performance that is considerably better than the worst-case. To study multi-valued decision diagrams, we propose to use benchmarks that represent functions occurring in model checking, i.e., those similar in structure to classical models but with some non-boolean inputs and outputs. We anticipate two principal ways in which multi-valued models arise in model-checking. The first is from *inconsistency*: we may want to combine several slightly different classical models while preserving all of their differences. A few examples of such systems were given in [32, 53]. The second cause of multi-valued functions is *partial* systems. In a fully-specified system, each transition is either present or absent, whereas in a partial system, some transitions can be marked as *unknown*, indicated by the value M of algebra **3**. Finite total orders with more values may be used for different levels of partiality [12, 48].

As a source of classical models, we adopt benchmarks used in a previous performance study of symbolic model-checking [60]; these represent *transition relations* of several NuSMV [24] models and are expressed in the modeling language of NuSMV. The list of the models we use is given in Table 4. For brevity, in this paper we present partial experimental results for models **abp4** and **dme2-16**. Results for all other models are identical, except for the absolute size of the decision diagram representation.

After constructing multi-valued models, we represent them using decision diagrams and then conduct experiments, measuring the size of the representations in terms of the *number of nodes*. This metric can be easily obtained; further, storage requirements can be computed by simply multiplying the branching factor of a given decision diagram by its number of nodes. Finally, the number of nodes allows us to compare running times for different mv-set operations, as described in Section 9.4.

As shown in Table 2 in Section 5, the four representations of mv-sets vary along two axes: the branching factor and the number of terminal nodes. We are interested in exploring the independent effects of these two axes of variation. We therefore want to design experiments that allow us to measure the effect of varying the number of terminal nodes, i.e. one row of Table 2 (see Section 9.2) and the branching factor, i.e., one column of Table 2 (see Section 9.3).

Finally, we perform each experiment on a number of multi-valued algebras. We restrict ourselves to algebras with at most 16 elements (to ensure tractability of the experiments). The algebras we use are listed in Table 5. The minimum number of join-irreducible elements for an algebra L is $\lceil \log_2(|L|) \rceil$. For example, algebra **2x2x2x2** has 16 elements and 4 join-irreducibles. The maximum number of join-irreducibles for an algebra L is $|L| - 1$. This number is attained when the elements of the algebra form an FTO – a finite total order², such as in algebra **16** which has 15 join-irreducibles. For our experiments, we pick algebras from each side of the spectrum and some from the middle, making sure that we have several algebras of different sizes with the same number of join-irreducible elements, e.g. **5**, **6C** and **3x3**, where each algebra has 4 join-irreducibles.

9.2 Experiments 1 & 2: Number of Terminal Nodes

The goal of these experiments is to study the effect of the number of terminal nodes on the size of the resulting decision diagram. Thus, we hold the branching factor of our diagrams constant and vary the

²An FTO is a lattice $(\mathcal{L}, \sqsubseteq)$, where \sqsubseteq is a *total order*. Algebra **5**, given in Figure 18(b), is an FTO.




Figure 18: Some algebras used in the experiments: (a) $\mathbf{6C}$; (b) $\mathbf{5}$. Join-irreducible elements are indicated by filled circles.

Name	Size	No. of join-irreducibles	Notes
$\mathbf{2}$	2	1	Classical logic, see Figure 1(a)
$\mathbf{3}$	3	2	Models uncertainty, see Figure 1(b)
$\mathbf{2x2}$	4	2	Models disagreement, see Figure 1(c)
$\mathbf{4}$	4	3	Finite total order with 4 elements
$\mathbf{5}$	5	4	Finite total order with 5 elements, see Figure 18(b)
$\mathbf{2x3}$	6	3	Product of $\mathbf{2}$ and $\mathbf{3}$
$\mathbf{6C}$	6	4	See Figure 18(a), used in [19]
$\mathbf{6}$	6	5	Finite total order with 6 elements
$\mathbf{2x2x2}$	8	3	Product of three instances of $\mathbf{2}$
$\mathbf{3x3}$	9	4	Product of two three-valued algebras ($\mathbf{3}$), see Figure 1(e)
$\mathbf{2x2x2x2}$	16	4	Product of four instances of $\mathbf{2}$
$\mathbf{16}$	16	15	Finite total order with 16 elements

Table 5: Algebras used in our experiments.

number of terminal nodes. In particular, we let the branching factor be equal to the size of the algebra, and thus study the trade-offs between the MDD and MBTDD implementations of mv-sets. Alternatively, we could have set the branching factor to 2 and studied the trade-offs between ADDs and BDDs.

9.2.1 Experiment Design

All model-checking benchmarks are based on boolean, i.e., crisp, functions; however, MDDs and MBTDD-vectors yield exactly the same sizes for such functions; further, the size does not change even when we allow the input to be multi-valued. This indicates that the number of terminal nodes does not affect the representation of boolean-output functions. However, if we go back to boolean benchmarks and instead make the output multi-valued, the MDD and the MBTDD-vector implementations exhibit different behaviour, which we study in this section. Functions with crisp input and multi-valued output arise from models where disagreements or partialities occur *only on transitions*, and likely encompass the majority of multi-valued model-checking problems.

In order to construct models with multi-valued transition relation, we start with a classical model and then change some of the transition values. Depending on the algebra, this change may represent either a disagreement (when the algebra is a product of several classical ones), an underspecification (in the case of an FTO), or a combination of the two. A global transition relation is typically computed via synchronous

```

function DIS (Algebra  $L$ , int  $n$ )
  //  $f$  is a crisp function over  $L$ , defined globally.
   $g_0 = f$ 
  foreach  $1 \leq i < n$ 
    compute  $\rho$  – a minimal unambiguous partial assignment s.t.  $g_i|_{\rho} \in \{\top, \perp\}$ 
    pick some  $\ell \in (L \setminus \{\top, \perp\})$ 
    compute  $g_{i+1}$  such that for each total assignment  $\rho'$ ,
      
$$g_{i+1}|_{\rho'} = \begin{cases} \ell & \text{if } \rho \subseteq \rho' \\ g_i|_{\rho'} & \text{otherwise} \end{cases}$$

  return  $g_n$ 

```

Figure 19: Adding disagreements to a crisp function f .

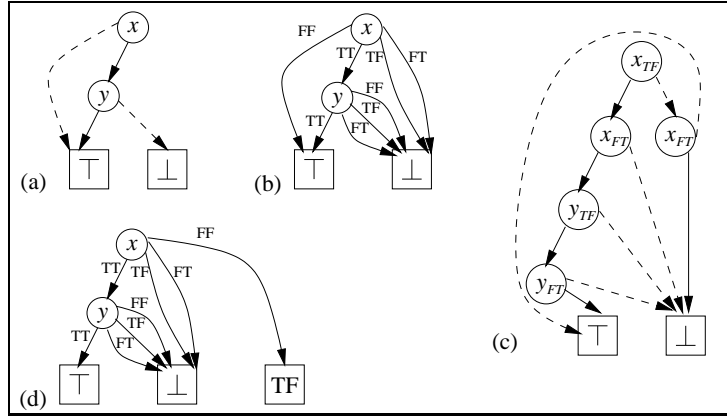


Figure 20: A boolean function $\neg x \vee y$: (a) as a BDD; (b) as an MBTDD-vector in $\mathbf{2x2}$; (c) as a BDD-vector in $\mathbf{2x2}$; (d) result of introducing a TF disagreement to (b).

or asynchronous parallel compositions of transition relations of smaller models. Thus, a disagreement on a single local transition may result in a large number of changes to the global transition relation. With this in mind, we propose the following method for creating multi-valued models.

Let δ be a function representing a transition relation of a classical model, and $f = \alpha_L(\delta)$ be the embedding of δ into L , defined via Definition 6. Multi-valued models are created by adding *disagreements* to f using a function DIS defined in Figure 19, parameterized by the algebra L and the number of disagreements n . The minimal unambiguous partial assignment, used in this algorithm, is defined in Section 6.1. The algorithm is declarative – the exact computation of g_{i+} is technical and is omitted from this presentation. For example, Figure 20(b) shows a representation of a boolean function f as an MBTDD-vector over algebra $\mathbf{2x2}$. The result of adding one disagreement where $\rho = \{(x, \perp)\}$ and $\ell = \text{TF}$ is a function g , shown in Figure 20(d). This function is identical to f except that the FF child of node x now points to TF.

For our experiments, we range the number of disagreements between 5 and 100, corresponding to an intuition that a model with *too* many disagreements is not yet ready for formal analysis. Given a algebra L and a number of disagreements n , let $\text{SAMPLE}(n, L)$ be a set of 100 different functions obtained by calling $\text{DIS}(n, L)$ on $\alpha_L(\delta)$. For each function $f_i \in \text{SAMPLE}(n, L)$, we generate an MDD u_i and an




Figure 21: Illustration of Hypothesis 1 with two disagreements: (a) the same output; (b) different outputs.

MBTDD-vector \hat{u}_i such that $f_i = f^{u_i} = f^{\hat{u}_i}$, thus creating two new sets of functions:

$$\begin{aligned} \text{MDDSAMPLE}(n, L) &\triangleq \{u_i \mid f^{u_i} = f_i\} \\ \text{MBTDDSAMPLE}(n, L) &\triangleq \{\hat{u}_i \mid f^{\hat{u}_i} = f_i\} \end{aligned}$$

We use these sets in the experiments discussed below.

9.2.2 Experimental Results

Before performing the experiments, we make an initial hypothesis, formulated below, that, all other things being equal, an MDD based on a larger algebra should contain more nodes than an MDD based on a smaller one:

Hypothesis 1 *Let $u \in \text{MDDSAMPLE}(n, L)$, $v \in \text{MDDSAMPLE}(n, L')$, and $|L| > |L'|$. Then we expect that $\text{size}(u) > \text{size}(v)$.*

For example, consider adding two disagreements ρ and ρ' with values ℓ and ℓ' , respectively. If $\ell = \ell'$, then the two paths corresponding to the disagreements must share a suffix (see Figure 21(a)). In the worst case, the shared suffix consists of just one terminal node corresponding to ℓ . On the other hand, if $\ell \neq \ell'$, the two paths cannot have a common suffix at all (see Figure 21(b)). Thus, the higher the chance that any two disagreements have the same value, the smaller the size of the resulting diagram. Clearly, the possibility of assigning the same value to disagreements depends on the number of possible values, which in our case is $|L| - 2$; therefore, the larger the algebra, the larger the size of the resulting MDD.

To validate this hypothesis, we conduct the following experiment:

Experiment 1(a):

For each algebra and each possible number of disagreements, we create 100 MDDs and measure their size. The sizes of MDDs are clustered within a narrow range, so we feel justified in extracting their mean value, denoted $\|\text{MDDSAMPLE}(n, L)\|$, and plotting it as a single point. Henceforward we shall simply call this value “the MDD size for n and L ”. Results of the experiment are shown in Figure 22(a). For each algebra, we plot the increase of the mean size of MDDs as the number of disagreements grows, for model **abp4**. Comparing the MDD size at a given number of disagreements for any two algebras validates Hypothesis 1. For example, MDDs over **3** have the smallest size, whereas those over **16** and **2x2x2x2x2** have the largest; further, curves corresponding to those two algebras (both with 16 elements) coincide.

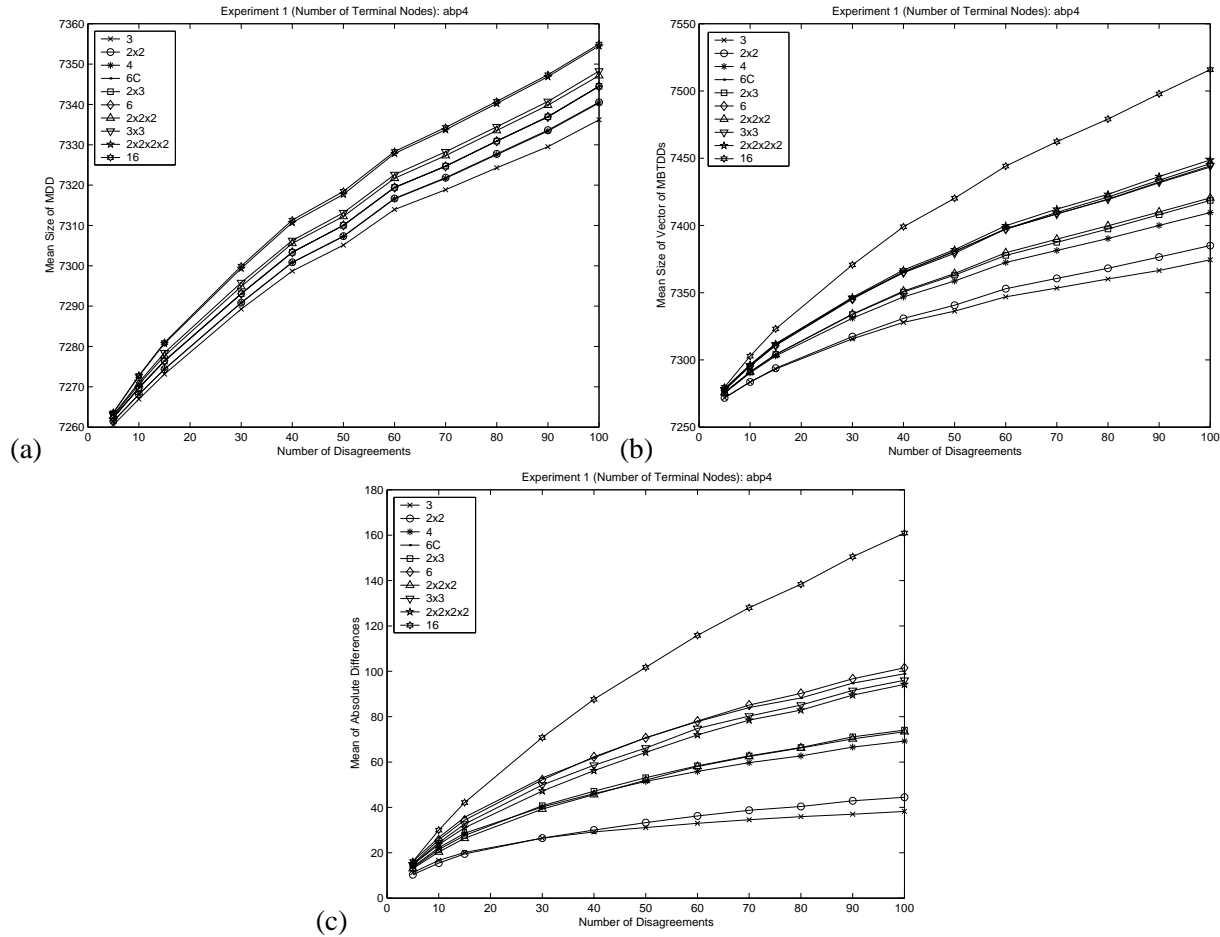


Figure 22: Graphs for Experiment 1: (a) MDD; (b) MBTDD-vector (note that y -axis is labeled differently here than in (a)); (c) The difference between the MDD and MBTDD-vector representations.

From the shape of the curves in the graph in Figure 22(a) we conclude that, as more disagreements are added, the marginal impact of each subsequent disagreement is lessened. That is, for any i, j, k with $i < j < k \leq 100$ and $k - j = j - i$,

$$\|\text{MDDSAMPLE}(j, L)\| - \|\text{MDDSAMPLE}(i, L)\| \geq \|\text{MDDSAMPLE}(k, L)\| - \|\text{MDDSAMPLE}(j, L)\|$$

We have not explored the space of parameters beyond 100, where there may be some unexpected phenomena. \square

We now turn our attention to measuring the size of MBTDD-vectors. Note that the observation about mean sizes of MDDs, made above, holds for MBTDD-vectors as well; this justifies computing $\|\text{MBTDDSAMPLE}(n, L)\|$ and referring to it as “MBTDD-vector size”. Since each MBTDD is also an MDD, we also expect subsequent changes to have a diminishing impact on the individual MBTDDs of a vector, just as they do for single MDDs, formulated in the following hypothesis:

Hypothesis 2 *The diminishing impact of disagreements observed for MDDs also holds for MBTDD-vectors: for any i, j, k with $i < j < k \leq 100$ and $k - j = j - i$,*

$$\begin{aligned} & \|\text{MBTDDSAMPLE}(j, L)\| - \|\text{MBTDDSAMPLE}(i, L)\| \\ & \geq \|\text{MBTDDSAMPLE}(k, L)\| - \|\text{MBTDDSAMPLE}(j, L)\| \end{aligned}$$

The number of terminal nodes in MBTDD-vectors remains set at 2, regardless of the number of algebraic values. Therefore, we do not anticipate the same dependence on the number of possible outputs as we observed with MDDs. However, we do predict dependence on the number of join-irreducibles, since it corresponds to the number of diagrams in the vector. For example, consider an MBTDD-vector $\hat{u} \in \text{MBTDDSAMPLE}(n, L)$ with $|\mathcal{J}(L)| = 2$. The two individual MBTDDs comprising this vector share some nodes, but each contains some nodes not appearing in the other. If we compare this case with the one of $\hat{v} \in \text{MBTDDSAMPLE}(n, L')$, with $|\mathcal{J}(L')| = 3$, we note that the shared size of any *two* individual MBTDDs in this vector is close to $\text{size}(\hat{u})$. However, the third MBTDD is likely to have some nodes not shared with the first two and thus contributes a non-zero number of unshared nodes to the overall size of the MBTDD-vector. We formulate this intuition in the following hypothesis:

Hypothesis 3 *Let $u \in \text{MBTDDSAMPLE}(n, L)$, $v \in \text{MBTDDSAMPLE}(n, L')$, and $|\mathcal{J}(L)| > |\mathcal{J}(L')|$. Then we expect that $\text{size}(\hat{u}) > \text{size}(\hat{v})$.*

Hypothesis 4 *Introducing disagreements destroys most of the sharing which holds within an MBTDD-vector.*

Experiment 1(b):

For each algebra and each possible number of disagreements, we create 100 MBTDD-vectors and measure their size. Results are given in Figure 22(b). These results allow us to make the following observations: (a) we see the same diminishing impact of subsequent disagreements in Figure 22(b), confirming Hypothesis 2; (b) comparing results of running experiments on two algebras with different numbers of join-irreducibles, we see that Hypothesis 3 holds in general. For example, the curve for **16**, with 15 join-irreducibles, is above the one for **2x2x2x2**, with 4 join-irreducibles. However, when two algebras have the same number of join-irreducibles, e.g., **3** and **2x2**, we see that their curves are clustered together, but do not coincide.

We now compare results in Figure 22(a) and 22(b). First, we note that for the same algebra L and the same number of disagreements n , the MDD size is less than the MBTDD-vector size. We can show that this is not *necessarily* the case. Consider an example in Figure 23. Figure 23(a) shows an MDD for some crisp-input function h with outputs from algebra **2x2**. This function requires 7 nodes (4 terminal and 3 non-terminal) to be represented as an MDD under the variable ordering $x < y$. In contrast, the same function represented as a vector of two MBTDDs (see Figure 23(b)) requires just 4 nodes: 2 terminal and 2 non-terminal. As can be seen from this simple example, the relationship we observe in the experimental data between the MDD size and the MBTDD-vector size is not universal, but only particular to our domain of interest.

Second, even though the MBTDD-vector size is greater than the MDD size, their ratio never exceeds 1.1. Without disagreements, each MBTDD is the same size as the original MDD. Were Hypothesis 4 to hold, we would expect to see the size of the MBTDD-vector to be close to $|\mathcal{J}(L)|$ times the size of the




Figure 23: Representation of a crisp-input function h in x and y : (a) as an MDD (to reduce clutter, only outgoing edges for TT (solid) and FF (dashed) are shown; all others lead directly to FF); (b) as an MBTDD-vector.

MDD: with minimal sharing, each pair of MBTDDs is mostly disjoint, and thus the size of the MBTDD-vector is roughly the sum of the sizes of its elements. We therefore conclude that this Hypothesis does *not* hold, and the addition of disagreements leaves most sharing between elements of the MBTDD-vector intact.

Third, to get a clearer picture of the differences in size, we also graph the absolute difference between mean MDD size and mean MBTDD-vector size on a single grid. This graph, shown in Figure 22(c), allows us to observe that (a) the impact of disagreements on the *difference* in size also diminishes; and (b) for any two algebras L, L' with the same number of join-irreducibles and any number of disagreements n , if $|L'| > |L|$, then the difference between MDD size and MBTDD-vector size is *smaller* for L' than for L : that is, the *larger* the algebra, the *less* is the difference in the representation size. For example, comparing results for two algebras with 4 join-irreducible elements, $\mathbf{3x3}$ with size 9, and $\mathbf{2x2x2x2}$ with size 16, we note that the curve corresponding to the former is above the one corresponding to the latter. □

Results of Experiment 1(b) allow us to make a yet another conjecture: given two algebras with the same number of join-irreducibles, the size of the MBTDD-vector also depends on the *size of the algebra*. We illustrate this point on a concrete example. Consider two algebras with the same number of join-irreducibles, $\mathbf{3}$ and $\mathbf{2x2}$. However, algebra $\mathbf{3}$ has only one “non-crisp” join-irreducible (M) whereas $\mathbf{2x2}$ has two (TF and FT). Thus, each disagreement, i.e., a change of an output from \perp to some other value of the algebra, requires a change to exactly one MBTDD, the M-cut, leaving the T-cut intact. The same disagreement added to a function in $\mathbf{2x2}$ causes a change *either* to the FT-cut or to the TF-cut; thus, as disagreements accumulate, both decision diagrams are changed in different ways, diminishing the sharing between them. Further, we observe the diminishing effect of disagreements: changing the M-cut twice has less effect than changing each of TF- and FT-cuts once. Thus, we formulate the following hypothesis:

Hypothesis 5 *If the same value ℓ is used for adding disagreements, then the size of an MBTDD-vector after adding disagreements is the same in algebra $\mathbf{2x2}$ as it is in algebra $\mathbf{3}$, despite the differences in the structure of the two algebras.*

In order to test our hypothesis about this dependency on structure, we perform an additional experiment:

Experiment 2:




Figure 24: Graph for Experiment 2 with algebras **3** and **2x2**.

In this experiment, we compare sizes of MBTDDs for algebras **3** and **2x2**. The method for adding disagreements with **3** remains unchanged, but in the case of **2x2**, we choose ℓ to be just TF, instead of varying it randomly between TF and FT. We now expect to see no dependence on the size of the algebra, apart from random noise, since our sample size may not be sufficiently large to even it out. Figure 24 gives the results of performing this experiment, validating Hypothesis 5 and showing the reduction in size for algebra **2x2** with the new treatment of disagreements over the original one.

□

9.3 Experiment 3: Branching Factor

The goal of this experiment is to study the effect of the branching factor on the different mv-set representations. Thus, we hold the number of terminal nodes in our diagrams constant and vary the branching factor. In particular, we fix the number of terminal nodes at 2 and study the trade-offs between the representation of mv-sets as MBTDD-vectors and as BDD-vectors. Equivalently, it is possible to keep the number of terminal nodes equal to the size of the algebra and study the trade-offs between MDDs and ADDs.

9.3.1 Experiment Design

At a first glance, a representation of any given function as an MBTDD-vector requires fewer nodes than its representation as a BDD-vector. Intuitively, this follows from the fact that each multi-valued node in an MBTDD is simulated by several boolean nodes in a BDD. However, replacing each multi-valued variable by a collection of boolean variables enables additional freedom for reordering; therefore, it is possible for a BDD-vector of a given function to be smaller than its MBTDD-vector representation.

Every multi-valued node in an MBTDD can be simulated by a BDD with at most $2^{d+1} - 1$ nodes, where d is the maximum number of bits required to encode the branching factor of the MBTDD. Thus, in the optimal case $d = \lceil \log_2(|L|) \rceil$, and if the join-irreducible encoding, as proposed in Section 6.2.3, is used, then $d = |\mathcal{J}(L)|$. However, this only provides an upper bound on the representation. A better estimate depends on the actual encoding used, and the variable orderings enabled by it.

Algebras	Number of nodes			Growth factor
	MBTDD-vector	BDD-vector	BDD-vector (reordered)	
2	141840	141840	141839	1.00
3	141840	314495	314495	2.22
4	141840	487150	467704	3.43
2x2	141840	314495	314495	2.22
2x3	141840	487150	467704	3.43
6C	141840	659805	613704	4.65
6	141840	832460	801850	5.87
2x2x2	141840	487150	467704	3.43
3x3	141840	659805	613704	4.65
2x2x2x2	141840	659805	613704	4.65

Table 6: Results for Experiment 3.

Instead, let us consider an embedding of a boolean function into some algebra L . For simplicity, assume that the boolean function is given by a BDD (see Figure 20(a) where \perp -edges are indicated by dashed lines). Clearly, the number of nodes in an MBTDD-vector for this function is independent of the algebra used. Every BDD node is simply replaced by a corresponding multi-valued MBTDD node with additional children pointing to \perp (see Figure 20(b)). Thus, the result is an MBTDD-vector where (a) all MBTDDs are equivalent; (b) every node has at least two and at most three distinct children; and (c) $\forall c \in L \setminus \{\top, \perp\} \cdot \text{child}_c(u) = \perp$. The BDD-vector is then obtained by replacing each MBTDD node that has three distinct children by $2 \times |\mathcal{J}(L)| - 1$ BDD nodes, and replacing each node that has two distinct children by $|\mathcal{J}(L)|$ BDD nodes. The result is shown in Figure 20(c). Comparing this figure with Figure 20(b), we note that node x is replaced by 3 nodes, whereas node y is replaced by 2 nodes. Thus, an embedding of a boolean function given by an MBTDD-vector with n nodes can be encoded by a BDD-vector with $k \times n$ nodes, where k ranges between $|\mathcal{J}(L)|$ and $2 \times |\mathcal{J}(L)| - 1$. We measure k empirically in the experiments described below.

9.3.2 Experimental Results

As shown in the above discussion, we expect the growth factor k to be at least $|\mathcal{J}(L)|$. However, we know that it depends on the number of nodes that do not have \perp children in the original BDD representation of the transition relation. Since the transition relation of a typical model is rather sparse, the reachable state space is much smaller than the total state space, so we expect k to be closer to its lower bound.

Hypothesis 6 *Let L be an algebra, f be a transition relation of some model, \hat{u} be an MBTDD-vector representation of $\alpha_L(f)$, and \hat{v} be a BDD-vector representation of $\alpha_L(f)$. Then $k = \text{size}(\hat{u})/\text{size}(\hat{v})$ is close to $|\mathcal{J}(L)|$.*

Experiment 3(a):

Operation	Complexity	Complexity relative to number of nodes
Unary (\neg)	$O(G)$	$O(b \times n)$
Binary (\sqcup, \sqcap)	$O(G \times F)$	$O(b^2 \times n_G \times n_F)$

Table 7: Complexity of operations on decision diagrams.

The goal of this experiment is to empirically measure the growth factor k . In this experiment we embed the transition relation of each model of Table 4 into a quasi-boolean algebra and then measure its size as an MBTDD-vector and as a BDD-vector.

Table 6 summarizes the results of our experiments for the **dme2-16** model. As expected, the size of the MBTDD-vector remains constant, and the size of the BDD-vector increases dramatically. Notice that the size of the BDD-vector does not depend on the size of the algebra, but rather on the number of join-irreducible elements in it. For example, the size of the BDD-vector for algebras **2x2x2x2** and **3x3** is exactly the same, although the algebras have different number of elements. The actual growth factor k is given by the last column of the table. As suggested by Hypothesis 6, the growth factor is very close to the number of join-irreducible elements in the algebra. For example, for **2x2x2x2** that has 4 join-irreducible elements, the growth factor is just 4.65.

□

Experiment 3(b):

In this experiment, we repeat Experiment 3(a) while trying to achieve the best ordering for the BDD-vector representation, using the *sifting* reordering heuristic of CUDD [57]. The results of this experiment are summarized in the Table 6 under the heading *BDD-vector (reordered)*. As can be seen from the table, the reordering heuristic reduced the size of the BDD-vector only marginally and did not affect the size of the BDD-vector at all in many of our experiments on smaller models. Since it is hard to predict an effect of a given reordering heuristic on an arbitrary decision diagram, we have not been able to explain this fact. We conjecture that a more aggressive heuristic should be able to significantly reduce the BDD-vector size.

□

9.4 Complexity

In this section we analyze complexity of different operations on decision diagrams. It is well known that using a different branching factor, or representing a function by a vector of diagrams, has no effect on the complexity of the operations used during model-checking [59]. The middle column of Table 7 summarizes these complexities of the operations from the left column with respect to the size of the graph representing the diagram. Note that even though we can think of representing an mv-set using a vector of diagrams, the underlying implementation constructs a single directed acyclic graph. Moreover, since the underlying graph is connected, we can express the complexity of operations relative to the number of nodes in this graph. These complexities are given in the right column of Table 7, where n is the number of nodes and b is the branching factor of the decision diagram. Using this representation of complexity, we infer the expected running time based on the empirical evidence on the sizes of different decision diagrams.

9.4.1 Worst-case Complexity

As can be seen from the experiments, representing an mv-set using a vector of diagrams incurs only a small overhead in the number of nodes. Thus, since the same branching factor is used to represent a given function as a single MDD or as an MBTDD-vector, the expected worst-case time complexity of an MBTDD-vector is only slightly worse than that of a single MDD. On the other hand, representing the same function using decision diagrams with a fixed branching factor of 2 (either as a BDD-vector, or as a single ADD), incurs a considerable overhead, since each multi-valued node is replaced by several boolean nodes. Let us assume that this size increase is given by k , that is, if a function represented by an MBTDD-vector has n nodes, then the same function represented by a BDD-vector has $k \times n$ nodes. Since BDDs enjoy a constant branching factor of 2, it follows that for a fixed logic L , the BDD-vector representation has lower worst-case time complexity than an MBTDD-vector representation if and only if $2k < |L|$.

9.4.2 Special Cases

Unfortunately, as outlined in Section 9.3, it is hard to obtain a good bound on the growth factor k , given an arbitrary function f , without constructing its decision diagram representation first. Thus, in general, this analysis is not helpful in deciding when a given representation is best suited for the task. However, there are some special cases. For example, the growth factor for a boolean function embedded into $2 \times 2 \times 2 \times 2$ ranges between 4 and 7, whereas the branching factor of an MBTDD in this logic is 16. Moreover, our empirical evidence suggests that for functions occurring in model-checking, the growth factor is actually close to its lower bound. Therefore, we expect that, in this case, a BDD-vector representation has a significantly lower time complexity than that of an MBTDD-vector (and similarly a single MDD) representation.

The above analysis suggests that for a fixed branching factor there is no big distinction, in terms of the time complexity, between representing an mv-set using a single decision diagram, or a vector of decision diagrams. However, it hides some of the advantages of the vector representation. Let f be a function and let u and \hat{u} be its representations as an MDD and as an MBTDD-vector, respectively. Suppose further that we are trying to perform a meet (or join) of f and a crisp constant $c \in \{\top, \perp\}$, referred to as $f^u \sqcap c$ or $f^{\hat{u}} \sqcap c$. The complexity of this operation is *constant* for an MDD operation and is linear in the *number of join-irreducible elements* of the logic, i.e. $O(|\mathcal{J}(\mathcal{L})|)$, for an MBTDD-vector operation. On the other hand, if c is a non-crisp value, the complexity of the MDD operation is linear in the *size of the graph representing u* , whereas it remains linear in the *number of join-irreducibles* in the case of the MBTDD-vector operation. Thus, if we expect that a large number of operations to be performed involve non-crisp values from the logic, then we should represent mv-sets by a vector of decision diagrams. Notice that the exact same argument illustrates the differences between an ADD and a BDD-vector representations of mv-sets.

Another hidden benefit of using a vector of decision diagrams to represent an mv-set is seen during the implementation of the merge procedure for several viewpoints (see Section 3). Assume that the difference between the viewpoints is restricted to a transition relation, that is, all viewpoints are defined on a common state space S . Under this assumption, the merge problem is formulated as follows: given a collection of boolean functions $\{f_i\}$ representing a transition relation of each viewpoint ($f_i : S \times S \rightarrow \{\top, \perp\}$), construct a crisp-input function $f : S \times S \rightarrow L$ such that $f(y, y') = (x_1, x_2, \dots) \Rightarrow \forall f_i \cdot f_i(y, y') = x_i$, where L is the corresponding product logic. It is easy to show that if the result of the merge operation is to be represented as a single decision diagram (either an MDD or an ADD), the time

complexity is linear in the *size* of each viewpoint. At the same time, if the result is to be represented as a vector of decision diagrams (either MBTDD-vector, or BDD-vector), the complexity is linear in the *number* of viewpoints being merged.

The above analysis shows the difference in the worst-case time complexity between different implementations of mv-sets. However, it is not clear what it says about the running time of a model-checker that uses these implementations. The problem lies in the fact that different implementations of mv-sets dictate different encodings of *model-checking properties*, and thus different sizes of decision diagrams used during the analysis. We study the observed effects of different representations using several complete model-checking case studies in the next section.

10 Conclusion

In this section we summarize the contribution of this paper and outline venues for future research.

10.1 Summary

In this paper, we cast the problem of model-checking a class of multi-valued logics, defined over quasi-boolean algebras, in terms of operations on mv-sets. We also study two major choices for representing a membership function in an mv-set and its implementation using decision diagrams: either as one multi-valued characteristic function, implemented using MDDs or ADDs, or as a set of boolean characteristic functions, implemented using an MBTDD-vector or a BDD-vector. The naive implementation of the latter choice can be improved if we represent mv-sets as operations on join-irreducible elements of the lattices underlying the algebra.

The main result of the paper is a series of experiments with the four decision diagram representations. As is usual with decision diagrams, no one representation is a clear winner in all situations. However, we did characterize, both analytically and using empirical data, some situations where one representation performs better than the others. We further confirmed these results on a number of real case studies.

We currently have a prototype implementation of χ Chek [22] – a symbolic multi-valued model-checker that uses any of the four decision diagram varieties as the representation of mv-sets. χ Chek is available upon request.

10.2 Related Work

Work on multi-valued model-checking was surveyed in Section 1. Here we look at the work related to decision diagrams: representation of multi-valued functions and their use in non-classical model-checking.

Some of the questions relating to decision-diagram representations of finite functions, discussed in this paper, have also arisen in different domains of application. For instance, incompletely-specified boolean functions [59] have been represented as single diagrams with three outputs (true (T), false (F), don't care (DC)), i.e., via 3-valued MDD diagrams. Alternatively, each incompletely-specified function f can be represented by three diagrams, f_T , f_F , f_{DC} , with each diagram representing a characteristic function of the respective output. For example, $f_{DC}(x) = 1$ if and only if $f(x)$ is undefined. Since the value of the function f is exactly one of T , F or DC , storing two characteristic functions is sufficient. A similar encoding is used in STE to represent a four-valued Belnap function by a pair of characteristic functions [40]. These are specific cases of the general problem which we have examined in our domain

of interest: that is, given a function range of size greater than 2, what are the tradeoffs between using a single diagram with multiple terminal nodes, and a vector of diagrams with only 2 terminal nodes. Sasao et al. [54, 2] have explored these alternatives experimentally in the domain of arithmetic circuits. Orthogonal to this issue is that of the arity of *input* variables. Sasao’s experiments of [54] also deal with this question, but it has been explored considerably less. The above experiments typically concluded that there is no clear winner, and the best solution depends on the problem at hand. In this paper we characterized, in the domain of model-checking, some of the circumstances under which different implementations are more preferable.

Decision diagrams have been used as underlying data structures in several other fixpoint-based approaches over labeled graphs. For example, Kwiatkowska and her group [45] use MTBDDs (multi-terminal boolean diagrams) for model-checking probabilistic systems. These diagrams are roughly equivalent to ADDs. MTBDDs are also used by Baier and Clarke [4], for solving fixed-point problems such as shortest-path computations. These problems are defined via a numeric-valued generalization of μ -calculus, called *algebraic μ -calculus*. In both cases, the algorithms are quite similar to ours. However, these problems do not enjoy Knaster-Tarski convergence conditions; therefore, exact solutions often cannot be obtained, and approximations are used instead. Another example are Multiway Decision Graphs (MDGs) [28] that are similar in name and structure to MDDs. However, MDGs are used to represent first-order formulas and allow edges to be labeled with first-order terms that are not necessarily disjoint. This makes their operations significantly different from those of MDDs.

In our previous work, we made several attempts to implement a symbolic multi-valued model-checker [23, 19, 21, 22]. The first presentation, [23], represented each mv-set by associating each value of the algebra with a corresponding classical set, using explicit tables to compute mv-set operations and their inverses (see discussion in the beginning of Section 7). The second, [19] used MDDs, and part of our Section 6 is based on this work. Finally, [21] defined an encoding of mv-set operations using join-irreducibles, allowing us to use vectors of BDDs and forming the core of Section 7. In the current article, we unified these approaches as alternatives for implementing an mv-set library and empirically compared the effectiveness of the latter two approaches, as well as their natural extensions: ADDs and MBTDD-vectors.

10.3 Future Work

In the future, we plan to concentrate on implementing and evaluating a number of optimizations of the approach presented in this paper.

Here, we have assumed that the transition relation of a system is given by a single monolithic decision diagram. However, in many classical problems, the transition relation is too large to be represented by a single diagram. An alternative solution is to partition the transition relation into several decision diagrams [17]. The partitioning can be property-based [55], or model-based, as is currently routinely done in software model-checking [6]. Such partitioning may reduce the overall size of the representation of the transition relation, and thus result in a significant increase in the size of the systems that can be analyzed. We are interested in extending the idea of partitioning to multi-valued domains.

In our current implementation we utilize only a fraction of all possible optimizations that are currently available in the state-of-the-art BDD libraries such as CUDD [57]. Many of these optimizations, such as reordering heuristics, complement (shift) edges, etc., have a natural extension to multi-valued domains. We are interested in exploring their effect and the resulting tradeoffs between different mv-set representations.

Throughout this paper, we maintain a very strong assumption about the branching factor and the number of terminal nodes in the diagrams. These parameters are always fixed either at 2 or the size of the algebra. However, it is interesting to explore the middle ground. For example, instead of representing functions in a given algebra L by $\mathcal{J}(L)$ diagrams with 2 terminal nodes each, we can reduce the number of diagrams by allowing each diagram to have more terminal nodes. Similarly, it is possible to use any branching factor between 2 and the size of the algebra. However, a more interesting approach is to allow for the branching factor to vary throughout the diagram. The problem with using a fixed branching factor becomes evident when we consider functions that are input-crisp in almost all inputs. For example, if the branching factor is fixed at the size of the algebra, each crisp input variable must be represented by a node with a large branching factor, which results in performance penalty since the worst-time complexity is linear in the branching factor. Alternatively, with varying branching factor diagrams each crisp input can be represented by a node with a branching factor of 2, resulting in no significant loss in performance. Essentially, this is equivalent to adding types to decision diagrams, where a node in the diagram has a branching factor corresponding to the *type* of the variable it represents (i.e. boolean variables are represented by nodes with branching factor of 2, ternary with branching factor of 3, etc.). Diagrams with varying branching factor are also good candidates for representing enumerated types in both classical and multi-valued model-checking. Since, the branching factor is not fixed, there is an additional freedom in deciding how to encode a variable of an enumerated type in the diagram. For example, a variable on an enumerated type with 3 values can be encoded by a single node with a branching factor of 3, or potentially a ternary decomposition can be used for all enumerated types in the model independently of their size. We would like to further study the effects of the number of terminal nodes and the various branching factors on the performance of model-checking, and hope to develop different heuristics for finding the optimal parameters to use.

The success of model-checking depends on the ability to construct effective abstractions of a model. In recent work, we have developed a theoretical foundation for abstracting multi-valued models [39]. In practice, an abstraction is performed either syntactically, i.e., on the level of the modeling language, or on the level of decision diagrams, after the final transition relation has been computed. The second approach is facilitated by the technique of *abstract binary decision diagrams (aBDD)* [26]. We are interested in exploring the applicability of the aBDD approach to multi-valued domains to enable abstraction of a merge of different viewpoints for which a syntactic description is typically unavailable.

Finally, we are interested in carrying out more case studies of the use of χ Chek. We hope to form a better understanding of its domains of applicability and factors determining its performance.

Acknowledgments

We thank Wendy MacCaull, Christopher Thompson-Walsh and Victor Petrovykh for many interesting discussions and for their help implementing the model-checker. We are also grateful to the members of the University of Toronto formal methods reading group and the anonymous referees for helping us refine the ideas presented in this paper and improve the clarity of the presentation. This work was financially supported by NSERC and CITO.

References

- [1] S. Akers. “Binary Decision Diagrams”. *IEEE Trans. Computers*, C-27:509–516, 1978.
- [2] H.H. Babu and T. Sasao. “Representations of Multiple-Output Functions Using Binary Decision Diagrams for Characteristic Functions”. *IEICE Transactions on Fundamentals*, E82-A(11):2398–2406, 1999.
- [3] R.I. Bahar, E.A. Frohm, C.M. Gaona, G.D. Hachtel, E. Macii, A. Pardo, and F. Somenzi. “Algebraic Decision Diagrams and Their Applications”. In *IEEE/ACM International Conference on Computer-Aided Design (ICCAD’93)*, pages 188–191, Santa Clara, California, 1993. IEEE Computer Society Press.
- [4] C. Baier and E. M. Clarke. “The Algebraic Mu-Calculus and MTBDDs”. In *Proceedings of the 5th Workshop on Logic, Language, Information and Computation, (WoLLIC’98)*, pages 27–38, 1998.
- [5] T. Ball, A. Podelski, and S. Rajamani. “Boolean and Cartesian Abstraction for Model Checking C Programs”. In *Proceedings of 7th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS’01)*, volume 2031 of *LNCS*, pages 268–283, April 2001.
- [6] T. Ball and S. Rajamani. “The SLAM Toolkit”. In *Proceedings of 13th International Conference on Computer-Aided Verification (CAV’01)*, volume 2102 of *LNCS*, pages 260–264, July 2001.
- [7] I. Beer, S. Ben-David, C. Eisner, and Y. Rodeh. “Efficient Detection of Vacuity in Temporal Model Checking”. *Formal Methods in System Design*, 18(2):141–163, March 2001.
- [8] N.D. Belnap. “A Useful Four-Valued Logic”. In Dunn and Epstein, editors, *Modern Uses of Multiple-Valued Logic*, pages 30–56. Reidel, 1977.
- [9] G.C. Berney and S.M. dos Santos. *Elevator Analysis, Design and Control*. IEE Control Engineering Series 2. Peter Peregrinus Ltd., 1985.
- [10] G. Birkhoff. *Lattice Theory*. American Mathematical Society, Providence, RI, 3 edition, 1967.
- [11] L. Bolc and P. Borowik. *Many-Valued Logics*. Springer-Verlag, 1992.
- [12] G. Bruns and P. Godefroid. “Generalized Model Checking: Reasoning about Partial State Spaces”. In *Proceedings of 11th International Conference on Concurrency Theory (CONCUR’00)*, volume 1877 of *LNCS*, pages 168–182. Springer, August 2000.
- [13] G. Bruns and P. Godefroid. “Temporal Logic Query-Checking”. In *Proceedings of 16th Annual IEEE Symposium on Logic in Computer Science (LICS’01)*, pages 409–417, Boston, MA, USA, June 2001. IEEE Computer Society.
- [14] R. E. Bryant. “Graph-based algorithms for boolean function manipulation.”. *Transactions on Computers*, 8(C-35):677–691, 1986.
- [15] R. E. Bryant. “Symbolic Boolean Manipulation with Ordered Binary-Decision diagrams”. *Computing Surveys*, 24(3):293–318, September 1992.
- [16] T. Bultan, R. Gerber, and C. League. “Composite Model Checking: Verification with Type-Specific Symbolic Representations”. *ACM Transactions on Software Engineering and Methodology*, 9(1):3–50, January 2000.
- [17] J.R. Burch, E.M. Clarke, K.L. McMillan, D.L. Dill, and L.J. Hwang. “Symbolic Model Checking for Sequential Circuit Verification”. *IEEE Transactions on Computer-Aided Design of Integrated Circuits*, 13(4):401–424, June 1992.
- [18] W. Chan. “Temporal-Logic Queries”. In *Proceedings of the 12th Conference on Computer Aided Verification (CAV’00)*, volume 1855 of *LNCS*, pages 450–463, Chicago, IL, USA, July 2000. Springer.
- [19] M. Chechik, B. Devereux, and S. Easterbrook. “Implementing a Multi-Valued Symbolic Model-Checker”. In *Proceedings of TACAS’01*, volume 2031 of *LNCS*, pages 404–419. Springer, April 2001.

- [20] M. Chechik, B. Devereux, S. Easterbrook, and A. Gurfinkel. “Multi-Valued Symbolic Model-Checking”. *ACM Transactions on Software Engineering and Methodology*, 12(4):1–38, October 2003.
- [21] M. Chechik, B. Devereux, S. Easterbrook, A. Lai, and V. Petrovykh. “Efficient Multiple-Valued Model-Checking Using Lattice Representations”. In *Proceedings of 12th International Conference on Concurrency Theory (CONCUR’01)*, volume 2154 of *LNCS*, pages 451–465, Aalborg, Denmark, August 2001.
- [22] M. Chechik, B. Devereux, and A. Gurfinkel. “ χ Chek: A Multi-Valued Model-Checker”. In *Proceedings of 14th International Conference on Computer-Aided Verification (CAV’02)*, volume 2404 of *LNCS*, pages 505–509, Copenhagen, Denmark, July 2002. Springer.
- [23] M. Chechik, S. Easterbrook, and V. Petrovykh. “Model-Checking Over Multi-Valued Logics”. In *Proceedings of Formal Methods Europe (FME’01)*, volume 2021 of *LNCS*, pages 72–98. Springer, March 2001.
- [24] A. Cimatti, E.M. Clarke, F. Giunchiglia, and M. Roveri. “NUSMV: a new Symbolic Model Verifier”. In N. Halbwachs and D. Peled, editors, *Proceedings of 11th Conference on Computer-Aided Verification (CAV’99)*, volume 1633 of *LNCS*, pages 495–499, Trento, Italy, July 1999. Springer.
- [25] E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 1999.
- [26] E. Clarke, S. Jha, Y. Lu, and D. Wang. “Abstract BDDs: A Technique for Using Abstraction in Model Checking”. In *Proceedings of 10th IFIP WG 10.5 Advanced Research Working Conference on Correct Hardware Design and Verification Methods (CHARME’99)*, volume 1703 of *LNCS*, pages 172–186, Germany, September 1999. Springer.
- [27] E.M. Clarke, E.A. Emerson, and A.P. Sistla. “Automatic Verification of Finite-State Concurrent Systems Using Temporal Logic Specifications”. *ACM Transactions on Programming Languages and Systems*, 8(2):244–263, April 1986.
- [28] F. Corella, Z. Zhou, X. Song, M. Langevin, and E. Cerny. “Multiway Decision Graps for Automated Hardware Verification”. *Formal Methods in System Design*, 10(1):7–46, 1997.
- [29] B.A. Davey and H.A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, 1990.
- [30] R. Drechsler and B. Becker. *Binary Decision Diagrams: Theory and Implementation*. Kluwer Academic, 1998.
- [31] J.M. Dunn. “A Comparative Study of Various Model-Theoretic Treatments of Negation: A History of Formal Negation”. In Dov Gabbay and Heinrich Wansing, editors, *What is Negation*. Kluwer Academic Publishers, 1999.
- [32] S. Easterbrook and M. Chechik. “A Framework for Multi-Valued Reasoning over Inconsistent Viewpoints”. In *Proceedings of International Conference on Software Engineering (ICSE’01)*, pages 411–420, Toronto, Canada, May 2001. IEEE Computer Society Press.
- [33] M. Fujita, P. C. McGeer, and J. C.-Y. Yang. “Multi-Terminal Binary Decision Diagrams: An Efficient Data Structure for Matrix Representation”. *Formal Methods in System Design: An International Journal*, 10(2/3):149–169, April 1997.
- [34] P. Godefroid and R. Jagadeesan. “Automatic Abstraction Using Generalized Model-Checking”. In *Proceedings of 14th International Conference on Computer-Aided Verification (CAV’02)*, volume 2404 of *LNCS*, pages 137–150, July 2002.
- [35] A. Gurfinkel and M. Chechik. “Generating Counterexamples for Multi-Valued Model-Checking”. In *Proceedings of Formal Methods Europe (FME’03)*, volume 2805 of *LNCS*, pages 503–521, Pisa, Italy, September 2003. Springer.

- [36] A. Gurfinkel and M. Chechik. “Multi-Valued Model-Checking via Classical Model-Checking”. In *Proceedings of 14th International Conference on Concurrency Theory (CONCUR’03)*, volume 2761 of *LNCS*, pages 263–277. Springer, September 2003.
- [37] A. Gurfinkel and M. Chechik. “How Vacuous Is Vacuous?”. In *Proceedings of 10th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS’04)*, volume 2988 of *LNCS*, pages 451–466, Barcelona, Spain, March 2004. Springer.
- [38] A. Gurfinkel, M. Chechik, and B. Devereux. “Temporal Logic Query Checking: A Tool for Model Exploration”. *IEEE Transactions on Software Engineering*, 29(10):898–914, October 2003.
- [39] A. Gurfinkel, O. Wei, and M. Chechik. “Systematic Construction of Abstractions for Model-Checking”. In *Proceedings of 7th International Conference on Verification, Model-Checking, and Abstract Interpretation (VMCAI’06)*, volume 3855 of *LNCS*, pages 381–397, Charleston, SC, January 2006. Springer.
- [40] S. Hazelhurst and C. H. Seger. “Model Checking Lattices: Using and Reasoning about Information Orders for Abstraction”. *Logic Journal of the IGPL*, 7(3):375–411, May 1999.
- [41] M. Huth and M. Ryan. *Logic in Computer Science: Modeling and Reasoning About Systems*. Cambridge University Press, 2000.
- [42] S. C. Kleene. *Introduction to Metamathematics*. New York: Van Nostrand, 1952.
- [43] T. Kropf. *An Introduction to Formal Verification*. Springer-Verlag, 1999.
- [44] O. Kupferman and M. Vardi. “Vacuity Detection in Temporal Model Checking”. *International Journal on Software Tools for Technology Transfer (STTT)*, 4(2):224–233, February 2003.
- [45] M.Z. Kwiatkowska, G. Norman, D.A. Parker, and R. Segala. “Symbolic Model Checking of Probabilistic Processes Using MTBDDs and the Kronecker Representation”. In *Proceedings of TACAS 2000*, volume 1587 of *LNCS*. Springer-Verlag, 2000.
- [46] K.L. McMillan. *Symbolic Model Checking*. Kluwer Academic, 1993.
- [47] S. Minato. “Arithmetic Boolean Expression Manipulator using BDDs”. *Formal Methods in System Design*, 10:221–242, 1997.
- [48] J.J. Pazos-Arias and J.G. Duque. “SCTL-MUS: A Formal Methodology for Software Development of Distributed Systems. A Case Study”. *Formal Aspects of Computing*, 13:50–91, 2001.
- [49] M.C. Plath and M.D. Ryan. A semantics of a feature construct for SMV: A case study in non-monotonic composition. Technical report, School of Computer Science, University of Birmingham, 1999. Available as <ftp://ftp.cs.bham.ac.uk/pub/tech-reports/1999/CSR-99-10.ps.gz>.
- [50] M.C. Plath and M.D. Ryan. “SFI: A Feature Integration Tool”. In R. Berghammer and Y. Lakhnech, editors, *Tool Support for System Specification, Development and Verification*, Advances in Computer Science, pages 201–216. Springer, 1999.
- [51] H. Rasiowa. *An Algebraic Approach to Non-Classical Logics. Studies in Logic and the Foundations of Mathematics*. Amsterdam: North-Holland, 1978.
- [52] T.W. Reps, M. Sagiv, and R. Wilhelm. “Static Program Analysis via 3-Valued Logic”. In *Proceedings of 16th International Conference on Computer-Aided Verification (CAV’04)*, volume 3114 of *LNCS*, pages 15–30, 2004.
- [53] M. Sabetzadeh and S.M. Easterbrook. “Analysis of Inconsistency in Graph-Based Viewpoints: A Category-Theoretic Approach”. In *Proceedings of 18th IEEE International Conference on Automated Software Engineering (ASE’03)*, pages 12–21. IEEE Computer Society, October 2003.

- [54] T. Sasao and J.T. Butler. “A Method to Represent Multiple-Output Switching Functions Using Multi-Valued Decision Diagrams”. In *Proceedings of IEEE International Symposium on Multiple-Valued Logic (ISMVL'96)*, pages 248–254, Santiago de Compostela, Spain, 1996.
- [55] R. Sebastini, E. Singerman, S. Tonetta, and M.Y. Vardi. “GSTE Is Partitioned Model Checking”. In *Proceedings of 16th International Conference on Computer-Aided Verification (CAV'04)*, volume 3114 of *LNCS*, pages 229–241, 2004.
- [56] S. Shoham and O. Grumberg. “A Game-Based Framework for CTL Counter-Examples and 3-Valued Abstraction-Refinement”. In *Proceedings of the 15th Conference on Computer Aided Verification (CAV'03)*, volume 2725 of *Lecture Notes in Computer Science*, pages 275–287. Springer, July 2003.
- [57] F. Somenzi. “Binary Decision Diagrams”. In Manfred Broy and Ralf Steinbrüggen, editors, *Calculational System Design*, volume 173 of *NATO Science Series F: Computer and Systems Sciences*, pages 303–366. IOS Press, 1999.
- [58] A. Srinivasan, T. Kam, S. Malik, and R.E. Brayton. “Algorithms for Discrete Function Manipulation”. In *IEEE/ACM International Conference on Computer-Aided Design (ICCAD'90)*, pages 92–95, Santa Clara, CA, USA, November 1990. IEEE Computer Society.
- [59] I. Wegener. “*Branching Programs and Binary Decision Diagrams: Theory and Applications*”. Monographs on Discrete Mathematics and Applications. SIAM, 2000.
- [60] B. Yang, R.E. Bryant, D. R. O'Hallaron, A. Biere, O. Coudert, G. Janssen, R. K. Ranjan, and F. Somenzi. “A Performance Study of BDD-Based Model Checking”. In *Formal Methods in Computer-Aided Design*, pages 255–289, 1998.
- [61] L.A. Zadeh. “Fuzzy Sets”. In R. R. Yager, S. Ovchinnikov, R. M. Tong, and H. T. Nguyen, editors, *Fuzzy Sets and Applications: Selected Papers by L.A. Zadeh*, pages 29–44, New York, 1987. John Wiley & Sons, Inc.