

NASA/CR-2010-216710



# Model-Checking with Edge-Valued Decision Diagrams

*Pierre Roux*  
*Ecole Normale Supérieure de Lyon, France*

*Radu I. Siminiceanu*  
*National Institute of Aerospace, Hampton, Virginia*

## NASA STI Program . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NASA Aeronautics and Space Database and its public interface, the NASA Technical Report Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- **TECHNICAL PUBLICATION.** Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.
- **TECHNICAL MEMORANDUM.** Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.
- **CONTRACTOR REPORT.** Scientific and technical findings by NASA-sponsored contractors and grantees.

- **CONFERENCE PUBLICATION.** Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.
- **SPECIAL PUBLICATION.** Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.
- **TECHNICAL TRANSLATION.** English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include creating custom thesauri, building customized databases, and organizing and publishing research results.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at <http://www.sti.nasa.gov>
- E-mail your question via the Internet to [help@sti.nasa.gov](mailto:help@sti.nasa.gov)
- Fax your question to the NASA STI Help Desk at 443-757-5803
- Phone the NASA STI Help Desk at 443-757-5802
- Write to:  
NASA STI Help Desk  
NASA Center for AeroSpace Information  
7115 Standard Drive  
Hanover, MD 21076-1320

NASA/CR-2010-216710



# Model-Checking with Edge-Valued Decision Diagrams

*Pierre Roux*  
*Ecole Normale Supérieure de Lyon, France*

*Radu I. Siminiceanu*  
*National Institute of Aerospace, Hampton, Virginia*

National Aeronautics and  
Space Administration

Langley Research Center  
Hampton, Virginia 23681-2199

Prepared for Langley Research Center  
under Cooperative Agreement NNX08AC59A

June 2010

## Acknowledgments

This work has been supported by NASA Cooperative Agreement NNX08AC59A, subagreement number 27-001310.

Trade names and trademarks are used in this report for identification only. Their usage does not constitute an official endorsement, either expressed or implied, by the National Aeronautics and Space Administration.

Available from:

NASA Center for AeroSpace Information  
7115 Standard Drive  
Hanover, MD 21076-1320  
443-757-5802

# Abstract

We describe an algebra of Edge-Valued Decision Diagrams (EVMDDs) to encode arithmetic functions and its implementation in a model checking library along with state-of-the-art algorithms for building the transition relation and the state space of discrete state systems.

We provide efficient algorithms for manipulating EVMDDs and give upper bounds of the theoretical time complexity of these algorithms for all basic arithmetic and relational operators. We also demonstrate that the time complexity of the generic recursive algorithm for applying a binary operator on EVMDDs is no worse than that of Multi-Terminal Decision Diagrams.

We have implemented a new symbolic model checker with the intention to represent in one formalism the best techniques available at the moment across a spectrum of existing tools: EVMDDs for encoding arithmetic expressions, identity-reduced MDDs for representing the transition relation, and the saturation algorithm for reachability analysis. We compare our new symbolic model checking EVMDD library with the widely used CUDD package and show that, in many cases, our tool is several orders of magnitude faster than CUDD.

## 1 Introduction

Binary decision diagrams (BDD) [4] have revolutionized the reachability analysis and model checking technology. Arithmetic decision diagrams [3], also called Multi Terminal Binary Decision Diagrams (MTBDD) [9] are the natural extension of regular BDDs to arithmetic functions. They take advantage of the symbolic encoding scheme of BDDs, but functions with large co-domains do not usually have a very compact representation because there are fewer chances for suffixes to be shared.

Edge-valued decision diagrams have been previously introduced, but only scarcely used. An early version, the edge valued binary decision diagrams (EVBDD) [13, 14], is particularly useful when representing both arithmetic and logic functions, which is the case for discrete state model checking. However, EVBDDs have only been applied to rather obscure applications, such as computing the probability spectrum and the Reed-Muller spectrum of (pseudo)-Boolean functions.

Binary Moment Diagrams [5] were designed to overcome the limitations of BDDs when encoding multiplier functions. However, their efficiency seems to be limited only to this particular type of functions. A new canonization rule for edge-valued decision diagrams enabling them to encode functions in  $\mathbb{Z} \cup \{+\infty\}$  was introduced in [7] along with an extension to multi-way diagrams (MDD) [12], but, again, this was applied to a very specific task, of finding minimum length counterexamples for safety properties. Later, EVMDDs have been also used for partial reachability analysis.

In this paper we first present a theoretical comparison between EVMDDs and MTMDDs for building the transition relation of discrete state systems before dealing with an implementation in a model checker along with state-of-the-art algorithms for state space construction.

## 2 Background

### 2.1 Discrete-state Systems

A discrete-state model is a triple  $(S, S_0, T)$ , where the discrete set  $S$  is the *potential state space* of the model; the set  $S_0 \subseteq S$  contains the *initial states*; and  $T : S \rightarrow 2^S$  is the *transition function* specifying which states can be reached from a given state in one step, which we extend to sets:  $T(X) = \bigcup_{i \in X} T(i)$ . We consider structured systems modeled as a collection of  $K$  *submodels*. A (global) system state  $i$  is then a  $K$ -tuple  $(i_K, \dots, i_1)$ , where  $i_k$  is the *local state* for submodel  $k$ , for  $K \geq k \geq 1$ , and  $S$  is defined as  $S_K \times \dots \times S_1$ , the cross-product of  $K$  local state spaces  $S_k$ , which

we abstract to  $\{0, \dots, n_k - 1\}$ . The (*reachable*) *state space*  $R \subseteq S$  is the smallest set containing  $S_0$  and closed with respect to  $T$ , i.e.

$$R = S_0 \cup T(S_0) \cup T(T(S_0)) \cup \dots = T^*(S_0).$$

Thus,  $R$  is least fixpoint of function  $X \mapsto S_0 \cup T(X)$ .

## 2.2 Symbolic State–space Generation: Breadth–first vs. Saturation

The traditional approach to generate the reachable states of a system is based on a breadth–first traversal, as derived from classical fixed–point theory, and applies a monolithic  $T$  (even when encoded as  $\bigcup_{e \in E} T_e$ ). After  $d$  iterations, the currently–known state space contains all states whose distance from any state in  $S_0$  is at most  $d$ . However, recent advances have shown that non–BFS, guided (or chaotic) exploration can result in a better iteration strategy.

An example is the *saturation* algorithm introduced in [6], which exhaustively *fires* all events of  $E_k$  in an MDD node at level  $k^1$ , thereby greedily bringing the node to its final “saturated” form.

## 2.3 Decision Diagrams

We assume implicitly that the decision diagrams are *ordered*, i.e. the variables labeling nodes along any path from the root must follow the order  $x_K, \dots, x_1$ . Ordered DDs can be either *reduced* (no duplicate nodes and no node with all edges pointing to the same node, but edges possibly spanning multiple levels) or *quasi–reduced* (no duplicate nodes, and all edges spanning exactly one level), either form being *canonical*.

We also adopt the extension of BDDs to integer variables, i.e., multi–valued decision diagrams (MDDs) [12]. MDDs are more naturally suited to represent the state space of arbitrary discrete systems than BDDs, since no binary encoding must be used to represent the local states for level  $k$  when  $n_k > 2$ . An even more important reason to use MDDs is that they allow us to better exploit the *event locality* present in systems exhibiting a globally–asynchronous locally–synchronous behavior.

# 3 EVMDDs

## 3.1 Definition

**Definition 3.1** An *EVMDD* on a group  $(G, *)$ , is a pair  $A = \langle v, n \rangle$  where  $v \in G$  will also be noted  $A.\text{val}$  and  $n$ , also noted  $A.\text{node}$ , is a node.

A node  $n$  is either the unique terminal node  $\langle 0, e \rangle$  where  $e$  is the identity element of  $G$  or a pair  $\langle k, p \rangle$  where  $1 \leq k \leq K$  and  $p$  is an array of edges of size  $n_k$ . The first element of the pair will be denoted  $k = n.\text{level}$  and, when relevant, the element of index  $i$  in the array will be denoted by  $n[i]$ .

Additionally, the notation  $n[i_k, \dots, i_{k'}]$  is used as a shortcut for  $n[i_k] \dots [i_{k'}].\text{node}$ .

**Definition 3.2** An ordered *EVMDD* is an *EVMDD* in which every node  $n$  satisfies

$$\forall i \in S_{n.\text{level}} \cdot n[i].\text{node.level} < n.\text{level}$$

As already mentioned, we only consider ordered EVMDDs. The canonicity of unordered EVMDD is significantly more complex to establish.

<sup>1</sup> $T$  is then encoded as a disjunction  $\bigcup_{e \in E} T_e$  of events  $e$  and  $E$  is then divided in  $\bigcup_{1 \leq k \leq K} E_k$  with each  $E_k$  grouping events not depending on submodels above  $k$  and not affecting them.

**Example 3.1** *Graphs are a convenient representation for EVMDDs. For ordered EVMDDs, the graph directed by node levels is acyclic. Graphically, the terminal node is represented by a circle at bottom and internal nodes are drawn above according to their level, with all edges pointing down to children nodes. Examples of graph representation of EVMDDs are given in Figure 1.*




Figure 1. Three EVMDDs on group  $(\mathbb{Z}, +)$  representing the same function  $f : \{0, 1\}^2 \rightarrow \mathbb{Z}, (x_2, x_1) \mapsto x_2 \times x_1$ .

**Definition 3.3** *Given a node  $n$  with  $n.\text{level} = k$  and  $i_k, \dots, i_{k'} \in S_k \times \dots \times S_{k'}$ , we define  $n(i_k, \dots, i_{k'})$*

$$n(i_k, \dots, i_{k'}) = \begin{cases} n[i_k].\text{val} & \text{if } n[i_k].\text{node.level} < k' \\ n[i_k].\text{val} * n[i_k].\text{node}(i_{n[i_k].\text{node.level}}, \dots, i_{k'}) & \text{if } n[i_k].\text{node.level} \geq k' \end{cases}$$

This allows the definition of the *represented function* for any EVMDD  $A$  as

$$f : \begin{array}{ccc} S & \rightarrow & G \\ (i_K, \dots, i_1) & \mapsto & A.\text{val} * A.\text{node}(i_{A.\text{node.level}}, \dots, i_1) \end{array}$$

In other words,  $n(i_k, \dots, i_{k'})$  is the repetitive application of law  $*$  on edge values along the path going down from node  $n$  and following directions given by  $(i_k, \dots, i_{k'})$ . Hence  $f(i_K, \dots, i_1)$  is  $*$  on a path from the root to the terminal node of  $A$ .

In this setup, every EVMDD  $A$  represents a function  $f : S \rightarrow G$ . The reciprocal is also true: given a function  $f$ , an EVMDD  $A$  representing  $f$  can be built by setting the values of all edges of  $A$  to the identity element  $e$  of  $G$ , except those pointing to the terminal node, which take proper values  $f(i_k, \dots, i_1)$  according to the incoming path leading to it.

**Example 3.2** *In Figure 1, the EVMDD  $B$  is built from  $f : \{0, 1\}^2 \rightarrow \mathbb{Z}, (x_2, x_1) \mapsto x_2 \times x_1$  according to the method explained above.*

**Definition 3.4** *A redundant node  $n$  has all outgoing edges identical*

$$\forall i, j \in S_{n.\text{level}}. n[i] = n[j]$$

*A reduced EVMDD contains no duplicate or redundant nodes.*

**Definition 3.5** A quasi-reduced *EVMDD* contains no duplicate nodes and all internal nodes  $n$  have all the descendants on the level below

$$\forall i \in S_{n.\text{level}} . n[i].\text{node.level} = n.\text{level} - 1$$

From any *EVMDD*  $A$ , we can build a reduced *EVMDD* representing the same function by just deleting nodes with all children identical and redirecting the incoming edges to its unique descendant. Similarly, from any *EVMDD*  $A$ , a quasi reduced *EVMDD* can be built by adding nodes with all children identical on edges spanning multiple levels.

**Example 3.3** In Figure 1 on the preceding page  $A$  is reduced, whereas  $B$  and  $C$  are quasi-reduced.

For the sake of simplicity we will only consider quasi-reduced *EVMDDs* in the following discussion. However, proofs and algorithms for the reduced version are very similar, only slightly more evolved in order to deal with edges skipping levels. We will turn back to reduced *EVMDDs* only for implementation since they are never larger in size than their quasi-reduced counterpart, hence could have more efficient algorithms.

As can be seen in the previous example, even when restricting to quasi-reduced diagrams, the *EVMDD* representation of a function  $f$  may not be uniquely defined.

**Definition 3.6** A canonical node is either the terminal node or a node  $n$  such that  $n[0].\text{val} = e$ .

A canonical *EVMDD* is an *EVMDD* in which all nodes are canonical.

It can be proved that for every function  $f$ , there exists a unique canonical *EVMDD* representing it [7].

In the following, *EVMDDs* are assumed to be canonical.

## 3.2 Extensions

*EVMDDs* can be used even when the algebraic structure is not a group. [7] offers a canonization rule for  $\mathbb{N} \cup \{+\infty\}$  and  $(\mathbb{Z}, \times)$  can be handled with the canonization rule “ $\text{gcd}\{n[i].\text{val} \mid i \in S_{n.\text{level}}\} = 1$  and  $(n[0].\text{val}, \dots, n[n_{n.\text{level}}].\text{val}) \geq_{\text{lex}} 0$ ”.

It is also interesting to notice that *EVMDDs* are just a generalization of binary decision diagrams with complemented edges, as presented for example in [11]. Indeed, they are edge-valued diagrams on  $G = \mathbb{Z}/2\mathbb{Z}$  and complemented (respectively not complemented) edges corresponding to value 1 (respectively 0).

## 4 *EVMDDs* compared to *MTMDDs*

*MTBDDs* are commonly used in model checking to build the transition relation of discrete-state systems. In this section we show that *EVMDDs* are at least as suited for that purpose and oftentimes significantly better. In the remainder of this section, we pick, often without loss of generality,  $(G, *) = (\mathbb{Z}, +)$ .

### 4.1 Space Complexity

As stated in section 2.2.6 of [13].

**Theorem 4.1** For any function  $f$ , the number of nodes of the *EVMDD* representing  $f$  is at most the number of nodes of the *MTMDD* representing the same function  $f$ .

**Proof.** Let  $A$  be the MTMDD representing  $f$ . From  $A$  we construct the EVMDD (not in canonical form)  $A_0$  by replacing each edge from level 1 to a terminal with value  $v$  with an edge with value  $v$  to the unique terminal node 0 and associating value 0 to all other edges (see Figure 2 for an example)<sup>2</sup>. Then, we iteratively compute the EVMDD  $A_k$ , for each  $k$  from 1 to  $n$ , through the following process:




Figure 2. Building the EVMDD  $A_0$  (right) from MTMDD  $A$  (left)

- for each node  $n$  at level  $k$ , subtract  $n[0].val$  from all outgoing edges and add this value to all incoming edges;
- merge all duplicate nodes at level  $k$  (by duplicate nodes we mean two nodes having edges  $x_i$  holding same value and pointing to same children for each  $i$  in the range of variable  $x_k$ ).

See Figure 3 on the next page for an example.

To prove that  $A_k$  and  $A$  represent the same function, it is sufficient to see that  $A$  and  $A_0$  represent the same function and that the iterative transformation preserves the sum of values on any path  $(i_K, \dots, i_1)$  from the root of the diagram to the unique terminal node (plus the value of the root's incoming edge)

$$A_k.val + A_k.node(i_K, \dots, i_1) = A_{k-1}.val + A_{k-1}.node(i_K, \dots, i_1)$$

Since  $A_K$  is in canonical form and since for each  $k$ , the number of nodes of  $A_k$  is at most the number of nodes of  $A_{k-1}$ , we can conclude that the size of an EVMDD is never larger than that of the corresponding MTMDD.  $\square$

This doesn't prove that EVMDDs always require less memory than MTMDDs since they need extra space to store the edge values, but no worse than up to a small factor<sup>3</sup>. On the other hand, EVMDDs can be exponentially better than MTMDD in some cases. For example, the function

$$\begin{aligned} \{0, B-1\}^K &\rightarrow \mathbb{Z} \\ (i_K, \dots, i_1) &\mapsto \sum_{k=1}^K i_k B^{k-1} \end{aligned}$$

<sup>2</sup>This process is similar to the one used in section 3.1 on page 2 to prove that every function can be represented by an EVMDD.

<sup>3</sup>Usually 2, assuming that edge values are as big as node pointers.




Figure 3. Constructing EVMDD  $A_1$  (right) from EVMDD  $A_0$  (left)

where ( $B \geq 2$ ), requires  $\frac{B^{K+1} - 1}{B - 1}$  nodes in its MTMDD representation<sup>4</sup> whereas it can be represented as an EVMDD with only  $K + 1$  nodes<sup>5</sup>.

## 4.2 Time complexity

What makes decision diagrams a useful data structure for symbolic model-checking is not only their space efficiency but the ability to efficiently compute common operations.

Section 2 of [9] gives an algorithm to compute any binary operation on MTBDDs. The *apply* algorithm can be easily generalized to MDDs for any  $n$ -ary operator  $\square_n$ . It computes its result in time  $O\left(\prod_{i=1}^n |f_i|\right)$ , where  $|f_i|$  is the size of MTMDD representing operand  $i$  (in nodes).

Section 2.2 of [13] gives the equivalent Algorithm 1 on the facing page for edge-valued decision diagrams. This is the binary version, but  $n$ -ary one is very similar for any  $n \in \mathbb{N}^*$ .

As stated in section 2.2.6 of [13]

**Theorem 4.2** *The number of recursive calls of the above apply algorithm is the same for MTMDDs and for EVMDDs representing the same functions.*

**Lemma 4.1** *Two paths  $(i_K, \dots, i_{k+1})$  and  $(j_K, \dots, j_{k+1})$  lead to the same node in  $A$*

$$A.\text{node}[i_K, \dots, i_{k+1}] = A.\text{node}[j_K, \dots, j_{k+1}]$$

*if and only if they lead to the same node in  $A_{k-1}$*

$$A_{k-1}.\text{node}[i_K, \dots, i_{k+1}] = A_{k-1}.\text{node}[j_K, \dots, j_{k+1}]$$

**Proof.**  $A$  and  $A_{k-1}$  are identical from level  $k + 1$  to level  $K$ .  $\square$

<sup>4</sup>Since all terminal values are distinct.

<sup>5</sup>More generally, any linear function requires only one node per level in its EVMDD representation.

---

**Algorithm 1** computes any binary operation  $\square_2$  on EVMDDs  $\langle v, n \rangle$  and  $\langle v', n' \rangle$

---

apply( $\square_2$  : edge \* edge  $\rightarrow$  edge,  $\langle v, n \rangle$  : edge,  $\langle v', n' \rangle$  : edge) : edge

$k \leftarrow n.\text{level} \ // = n'.\text{level}$  since EVMDDs are quasi-reduced

// base case

**if**  $k = 0$  **then**

**return**  $\langle v \square_2 v', t \rangle$  //  $t$  is the unique terminal node

// lookup in cache

**if** CacheFind( $\square_2$ ,  $\langle v, n \rangle$ ,  $\langle v', n' \rangle$ ,  $\langle m, r \rangle$ ) **then**

**return**  $\langle m, r \rangle$

$r \leftarrow \text{NewNode}(k)$

**for**  $i = 0$  **to**  $n_k - 1$  **do**

$r[i] \leftarrow \text{apply}(\square_2, \langle v + n[i].\text{val}, n[i].\text{node} \rangle, \langle v' + n'[i].\text{val}, n'[i].\text{node} \rangle)$

$m \leftarrow r[0].\text{val}$

**for**  $i = 0$  **to**  $n_k - 1$  **do**

$r[i].\text{val} \leftarrow r[i].\text{val} - m$

// check if a node identical to  $r$  already exists

$r \leftarrow \text{FindOrAdd}(r)$

// save result in cache

CacheInsert( $\square_2$ ,  $\langle v, n \rangle$ ,  $\langle v', n' \rangle$ ,  $\langle m, r \rangle$ )

**return**  $\langle m, r \rangle$

---

**Lemma 4.2** Two paths  $(i_K, \dots, i_{k+1})$  and  $(j_K, \dots, j_{k+1})$  lead to the same node in  $A_K$  with the same value

$$\begin{aligned} & A_K.\text{node}[i_K, \dots, i_{k+1}] = A_K.\text{node}[j_K, \dots, j_{k+1}] \\ \text{and} \quad & A_K.\text{node}(i_K, \dots, i_{k+1}) = A_K.\text{node}(j_K, \dots, j_{k+1}) \end{aligned}$$

if and only if they lead to the same node in  $A_k$  with the same value

$$\begin{aligned} & A_k.\text{node}[i_K, \dots, i_{k+1}] = A_k.\text{node}[j_K, \dots, j_{k+1}] \\ \text{and} \quad & A_k.\text{node}(i_K, \dots, i_{k+1}) = A_k.\text{node}(j_K, \dots, j_{k+1}) \end{aligned}$$

**Proof.** First, let us prove by induction on  $l$ , from  $k$  to  $K$ , that:

$$\begin{aligned} & A_k.\text{node}[i_K, \dots, i_{k+1}] = A_l.\text{node}[i_K, \dots, i_{k+1}] \\ \text{and} \quad & A_k.\text{val} + A_k.\text{node}(i_K, \dots, i_{k+1}) = A_l.\text{val} + A_l.\text{node}(i_K, \dots, i_{k+1}) \end{aligned}$$

- for  $l = k$ , the property trivially holds;
- if the property holds for a value of  $l$  between  $k$  and  $K - 1$ , it still holds for  $l + 1$ , since the computation of  $A_{l+1}$  from  $A_l$  can only merge duplicate nodes at level  $l + 1 \geq k + 1$ .

□

**Proof.** [Theorem 4.2] We show the proof for the unary version of the algorithm. Proofs for other versions are similar, only a bit more verbose.

If we do not take into consideration the caches, the two algorithms are obviously equivalent so what we need to prove is that a cache hit occurs in the EVMDD *apply* algorithm with diagram  $A_K$  if and only if it occurs in the MTMDD algorithm with diagram  $A$ . In other words, we have to prove for every two paths  $(i_K, \dots, i_{k+1})$  and  $(j_K, \dots, j_{k+1})$  that they lead to the same node in  $A$

$$A.\text{node}[i_K, \dots, i_{k+1}] = A.\text{node}[j_K, \dots, j_{k+1}]$$

if and only if they reach the same node in  $A_K$  with the same value:

$$\begin{aligned} & A_K.\text{node}[i_K, \dots, i_{k+1}] = A_K.\text{node}[j_K, \dots, j_{k+1}] \\ \text{and} \quad & A_K.\text{node}(i_K, \dots, i_{k+1}) = A_K.\text{node}(j_K, \dots, j_{k+1}) \end{aligned}$$

Therefore, from lemmas 4.1 on page 6 and 4.2 on the preceding page it remains to prove that two paths  $(i_K, \dots, i_{k+1})$  and  $(j_K, \dots, j_{k+1})$  lead to the same node in  $A_{k-1}$

$$A_{k-1}.\text{node}[i_K, \dots, i_{k+1}] = A_{k-1}.\text{node}[j_K, \dots, j_{k+1}] \quad (1)$$

if and only if they reach the same node in  $A_k$  with the same value:

$$\begin{aligned} & A_k.\text{node}[i_K, \dots, i_{k+1}] = A_k.\text{node}[j_K, \dots, j_{k+1}] \\ \text{and} \quad & A_k.\text{node}(i_K, \dots, i_{k+1}) = A_k.\text{node}(j_K, \dots, j_{k+1}) \end{aligned} \quad (2)$$

- Let us prove that (1) implies (2). The first part of (2) is obvious: if two paths lead to the same node before merging, this still holds after merging. Second part is slightly more complex. Indeed,  $A_k.\text{node}(i_K, \dots, i_{k+1})$  is just the value attached to edge  $i_{k+1}$  of node  $A_k.\text{node}[i_K, \dots, i_{k+2}]$  since all other edges hold value 0. This value is the one reported up when constructing  $A_k$  from  $A_{k-1}$  for node  $A_k.\text{node}[i_K, \dots, i_{k+1}] = A_k.\text{node}[j_K, \dots, j_{k+1}]$  so it is the same for  $A_k.\text{node}(j_K, \dots, j_{k+1})$ , which is the second part of (2).
- In the other direction: (2) implies (1).  
If (2) holds then  $A_{k-1}.\text{node}[i_K, \dots, i_{k+1}]$  and  $A_{k-1}.\text{node}[j_K, \dots, j_{k+1}]$  are identical (for if they were not, then (2) does not hold on  $A_k$  computed from  $A_{k-1}$ ). Hence (1), since  $A_{k-1}$  doesn't contain any duplicate at level  $k$ .

□

In conclusion, EVMDD computations are at least not slower than the MTMDD equivalent. However, particular operators  $\square_n$  may enable much better algorithms on EVMDDs.

#### 4.2.1 Addition of Constant: EVMDD + c

Adding a constant  $c$  to an EVMDD  $\langle v, n \rangle$  is just computing  $\langle c+v, n \rangle$ , which can be done in constant time.

#### 4.2.2 Multiplication with Scalar: EVMDD $\times c$

As stated in section 2.2.2 of [13], computing  $f \times c$  is just multiplying each edge value of EVMDD representing  $f$  by  $c$ , which can be done in time  $O(|f|)$ .

#### 4.2.3 Addition: EVMDD + EVMDD

As stated in section 2.2.2 of [13], addition satisfies the property

$$\langle v, n \rangle + \langle v', n' \rangle = (\langle 0, n \rangle + \langle 0, n' \rangle) + (v + v')$$

allowing to cache only edges with value 0, therefore leading to a slightly modified *apply*, Algorithm 2 on the next page The complexity of the algorithm to compute  $f + g$  is  $O(|f| |g|)$ .

It's also interesting to notice that this algorithm offers a simple upper bound to the size of the result of an addition

$$|f + g| \leq |f| |g|$$

#### 4.2.4 Remainder and Euclidean Division: EVMDD $\%c$ and EVMDD $/c$

As stated in section 2.2.4 of [13], there is no need to cache values equal modulo  $c$ , hence the complexity of these algorithms is  $O(|f|c)$ .

---

**Algorithm 2** computes sum of EVMDDs  $\langle v, n \rangle$  and  $\langle v', n' \rangle$

---

```

plus( $\langle v, n \rangle$  : edge,  $\langle v', n' \rangle$  : edge) : edge
   $k \leftarrow n.\text{level} \ // = n'.\text{level}$  since EVMDDs are quasi-reduced

  // base case
  if  $k = 0$  then
    return  $\langle v + v', t \rangle$  //  $t$  is the unique terminal node

  // lookup in cache
  if CacheFind(+,  $\langle 0, n \rangle$ ,  $\langle 0, n' \rangle$ ,  $r$ ) then
    return  $\langle v + v', r \rangle$ 

   $r \leftarrow \text{NewNode}(k)$ 
  for  $i = 0$  to  $n_k - 1$  do
     $r[i] \leftarrow \text{plus}(n[i], n'[i])$ 

  // check if a node identical to  $r$  already exists
   $r \leftarrow \text{FindOrAdd}(r)$ 

  // save result in cache
  CacheInsert(+,  $\langle v, n \rangle$ ,  $\langle v', n' \rangle$ ,  $\langle m, r \rangle$ )

  return  $\langle v + v', r \rangle$ 

```

---

#### 4.2.5 Minimum and Maximum

Per section 2.2.3 of [13],  $\max \{f(x) \mid x \in S\}$  and  $\min \{f(x) \mid x \in S\}$  can be easily computed by traversing the graph and caching the result for each node, hence the complexity is  $O(|f|)$ .

#### 4.2.6 Multiplication: EVMDD $\times$ EVMDD

As stated in section 2.2.5 of [13], the result of a multiplication can have an EVMDD representation of exponential size in terms of the operands. For example, let  $S$  be  $\{0, 1\}^K$ ,  $f : (x_K, \dots, x_1) \mapsto \sum_{k=2}^K x_k 2^{k-2}$  and  $g : (x_K, \dots, x_1) \mapsto x_1$ ,  $f$  and  $g$  have both an EVMDD representation with  $K + 1$  nodes whereas  $fg$  needs  $2^K$  nodes (here again, we see the importance of variable ordering, putting  $x_1$  at top allows  $fg$  to be represented with  $2K$  nodes). Therefore, we cannot expect to find an algorithm with better worst-case complexity. However, the following equation

$$\langle v, n \rangle \times \langle v', n' \rangle = vv' + v\langle 0, n' \rangle + v'\langle 0, n \rangle + \langle 0, n \rangle \times \langle 0, n' \rangle$$

suggests the alternative Algorithm 3 on the following page. The first product is an integer multiplication done in constant time. The next two are multiplications by a constant done in  $O(|f|)$  and  $O(|g|)$ , respectively. The last one is done through recursive calls. The first addition takes constant time, the second one takes  $O(|f| |g|)$  and produce a result of size at most  $|f| |g|$ , hence a cost of  $O(|f| |g| |fg|)$  for the last addition. The function `times` is called  $O(|f| |g|)$  times, hence a final complexity of  $O(|f|^2 |g|^2 |fg|)$ .

Although we were unable to theoretically compare this algorithm to the generic Algorithm 1 on page 7, it seems to perform far better in practice when the size of the result is moderate.

#### 4.2.7 Relational Operators: EVMDD $< c$

Relational operators ( $<$ ,  $>$ ,  $\leq$ ,  $\geq$ ,  $=$  and  $\neq$ ) are used in building transition relations. Unfortunately, the corresponding operations remain somewhat expensive. But, as stated in section 2.2.3 of [13],

---

**Algorithm 3** computes product of EVMDDs  $\langle v, n \rangle$  and  $\langle v', n' \rangle$

---

```

times( $\langle v, n \rangle$  : edge,  $\langle v', n' \rangle$  : edge) : edge
   $k \leftarrow n.\text{level}$  // =  $n'.$ level since EVMDDs are quasi-reduced

  // base case
  if  $k = 0$  then
    return  $\langle v \times v', t \rangle$  //  $t$  is the unique terminal node

  // lookup in cache
  if CacheFind( $\times$ ,  $\langle 0, n \rangle$ ,  $\langle 0, n' \rangle$ ,  $r$ ) then
    return plusc(plus(plus(timesc( $\langle 0, n \rangle$ ,  $v'$ ), timesc( $\langle 0, n' \rangle$ ,  $v$ )),  $r$ ),  $v \times v'$ )

   $r \leftarrow \text{NewNode}(k)$ 
  for  $i = 0$  to  $n_k - 1$  do
     $r[i] \leftarrow \text{times}(n[i], n'[i])$ 

  // check if a node identical to  $r$  already exists
   $r \leftarrow \text{FindOrAdd}(r)$ 

  // save result in cache
  CacheInsert( $\times$ ,  $\langle v, n \rangle$ ,  $\langle v', n' \rangle$ ,  $\langle m, r \rangle$ )

  return plusc(plus(plus(timesc( $\langle 0, n \rangle$ ,  $v'$ ), timesc( $\langle 0, n' \rangle$ ,  $v$ )),  $r$ ),  $v \times v'$ )

```

---

$f < g$  can be computed as  $f - g < 0$  and the operation “less than constant” can be greatly improved as shown in Algorithm 4 on the facing page through the use of min and max, which are easy to compute from section 4.2.5 on the previous page.

We also developed Algorithm 5 on page 12, which initially was deemed as “optimized”. The idea behind this algorithm is that for a call  $\text{lt}((-, p), -)$  returning  $(lb, ub, q)$ , the interval  $[lb, ub]$  is the set  $\{c \mid \forall x. p(x) < c \Leftrightarrow q(x) = 1\}$ . In other words,  $[lb, ub]$  is the exact interval for which the relation  $p < .$  is given by  $q$ .

The later algorithm is theoretically better than the former in that more cache hits occur. It can even be viewed as optimal in the sense that it never creates duplicate nodes from a same node of the input EVMDD. However, in practice, this advantage appears quite small in most cases. Moreover, in order to implement it, it requires an ordered cache, such as a search tree, instead of a simple hash table and the resulting log factor seems to overcome any advantage.

Algorithms for other relational operators analogous to  $<$ .

### 4.3 Versatility

The advantages of EVMDDs over MTMDDs come at a price of a slightly reduced versatility. Both EVMDDs and MTMDDs can be used to represent functions more general than arithmetic functions over integers. For MTMDDs, the only restriction is that the domain and co-domain of the encoded function have to be finite. However EVMDDs, despite the extensions of section 3.2 on page 4, also require some algebraic structure. For example, they cannot be directly used with floats since the sum of two floating point values has to be rounded to be represented as a floating point value. In computer arithmetic, [15] encodes floating point functions through integer EVMDDs of their binary (bit-level) representation. Even though space efficient, floating point computations on this representation remain expensive.

Another problem affecting EVMDDs but not MTMDDs is *false overflow*. Indeed, computers do not operate over  $\mathbb{Z}$  proper, but on  $\mathbb{Z}/2^{32}\mathbb{Z}$  (for 32bit platforms, or the equivalent structure for the appropriate size of integers – 16, 32, 64 bits). Integer overflow is possible when adding edge values, even when the represented function may fit within 32 bits. An example of such an internal

---

**Algorithm 4** computes  $\langle v, n \rangle < c$  for EVMDD  $\langle v, n \rangle$  and integer  $c$

---

```

lt( $\langle v, n \rangle$  : edge,  $c$  : int) : edge
   $k \leftarrow n.$ level

  // base cases
  if  $c - v \leq \min(n)$  then
    return  $\langle 0, t \rangle$  //  $t$  is the unique terminal node
  if  $c - v > \max(n)$  then
    return  $\langle 1, t \rangle$  //  $t$  is the unique terminal node

  // lookup in cache
  if CacheFind( $\langle, \langle 0, n \rangle, c - v, \langle m, r \rangle$ ) then
    return  $\langle m, r \rangle$ 

   $r \leftarrow$  NewNode( $k$ )
  for  $i = 0$  to  $n_k - 1$  do
     $r[i] \leftarrow$  lt( $n[i]$ ,  $c - v$ )
   $m \leftarrow r[0].$ val
  for  $i = 0$  to  $n_k - 1$  do
     $r[i].$ val  $\leftarrow$   $r[i].$ val  $- m$ 

  // check if a node identical to  $r$  already exists
   $r \leftarrow$  FindOrAdd( $r$ )

  // save result in cache
  CacheInsert( $\langle, \langle 0, n \rangle, c - v, \langle m, r \rangle$ )

  return  $\langle m, r \rangle$ 

```

---

overflow can be seen in figure 4 on the following page. Since  $\mathbb{Z}/2^{32}\mathbb{Z}$  is still an additive group, these overflows are harmless for addition and multiplication by a constant. The general algorithm from section 4.2 on page 6 is not affected either, as it applies operations only on leaf nodes, but this observation no longer holds on the other algorithms of this section. A necessary and sufficient condition on a function  $f$  for not having false overflows in the corresponding EVMDD is

$$\forall k \in \{1, \dots, K\}. \forall (i_K, \dots, i_{k+1}) \in S_K \times \dots \times S_{k+1}. \forall i_k \in S_k. \\ f(i_K, \dots, i_{k+1}, i_k, 0, \dots, 0) - f(i_K, \dots, i_{k+1}, 0, 0, \dots, 0) \text{ does not overflow}$$

## 5 Implementation

Symbolic model checkers such as (Nu)SMV [1] or SAL [2] are based on the library CUDD [10] which offers an efficient implementation of BDDs and MTBDDs. For state space generation, they use a plain breadth first search (BFS) algorithm.

Our goal was to implement a new symbolic model checking library featuring EVMDDs for transition relation construction and saturation [6] for state space generation. We also developed a basic model checking front-end to test the library and compare it to CUDD. Binaries for both and model checker sources are available at <http://research.nianet.org/~radu/evmdd/>.

In this section, we discuss some implementation details before presenting experimental results.

### 5.1 Memory Management

It is well known that model checking can be memory intensive, making memory management a critical issue when implementing a decision diagrams library.

---

**Algorithm 5** computes  $\langle v, n \rangle < c$  for EVMDD  $\langle v, n \rangle$  and integer  $c$

---

```

It( $\langle v, n \rangle$  : edge,  $c$  : int) : int * int * edge
   $k \leftarrow n.level$ 

  // base cases
  if  $c - v \leq \min(n)$  then
    return  $(-\infty, \min(n), \langle 0, t \rangle)$  //  $t$  is the unique terminal node
  if  $c - v > \max(n)$  then
    return  $(\max(n) + 1, +\infty, \langle 1, t \rangle)$  //  $t$  is the unique terminal node

  // lookup in cache for an interval  $[lb, ub]$  containing  $c - v$ 
  if CacheFind( $\langle, \langle 0, n \rangle, c - v, (lb, ub, \langle m, r \rangle)$ ) then
    return  $(lb, ub, \langle m, r \rangle)$ 

   $r \leftarrow \text{NewNode}(k)$ 
   $lb \leftarrow -\infty$ 
   $ub \leftarrow +\infty$ 
  for  $i = 0$  to  $n_k - 1$  do
     $lb', ub', r[i] \leftarrow \text{It}(n[i], c - v)$ 
     $lb \leftarrow \max(lb, lb' + n[i].val)$ 
     $ub \leftarrow \min(ub, ub' + n[i].val)$ 
   $m \leftarrow r[0].val$ 
  for  $i = 0$  to  $n_k - 1$  do
     $r[i].val \leftarrow r[i].val - m$ 

  // check if a node identical to  $r$  already exists
   $r \leftarrow \text{FindOrAdd}(r)$ 

  // save result in cache
  CacheInsert( $\langle, \langle 0, n \rangle, [lb, ub], (lb, ub, \langle m, r \rangle)$ )

  return  $(lb, ub, \langle m, r \rangle)$ 

```

---




Figure 4. MTMDD (left) and EVMDD (right) encoding the same function over 32-bit integers,  $f : x \in \{0, 1\} \mapsto 2,000,000,000 \times (2x - 1)$ .

The usual addressing scheme of BDD nodes is through pointers. Here, with MDD, nodes at different levels with different variable ranges are of different sizes. For that reason, we chose an index-based addressing scheme  $(level, index)$ , allowing to allocate nodes independently at each level.

Decision diagrams algorithms rely heavily on several cache structures. They are usually implemented as lossy hash tables: upon collision, older entries are simply discarded. We avoided the loss of information by choosing dynamically resizing hash tables with chaining. The non-lossy caches give slightly faster algorithms at the price of some additional memory.

Different diagrams often share subgraphs, making explicit freeing of disconnected sub-diagrams very difficult. Moreover, symbolic model checking algorithms often produce a large number of temporary nodes which rapidly become disconnected (also called garbage). The easiest way to reclaim memory is to use some garbage collection procedure. The most commonly used technique is via reference counting. This is perfectly suited since diagrams are DAGs, hence avoiding circular reference, the main drawback of reference counting. Instead, we chose to use a simple *mark & sweep* algorithm that proved to be both simple and efficient, because it requires minimal book-keeping. We keep reference counting only as an interface with library users.

## 5.2 Encoding the Transition Relation

We represent the transition relation  $T$  as a disjunction  $\bigcup_{e \in E} T_e$  of events  $e$ . Each event  $T_e$  is represented as an MDD of high  $2K$ , with level  $2k$  encoding variable  $x_k$  before the transition and level  $2k - 1$  same variable  $x'_k$  but after the transition<sup>6</sup>.

This disjunction representation is well suited for globally-asynchronous locally-synchronous systems, where each event encodes some local transition. However, we could end up with many events that are just the identity relation for most variables, The numerous “identity patterns” in the MDD are very expensive to deal with, both in terms of memory usage and computation time. To avoid this problem, we chose yet another reduction rule of nodes, different than the two already presented in definitions 3.4 on page 3 and 3.5 on page 4: the *full-identity reduction* from [8].




Figure 5. An identity pattern in a reduced MDD (left) and its identity reduced equivalent (right).

An example of such an identity pattern and its full-identity reduction is shown in Figure 5. In this figure, edges leading to terminal 0 are omitted for clarity.

<sup>6</sup>That is, we consider  $T_e$  as a part of  $S \times S = S_K \times S_K \times \dots \times S_1 \times S_1$ . It is interesting to notice, that  $T_e$  does not need to be a function  $S \rightarrow S$ , allowing to model non deterministic behaviors.

Table 1 shows the execution time with standard reduction and full-identity reduction rules for the classical model of dining philosophers.

Model size	standard (sec)	full-identity (sec)
50	0.36	0.02
100	1.68	0.04
200	7.85	0.07
400	36.72	0.20
1000	—	0.91
7000	—	37.47

Table 1. Execution time for building the transition relation using standard and fully-identity reduced MDDs (“—” means “out of memory”).

### 5.3 State Space Construction

For state space construction, we use the Saturation algorithm [6] instead of the classical breadth first search exploration. This heuristic often gives spectacular improvements when building the state spaces of globally-asynchronous locally-synchronous systems. This is certainly the major source of improvement of our implementation over existing BDD libraries.

As advised in [8], we chose to merge all events  $e$  with same topmost affected level<sup>7</sup> in the same MDD. The cost of the unions slows down the generation of the transition relation but generally makes the state space construction up to several times faster.

### 5.4 Experimental Results

Our new model checker comprises 7000 lines of ANSI-C code for the library and 4000 lines for the simple model checker that provides a common interface to our library and CUDD. Table 2 on the facing page shows execution times for finding deadlocks on a suite of classical models. The search for deadlocks in a deadlock-free system equates to building state space. Programs to generate all models can be found in the `examples` directory of our simple model checker source distribution, available at <http://research.nianet.org/~radu/evmdd/>. We collected the results on a Linux machine with Intel Core 2 processor, 1.2GHz, 1.5GB of memory.

Compared to the first implementation of saturation algorithm [6] in the tool SMART, our new implementation is always several (up to a few dozens) times faster. This is due to both the encoding of the transition relation and our simple C implementation in comparison to the object-oriented C++ version.

## 6 Conclusions and Future Work

We studied the advantages of the EVMDD data structure over the widely used MTBDDs for the construction of transition relation of finite state systems and implemented them in a library, along with state-of-the-art algorithms for state space generation. We obtained execution times several orders of magnitude faster than the CUDD library and classical algorithms, with a reduced memory usage enabling to handle extremely large systems. Future work should focus primarily on integrating our library into the SAL model checker.

<sup>7</sup>i.e. same value of  $\max \{k \mid \exists i, j \in S_k . i \neq j \wedge (i, j) \in T_e\}$

Model size	Reachable states	CUDD (sec)	EVMDD (sec)	Model size	Reachable states	CUDD (sec)	EVMDD (sec)
Dining philosophers				Kanban assembly line			
100	$4 \times 10^{62}$	5.15	0.04	10	$1 \times 10^9$	0.84	0.01
200	$2 \times 10^{125}$	1493.36	0.10	20	$8 \times 10^{11}$	811.89	0.04
1000	$9 \times 10^{626}$	—	1.10	100	$1 \times 10^{19}$	—	3.25
10000	$4 \times 10^{6269}$	—	77.77	200	$3 \times 10^{22}$	—	32.31
16000	$2 \times 10^{10031}$	—	196.84	400	$6 \times 10^{25}$	—	697.90
Round robin mutual exclusion protocol				Knights problem			
40	$9 \times 10^{13}$	4.04	0.46	5	$6 \times 10^7$	921.95	0.30
60	$1 \times 10^{20}$	15.59	1.54	7	$1 \times 10^{15}$	—	3.73
100	$2 \times 10^{32}$	599.16	7.42	9	$8 \times 10^{24}$	—	47.80
200	$7 \times 10^{62}$	—	75.94	Randomized leader election protocol			
Slotted ring protocol				6	$2 \times 10^6$	4.62	1.20
10	$8 \times 10^9$	1.07	0.01	7	$2 \times 10^7$	23.87	3.69
20	$2 \times 10^{20}$	1804.48	0.04	8	$3 \times 10^8$	176.16	10.04
100	$2 \times 10^{105}$	—	3.74	9	$5 \times 10^9$	810.36	24.86
300	$3 \times 10^{318}$	—	111.29	10	$6 \times 10^{10}$	—	85.54
500	$5 \times 10^{531}$	—	505.04	11	$9 \times 10^{11}$	—	423.41

Table 2. Execution times for finding deadlocks using our library or CUDD (“—” means “> 1hour”).

Our results show that, as old fashioned as it may seem, symbolic model checking remains an efficient technique for analyzing globally-asynchronous locally-synchronous systems and significant improvements are still possible.

## References

1. NuSMV model checker. <http://nusmv.iirst.itc.it/>.
2. SAL model checker. <http://sal.csl.sri.com/>.
3. R. Iris Bahar, Erica A. Frohm, Charles M. Gaona, Gary D. Hachtel, Enrico Macii, Abelardo Pardo, and Fabio Somenzi. Algebraic decision diagrams and their applications. In *ICCAD*, pages 188–191, 1993.
4. Randal E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Transactions on Computers*, 35:677–691, 1986.
5. Randal E. Bryant and Yirng-An Chen. Verification of arithmetic functions with binary moment diagrams. Technical Report CS-94-160, CMU, 1994.
6. Gianfranco Ciardo, Gerald Lüttgen, and Radu Siminiceanu. Saturation: An efficient iteration strategy for symbolic state space generation. In Tiziana Margaria and Wang Yi, editors, *Proc. Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, LNCS 2031, pages 328–342, Genova, Italy, April 2001. Springer.
7. Gianfranco Ciardo and Radu Siminiceanu. Using edge-valued decision diagrams for symbolic generation of shortest paths. In Mark D. Aagaard and John W. O’Leary, editors, *Proc. Fourth International Conference on Formal Methods in Computer-Aided Design (FMCAD)*, LNCS 2517, pages 256–273, Portland, OR, USA, November 2002. Springer.

8. Ciardo, G. and Yu, J. Saturation-based symbolic reachability analysis using conjunctive and disjunctive partitioning. In *Proc. CHARME*, October 2005.
9. Edmund Clarke, M. Fujita, and X. Zhao. Application of multi-terminal binary decision diagrams. Technical report, IFIP WG 10.5 Workshop on Applications of the Reed-Muller Expansion in Circuit Design, 1995.
10. Fabio Somenzi. CUDD library, Colorado University Decision Diagram. <http://vlsi.colorado.edu/~fabio/CUDD/>.
11. K. S. Brace and R. L. Rudell and R. E. Bryant. Efficient implementation of a BDD package. In *Proceedings of the 27th Design Automation Conference*, pages 40–45, June 1990.
12. T. Kam, T. Villa, R.K. Brayton, and Alberto Sangiovanni-Vincentelli. Multi-valued decision diagrams: theory and applications. *Multiple-Valued Logic*, 4(1–2):9–62, 1998.
13. Y.-T. Lai, M. Pedram, and B. K. Vrudhula. Formal verification using edge-valued binary decision diagrams. *IEEE Transactions on Computers*, 45:247–255, 1996.
14. Yung-Te Lai and Sarma Sastry. Edge-valued binary decision diagrams for multi-level hierarchical verification. In *DAC*, pages 608–613, 1992.
15. S. Nagayama, and T. Sasao, and J. T. Butler. Floating-point numerical function generators using EVMDDs for monotone elementary functions. In *39th International Symposium on Multiple-Valued Logic (ISMVL 2009)*, pages 349–355, 2009.

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188		
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p><b>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</b></p>					
1. REPORT DATE (DD-MM-YYYY) 01-06-2010		2. REPORT TYPE Contractor Report		3. DATES COVERED (From - To)	
4. TITLE AND SUBTITLE Model-Checking with Edge-Valued Decision Diagrams			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER NNX08AC59A		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Roux, Pierre; Siminiceanu, Radu I.			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER 645846.02.07.07.15.03		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) NASA Langley Research Center Hampton, VA 23681-2199			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Washington, DC 20546-0001			10. SPONSOR/MONITOR'S ACRONYM(S) NASA		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) NASA/CR-2010-216710		
12. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified - Unlimited Subject Category 64 Availability: NASA CASI (443) 757-5802					
13. SUPPLEMENTARY NOTES Langley Technical Monitor: Benedetto L. Di Vito					
14. ABSTRACT We describe an algebra of Edge-Valued Decision Diagrams (EVMDDs) to encode arithmetic functions and its implementation in a model checking library along with state-of-the-art algorithms for building the transition relation and the state space of discrete state systems. We provide efficient algorithms for manipulating EVMDDs and give upper bounds of the theoretical time complexity of these algorithms for all basic arithmetic and relational operators. We also demonstrate that the time complexity of the generic recursive algorithm for applying a binary operator on EVMDDs is no worse than that of Multi-Terminal Decision Diagrams. We have implemented a new symbolic model checker with the intention to represent in one formalism the best techniques available at the moment across a spectrum of existing tools: EVMDDs for encoding arithmetic expressions, identity-reduced MDDs for representing the transition relation, and the saturation algorithm for reachability analysis. We compare our new symbolic model checking EVMDD library with the widely used CUDD package and show that, in many cases, our tool is several orders of magnitude faster than CUDD.					
15. SUBJECT TERMS Formal methods; Hybrid abstraction; Model checking; Decision diagrams					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			STI Help Desk (email: help@sti.nasa.gov)
U	U	U	UU	23	19b. TELEPHONE NUMBER (Include area code) (443) 757-5802



