

Model Checking I

Binary Decision Diagrams

Edmund M. Clarke, Jr.
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Binary Decision Diagrams

Ordered binary decision diagrams (OBDDs) are a canonical form for boolean formulas.

OBDDs are often substantially more compact than traditional normal forms.

Moreover, they can be manipulated very efficiently.

- ▶ R. E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Transactions on Computers*, C-35(8), 1986.

Binary Decision Trees

To motivate our discussion of binary decision diagrams, we first consider *binary decision trees*.

A binary decision tree is a rooted, directed tree with two types of vertices, *terminal vertices* and *nonterminal vertices*.

Each nonterminal vertex v is labeled by a variable $var(v)$ and has two successors:

- ▶ $low(v)$ corresponding to the case where the variable v is assigned 0, and
- ▶ $high(v)$ corresponding to the case where the variable v is assigned 1.


Each terminal vertex v is labeled by $value(v)$ which is either 0 or 1.

Binary Decision Trees (Cont.)

A BDD for the two-bit comparator given by the formula

$$f(a_1, a_2, b_1, b_2) = (a_1 \leftrightarrow b_1) \wedge (a_2 \leftrightarrow b_2),$$

is shown in the figure below:



Binary Decision Trees (Cont.)

We can decide if a truth assignment satisfies the formula as follows:

- ▶ Traverse the tree from the root to a terminal vertex.
- ▶ If variable v is assigned 0, the next vertex on the path will be $low(v)$.
- ▶ If variable v is assigned 1, the next vertex on the path will be $high(v)$.
- ▶ The value that labels the terminal vertex will be the value of the function for this assignment.

In the comparator example, the assignment

$$\langle a_1 \leftarrow 1, a_2 \leftarrow 0, b_1 \leftarrow 1, b_2 \leftarrow 1 \rangle$$

leads to a leaf vertex labeled 0, so the formula is false for this assignment.

A More Concise Representation

Binary decision trees do not provide a very concise representation for boolean functions.

However, there is usually a lot of redundancy in such trees.

In the comparator example, there are eight subtrees with roots labeled by b_2 , but only three are distinct.

Thus, we can obtain a more concise representation by merging isomorphic subtrees.

This results in a directed acyclic graph (DAG) called a *binary decision diagram*.

Binary Decision Diagrams

More precisely, a *binary decision diagram* is a rooted, directed acyclic graph with two types of vertices, *terminal vertices* and *nonterminal vertices*.

Each nonterminal vertex v is labeled by a variable $var(v)$ and has two successors, $low(v)$ and $high(v)$.

Each terminal vertex is labeled by either 0 or 1.

Binary Decision Diagrams (cont'd)

A binary decision diagram with root v determines a boolean function $f_v(x_1, \dots, x_n)$ in the following manner:

1. If v is a terminal vertex:
 - 1.1 If $value(v) = 1$ then $f_v(x_1, \dots, x_n) = 1$.
 - 1.2 If $value(v) = 0$ then $f_v(x_1, \dots, x_n) = 0$.
2. If v is a nonterminal vertex with $var(v) = x_i$ then $f_v(x_1, \dots, x_n)$ is given by

$$\bar{x}_i \cdot f_{low(v)}(x_1, \dots, x_n) + x_i \cdot f_{high(v)}(x_1, \dots, x_n)$$

Canonical Form Property

In practical applications, it is desirable to have a *canonical representation* for boolean functions.

This simplifies tasks like checking equivalence of two formulas and deciding if a given formula is satisfiable or not.

Such a representation must guarantee that two boolean functions are logically equivalent if and only if they have isomorphic representations.

Canonical Form Property (Cont.)

Two binary decision diagrams are *isomorphic* if there exists a bijection h between the graphs such that

- ▶ terminals are mapped to terminals and nonterminals are mapped to nonterminals,
- ▶ for every terminal vertex v , $value(v) = value(h(v))$, and
- ▶ for every nonterminal vertex v :
 - ▶ $var(v) = var(h(v))$,
 - ▶ $h(low(v)) = low(h(v))$, and
 - ▶ $h(high(v)) = high(h(v))$.

Canonical Form Property (Cont.)

Bryant showed how to obtain a canonical representation for boolean functions by placing two restrictions on binary decision diagrams:

- ▶ First, the variables should appear in the same order along each path from the root to a terminal.
- ▶ Second, there should be no isomorphic subtrees or redundant vertices in the diagram.

Canonical Form Property (Cont.)

The first requirement is easy to achieve:

- ▶ We impose total ordering $<$ on the variables in the formula.
- ▶ We require that if vertex u has a nonterminal successor v , then $var(u) < var(v)$.

Canonical Form Property (Cont.)

The second requirement is achieved by repeatedly applying three transformation rules that do not alter the function represented by the diagram:

Remove duplicate terminals: Eliminate all but one terminal vertex with a given label and redirect all arcs to the eliminated vertices to the remaining one.

Remove duplicate nonterminals: If nonterminals u and v have $var(u) = var(v)$, $low(u) = low(v)$ and $high(u) = high(v)$, then eliminate one of the two vertices and redirect all incoming arcs to the other vertex.

Remove redundant tests: If nonterminal vertex v has $low(v) = high(v)$, then eliminate v and redirect all incoming arcs to $low(v)$.

Canonical Form Property (Cont.)

The canonical form may be obtained by applying the transformation rules until the size of the diagram can no longer be reduced.

Bryant shows how this can be done by a procedure called *Reduce* in linear time.

Ordered Binary Decision Diagrams


The term *ordered binary decision diagram* (OBDD) will be used to refer to the graph obtained in this manner.

If OBDDs are used as a canonical form for boolean functions, then

- ▶ checking equivalence is reduced to checking isomorphism between OBDDs, and
- ▶ satisfiability can be determined by checking equivalence with the trivial OBDD that consists of only one terminal labeled by 0.

OBDD for Comparator Example


If we use the ordering $a_1 < b_1 < a_2 < b_2$ for the comparator function, we obtain the OBDD below:



Variable Ordering Problem

The size of an OBDD depends critically on the variable ordering.

If we use the ordering $a_1 < a_2 < b_1 < b_2$ for the comparator function, we get the OBDD below:



Variable Ordering Problem (Cont.)

For an n -bit comparator:

- ▶ if we use the ordering $a_1 < b_1 < \dots < a_n < b_n$, the number of vertices will be $3n + 2$.
- ▶ if we use the ordering $a_1 < \dots < a_n < b_1 \dots < b_n$, the number of vertices is $3 \cdot 2^n - 1$.

In general, finding an optimal ordering is known to be NP-complete.

Moreover, there are boolean functions that have exponential size OBDDs for any variable ordering.

An example is the middle output (n^{th} output) of a combinational circuit to multiply two n bit integers.

Heuristics for Variable Ordering

Heuristics have been developed for finding a good variable ordering when such an ordering exists.

The intuition for these heuristics comes from the observation that OBDDs tend to be small when related variables are close together in the ordering.

The variables appearing in a subcircuit are related in that they determine the subcircuit's output.

Hence, these variables should usually be grouped together in the ordering.

This may be accomplished by placing the variables in the order in which they are encountered during a depth-first traversal of the circuit diagram.

Dynamic Variable Ordering

A technique, called *dynamic reordering* appears to be useful if no obvious ordering heuristic applies.

When this technique is used, the OBDD package internally reorders the variables periodically to reduce the total number of vertices in use.

Logical Operations on OBDDs

We begin with the function that restricts some argument x_i of the boolean function f to a constant value b .

This function is denoted by $f|_{x_i \leftarrow b}$ and satisfies the identity

$$f|_{x_i \leftarrow b}(x_1, \dots, x_n) = f(x_1, \dots, x_{i-1}, b, x_{i+1}, \dots, x_n).$$

If f is represented as an OBDD, the OBDD for the restriction $f|_{x_i \leftarrow b}$ is computed by a depth-first traversal of the OBDD.

For any vertex v which has a pointer to a vertex w such that $\text{var}(w) = x_i$, we replace the pointer by $\text{low}(w)$ if b is 0 and by $\text{high}(w)$ if b is 1.

When the graph is not in canonical form, we apply *Reduce* to obtain the OBDD for $f|_{x_i \leftarrow b}$.

Logical Operations (Cont.)

All 16 two-argument logical operations can be implemented efficiently on boolean functions that are represented as OBDDs.

In fact, the complexity of these operations is linear in the size of the argument OBDDs.

The key idea for efficient implementation of these operations is the *Shannon expansion*

$$f = \bar{x} \cdot f|_{x \leftarrow 0} + x \cdot f|_{x \leftarrow 1} .$$

Logical Operations (Cont.)

Bryant gives a uniform algorithm called *Apply* for computing all 16 logical operations.

Let \star be an arbitrary two argument logical operation, and let f and f' be two boolean functions.

To simplify the explanation of the algorithm we introduce the following notation:

- ▶ v and v' are the roots of the OBDDs for f and f' .
- ▶ $x = \text{var}(v)$ and $x' = \text{var}(v')$.

Logical Operations OBDDs (Cont.)

We consider several cases depending on the relationship between v and v' .

- ▶ If v and v' are both terminal vertices, then $f \star f' = \text{value}(v) \star \text{value}(v')$.
- ▶ If $x = x'$, then we use the Shannon expansion

$$f \star f' = \bar{x} \cdot (f|_{x \leftarrow 0} \star f'|_{x \leftarrow 0}) + x \cdot (f|_{x \leftarrow 1} \star f'|_{x \leftarrow 1})$$

to break the problem into two subproblems. The subproblems are solved recursively.

The root of the resulting OBDD will be v with $\text{var}(v) = x$.

$\text{Low}(v)$ will be the OBDD for $(f|_{x \leftarrow 0} \star f'|_{x \leftarrow 0})$.

$\text{High}(v)$ will be the OBDD for $(f|_{x \leftarrow 1} \star f'|_{x \leftarrow 1})$.

Logical Operations OBDDs (Cont.)

- ▶ If $x < x'$, then $f' |_{x \leftarrow 0} = f' |_{x \leftarrow 1} = f'$ since f' does not depend on x .

In this case the Shannon Expansion simplifies to

$$f \star f' = \bar{x} \cdot (f |_{x \leftarrow 0} \star f') + x \cdot (f |_{x \leftarrow 1} \star f')$$

and the OBDD for $f \star f'$ is computed recursively as in the second case.

- ▶ If $x' < x$, then the required computation is similar to the previous case.

Logical Operations (Cont.)

By using dynamic programming, it is possible to make the algorithm polynomial.

- ▶ A hash table is used to record all previously computed subproblems.
- ▶ Before any recursive call, the table is checked to see if the subproblem has been solved.
- ▶ If it has, the result is obtained from the table; otherwise, the recursive call is performed.
- ▶ The result must be reduced to ensure that it is in canonical form.

A single OBDD can be used to represent a collection of boolean functions:

- ▶ The same variable ordering is used for whole the collection.
- ▶ As before, no isomorphic subgraphs or redundant vertices.

Two functions in the collection are identical if and only if they have the same root.

Consequently, checking whether two functions are equal can be implemented in constant time.

OBDD Extensions (Cont.)

A possible extension is to add labels to the arcs in the graph to denote boolean negation.

This makes it unnecessary to use different subgraphs to represent a formula and its negation.

Canonicity?

OBDDs and Finite Automata

OBDDs can also be viewed as *deterministic finite automata* (DFAs).

An n -argument boolean function can be identified with the set of strings in $\{0, 1\}^n$ that evaluate to 1.

This is a finite language. Finite languages are regular. Hence, there is a minimal DFA that accepts the language.

The DFA provides a canonical form for the original boolean function.

Logical operations on boolean functions can be implemented by standard constructions from elementary automata theory.

Representing Finite Relations

OBDDs are extremely useful for obtaining concise representations of relations over finite domains.

If R is an n -ary relation over $\{0, 1\}$ then R can be represented by the OBDD for its *characteristic function*

$$f_R(x_1, \dots, x_n) = 1 \text{ iff } R(x_1, \dots, x_n).$$

- ▶ J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang. Symbolic model checking: 10^{20} states and beyond. *Information and Computation*, 98(2):142–170, June 1992.
- ▶ J. R. Burch, E. M. Clarke, D. E. Long, K. L. McMillan, and D. L. Dill. Symbolic model checking for sequential circuit verification. *IEEE Transactions on Computer-Aided Design of Integrated Circuits*, 13(4):401–424, 1994.

Representing Relations (Cont.)

Suppose R is an n -ary relation over the domain D , where D has 2^m elements for some $m > 1$.

To represent R as an OBDD, we simply encode elements of D using a binary representation. Formally, we use a bijection $\phi : \{0, 1\}^m \rightarrow D$ that maps each boolean vector of length m to an element of D .

We construct a new boolean relation R' of arity $m \times n$ according to the following rule:

$$R'(\bar{x}_1, \dots, \bar{x}_n) = R(\phi(\bar{x}_1), \dots, \phi(\bar{x}_n))$$

where \bar{x}_i is a vector of m boolean variables which encodes the variable x_i that takes values in D .

R can now be represented as the OBDD for the characteristic function $f_{R'}$ of R' .

Representing Relations (Cont.)

This technique can be easily extended to relations over different domains, D_1, \dots, D_n .

Since sets can be viewed as unary relations, the same technique can be used to represent sets as OBDDs.

Quantified Boolean Formulas (QBF)

Sometimes it is convenient to use existential or universal quantification over boolean variables.

The resulting logic is the logic of *Quantified Boolean Formulas* (QBF).

Suppose x is a propositional variable.

- ▶ The QBF formula $\exists x.f$ is true iff: $f|_{x \leftarrow 0}$ is true or $f|_{x \leftarrow 1}$ is true.
- ▶ The QBF formula $\forall x.f$ is true iff: $f|_{x \leftarrow 0}$ is true and $f|_{x \leftarrow 1}$ is true.

(We identify 1 with true and identify 0 with false.)

Formally:

Given a set of propositional variables $V = \{v_1, \dots, v_n\}$, $QBF(V)$ is the smallest set of formulas such that

- ▶ every variable in V is a formula,
- ▶ if f and g are formulas, then $\neg f$, $f \vee g$, and $f \wedge g$ are formulas, and
- ▶ if f is a formula and $v \in V$, then $\exists v f$ and $\forall v f$ are formulas.

Truth Assignments

A *truth assignment* for $QBF(V)$ is a function $\sigma : V \rightarrow \{0, 1\}$.

If $a \in \{0, 1\}$, then we will use the notation $\sigma\langle v \leftarrow a \rangle$ for the truth assignment defined by

$$\sigma\langle v \leftarrow a \rangle(w) = \begin{cases} a & \text{if } v = w \\ \sigma(w) & \text{otherwise.} \end{cases}$$

If f is a formula in $QBF(V)$ and σ is a truth assignment, we will write $\sigma \models f$ when f is true under the assignment σ .

The relation \models is defined recursively in the obvious manner:

- ▶ $\sigma \models v$, iff $\sigma(v) = 1$,
- ▶ $\sigma \models \neg f$, iff $\sigma \not\models f$,
- ▶ $\sigma \models f \vee g$, iff $\sigma \models f$, or $\sigma \models g$,
- ▶ $\sigma \models f \wedge g$, iff $\sigma \models f$, and $\sigma \models g$,
- ▶ $\sigma \models \exists v f$, iff $\sigma\langle v \leftarrow 0 \rangle \models f$, or $\sigma\langle v \leftarrow 1 \rangle \models f$,
- ▶ $\sigma \models \forall v f$, iff $\sigma\langle v \leftarrow 0 \rangle \models f$, and $\sigma\langle v \leftarrow 1 \rangle \models f$

QBF formulas have the same expressive power as ordinary propositional formulas; however, they are sometimes much more concise.

Every QBF formula determines an n -ary boolean relation consisting of those truth assignments for the variables in V that make the formula true.

We will identify each formula with the boolean relation that it determines.

QBF formulas and Relations

Previously, we showed how to associate an OBDD with each formula of propositional logic.

In principle, it is easy to construct OBDDs for $\exists v f$ and $\forall v f$ when f is given as an OBDD.

- ▶ $\exists x f = f|_{x \leftarrow 0} + f|_{x \leftarrow 1}$
- ▶ $\forall x f = f|_{x \leftarrow 0} \cdot f|_{x \leftarrow 1}$

In practice, however, special algorithms are needed to handle quantifiers efficiently.

Relational Products

In model checking, quantifiers occur most frequently in *relational products*

$$\exists \bar{v} [f(\bar{v}) \wedge g(\bar{v})].$$

We give an algorithm that performs this computation in one pass over the OBDDs $f(\bar{v})$ and $g(\bar{v})$.

This is important since we avoid constructing the OBDD for $f(\bar{v}) \wedge g(\bar{v})$.

Before giving the algorithm for the relational product

$$\exists \bar{v} [f(\bar{v}) \wedge g(\bar{v})]$$

let's review the algorithm for normal conjunction,

$$f(\bar{v}) \wedge g(\bar{v})$$

Conjunction

Conjoin(f, g : OBDD)

if $f = \text{false} \vee g = \text{false}$ **then return** *false*

else if $f = \text{true} \wedge g = \text{true}$ **then return** *true*

else if $\text{cache}[(f, g)] \neq \text{null}$ **then return** $\text{cache}[(f, g)]$

else

 let x be the top variable of f

 let y be the top variable of g

 let z be the topmost of x and y

$h_0 := \text{Conjoin}(f|_{z=0}, g|_{z=0})$

$h_1 := \text{Conjoin}(f|_{z=1}, g|_{z=1})$

$h := \text{IfThenElse}(z, h_1, h_0)$ /* BDD for $(z \wedge h_1) \vee (\neg z \wedge h_0)$ */

$\text{cache}[(f, g)] := h$

return h

endif

Relational Products (Cont.)

RelProd(f, g : OBDD, E : set of variables)

if $f = \text{false} \vee g = \text{false}$ **then return** *false*

else if $f = \text{true} \wedge g = \text{true}$ **then return** *true*

else if $\text{cache}[(f, g, E)] \neq \text{null}$ **then return** $\text{cache}[(f, g, E)]$

else

 let x be the top variable of f

 let y be the top variable of g

 let z be the topmost of x and y

$h_0 := \text{RelProd}(f|_{z=0}, g|_{z=0}, E)$

$h_1 := \text{RelProd}(f|_{z=1}, g|_{z=1}, E)$

if $z \in E$ **then** $h := \text{Or}(h_0, h_1)$ /* BDD for $h_0 \vee h_1$ */

else $h := \text{IfThenElse}(z, h_1, h_0)$ /* BDD for $(z \wedge h_1) \vee (\neg z \wedge h_0)$ */

endif

$\text{cache}[(f, g, E)] := h$

return h

endif

Relational Products (Cont.)

Like many OBDD algorithms, *RelProd* uses a result cache.

In this case, entries in the cache have the form $((f, g, E), h)$, where E is a set of variables that are quantified out and f , g and h are OBDDs.

If such an entry is in the cache, then a previous call to $RelProd(f, g, E)$ returned h as its result.

Although the algorithm works well in practice, it has exponential complexity in the worst case.