

AD-A281 028




Computer Science

Verification of Arithmetic Functions
with Binary Moment Diagrams

Randal E. Bryant, Yirng-An Chen

May 31, 1994

CMU-CS-94-160



DTIC
ELECTE
JUL 05 1994
S G D

DTIC QUALITY INSPECTED 2

Carnegie
Mellon

94-20298



4098

94 7 1 058

0

**Verification of Arithmetic Functions
with Binary Moment Diagrams**

Randal E. Bryant, Yirng-An Chen
May 31, 1994
CMU-CS-94-160

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3890

This research is sponsored by the Wright Laboratory, Aeronautical Systems Center, Air Force Materiel Command, USAF, and the Advanced Research Projects Agency (ARPA) under grant F33615-93-1-1330. The US Government is authorized to reproduce and distribute reprints for Government purposes, notwithstanding any copyright notation thereon. Views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of Wright Laboratory or the United States Government.



Keywords: Formal verification, binary decision diagrams, arithmetic circuits, multipliers



Abstract

Binary Moment Diagrams (BMDs) provide a canonical representations for linear functions similar to the way Binary Decision Diagrams (BDDs) represent Boolean functions. Within the class of linear functions, we can embed arbitrary functions from Boolean variables to real, rational, or integer values. BMDs can thus model the functionality of data path circuits operating over word level data. Many important functions, including integer multiplication, that cannot be represented efficiently at the bit level with BDDs have simple representations at the word level with BMDs. Furthermore, BMDs can represent Boolean functions with around the same complexity as BDDs.

We propose a hierarchical approach to verifying arithmetic circuits, here basic building blocks are first shown to implement a word-level specification. The overall circuit functionality is then verified at the word level. Multipliers with word sizes of up to 62 bits have been verified by this technique.

Accession For	
NTIS CRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification _____	
By _____	
Distribution / _____	
Availability Codes	
Dist	Avail and/or Special
A-1	

Verification of Arithmetic Functions with Binary Moment Diagrams*

Abstract

Binary Moment Diagrams (BMDs) provide a canonical representations for linear functions similar to the way Binary Decision Diagrams (BDDs) represent Boolean functions. Within the class of linear functions, we can embed arbitrary functions from Boolean variables to real, rational, or integer values. BMDs can thus model the functionality of data path circuits operating over word level data. Many important functions, including integer multiplication, that cannot be represented efficiently at the bit level with BDDs have simple representations at the word level with BMDs. Furthermore, BMDs can represent Boolean functions with around the same complexity as BDDs.

We propose a hierarchical approach to verifying arithmetic circuits, here basic building blocks are first shown to implement a word-level specification. The overall circuit functionality is then verified at the word level. Multipliers with word sizes of up to 62 bits have been verified by this technique.

Keywords: Formal verification, binary decision diagrams, arithmetic circuits, multipliers

1. Introduction

Binary Decision Diagrams (BDDs) have proved successful for representing and manipulating Boolean functions symbolically [4] in a variety of application domains. Building on this success, there have been several efforts to extend the BDD concept to represent functions over Boolean variables, but having non-Boolean ranges, such as integers or real numbers

*This research is sponsored by the Wright Laboratory, Aeronautical Systems Center, Air Force Materiel Command, USAF, and the Advanced Research Projects Agency (ARPA) under grant number F33615-93-1-1330.

	Boolean	Numeric	
		Terminal	Edge-Weighted
Pointwise Moment	BDD	MTBDD, ADD	EVBDD
	FDD	BMD	*BMD

Table 1: Categorization of Graphical Function Representations.

[1, 7, 8, 15, 17]. This class of functions is sometimes termed “pseudo-Boolean” [12]. Many tasks can be expressed in terms of operations on such functions, including integer linear programming, matrix manipulation, spectral transforms, and word-level digital system analysis. To date, the proposed representations for these functions have proved too fragile for routine application—too often the data structures grow exponentially in the number of variables.

In this paper we propose a new representation called Multiplicative Binary Moment Diagrams (*BMDs) that improve on previous methods. *BMDs incorporate two novel features: they are based on a decomposition of a linear function in terms of its “moments,” and they have weights associated with their edges which are combined multiplicatively. These features have as heritage ideas found in previous function representations, namely the Reed-Muller decomposition used by Functional Decision Diagrams (FDDs) [9, 14], and the additive edge weights found in Edge-Valued Binary Decision Diagrams (EVBDDs) [15]. The relations between the various representations are described more fully below.

*BMDs are particularly effective for representing digital systems at the word level, where sets of binary signals are interpreted as encoding integer (fixed point) or rational (floating point) values. Common integer and floating point encodings have efficient representations as *BMDs, as do operations such as addition and multiplication. *BMDs can also represent Boolean functions as a special case, with size comparable to BDDs.

*BMDs can serve as the basis for a hierarchical methodology for verifying circuits such as multipliers. At the low level, we have a set of building blocks such as add steppers, Booth steppers, and carry save adders described at both the bit level (as combinational circuits) and at the word level (as algebraic expressions). Using a methodology proposed by Lai and Sastry [15], we verify that the bit-level implementation of each block implements its word-level specification. At the higher level (or levels), a system is described as an interconnection of blocks having word-level representations, and the specification is also given at the word-level. We then verify that the composition of the block functions corresponds to the system specification. By this technique we can verify systems, such as multipliers [5], that cannot be represented efficiently at the bit level. We also can handle a more abstract level of specification than can methodologies that work entirely at the bit level.


2. Graphical Function Representations

Methods related to ordered BDDs for representing functions as graphs can be categorized as shown in Table 1. First, the range of a function can be either Boolean or numeric, e.g., integer, rational, or real. Second, we will consider two methods of decomposing a function with respect to a Boolean variable x : in terms of its value at $x = 1$ and $x = 0$ (pointwise

Pointwise Decomposition

y	z	F
0	0	8
0	1	-12
1	0	10
1	1	-6

MTBDD



Linear Decomposition

$$\begin{aligned}
 &8 \quad + \\
 &-20z \quad + \\
 &2y \quad + \\
 &4yz
 \end{aligned}$$

BMD




Figure 1: Example Function Decompositions. MTBDDs are based on a pointwise decomposition (left), while BMDs are based on a linear decomposition (right).

decomposition), or its “moments,” i.e., its value at $x = 0$ and how this value changes as x changes to 1. Finally, the values of a numeric function can be expressed in terms of values associated with the leaves or with the edges. Note that in all cases we assume a total ordering of the variables and that variables are tested according to this ordering along any path from the root to a leaf.

To illustrate the two ways of decomposing a function, consider the function F over a set of Boolean variables y and z , yielding the integer values shown in the table of Figure 1. A pointwise decomposition characterizes a function by its value for every possible set of argument values. By extending BDDs to allow numeric leaf values, the pointwise decomposition leads to a “Multi-Terminal” BDD (MTBDD) representation of a function [7, 8] (also called “ADD” [1]), as shown on the left side of Figure 1. In our drawings of graphs based on a pointwise decomposition, the dashed line from a vertex denotes the case where the vertex variable is 0, and the solid line denotes the case where the variable is 1. Observe that the leaf values correspond directly to the entries in the function table.

Exploiting the fact that the function variables take on only the values 0 and 1, we can write a linear expression for function F directly from the function table. For variable y , the assignment $y = 1$ is encoded as y , and the assignment $y = 0$ is encoded as $1 - y$:

$$\begin{aligned}
 F(x, y) = & \begin{array}{r} 8 \quad (1 - y) \quad (1 - z) \quad + \\ -12 \quad (1 - y) \quad z \quad + \\ 10 \quad y \quad (1 - z) \quad + \\ -6 \quad y \quad z \end{array}
 \end{aligned}$$

Expanding this expression and combining common terms yields the expression:

$$F(x, y) = 8 - 20z + 2y + 4yz$$

$$= 8y^0z^0 + -20y^0z^1 + 2y^1z^0 + 4y^1z^1$$

This representation is called the "monomial expansion" of F . It represents the function as a sum of terms $\alpha y^{b_y} z^{b_z}$ where α is a numeric coefficient and both b_y and b_z are either 0 or 1. This expansion leads to the BMD representation of a function, as shown on the right side of Figure 1. In our drawings of graphs based on a moment decomposition, the dashed line from a vertex indicates the case where the function is independent of the vertex variable x ($b_x = 0$), while the solid line indicates the case where the function varies linearly ($b_x = 1$).

2.1. Recursive Decompositions of Functions

The graph representations of functions we consider expand a function one variable at a time, rather than in terms of all the variables, as do the tabular form and the monomial expansions of Figure 1. Better insight can be gained by considering recursive decompositions of the function, where a function is decomposed in terms of a variable into two subfunctions. In our graphical representation, each vertex denotes a function. The outgoing branches from the vertex indicate the subfunctions resulting from the decomposition with respect to the vertex variable.

For function f over a set of Boolean variables, let f_x (respectively, $f_{\bar{x}}$) denote the positive (resp., negative) cofactor of f with respect to variable x , i.e., the function resulting when constant 1, (resp., 0) is substituted for x . BDDs are based on a pointwise decomposition, where the function is characterized with respect to some variable x in terms of its cofactors. Function f can be expressed in terms of an expansion (variously credited to Shannon and to Boole):

$$f = \bar{x} \wedge f_{\bar{x}} \vee x \wedge f_x$$

In this equation we use \wedge and \vee to represent Boolean sum and product, and overline to represent Boolean complement.

For expressing functions having numeric range, the Boole-Shannon expansion can be generalized as:

$$f = (1 - x) \cdot f_{\bar{x}} + x \cdot f_x \quad (1)$$

where \cdot , $+$, and $-$ denote multiplication, addition, and subtraction, respectively. Note that this expansion relies on the assumption that variable x is Boolean, i.e., it will evaluate to either 0 or 1. Both MTBDDs and EVBDDs [15, 17] are based on such a pointwise decomposition. As with BDDs, each vertex v describes a function f in terms of its decomposition with respect to variable $x = \text{Var}(v)$. The two outgoing arcs: $\text{Lo}(v)$ and $\text{Hi}(v)$ denote functions $f_{\bar{x}}$ and f_x , respectively. A leaf vertex v in an MTBDD has an associated value $\text{Val}(v)$.

The moment decomposition of a function is obtained by rearranging the terms of Equation 1:

$$\begin{aligned} f &= f_{\bar{x}} + x \cdot (f_x - f_{\bar{x}}) \\ &= f_{\bar{x}} + x \cdot f_{\dot{x}} \end{aligned} \quad (2)$$

where $f_{\dot{x}} = f_x - f_{\bar{x}}$ is called the *linear moment* of f with respect to x . This terminology arises by viewing f as being a linear function with respect to its variables, and thus $f_{\dot{x}}$ is the

partial derivative of f with respect to x . Since we are interested in the value of the function for only two values of x , we can always extend it to a linear form. The negative cofactor will be termed the *constant moment*, i.e., it denotes the portion of function f that remains constant with respect to x , while f_x denotes the portion that varies linearly. Relating to the monomial expansion presented earlier, the two moments of function f partition the set of monomial terms into those that are independent of x , i.e., $b_x = 0$ ($f_{\bar{x}}$), and those that vary linearly with x , i.e., $b_x = 1$ (f_x).

We will define two forms of graphs representing functions according to a moment decomposition. In both cases, vertex v denoting function f is labeled by a variable $x = \text{Var}(v)$, and has two outgoing arcs: $\text{Lo}(v)$ denoting function $f_{\bar{x}}$ and $\text{Hi}(v)$ denoting function f_x . We will term graphs of this form "Moment" Diagrams (MDs) as opposed to "Decision" Diagrams (DDs). The distinction is based on the rules used to evaluate a function for some valuation of the variables. In a decision diagram one simply traverses the unique path from the root to a leaf determined by the variable values, possibly accumulating edge weights. For example, consider the evaluation of a MTBDD for Boolean variable assignment ϕ . That is, ϕ denotes a function that for each variable x assigns a value $\phi(x)$ equal to either 0 or to 1. The evaluation starting at vertex v can be defined as:

$$MTBDDeval(v, \phi) = \begin{cases} \text{Val}(v), & v \text{ is leaf} \\ MTBDDeval(\text{Lo}(v), \phi), & \phi(\text{Var}(v)) = 0 \\ MTBDDeval(\text{Hi}(v), \phi), & \phi(\text{Var}(v)) = 1 \end{cases} \quad (3)$$

In a moment diagram, evaluation requires consideration of multiple paths in the graph. For every vertex v labeled by a variable x that evaluates to 1, subgraphs $\text{Lo}(v)$ and $\text{Hi}(v)$ must both be evaluated and their results summed. The evaluation of BMD for Boolean variable assignment ϕ starting at vertex v can be defined as:

$$BoolEval(v, \phi) = \begin{cases} \text{Val}(v), & v \text{ is leaf} \\ BoolEval(\text{Lo}(v), \phi), & \phi(\text{Var}(v)) = 0 \\ BoolEval(\text{Lo}(v), \phi) + BoolEval(\text{Hi}(v), \phi), & \phi(\text{Var}(v)) = 1 \end{cases} \quad (4)$$

In return for the more complex evaluation rule of moment diagrams, we obtain graphs that are potentially much more compact.

By way of comparison, the moment decomposition of Equation 2 is analogous to the Reed-Muller expansion (also called the positive Davio expansion [9]) for Boolean functions:

$$f = f_{\bar{x}} \oplus x \wedge (f_x \oplus f_{\bar{x}})$$

The expression $f_x \oplus f_{\bar{x}}$ is referred to as the *Boolean difference* of f with respect to x [21], and in many ways is analogous to our linear moment. Other researchers [9, 14] have explored the use of graphs for Boolean functions based on this expansion, calling them Functional Decision Diagrams (FDDs). By our terminology, we would refer to such a graph as a "moment" diagram rather than a "decision" diagram.

2.2. Edge Versus Terminal Weights




Figure 2: Different Representations for Binary-Weighted Bits. All represent the function $X = 4x_2 + 2x_1 + x_0$.

One method to represent functions yielding numeric values, used by MTBDDs and by BMDs, is to simply introduce a distinct leaf vertex for each constant value needed. This approach has the drawback, however, that many leaves may be required, often exponential in the number of variables. Figure 2 illustrates the complexity of the function mapping a vector of Boolean variables: x_{n-1}, \dots, x_1, x_0 to an integer value according to its interpretation as an unsigned binary number. As can be seen, the MTBDD representation will grow exponentially with the word size, since there are 2^n different values for the function.

A second method for defining function values is to associate weights with the edges. This idea was originated by Lai, *et al* in their definition of EVBDDs. In their case, edge weights are combined additively, i.e., the value of a function is determined by following a path from a root to a leaf, summing the edge weights encountered. As shown on the right side of Figure 2, the edge weights of EVBDDs can lead to a much more compact representation than with MTBDDs. In our drawings of EVBDDs, edge weights are shown in square boxes, where an edge without a box has weight 0. For representing a sum of weighted bits, this representation achieves a linear complexity. Various schemes can be used for "normalizing" edge weights so that the resulting graph provides a canonical form for the function. For example, the standard formulation of EVBDDs requires that edge $Lo(v)$ for any vertex v have weight 0.

The bottom of Figure 2 shows the BMD representation of the same function. Observe that the graph for this function grows linearly with word size. In our drawings for BMDs, the solid line leaving vertex v indicates $Hi(v)$, the linear moment. The linear moment of X with respect to any variable x_i is simply its binary weight 2^i , giving rise to the simple

linear structure shown. Thus, the moment decomposition is sufficient for simplifying the representation of this function.

*BMDs also have edge weights, although the weights combine multiplicatively rather than additively. Although not the case for Figure 3, edge weighting can lead to a much more concise representation of a function. As an illustration, Figure 3 shows three representations of the function $8 - 20z + 2y + 4yz + 12x + 24xz + 15xy$. The upper graph is a BMD, with the leaf values corresponding to the coefficients in the monomial expansion. As the figure shows, the BMD data structure misses some opportunities for sharing of common subexpressions. For example, the terms $2y + 4yz$ and $12x + 24xz$ can be factored as $2y(1 + 2z)$ and $12x(1 + 2z)$, respectively. The representation could therefore save space by sharing the subexpression $1 + 2z$. For more complex functions, one might expect more opportunities for such sharing.

The two forms of *BMDs, shown at the bottom of Figure 3 indicate how *BMDs are able to exploit the sharing of common subexpressions. In our drawings of *BMDs, we indicate the weight of an edge in a square box. Unlabeled edges have weight 1. In evaluating the function for a set of arguments, the weights are multiplied together when traversing downward. There are a variety of different rules for manipulating edge weights, resulting in different representations. We will describe two different sets of rules—one that results in rational weights, even when manipulating integer functions (left), and one that yields integer weights, but is only applicable for integer functions (right). Observe that these two rules yield graphs with identical branching structure, but differing in edge weights.

For the remainder of the presentation we will consider mainly *BMDs. The effort required to implement weighted edges is justified by the savings in graph sizes. For functions with integer ranges, we will use integer edge weights. Keeping edge weights as integers is easier than maintaining rational numbers. If we approximate rational numbers with floating point representations, the vagaries of the rounding behavior could greatly complicate the use of *BMDs in formal verification.

2.3. Algebraic Structure

Although we have presented BMDs and *BMDs as methods for representing functions over Boolean variables, they can also be viewed as representing arbitrary linear functions. For example, the BMD of Figure 1 can be viewed as representing the function $F(x, y) = 8 - 20z + 2y + 4yz$ for arbitrary values of y and z . The rule for evaluating a graph given a numeric variable assignment ϕ then becomes:

$$LinEval(v, \phi) = \begin{cases} Val(v), & v \text{ is leaf} \\ LinEval(Lo(v), \phi) + \phi(Var(v)) \cdot LinEval(Hi(v), \phi) & \text{otherwise} \end{cases} \quad (5)$$

The class of linear functions can be defined as either those that can be expressed as a sum of monomial terms, or as those functions that obey Equation 2 for all variables.

An algebraic structure for linear functions provides further insight into our representation. Let L denote the set of linear functions, and for a variable assignment ϕ let $f(\phi)$ denote the result of evaluating linear function f according to this assignment. We can define addition of linear functions in the usual way, i.e., the sum of two functions $f + g$ is a function h such




Figure 3: **Examples of BMD and *BMDs.** All represent the function $8 - 20z + 2y + 4yz + 12x + 24xz + 15xy$. *BMDs have weights on the edges that combine multiplicatively.

that $h(\phi) = f(\phi) + g(\phi)$. It can be seen that the algebraic structure $\langle L, + \rangle$ forms a group, having as identity element the function that always evaluates to 0.

We could define a multiplication over functions in a similar fashion, but then the class of linear functions would not be closed under this operation. The product of two linear functions could yield a quadratic function. In particular, the product of functions f and g , denoted $f \cdot g$ can be defined recursively as follows. If these functions evaluate to constants a and b , respectively, then their product is simply $f \cdot g = a \cdot b$. Otherwise assume the functions are given by their moment expansions (Equation 2) with respect to some variable x . The product of the functions can then be defined as:

$$f \cdot g = f_{\bar{x}} \cdot g_{\bar{x}} + x(f_{\bar{x}} \cdot g_x + f_x \cdot g_{\bar{x}}) + x^2 f_x \cdot g_x \quad (6)$$

One can readily show that this definition is unambiguous—the result is independent of the ordering of the variables in the successive decompositions.

Instead of conventional multiplication, we can define an operation $\hat{\cdot}$ with similar properties, except that it preserves linearity. This involves “demoting” the quadratic term in the equation for conventional multiplication to a linear term. The *linear product* of functions f and g , denoted $f \hat{\cdot} g$, is defined recursively as follows. If these functions evaluate to constants a and b , respectively, then their linear product is simply their product: $f \hat{\cdot} g = a \cdot b$. Otherwise assume the functions are given by their moment expansions (Equation 2) with respect to some variable x . Their linear product is defined as

$$f \hat{\cdot} g = f_{\bar{x}} \hat{\cdot} g_{\bar{x}} + x(f_{\bar{x}} \hat{\cdot} g_x + f_x \hat{\cdot} g_{\bar{x}} + f_x \hat{\cdot} g_x) \quad (7)$$

One can show that the definition is independent of the ordering in the decomposition. The algebraic structure $\langle L, +, \hat{\cdot} \rangle$ forms a ring. That is, $\hat{\cdot}$ is associative, and it distributes over $+$. Furthermore, the function that always yields 1 serves as a unit for this ring.

Although the linear product operation is not the same as conventional multiplication, there are two important cases where we can safely use $f \hat{\cdot} g$ as a replacement for $f \cdot g$. First, under the *Boolean domain restriction*, i.e., considering only variable assignments ϕ such that $\phi(x) \in \{0, 1\}$, we are guaranteed that $[f \cdot g](\phi) = [f \hat{\cdot} g](\phi)$. Second, define the *support* of a function f as those variables x such that $f_x \neq 0$. Under the *independent support assumption*, where functions f and g have disjoint support sets, we have that $f \cdot g = f \hat{\cdot} g$ for any variable assignment. In particular, for any variable x we must have that either f_x or g_x is identically 0, and hence the quadratic term of Equation 6 drops out.

In general, we can “linearize” any operation op to create an operation \hat{op} such that for any Boolean variable assignment ϕ , we have $[f \hat{op} g](\phi) = f(\phi) op g(\phi)$. This involves generating moments with respect to each variable x as:

$$[f \hat{op} g]_{\bar{x}} = f_{\bar{x}} \hat{op} g_{\bar{x}} \quad (8)$$

$$\begin{aligned} [f \hat{op} g]_x &= [f \hat{op} g]_x - [f \hat{op} g]_{\bar{x}} \\ &= [f_x \hat{op} g_x] - [f_{\bar{x}} \hat{op} g_{\bar{x}}] \\ &= [(f_{\bar{x}} + f_x) \hat{op} (g_{\bar{x}} + g_x)] - [f_{\bar{x}} \hat{op} g_{\bar{x}}] \end{aligned} \quad (9)$$

As before, the definition is independent of the variable ordering. In general, this linearization would not yield valid results for non-Boolean variable assignments, whether or not the




Figure 4: Representations of Signed Integers. All commonly used encodings can be represented easily.

arguments have independent support. For example, the linearized form of exponentiation would convert $(x + 2)^y$ into $1 + y + xy$.

3. Representation of Numeric Functions

*BMDs provide a concise representation of functions defined over "words" of data, i.e., vectors of bits having a numeric interpretation. Let \vec{x} represent a vector of Boolean variables: x_{n-1}, \dots, x_1, x_0 . These variables can be considered to encode an integer X according to some encoding, e.g., unsigned binary, two's complement, BCD, etc. As Figure 2 shows, the *BMD (as well as BMD) representations for X according to an unsigned binary encoding have linear complexity. Figure 4 illustrates the *BMD representations of several common encodings for signed integers, where x_{n-1} is the sign bit. The sign-magnitude encoding gives integer value $X = -1^{x_{n-1}}X'$, where X' is the unsigned integer encoded by the remaining bits. Observe that this can be expressed in the linear form $(1 - 2x_{n-1})X'$, yielding a graph structure where both moments for variable x_{n-1} point to the graph for X' , but having edge weights 1 and -2. As the other graphs in the figure illustrate, both two's complement and one's complement encodings can be viewed as sums of weighted bits, where the sign bit is weighted either -2^{n-1} (two's complement) or $1 - 2^{n-1}$ (one's complement) [18].

The conciseness of *BMDs arises from two important properties of typical encodings. First, many encodings are based on a sum of weighted bits. In terms of the monomial expansion, this implies that the terms are all of low degree. Second, the regularity of the encodings gives rise to many subexpressions differing only by multiplicative factors. This leads to sharing of subgraphs in the *BMD, with edge weights denoting the different factors.

3.1. Word-Level Operations

Form	X	$X + Y$	$X * Y$	X^2	c^X
MTBDD	exponential	exponential	exponential	exponential	exponential
EVBDD	linear	linear	exponential	exponential	exponential
BMD	linear	linear	quadratic	quadratic	exponential
*BMD	linear	linear	linear	quadratic	linear

Table 2: **Word-Level Operation Complexity.** Expressed in how the graph sizes grow relative to the word size.

Table 2 provides a comparative summary of the four function representations for a number of word-level operations on unsigned data. *BMD examples of these functions are included in this paper. As can be seen, MTBDDs are totally unsuited for this class of functions. The range of the functions to be represented is simply too large. EVBDDs yield better results for representing word-level data and for representing “additive” operations (e.g, addition and subtraction) at the word level. This capability was exploited by Lai and Sastry in verifying adder circuits against word-level specifications [15]. On the other hand, EVBDDs cannot efficiently represent more complex functions such as multiplication, squaring, and exponentiation. Thus, for example, they cannot be used for verifying multipliers. In fact, all published examples that can be handled efficiently at the word level using EVBDDs can be handled at the bit level using BDDs. Their utility in verifying circuits is mainly for providing a more abstract form of specification.

Both BMDs and *BMDs are much more effective for representing word-level operations. BMDs remain of polynomial (quadratic) size for both multiplication and for squaring, although they grow exponentially for exponentiation. *BMDs do even better, being quadratic for squaring and linear for all other operations listed. By verifying circuits at the word level with *BMDs, we can handle classes of systems that are beyond the capability of BDDs and other bit-level techniques.

Figure 5 illustrates the *BMD representations of addition and multiplication expressed at a word level. Observe that the sizes of the graphs grow only linearly with the word size n . Word-level addition can be viewed as summing a set of weighted bits, where bits x_i and y_i both have weight 2^i . Word-level multiplication can be viewed as summing a set of partial products of the form $x_i 2^i Y$.

As with BDDs, the representation of a function depends on the variable ordering. For example, Figure 6 shows the *BMDs for word-level multiplication under two additional variable orderings. Observe that although these graphs appear more complex than the one of Figure 5, their complexity still grows only linearly with n . In our experience, *BMDs are much less sensitive to variable ordering than are BDDs.

Figure 7 illustrates the *BMD representations of two unary operations on word-level data. For representing the function c^X (in this case $c = 2$), the *BMD has linear complexity. It expresses the function as a product of factors of the form $c^{2^i x_i} = (c^{2^i})^{x_i}$. Since x_i evaluates to either 0 or to 1, the exponentiation can be linearized as: $a^{x_i} = 1 + (a - 1)x_i$. In the graph, a vertex labeled by variable x_i has outgoing edges with weights 1 and $c^{2^i} - 1$ both leading to a common vertex denoting the product of the remaining factors.




Figure 5: Representations of Word-Level Sum and Product. The graphs grow linearly with word size.

For representing the function X^2 , both the BMD and the *BMD have quadratic complexity. The representation can be seen to follow a recursive expansion of the function based on the decomposition: $X = X_n = 2^{n-1}x_{n-1} + X_{n-1}$, where X_k denotes the weighted sum of variables x_0 through x_{k-1} . In terms of this decomposition we have:

$$X_n^2 = (2^{n-1}x_{n-1} + X_{n-1})^2$$

$$2^{2n-2}x_{n-1}^2 + 2^n x_{n-1}X_{n-1} + X_{n-1}^2$$

Since x_{n-1} is Boolean-valued, we can “demote” the quadratic term x_{n-1}^2 to a linear term x_{n-1} . Thus, the constant moment for the function is X_{n-1}^2 , while the linear moment is $2^{2n-2} + 2^n X_{n-1} = 2^n(X_{n-1} + 2^{n-2})$. In our example with $n = 4$, the left subgraph represents the function X_3^2 , while the right side represents the subgraph $16(X_3 + 4)$. Observe that the different constant offsets for each bit cause the growth of the graph to be quadratic rather than linear. That is, there is no sharing between the graphs for the terms $X_{i-1} + 2^{i-2}$ for different values of i . For many applications, this quadratic complexity is acceptable. For example, we could represent the square of a 32-bit number by a graph of around 530 vertices.

4. Representation of Boolean Functions

Boolean functions are just a special case of numeric functions having a restricted range. Therefore such functions can be represented as BMDs or *BMDs. The algebraic structure introduced in Section 2.3 provides a convenient notation for translating Boolean operations




Figure 6: Representations of Word-Level Product for Other Variable Orderings. The graphs grow linearly with the word size regardless of the variable ordering.




Figure 7: Representations of Unary Operations at Word-Level. The graph for 2^X grows linearly with the word size, while that for X^2 grows quadratically.




Figure 8: Representations of Boolean Functions. Representations as *BMDs are comparable in size to BDDs.

into operations on linear functions. In particular, let f and g denote functions have Boolean ranges. Then we can define the standard Boolean operations as:

$$\begin{aligned}
 \bar{f} &= 1 - f \\
 f \wedge g &= f \hat{\cdot} g \\
 f \vee g &= f + g - (f \hat{\cdot} g) \\
 f \oplus g &= f + g - 2(f \hat{\cdot} g)
 \end{aligned}
 \tag{10}$$

Figure 8 illustrates the *BMD representations of several common Boolean functions over multiple variables, namely their Boolean product and sum, as well as their exclusive-or sum. As this figure shows, the *BMD of Boolean functions may have values other than 0 or 1 for edge weights and leaf values. Under all variable assignments, however, the function will evaluate to 0 or to 1. As can be seen in the figure, these functions all have representations that grow linearly with the number of variables, as is the case for their BDD representations. The representation for AND follows due to the parallel between Boolean and linear products. The representation for OR can be seen to follow an iterative structure. In particular, let F_n denote the OR of variables x_1, x_2, \dots, x_n , and G_n denote their NOR, i.e., $G_n = 1 - F_n$. Function F_n can be rewritten as:

$$\begin{aligned}
 F_n &= x_n \vee F_{n-1} \\
 &= x_n + F_{n-1} - (x_n \hat{\cdot} F_{n-1}) \\
 &= F_{n-1} + x_n(1 - F_{n-1}) \\
 &= F_{n-1} + x_n G_{n-1}
 \end{aligned}$$

Thus, the moments of function F_n with respect to variable x_n are F_{n-1} and G_{n-1} . Based on this result, function G_n can be rewritten as:

$$\begin{aligned}
 G_n &= 1 - F_n \\
 &= 1 - F_{n-1} - x_n G_{n-1} \\
 &= G_{n-1} + x_n(-G_{n-1})
 \end{aligned}$$

Thus, the moments of function G_n with respect to variable x_n are G_{n-1} and $-G_{n-1}$. In the center graph of Figure 8, the vertices on the left side denote the sequence of OR functions, while those on the right side denote the sequence of NOR functions.

The representation for EXCLUSIVE-OR follows a similar iterative structure. It can be generated by defining function F_n to be the EXCLUSIVE-OR of variables x_1, x_2, \dots, x_n , while letting G_n denote the function $G_n = 1 - 2F_n$. It can be shown that F_n has moments F_{n-1} and G_{n-1} , while G_n has moments G_{n-1} and $-2G_{n-1}$.

Figure 9 illustrates the similarity between BDDs and *BMDs when representing the Boolean functions describing an adder circuit at the bit level. Observe the relation between the word-level representation (Figure 5) and the bit-level representation of addition. Both are functions over variables representing the adder inputs, but the former is a single function yielding an integer value, while the latter is a set of Boolean functions: one for each output




Figure 9: Bit-Level Representations of Addition Functions. Each graph represents all four outputs of a 3-bit adder.

signal for the circuit. The relation between these two representations will be discussed more fully in our development of a verification methodology.

The BDD representation shown in Figure 9 employs two techniques to reduce its size [3]. First, it represents a set of functions by a single graph with multiple roots, allowing different functions to share common subgraphs. In fact, the set of functions is maintained in *strong canonical form*, where every function to be represented is denoted by a unique root vertex. The *BMD representation can also use this form of sharing and maintained in strong canonical form. Second, the BDD contains "negative edges" (indicated by dots on the edge) to denote Boolean complementation. The use of edge weights in *BMDs has a similar effect, although edge weights cannot be used to directly represent the complement operation: $\bar{f} = 1 - f$. Observe in any case that the *BMD representation for these functions has a similar structure to the BDD representation. Both grow linearly with the word size, with the *BMD requiring 7 vertices per bit position, and the BDD requiring 5.

In all of the examples shown, the *BMD representation of a Boolean function is of comparable size to its BDD representation. We conjecture, however, that this is not always the case. The two representations are based on different expansions of the function, and hence there would not seem to be any fundamental reason for them to be of similar complexity.

5. Factoring and Other Decision Properties

One powerful property of BDDs is that, given a BDD representation of a function f over a set of variables \vec{x} , one can easily find solutions to the equation $f(\vec{x}) = 0$ by tracing paths from the root to the leaf with value 0. This strength of BDDs is also a limitation. Since any problem that can be expressed as a function f having an efficient BDD representation is amenable to easy solution, this implies that BDDs cannot efficiently represent functions corresponding to intractable problems.

Imagine for example, that it were possible to construct the $2n$ BDDs giving a bit-level representation of multiplication over n -bit integers \vec{x} and \vec{y} . Then we could potentially factor a large number K , by solving the equation:

$$\bigwedge_{i=0}^{2n-1} P_i(\vec{x}, \vec{y}) \oplus k_i = 0$$

where P_i is the function representing bit i of the product, and k_i is the i th bit of K . Observe in this equation that the values k_i are constants, and therefore the computation involves forming the product of either true or complemented multiplier output functions. Experts consider factoring to be a "hard" problem. In fact, the RSA encryption algorithm [23] relies on the assumption that given the public key, one cannot derive the two prime factors of the key in a reasonable amount of time. Thus, one would expect that some step in the above scheme for factoring would break down. In the case of BDDs, the problem comes in trying to generate the BDD representations of the functions P_i . It can be shown that these graphs grow exponentially with the word size [5].

Define the task of "finding a zero for function f " as finding a (Boolean) variable assignment such that $f(x) = 0$. We will call a representation for functions "easily invertible" if it




Figure 10: **Representation of Factoring Problem.** Solving requires finding variable assignment that evaluates to 0—not an easy task.

is always possible to find a zero for the function in time polynomial in the size of the representation. Both BDDs and MTBDDs have this property—one simply finds a path to the leaf with value 0. One can also show that FDDs are easily invertible [2], even though evaluation does not involve simply following a single path in the graph.

On the other hand, EVBDDs are not easily invertible, assuming $P \neq NP$. The following argument shows that the problem of finding a zero of a function represented by an EVBDD is NP-complete. First, the problem is clearly NP, since given an assignment to the variables, one can evaluate an EVBDD and determine whether the function yields 0 for this assignment. Furthermore, any instance of the NP-complete Partition problem [10] can readily be translated into an EVBDD solution problem. This problem is defined as: given a set of n elements A , where each element i has a nonnegative integer “size” s_i , determine whether there exists a subset A' such that

$$\sum_{i \in A'} s_i = \sum_{i \in A - A'} s_i$$

To translate this into an equation solution problem, let $S = \sum_{i \in A} s_i$, and define the function f as:

$$f(x_1, \dots, x_n) = -S/2 + \sum_{i=1}^n x_i s_i \quad (11)$$

This function has an EVBDD with n nonterminal vertices. It is similar in structure to that of Figure 2, except that the outgoing solid arc from a vertex with variable x_i has weight s_i , and the root has weight $-S/2$. The challenge of solving this problem for EVBDDs can be seen to lie with the edge weights. One must find a path through the graph such that the edge weights encountered sum to 0.

By a similar argument, one can show that BMDs and *BMDs also do not form easily invertible representations. Both are clearly in NP, since evaluation can be performed in time linear in the graph sizes. Furthermore, both provide linear-sized representations of the function defined in Equation 11. For example, the BMD representation of this function has structure similar to that of Figure 2. The solid arc from a vertex with variable x_i points to a leaf with value s_i , while the dashed arc from the vertex with variable x_0 points to a leaf with value $-S/2$. The *BMD has similar structure, but possibly with weights moved up into the edges.

The challenge of finding a zero of a BMD or *BMD can be seen to lie with the evaluation rule, given by Equation 4—evaluation requires considering multiple paths in the graph. We can readily represent the factoring problem, as shown in Figure 10 by constructing a *BMD representation of the function $X \cdot Y - K$ (in this example $K = 35$). The BMD representation of this function is somewhat more complex, but still of size quadratic in n . The lack of an efficient inversion algorithm prevents one from factoring by this method.

The example of factoring illustrates the fact that the strengths and weaknesses of BDDs versus *BMDs are somewhat orthogonal. Tasks that can easily be performed on BDDs are much more difficult to perform on *BMDs. On the other hand, *BMDs can represent circuit functions that cause exponential blow up for BDDs or to their extensions as MTBDDs and EVBDDs.

6. Algorithms

In this section we describe key algorithms for constructing and manipulating *BMDs. The algorithms have a similar style to their counterparts for BDDs. Unlike operations on BDDs where the complexities are at worst polynomial in the argument sizes, most operations on *BMDs potentially have exponential complexity. We will show in the experimental results, however, that these exponential cases do not arise in our applications.

6.1. Representation of *BMDs

We will represent a function as a “weighted pair” of the form $\langle w, v \rangle$ where w is a numeric weight and v designates a graph vertex. Weights can either be maintained as integers or real numbers. Maintaining rational-valued weights follows the same rules as the real case. Vertex $v = \Lambda$ denotes a terminal leaf, in which case the weight denotes the leaf value. The weight w must be nonzero, except for the terminal case. Each vertex v has the following attributes:

$\text{Var}(v)$ The vertex variable.

$\text{Hi}(v)$ The pair designating the linear moment.

$\text{Lo}(v)$ The pair designating the constant moment.

$\text{Uid}(v)$ Unique identifier for vertex.

Observe that each edge in the graph is also represented as a weighted pair.

6.2. Maintaining Canonical Form

The functions to be represented are maintained as a single graph in *strong canonical form*. That is, pairs $\langle w_1, v_1 \rangle$ and $\langle w_2, v_2 \rangle$ denote the same function if and only if $w_1 = w_2$ and $v_1 = v_2$. We assume that the set of variables is totally ordered, and that all of the vertices constructed obey this ordering. That is, for any vertex v , its variable $\text{Var}(v)$ must be less than any variable appearing in the subgraphs $\text{Lo}(v)$ and $\text{Hi}(v)$.

Maintaining a canonical form requires obeying a set of conventions for vertex creation and for weight manipulation. These conventions are expressed by the pseudo-code shown in Figures 11 and 12. The *MakeBranch* algorithm provides the primary means of creating and reusing vertices in the graph. It is given as arguments a variable and two moments, each represented as weighted pairs. It returns a pair representing the function given by Equation 2. It assumes that the argument variable is less than any variable in the argument subgraphs. The steps performed by *MakeBranch* are illustrated in Figure 13. In this figure two moments are drawn as weighted pointers.

When the linear moment is the constant 0, we can simply return the constant moment as the result, since this function is independent of variable x . Observe that this rule differs from the reduction rule for a graph based on a pointwise decomposition such as BDDs. In

```

function MakeBranch(variable  $x$ , pair  $\langle w_l, v_l \rangle$ , pair  $\langle w_h, v_h \rangle$ ): pair
{ Create a branch, normalize weights. }
{ Assumes that  $x < \text{Var}(v_h)$  and  $x < \text{Var}(v_l)$  }
  if  $w_h = 0$  then return  $\langle w_l, v_l \rangle$ 
   $w \leftarrow \text{NormWeight}(w_l, w_h)$ 
   $w_l \leftarrow w_l/w$ 
   $w_h \leftarrow w_h/w$ 
   $v \leftarrow \text{UniqueVertex}(x, \langle w_l, v_l \rangle, \langle w_h, v_h \rangle)$ 
  return  $\langle w, v \rangle$ 

```

```

function UniqueVertex(variable  $x$ , pair  $\langle w_l, v_l \rangle$ , pair  $\langle w_h, v_h \rangle$ ): vertex
{ Maintain set of graph vertices such that no duplicates created }
   $key \leftarrow [x, w_l, \text{Uid}(v_l), w_h, \text{Uid}(v_h)]$ 
   $found, v \leftarrow \text{LookUp}(U\text{Table}, key)$ 
  if  $found$  then return  $v$ 
   $v \leftarrow \text{New}(\text{vertex})$ 
   $\text{Var}(v) \leftarrow x, \text{Uid}(v) \leftarrow \text{Unid}();$ 
   $\text{Lo}(v) \leftarrow \langle w_l, v_l \rangle; \text{Hi}(v) \leftarrow \langle w_h, v_h \rangle$ 
   $\text{Insert}(U\text{Table}, key, v)$ 
  return  $v$ 

```

```

function NormWeight(integer  $w_l$ , integer  $w_h$ ): integer
{ Normalization function, integer weights. }
   $w \leftarrow \text{gcd}(w_l, w_h)$ 
  if  $w_l < 0$ 
    then return  $-w$ 
  else return  $w$ 

```

```

function NormWeight(real  $w_l$ , real  $w_h$ ): real
{ Normalization function, real weights }
  if  $w_l = 0$ 
    then return  $w_h$ 
  else return  $w_l$ 

```

Figure 11: Algorithms for Maintaining *BMD. These algorithms preserve a strong canonical form.

```

function ApplyWeight(wtype  $w'$ , pair  $\langle w, v \rangle$ ): pair
{ Multiply function by constant }
  if  $w' = 0$  then return  $\langle 0, \Lambda \rangle$ 
  return  $\langle w' \cdot w, v \rangle$ 

```

Figure 12: Algorithm for Multiplying Function by Weight. This algorithm ensures that edge to a nonterminal vertex has weight 0.




Figure 13: Normalizing Transformations Made by *MakeBranch*. These transformations enforce the rules on branch weights.

such cases a vertex can be eliminated when both of its children are identical. This reflects the difference between the two different function decompositions. Our rule for *BMDs is similar to that for FDDs [9, 14].

For other values of the linear moment, the routine first factors out some weight w , adjusting the weights of the two arguments accordingly. We show two versions of a function *NormWeight* according to whether integer or real-valued weights are to be used. For the integer case, we want to extract any common factor while ensuring that all weights are integers. Hence we take the greatest common divisor (gcd) of the argument weights. In addition, we adopt the convention that the sign of the extracted weight matches that of the constant moment. This assumes that gcd always returns a nonnegative value. For real-valued weights we adopt the convention that the weighted pair designating the constant moment for a vertex always has weight 0 (only when this moment is the constant 0) or 1. In the former case the weight of the pair designating the first moment will have weight 1. Thus, normalizing real-valued weights involves moving one of the argument weights up and adjusting the other.

Once the weights have been normalized *MakeBranch* calls the function *UniqueVertex* to find an existing vertex or create a new one. This function maintains a table (typically a hash table) where each entry is indexed by a key formed from the variable and the two moments. Every vertex in the graph is stored according to such a key and hence duplicate vertices are never constructed.

Figure 12 shows the code for a function *ApplyWeight* to multiply a function, given as a weighted pair, by a constant value, given as a weight w' . This procedure simply adjusts the pair weight, detecting the special case where the multiplicative constant is 0.

As long as all vertices are created through calls to the *MakeBranch* function and all multiplications by constants are performed by calls to *ApplyWeight*, the graph will remain in strongly canonical form.

6.3. The Apply Operations

As with BDDs, *BMDs are constructed by starting with base functions corresponding to constants and single variables, and then building more complex functions by combining simpler functions according to some operation. In the case of BDDs this combination is expressed by a single algorithm that can apply an arbitrary Boolean operation to a pair of functions. In the case of *BMDs we require algorithms tailored to the characteristics of the individual operations. To simplify the presentation, we show only a few of these algorithms and attempt to do so in as uniform a style as possible. These algorithms are referred to collectively as "Apply" algorithms.

Figure 14 shows the fundamental algorithm for adding two functions. The function *PlusApply* takes two weighted pairs indicating the argument functions and returns a weighted pair indicating the result function. This algorithm can also be used for subtraction by first multiplying the second argument by weight -1 . This code closely follows the Apply algorithm for BDDs [3]. It utilizes a combination of recursive descent and "memoizing," where all computed results are stored in a table and reused whenever possible. The recursion is based on the property that taking moments of functions commutes with addition. That is, for

```

function PlusApply(pair  $\langle w_1, v_1 \rangle$ , pair  $\langle w_2, v_2 \rangle$ ): pair
{ Compute sum of two functions }
  done,  $\langle w, v \rangle \leftarrow$  TermCheck(+,  $\langle w_1, v_1 \rangle$ ,  $\langle w_2, v_2 \rangle$ )
  if done then return  $\langle w, v \rangle$ 

   $w', \langle w_1, v_1 \rangle, \langle w_2, v_2 \rangle \leftarrow$  Rearrange(+,  $\langle w_1, v_1 \rangle$ ,  $\langle w_2, v_2 \rangle$ )
  key  $\leftarrow$  [+ ,  $w_1$ , Uid( $v_1$ ),  $w_2$ , Uid( $v_2$ )]
  found,  $\langle w, v \rangle \leftarrow$  LookUp(OpTable, key)
  if found then return ApplyWeight( $w'$ ,  $\langle w, v \rangle$ )

   $x \leftarrow$  Min(Var( $v_1$ ), Var( $v_2$ ))

  { Begin recursive section }
   $\langle w_{1l}, v_{1l} \rangle \leftarrow$  SimpleMoment( $\langle w_1, v_1 \rangle$ ,  $x$ , 0)
   $\langle w_{2l}, v_{2l} \rangle \leftarrow$  SimpleMoment( $\langle w_2, v_2 \rangle$ ,  $x$ , 0)
   $\langle w_{1h}, v_{1h} \rangle \leftarrow$  SimpleMoment( $\langle w_1, v_1 \rangle$ ,  $x$ , 1)
   $\langle w_{2h}, v_{2h} \rangle \leftarrow$  SimpleMoment( $\langle w_2, v_2 \rangle$ ,  $x$ , 1)

   $\langle w_l, v_l \rangle \leftarrow$  PlusApply( $\langle w_{1l}, v_{1l} \rangle$ ,  $\langle w_{2l}, v_{2l} \rangle$ )
   $\langle w_h, v_h \rangle \leftarrow$  PlusApply( $\langle w_{1h}, v_{1h} \rangle$ ,  $\langle w_{2h}, v_{2h} \rangle$ )
  { End recursive section }

   $\langle w, v \rangle \leftarrow$  MakeBranch( $x$ ,  $\langle w_l, v_l \rangle$ ,  $\langle w_h, v_h \rangle$ )
  Insert(OpTable, key,  $\langle w, v \rangle$ )
  return ApplyWeight( $w'$ ,  $\langle w, v \rangle$ )

function SimpleMoment(pair  $\langle w, v \rangle$ , variable  $x$ , integer  $b$ ): pair
{ Find moment of function under special condition. }
{ Variable either at root vertex  $v$ , or not present in graph. }
{  $b = 0$  for constant moment,  $b = 1$  for linear }

  if Var( $v$ )  $\neq$   $x$ 
    if  $b = 0$ 
      then return  $\langle w, v \rangle$ 
      else return  $\langle 0, \Lambda \rangle$ 
    if  $b = 1$ 
      then return ApplyWeight( $w$ , Lo( $v$ ))
      else return ApplyWeight( $w$ , Hi( $v$ ))

```

Figure 14: Apply Algorithm for Adding Two Functions. The algorithm is similar to the counterpart for BDDs.

<i>op</i>	$\langle w_1, v_1 \rangle$	$\langle w_2, v_2 \rangle$	$\langle w, v \rangle$
+	$\langle 0, \Lambda \rangle$		$\langle w_2, v_2 \rangle$
+		$\langle 0, \Lambda \rangle$	$\langle w_1, v_1 \rangle$
+	$\langle w_1, v \rangle$	$\langle w_2, v \rangle$	<i>ApplyWeight</i> ($w_1 + w_2, \langle 1, v \rangle$)
*	$\langle w_1, \Lambda \rangle$		<i>ApplyWeight</i> ($w_1, \langle w_2, v_2 \rangle$)
*		$\langle w_2, \Lambda \rangle$	<i>ApplyWeight</i> ($w_2, \langle w_1, v_1 \rangle$)
\div		$\langle w_2, \Lambda \rangle$	<i>ApplyWeight</i> ($1/w_2, \langle w_1, v_1 \rangle$)

Table 3: Termination Cases for Apply Algorithms. Each line indicates an operation, a set of terminations, and the returned result.

<i>op</i>	Arguments		Results		
		Condition	w'	$\langle w_1, v_1 \rangle$	$\langle w_2, v_2 \rangle$
*		$\text{Uid}(v_1) > \text{Uid}(v_2)$	$w_1 \cdot w_2$	$\langle 1, v_1 \rangle$	$\langle 1, v_2 \rangle$
*		$\text{Uid}(v_1) \leq \text{Uid}(v_2)$	$w_1 \cdot w_2$	$\langle 1, v_2 \rangle$	$\langle 1, v_1 \rangle$
+		$ w_1 > w_2 $	<i>NormWeight</i> (w_1, w_2)	$\langle w_1/w', v_1 \rangle$	$\langle w_2/w', v_2 \rangle$
+		$ w_1 \leq w_2 $	<i>NormWeight</i> (w_2, w_1)	$\langle w_2/w', v_2 \rangle$	$\langle w_1/w', v_1 \rangle$
\div			w_1/w_2	$\langle 1, v_1 \rangle$	$\langle 1, v_2 \rangle$

Table 4: Rearrangements for Apply Algorithms. These rearrangements increase the likelihood of reusing a previously-computed result.

functions f and g and for variable x :

$$[f + g]_{\bar{x}} = f_{\bar{x}} + g_{\bar{x}}$$

$$[f + g]_{\dot{x}} = f_{\dot{x}} + g_{\dot{x}}$$

This routine, like the other Apply algorithms, first checks a set of termination conditions to determine whether it can return a result immediately. This test is indicated as a call to function *TermCheck* having as arguments the operation and the arguments of the operation. This function returns two values: a Boolean value *done* indicating whether immediate termination is possible, and a weighted pair indicating the result to return in the event of termination. Some sample termination conditions are shown in Table 3. For the case of addition, the algorithm can terminate if either argument represents the constant 0, or if the two arguments are multiples of each other, indicated by weighted pairs having the same vertex element.

Failing the termination test, the routine attempts to reuse a previously computed result. To maximize possible reuse it first rearranges the arguments and extracts a common weight w' . This process is indicated as a call to the function *Rearrange* having the same arguments as *TermCheck*. This function returns three values: the extracted weight and the modified arguments to the operation. Some sample rearrangements are shown in Table 4. For the case of addition rearranging involves normalizing the weights according to the same conditions used in *MakeBranch* and ordering the arguments so that the first has greater weight. For example, suppose at some point we compute $6y - 9z$. We will extract weight -3 (assuming integer weights) and rearrange the arguments as $3z$ and $-2y$. If we later attempt to compute $15z - 10y$, we will be able to reuse this previous result with extracted weight 5.

If the routine fails to find a previously computed result, it makes recursive calls to compute the sums of the two moments according to the minimum variable in its two arguments. In generating the arguments for the recursion, it calls a function *SimpleMoment* to compute the moments. This routine can only compute a moment with respect to a variable that either does not appear in the graph or is at its root, a condition that is guaranteed by the selection of x as the minimum variable in the two graphs. When the variable does not appear in the graph, the constant moment is simply the original function, while the linear moment is the constant 0. When the variable appears at the root, the result is the corresponding subgraph multiplied by the weight of the original argument. The final result of *PlusApply* is computed by calling *MakeBranch* to generate the appropriate function and multiplying this function by the constant extracted when rearranging the arguments.

Observe that the keys for table *OpTable* index prior computations by both the weights and the vertices of the (rearranged) arguments. In the worst case, the rearranging may not be effective at creating matches with previous computations. In this event, the weights on the arcs would be carried downward in the recursion, via the calls to *SimpleMoment*. In effect, we are dynamically generating BMD representations from the *BMD arguments. Thus, if functions f and g have BMD representations of size m_f and m_g , respectively, there would be no more than $m_f m_g$ calls to *PlusApply*, and hence the overall algorithm has worst case complexity $O(m_f m_g)$. As we have seen, many useful functions have polynomial BMD sizes, guaranteeing polynomial performance for *PlusApply*. On the other hand, some functions blow up exponentially in converting from a *BMD to a BMD representation, in which case

```

{ Begin recursive section }
⟨w1l, v1l⟩ ← SimpleMoment(⟨w1, v1⟩, x, 0)
⟨w2l, v2l⟩ ← SimpleMoment(⟨w2, v2⟩, x, 0)
⟨w1h, v1h⟩ ← PlusApply(SimpleMoment(⟨w1, v1⟩, x, 1), ⟨w1l, v1l⟩)
⟨w2h, v2h⟩ ← PlusApply(SimpleMoment(⟨w2, v2⟩, x, 1), ⟨w2l, v2l⟩)

⟨wl, vl⟩ ← BinApply(op, ⟨w1l, v1l⟩, ⟨w2l, v2l⟩)
⟨wh, vh⟩ ← PlusApply(BinApply(op, ⟨w1h, v1h⟩, ⟨w2h, v2h⟩), ⟨-wl, vl⟩)
{ End recursive section }

```

Figure 15: Recursive Section of Apply Algorithm for Arbitrary Binary Operation. This generic algorithm does not exploit particular properties of the operation.

the algorithm may have exponential complexity. We will see with the experimental results, however, that this exponential blow-up does not occur for the cases we have tried. The termination checks and rearrangements are very effective at stopping the recursion.

The Apply algorithms for other operations have a similar overall structure to that for addition, but differing in the recursive evaluation. Comments in the code of Figure 14 delimit the “recursive section” of the routine. In this section recursive calls are made to create a pair of weighted pointers $\langle w_l, v_l \rangle$ and $\langle w_h, v_h \rangle$ from which the returned result is constructed. For the remaining Apply algorithms we show only their recursive sections.

Figure 15 shows the recursive section for applying an arbitrary binary operation op to a pair of functions. This algorithm can be seen to implement the linearized form \hat{op} defined by Equations 9 and 9. At each recursive step of the computation in Figure 15, we must sum the moments of the arguments to generate their positive cofactors, recursively apply the operation to these cofactors, and then subtract the constant moment to obtain a linear moment. In effect we dynamically construct an MTBDD representation of the arguments. Thus, one would expect that this computation would perform poorly unless either the arguments have efficient MTBDD representations, or the termination checks and rearrangements can stop the recursion from expanding into a large number of cases.

Rather than resorting to the generic Apply algorithm of Figure 15, it is preferable to exploit properties of the operation so that the positive cofactors of the arguments do not need to be generated. Figure 16 shows how this can be done for multiplication, using the formulation of linear product given by Equation 7. Each call to *MultApply* requires four recursive calls, plus two calls to *PlusApply*. With the rearrangements shown in Table 4, we can always extract the weights from the arguments. Hence if the arguments have *BMD representations of m_f and m_g vertices, respectively, no more than $m_f m_g$ calls will be made to *MultApply*. Unfortunately, this bound on the calls does not suffice to show a polynomial bound on the complexity of the algorithm. The calls to *PlusApply* may blow up exponentially.

6.4. Affine Substitution

```

{ Begin recursive section }
⟨w1l, v1l⟩ ← SimpleMoment(⟨w1, v1⟩, x, 0)
⟨w2l, v2l⟩ ← SimpleMoment(⟨w2, v2⟩, x, 0)
⟨w1h, v1h⟩ ← SimpleMoment(⟨w1, v1⟩, x, 1)
⟨w2h, v2h⟩ ← SimpleMoment(⟨w2, v2⟩, x, 1)

⟨wl, vl⟩ ← MultApply(⟨w1l, v1l⟩, ⟨w2l, v2l⟩)
⟨whh, vhh⟩ ← MultApply(⟨w1h, v1h⟩, ⟨w2h, v2h⟩)
⟨whl, vhl⟩ ← MultApply(⟨w1h, v1h⟩, ⟨w2l, v2l⟩)
⟨wlh, vlh⟩ ← MultApply(⟨w1l, v1l⟩, ⟨w2h, v2h⟩)
⟨wh, vh⟩ ← PlusApply(⟨whh, vhh⟩, PlusApply(⟨whl, vhl⟩, ⟨wlh, vlh⟩))
{ End recursive section }

```

Figure 16: Recursive Section for Apply Operation for Multiplying Functions. This operation exploits the ring properties of linear product.

```

AffineSubst(pair ⟨w, v⟩, assignment μ, assignment β)
{ Replace each variable x in function by μ(x) · x + β(x) }
  if v = Λ then return ⟨w, v⟩
  Key ← [v, μ, β]
  found, ⟨wi, vi⟩ ← LookUp(SubstTable, key)
  if found then return ApplyWeight(w, ⟨wi, vi⟩)
  x ← Var(x)
  ⟨wl, vl⟩ ← AffineSubst(Lo(v), μ, β)
  ⟨wh, vh⟩ ← AffineSubst(Hi(v), μ, β)
  ⟨wl, vl⟩ ← PlusApply(⟨wl, vl⟩, ApplyWeight(β(x), ⟨wh, vh⟩))
  ⟨wh, vh⟩ ← ApplyWeight(μ(x), ⟨wh, vh⟩)
  ⟨wi, vi⟩ ← MakeBranch(x, ⟨wl, vl⟩, ⟨wh, vh⟩)
  Insert(SubstTable, key, ⟨wi, vi⟩)
  return ApplyWeight(w, ⟨wi, vi⟩)

```

Figure 17: Affine Substitution Algorithm. Each variable in the function is replaced by an affine transformation of the variable.




Figure 18: **Formulation of Verification Problem.** The goal of verification is to prove a correspondence between a bit-level circuit and a word-level specification

Figure 17 shows an algorithm for performing a very general form of function evaluation we will call *affine substitution*. The idea is to substitute for each variable x a function of the form $mx + b$. The result will be a function over the same set of variables, or possibly a subset of these variables. By selecting different values of m and b we can obtain many useful substitutions. For example, with $b = a$ and $m = 0$, we obtain the result of assigning value a to the variable. Thus, this operation generalizes the linear evaluation shown in Equation 5, including accounting for the edge weights. With $m = 1$ and $b = 0$, an identity substitution will be performed, and hence the algorithm can be used for partial evaluation, where some variables are assigned constants, while others are unchanged. With $m = -1$ and $b = 1$, we replace the variable by its Boolean complement.

The algorithm is shown as having functional arguments μ and β . When applied to a variable x , these “assignments” yield the constant factors to be used in the affine substitution. The algorithm follows from the linear expansion of function f with respect to each variable x . Given that $f = f_{\bar{x}} + xf_x$, substituting $mx + b$ for x yields:

$$f|_{x \leftarrow mx+b} = f_{\bar{x}} + bf_x + xm f_x$$

and hence this substitution yields a function with moments $f_{\bar{x}} + bf_x$ and $m f_x$.

The routine maintains a table of previously computed substitutions. Observe that for given assignments μ and β , recursive calls are generated from a vertex only once. The total number of calls to *AffineSubst* is therefore linear in the graph size. Of course, the resulting calls to *PlusApply* could cause the algorithm to blow up exponentially. For the special case of full evaluation, however, where $\mu(x) = 0$ for all variables x , each recursive call must return a constant function, and hence the overall complexity is linear.

7. Verification Methodology

Figure 18 illustrates schematically an approach to circuit verification originally formulated by Lai and Sastry [15] using EVBDDs. The overall goal is to prove a correspondence between

a combinational circuit, represented by a vector of Boolean functions \vec{f} , and the specification, represented by the word level function F . More precisely, assume that the circuit inputs are partitioned into vectors of binary signals $\vec{x}^1, \dots, \vec{x}^k$ (in the figure $k = 2$). For each set of signals \vec{x}^i , we are given an encoding function ENC_i describing a word level interpretation of the signals. This function will typically be a standard encoding, such as a 16-bit two's complement integer. The circuit implements a set of Boolean functions over the inputs, denoted by the vector of functions $\vec{f}(\vec{x}_1, \dots, \vec{x}_k)$. Typically this circuit is given in the form of a network of logic gates. Furthermore, we are given an encoding function ENC_o defining a word level interpretation of the output. Finally, we are given as specification a word-level function $F(X_1, \dots, X_k)$. The task of verification is then to prove the equivalence:

$$ENC_o(\vec{f}(\vec{x}_1, \dots, \vec{x}_k)) = F(ENC_1(\vec{x}_1), \dots, ENC_k(\vec{x}_k)) \quad (12)$$

That is, the circuit output, interpreted as a word should match the specification when applied to word interpretations of the circuit inputs.

*BMDs provide a suitable data structure for this form of verification, because they can represent both bit-level and word-level functions efficiently. EVBDDs can also be used for this purpose, but only for the limited class of circuit functions having efficient word-level representations as EVBDDs. By contrast, BDDs can only represent bit-level functions, and hence the specification must be expanded into bit-level form. While this can be done readily for standard functions such as binary addition, a more complex function such as binary to BCD conversion would be difficult to specify at the bit level.

7.1. Component Verification

For circuits that can be represented efficiently as *BMDs at both the bit and the word level, the test of Equation 12 can be implemented directly. As an example, consider an $n + m$ -Add-Stepper, illustrated in Figure 19 for $n = 3$ and $m = 2$. This circuit forms a basic building block for the class of multipliers we will verify. It has as inputs an $n + m$ -bit partial product input \vec{p} , split into high order elements h_{n-1}, \dots, h_0 , and low order elements l_{m-1}, \dots, l_0 . This naming convention is adopted to expedite the multiplier verification, as will be discussed shortly. The other inputs are an n -bit multiplicand x_{n-1}, \dots, x_0 , and a single bit multiplier y . It produces an $n + m + 1$ bit partial product output z_{n+m}, \dots, z_0 .

The bit-level structure for the circuit is shown in the figure, consisting of AND gates and full adders blocks. Each full adder has three inputs a , b , and c . It produces a sum output at the right hand side with function $a \oplus b \oplus c$. It has a carry output at the top, with function expressed in terms of linear operators as $a \wedge b + a \wedge c + b \wedge c - 2a \wedge b \wedge c$. From this representation we can use the algorithms *PlusApply* and *MultApply* to generate a *BMD representation of $f_i(\vec{p}, \vec{x}, y)$, the function at each output z_i for $0 \leq i \leq n + m$.

The word-level specification for an $n + m$ -Add-Stepper is simply $P + 2^m \cdot y \cdot X$, where P and X are the word-level interpretations of the partial product and multiplicand inputs. Both of these inputs are encoded as unsigned integers, as is the output. Verification therefore involves proving that the weighted sum of the bit-level output functions: $\sum_{i=0, n+m} 2^i f_i$ is equivalent to the word-level specification. As with BDDs, this process can be completely automated and works well even for more complex realizations such as carry-lookahead adders.




Figure 19: Bit-Level Representation of Add-Stepper. This circuit is a basic component of the Multiplier.

7.2. Hierarchical Verification

For larger scale circuits, representing the bit-level functionality becomes too cumbersome and hence the method described above cannot be applied directly. For example, attempting to construct the bit-level functions for a multiplier would cause exponential blow-up with *BMDs, just as it does with BDDs. Instead, we can follow a hierarchical approach in which the overall circuit is divided into components, each having a word-level specification. Verification then involves proving 1) that each component implements its word-level specification, and 2) that the composition of the word-level component functions matches the specification. This approach works well for circuits in which the components have simple word-level specifications. Such is the case for most arithmetic circuits.

Figure 20 illustrates this process for a 3-bit combinational multiplier. The bit level structure for this circuit is shown at the top. The first stage of this circuit is a Bit-multiplier (BM), containing just the AND gates of an Add-stepper. The remaining stages are Add-steppers with increasing values of m . At each stage i , input y_i serves as the multiplier bit. The justification for our hierarchical verification is shown by the progression from top to bottom in the figure. The verification of component AS 3 + 2, indicates an equivalence between the component output interpreted as a word, and its specification when applied to word interpretations of the circuit inputs (Figure 18). Thus, we can replace the final stage in the circuit by its specification, shifting the encoding operations to the component inputs (middle). Continuing, we can similarly replace the second to last stage by its specification, shifting the encoding operations to its inputs. Finally, we can replace the first stage by its specification, shifting the encoding operation to input \bar{x} (bottom). Observe that the multiplier inputs \bar{y} remain in bit-level form. In general this methodology can use word-level representations of some signals and bit-level representations of others.

As this figure illustrates, once we have verified all of the components, we can verify the overall circuit behavior by composing their word-level specifications. For the case of the multiplier this involves proving that a sequence of add steps implements multiplication. Note that in moving the encoding operations backward in the circuit, we require that the encoding function for a component input must match the output encoding of the component supplying that input.

7.3. Experimental Results

Table 5 indicates the results for verifying a number of multiplier circuits having the same structure as that of Figure 20. As can be seen, this approach remains practical for large word sizes. Our results are limited to a 62-bit word size only because our weight values are represented as 64-bit signed integers. We plan to extend our implementation to use arbitrary precision arithmetic, enabling us to go well beyond this limit.

The table also shows the time required to verify a single $n + n$ -Add-Stepper. One can see that this time grows linearly with the word size. Note also that the time to completely verify an $n \times n$ multiplier, including verifying all n Add-Steppers, is less than n times that of the final Add-Stepper. The reason for this is that much of the computation for the Add-Steppers can be reused. By the way we have named the partial product input variables \bar{p} , the bit-level




Figure 20: Hierarchical Verification of Multiplier. The bit-level representations of the circuit blocks are replaced by their word-level specifications.

Word Size	2-input gates	Mult. Time(sec)	Add-Stepper Time(sec)
4 × 4	100	0.37	0.25
8 × 8	456	2.25	0.68
12 × 12	1068	5.47	1.12
16 × 16	1936	11.53	1.68
20 × 20	3060	20.68	2.00
24 × 24	4440	25.28	2.50
28 × 28	6076	35.62	2.94
32 × 32	7968	49.17	3.22
40 × 40	12520	92.95	4.32
48 × 48	18096	152.65	5.08
56 × 56	24696	226.32	5.77
62 × 62	30318	217.17	6.87

Table 5: Verification Results for Combinational Multipliers.

outputs for the Add-Steppers hardly change. In our code, we run through the complete construction of all of the Add-Steppers, but many of the results are found in the various stored tables. Even so, the time for the multiplier verification grows slightly worse than quadratically in the word size. Given that the hardware complexity scales quadratically in the word size, this performance is reasonable, although we believe it could be further improved. We have no explanation why the verification of a 56-bit multiplier requires more time than a 62-bit one. The 56-bit result appears to be an outlier in the performance trend.

These results are especially appealing in light of prior results on multiplier verification. A brute force approach based on BDDs cannot get beyond even modest word sizes. Ochi *et al* [19] have successfully built the OBDDs for a 15-bit multiplier, requiring over 12 million vertices. Increasing the word size by one bit causes the number of vertices to increase by a factor of approximately 2.7, and hence even more powerful computers will not be able to get much beyond this point. Jain [13] verified the 16th output of circuit C6288, a 16 × 16 multiplier using a combination of BDDs, partial enumeration of the inputs, and probabilistic methods. The computation required 3-1/2 hours on a high performance workstation. Given the use of explicit enumeration, it is unlikely that this approach would scale well to larger word sizes. Burch [6] has implemented a BDD-based technique for verifying certain classes of multipliers. His method effectively creates multiple copies of the multiplier and multiplicand variables, leading to BDDs that grow cubically with the word size. This approach works for multipliers, such as ours, that form all possible product bits of the form $x_i \wedge y_j$ and then sum these bits. Burch reports verifying C6288 in 40 minutes on a Sun-3 using 12 MBytes of memory. The limiting factor in dealing with larger word sizes would be the cubic growth in memory requirement. Furthermore, this approach cannot handle multipliers that use multiplier recoding techniques, although Burch describes extensions to handle some forms of recoding.

Although we have only tried our methods on synthetically-generated multipliers based on add steps, we are confident that we can handle C6288, as well as multipliers using multiplier

recoding and other more advanced techniques.

8. Conclusions

*BMDS provide an efficient representation for functions mapping Boolean variables to numeric values. They can represent a number of word-level functions in a compact form. They also represent Boolean functions with complexity comparable to BDDs. They are therefore suitable for implementing a verification methodology in which bit-level circuits are compared to word-level specifications. By exploiting circuit hierarchy, we are able to verify circuits having functions that are intractable to represent at the bit level.

At this stage of research, there are many open problems regarding this representation. We need to characterize the behavior of *BMDs in representing Boolean functions. For all examples we have tried, their sizes are comparable to BDD representations. Either a formal relation should be established, such as has been done for EVBDDs [15], or a distinction should be proved, such as has been done for FDDs [2]. In addition, the performance of the Apply algorithms need to be characterized, indicating when they avoid exponential complexity.

Verification of multipliers and other arithmetic circuits using *BMDs seems quite promising, but these ideas must be tested and extended further. In developing a comprehensive verification system based on our hierarchical methodology, it would be good to have a "proof manager" that keeps track of what components have been verified, checks for compatibility between encodings, etc.

The hierarchical verification methodology described here extends to sequential circuits as well. For modeling such circuits, one could implement a form of symbolic simulator, where blocks of the circuit can be modeled at either the bit or the word level. For example, one could verify a sequential multiplier by first simulating a single cycle at the bit level to show it implements an add step, and then a series of cycles at the word level to show this implements multiplication.

Our method shows some promise for verifying floating point hardware, although difficult obstacles must be overcome. Using a version that supports rational numbers, we can efficiently represent the word level functions denoted by standard floating point formats. This fact follows from our ability to represent integer formats plus exponentials. Floating point hardware, however, only computes approximations of arithmetic functions. Thus, verification requires proving equivalence within some tolerance, rather than the strict equivalence of the current methodology. It is unclear whether such a test can be performed efficiently.

Many techniques developed for improving the efficiency and compactness of BDDs could be extended to *BMDs. Among these are dynamic variable reordering [20], and loosening the ordering requirement from a uniform total ordering to one in which variables may appear in different orders along different paths in the graphs [11, 22]. Our experience thus far has been that variable ordering is not as critical when representing functions at the word level as it is with bit-level representations. Nonetheless, these issues bear further investigation.

Some of the applications proposed for EVBDDs and MTBDDs may work well with *BMDs.

Among these are matrix operations and spectral transforms. Applications requiring efficient equation solving, such as integer linear programming, on the other hand, are probably not good candidates. In any case, the opportunities for further exploration seem limitless.

References

- [1] 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," *International Conference on Computer-Aided Design*, November, 1993, pp. 188-191.
- [2] B. Becker, R. Drechsler, and R. Werchner, "On the relation between BDDs and FDDs," Technical Report 12/93, University of Frankfurt, 1993.
- [3] K. S. Brace, R. L. Rudell, and R. E. Bryant, "Efficient implementation of a BDD package," *27th Design Automation Conference*, June, 1990, pp. 40-45.
- [4] R. E. Bryant, "Graph-based algorithms for Boolean function manipulation," *IEEE Transactions on Computers*, Vol. C-35, No. 8 (August, 1986), pp. 677-691.
- [5] R. E. Bryant, "On the complexity of VLSI implementations and graph representations of Boolean functions with application to integer multiplication," *IEEE Transactions on Computers*, Vol. 40, No. 2 (February, 1991), pp. 205-213.
- [6] J. R. Burch, "Using BDDs to verify multipliers," *28th Design Automation Conference*, June, 1991, pp. 408-412.
- [7] E. Clarke, K.L. McMillan, X. Zhao, M. Fujita, and J. C.-Y. Yang, "Spectral transforms for large Boolean functions with application to technology mapping," *30th ACM/IEEE Design Automation Conference*, Dallas, TX, June, 1993, pp. 54-60.
- [8] E. Clarke, M. Fujita, P. C. McGeer, K.L. McMillan, and J. C.-Y. Yang, "Multi-terminal binary decision diagrams: an efficient data structure for matrix representation," unpublished, 1993.
- [9] R. Drechsler, A. Sarabi, M. Theobald, B. Becker, and M. Perkowski, "Efficient representation and manipulation of switching functions based on ordered Kronecker function decision diagrams," *31st Design Automation Conference*, June, 1994, pp. 415-419.
- [10] M. R. Garey, and D. S. Johnson, *Computers and Intractability*, W. H. Freeman and Company, 1979.
- [11] J. Gergov, and C. Meinel, "Efficient analysis and manipulation of OBDDs can be extended to read-once-only branching programs," To appear in *IEEE Transactions on Computers*, 1994.
- [12] P. L. Hammer (Ivănescu), and S. Rudeanu, *Boolean Methods in Operations Research*, Springer-Verlag, 1968.

- [13] J. Jain, J. Bitner, J. A. Abraham, and D. S. Fussell, "Functional partitioning for verification and related problems," *Advanced Research in VLSI and Parallel Systems: Proceedings of the 1992 Brown/MIT Conference*, T. Knight, and J. Savage, eds., 1992, pp. 210-226.
- [14] U. Kebschull, E. Schubert, and W. Rosentiel, "Multilevel logic based on functional decision diagrams," *European Design Automation Conference*, 1992, pp. 43-47.
- [15] Y.-T. Lai, and S. Sastry, "Edge-valued binary decision diagrams for multi-level hierarchical verification," *29th Design Automation Conference*, June, 1992, pp. 608-613.
- [16] Y.-T. Lai, M. Pedram, and S. B. K. Vrudhula, "Edge-valued binary-decision diagrams," unpublished, 1993.
- [17] Y.-T. Lai, M. Pedram, and S. B. K. Vrudhula, "FGILP: An integer linear program solver based on function graphs," *International Conference on Computer-Aided Design*, November, 1993, pp. 685-689.
- [18] R. M. M. Oberman, *Digital Circuits for Binary Arithmetic*, John Wiley and Sons, 1979.
- [19] H. Ochi, K. Yasuoka, and S. Yajima, "Breadth-first manipulation of very large binary-decision diagrams," *International Conference on Computer-Aided Design*, November, 1993, pp. 48-55.
- [20] R. Rudell, "Dynamic variable ordering for ordered binary decision diagrams," *International Conference on Computer-Aided Design*, November, 1993, pp. 42-47.
- [21] F. F. Sellers, M. Y. Hsiao, and C. L. Bearnson, "Analyzing errors with the Boolean difference," *IEEE Transactions on Computers*, Vol. C-17 (1968), pp. 676-683.
- [22] D. Sieling, and I. Wegner, "Graph driven BDDs—a new data structure for Boolean functions," To appear in *Theoretical Computer Science*, 1994.
- [23] R. D. Silverman, "Massively distributed computing and factoring large integers," *Communications of the ACM*, Vol. 34, No. 11 (November, 1991), pp. 95-103.

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213-3890

Carnegie Mellon University does not discriminate and Carnegie Mellon University is required not to discriminate in admission, employment or administration of its programs on the basis of race, color, national origin, sex or handicap in violation of Title VI of the Civil Rights Act of 1964, Title IX of the Educational Amendments of 1972 and Section 504 of the Rehabilitation Act of 1973 or other federal, state or local laws, or executive orders.

In addition, Carnegie Mellon University does not discriminate in admission, employment or administration of its programs on the basis of religion, creed, ancestry, belief, age, veteran status, sexual orientation or in violation of federal, state or local laws, or executive orders.

Inquiries concerning application of these statements should be directed to the Provost, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, telephone (412) 668-6841 or to the President for Enrollment, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, telephone (412) 268-2056.