Advanced Research in VLSI Caltech 89

Using If-then-else DAGs for Multi-Level Logic Minimization

Kevin Karplus

Board of Studies in Computer Engineering
University of California, Santa Cruz
Santa Cruz, CA 95064

This article describes the use of if-then-else DAG*s for multi-level logic minimization.*

A new canonical form for if-then-else DAG*s, analogous to Bryant's canonical form for binary decision diagrams (*BDD*s), is introduced. The definitions of prime and irredundant expressions are extended to if-then-else* DAG*s. Expressions in Bryant's canonical form or in the new canonical form can be shown to be prime and irredundant.*

Objective functions for minimization are discussed, and estimators for predicting the area and delay of the circuit produced after technology mapping are proposed. A brief discussion of methods for applying don't-care information and for factoring expressions is included.

1 What is multi-level logic minimization?

Multi-level logic minimization is the transformation of a specification of a Boolean function into an equivalent representation that can be implemented as a circuit with better characteristics (smaller, faster, or more testable) than a circuit built from the original specification. The function usually has multiple outputs, and may be only partially specified.

Most previous work in multi-level logic synthesis is based on extensions of two-level (sum-of-products) minimization for PLAs [7, 6, 2, 14, 3]. A notable example is the misII multi-level minimization system [9], based on the espresso two-level minimizer [8].

Some subproblems of multi-level minimization may be easier in representations other than sum-of-products. For example, tautology checking, finding common subexpressions, and extracting XOR operations look more attractive in the if-then-else DAG form (see Section 2) than in sum-of-products form. We are investigating if-then-else DAGs

for multi-level logic minimization, and have some encouraging preliminary results.

Section 2 describes the if-then-else operator, which forms the basis for the representation used, and gives a quick introduction to binary decision diagrams and if-then-else DAGs. Section 3 introduces a new canonical form. Section 4 discusses ways to estimate the area and delay of a circuit in a technology-independent logic minimizer. Section 5 discusses conversion from networks of gates to if-then-else DAGs. Section 6 talks about ways to use don't-care information for simplifying if-then-else DAGs. Section 7 describes some crude factoring techniques that do surprisingly well.

2 Binary decision diagrams and if-then-else DAGs

The research described here is based on a single universal operator—the if-then-else operator.

Definition 1: *The* if-then-else *operator is a ternary Boolean function, with* (**if** a **then** b **else** c) *defined as* $ab + a'c$ *or, equivalently,* $(a + c')(a' + b)$.

All binary Boolean functions are easily defined with the if-then-else operator. For example,

- $ab =$ (**if** a **then** b **else** FALSE)
- $a + b =$ (**if** a **then** TRUE **else** b)
- $a \oplus b =$ (**if** a **then** b' **else** b).

If-then-else trees and DAGs have a long history [17, 1, 10]. We divide if-then-else representations into two classes: *binary decision diagrams*, in which the **if**-part is always a simple variable, and *if-then-else* DAG*s*, which may have arbitrary expressions in the **if**-part.

Definition 2: *A* binary decision diagram *is a binary directed acyclic graph with two leaves* TRUE *and* FALSE, *in which each non-leaf node is labeled with an atom and has two out-edges pointing to the* **then**-*part and the* **else**-*part. The meaning of a binary decision diagram is defined recursively as* (**if** *label(node)* **then** *meaning(***then**-*part)* **else** *meaning(***else**-*part)*).

Binary decision diagrams (BDDs) have been used often for logic verification work [12, 20, 21, 18]. They are attractive for such work as they are easy to manipulate and have a convenient canonical form (Bryant's canonical form) [10]. They have also been used

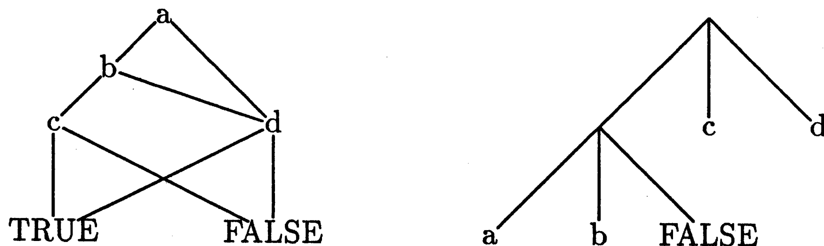

Figure 2.1: BDD and if-then-else DAG for $abc + a'd + b'd$, factored as (**if** ab **then** c **else** d).

for logic synthesis work, mainly for designing differential voltage-cascode switches, which implement BDDs directly [19, 13]. For logic minimization in other technologies, the mismatch between the BDD structure and the circuit structure has restricted their use.

If-then-else DAGs generalize binary decision diagrams by not restricting the **if**-parts to single variables:

Definition 3: *An* if-then-else DAG *is a ternary directed acyclic graph in which each leaf is labeled with* TRUE, FALSE *or a literal, and each internal node has three out-edges pointing to the* **if**-, **then**-, *and* **else**-*parts. The meaning of a leaf node is the label on the node, and the meaning of an internal node is defined recursively as* (**if** meaning(**if**-part) **then** meaning(**then**-part) **else** meaning(**else**-part)).

Negating an expression represented as an if-then-else DAG requires negating all the leaves. To simplify the negation operation, and to reduce the storage needed for representing expressions, we allow negation of an if-then-else DAG to be represented by flipping one flag bit, which we keep in the low-order bit of the pointer to the DAG.

There is a natural mapping from if-then-else DAGs to BDDs, which can be defined recursively. The **if**-, **then**- and **else**-parts are each mapped to the corresponding BDDs, and the common subBDDs of the **then**- and **else**-parts are merged. All pointers to TRUE in the image of the **if**-part are changed to pointers to the image of the **then**-part, and all pointers to FALSE are changed to pointers to the image of the **else**-part. The leaves of the if-then-else DAG are mapped to the obvious corresponding BDDs. Note that the mapping is many-to-one, with different if-then-else DAGs corresponding to different two-cuts in the BDD [15].

Figure 2.1 shows a BDD and an if-then-else DAG for the expression $abc+a'd+b'd$. Note that the if-then-else DAG represents the expression as (**if** ab **then** c **else** d), allowing the ab term to be shared with other expressions.

If-then-else DAGs offer several advantages over sum-of-products and Boolean decision diagram representations.

- If-then-else DAGs can be used to represent BDDs and sum-of-products expressions, but neither BDDs nor sum-of-products forms can represent if-then-else DAGs.
- Every 1- or 2-input gate can be represented as an if-then-else triple, so every acyclic network of gates can be represented by replacing each gate by the appropriate if-then-else triple. This means that circuits built out of arbitrary gates can be converted to if-then-else DAGs without losing any sharing of common subexpressions.
- If-then-else DAGs have two convenient canonical forms: Bryant's canonical form (as in BDD's) and a new canonical form presented in Section 3. Canonical forms are particularly valuable for tautology checking.
- The new canonical form for if-then-else DAGs allows more sharing of subexpressions than Bryant's canonical form for BDDs. Any shared subexpressions in Bryant's form have corresponding sharing in the new canonical form, but the new form allows more sharing in the **if**-part.

  For example, $ab(d+e)$ is (**if** (**if** a **then** b **else** FALSE) **then** (**if** d **then** TRUE **else** e) **else** FALSE), $c(d+e)$ is (**if** c **then** (**if** d **then** TRUE **else** e) **else** FALSE), and abd is (**if** (**if** a **then** b **else** FALSE) **then** d **else** FALSE). These three functions share the subexpressions ab =(**if** a **then** b **else** FALSE) and $(d+e)$ =(**if** d **then** TRUE **else** e), while the BDD representations would share only $(d+e)$.
- If-then-else DAGs are a more factored form than BDDs, providing for better printing and logic minimization. With the aid of the transformations described in Section 7, good factorings can often be found from the canonical forms or from arbitrary non-canonical expressions.
- Boolean operations can be computed as if-then-else triples, so that the symbol table used for storing canonical forms can be used for caching the results of operations as well.

3 A canonical form for if-then-else DAGs

A representation is *canonical* if any two expressions that are logically equivalent are identical. For example, if $ab + ab'$ is represented differently from a, then the representation is non-canonical.

Using canonical forms makes checking for equivalence easy—unfortunately, conversion from a non-canonical form to canonical form

may take a lot of time or memory. Because equivalence checking in canonical form is fast, but equivalence checking in a non-canonical form (such as clause form) is equivalent to the NP-complete problem SATISFIABILITY, we are essentially guaranteed that the conversion to any canonical form is exponential in the worst case. For most commonly used Boolean functions, however, a well-chosen canonical form can be small and easy to manipulate, and exponential blow-up is rare.

A recent paper by Randy Bryant shows that one common function, integer multiplication, requires exponentially many nodes to represent in his canonical form, no matter what ordering is used for the variables [11]. The same arguments can be applied to the new canonical form described here. Small if-then-else DAGs for integer multiplication are easy to construct from small circuits, but they all involve duplicating variables, and so are not canonical.

To make if-then-else DAGs canonical, we must place some restrictions on the expressions allowed in the **if**-, **then**-, and **else**-parts of the structure. Of the seven restrictions, the first three are slightly modified versions of the corresponding restrictions in Bryant's canonical form.

1. A total ordering is imposed on the atoms of the expression, and all the atoms in the **if**-part must be earlier in the order than all atoms in the **then**- and **else**-parts. A weaker restriction, that the variables of the **if**-part be disjoint from those of the **then**- and **else**-parts, would be enough to eliminate paths with duplicate variables in the corresponding BDD, but not enough to make the if-then-else DAG canonical. Non-canonical expressions using this weaker version of the restriction are useful for factoring.
2. The **then**- and **else**-parts of an expression must be distinct Boolean functions—exactly as in Bryant's canonical form.
3. A systematic choice must be made between the equivalent expressions (**if** a **then** b **else** c) and (**if** a' **then** c **else** b) and between (**if** a **then** b **else** c) and (**if** a **then** b' **else** c')$'$. We require that **if**- and **then**-parts of an expression be pure pointers, with negation allowed only for the **else**-part or the entire expression. This corresponds to Bryant's choice of atoms as node labels (never negations of atoms).
4. Triples of the form (**if** a **then** TRUE **else** FALSE) and (**if** a **then** FALSE **else** TRUE) are prohibited. The first triple should be represented simply as a, and the second one by a'.
5. Triples of the form (**if** TRUE **then** b **else** c) and (**if** FALSE **then** b **else** c) are prohibited, and should be replaced with b and c respectively.

6. In the triple (**if** a **then** b **else** c), b and c must not share both **then**- and **else**-parts. If b = (**if** b_a **then** b_b **else** c_c) and c = (**if** c_a **then** b_b **else** c_c), then the correct representation is (**if** (**if** a **then** b_a **else** c_a) **then** b_b **else** c_c). If b = (**if** b_a **then** b_b **else** b_c) and c = (**if** c_a **then** b_c **else** b_b), then use (**if** (**if** a **then** b_a **else** c'_a) **then** b_b **else** b_c).
7. In the triple (**if** a **then** b **else** c), b must not contain c as a **then**- or **else**-part. If b = (**if** b_1 **then** b_b **else** c) or b = (**if** b_2 **then** c **else** b_c), then the expression should be represented as (**if** (**if** a **then** b_1 **else** FALSE) **then** b_b **else** c) or (**if** (**if** a **then** b_2 **else** TRUE) **then** c **else** b_c). If c is one of the constants TRUE or FALSE, this restriction amounts to choosing left-associativity for commutative AND or OR operations. The symmetric test for c = (**if** c_1 **then** c_b **else** b) or c = (**if** c_2 **then** b **else** c_c) is also needed.

We can show that imposing the restrictions listed above defines a canonical form by exhibiting an isomorphism with Bryant's canonical form [15]. We can use essentially the same algorithm for converting to either Bryant's canonical form or the new form [15].

3.1 Two-cut canonical forms are prime and irredundant

Other researchers in multi-level minimization, working primarily with sum-of-products representations, have found the concepts of *primality* and *irredundancy* to be important, particularly for producing testable circuits [8, page 28], [7, page 202]. Both concepts have natural analogs in if-then-else DAG representations.

In sum-of-products form, an expression is said to be *prime* if no term can be modified by changing a literal to TRUE without changing the meaning of the expression. Similarly, an expression in sum-of-products form is *irredundant* if no term can be changed to FALSE without changing the meaning of the expression.

Definition 4: *An if-then-else* DAG *is* prime *if no pointer to a literal, sub*DAG*, or the constant* FALSE *could be replaced with a pointer to* TRUE *without changing the meaning of the expression. An if-then-else* DAG *is* irredundant *if no pointer to a literal, sub*DAG*, or the constant* TRUE *could be replaced with a pointer to* FALSE *without changing the meaning of the expression.*

The new definitions of "prime" and "irredundant" correspond to existing ones for sum-of-products and factored forms. For example, we can use an if-then-else tree (so that no sharing is done between terms)

to represent a sum-of-products expression by replacing each binary AND or OR operator by the corresponding if-then-else triple. If the if-then-else tree is prime or irredundant, then the sum-of-products expression must be, because the substitutions to be tested in the sum-of-products form are a subset of those tested in the if-then-else tree. The extra tests for primality and irredundancy in the if-then-else tree are easily satisfied for trees corresponding to sum-of-products representations [15].

Both Bryant's canonical form and the new canonical form presented in Section 3 can be shown to be prime and irredundant with the definition presented here [15].

4 Expression complexity, counting literals

When doing logic minimization, the first question is "what exactly is being minimized?" The goal of minimization is to reduce the area, power, or delay of the final circuit after technology mapping, but these costs are dependent on the technology used, and optimizations tied to a particular cell library quickly become obsolete.

For technology-independent minimization to work, we need measures that are not dependent on any particular cell library (or that are parameterized and easily tuned for different technologies), and that roughly approximate the cost or speed obtained by a technology mapper. Technology-independent *delay* estimates are hard to come up with, and so most research has concentrated on *size* minimization, leaving the delay minimization to the technology mapper.

The usual way to estimate the area for a network of gates is to estimate the cost for each gate and sum the estimates. The most popular gate area estimators are the number of literals in sum-of-products form and the number of literals in the factored form [7, page 235]. The literal count corresponds closely to the number of transistor pairs needed to implement the function as a static cMOS gate, and is an excellent area estimator if the mapper does not change the decomposition of the circuit. The estimate is not as good when the mapper splits or merges gates.

We have made some attempts to calibrate area estimators for misII's technology mapper on a collection of different designs, including the MCNC benchmarks. The mapper we attempted to calibrate was the command `map -m1; phase -g`. We looked at several different measures, including the ones reported by misII (number of nodes, sum of literals in sum-of-products form, sum of literals in factored form). The sum of literals in factored form is a good predictor, with the

ratio of actual area over predicted area having a standard deviation of 17.4% of the mean.

Many technology mappers do polarity assignment, adding or removing inverters to minimize delay or area. For such mappers, adding inverters in the input description usually does not increase the cost of the final solution, and so should have zero cost for the technology-independent minimizer. The standard cost functions do not have this property, and a cost estimator that hides such inverters should be a better estimator of final area. Subtracting the number of nodes in a network from the sum of literals makes inverters free, and has the added advantage of making the measure less sensitive to the size of the gates used in the decomposition. This measure is an excellent predictor, having a standard deviation of only 12.6% of the mean.

The standard measures described above are useful when a network has been decomposed into gates, but are not directly applicable to a network described as an if-then-else DAG with multiple roots. New measures are needed.

We have experimented with several estimators, including the following:

triples the number of if-then-else triples in the DAG,

size the number of triples plus the number of distinct variables,

opcount the number of n-input AND, n-input OR, and 2-input XOR gates produced by our decomposition algorithm,

height the longest path from a root to a leaf,

pcount the number of literals needed if each if-then-else triple were expanded to AND, OR, or XOR gates, and any shared nodes were duplicated.

count a recursively defined function that attempts to match the values of the estimator (literals(factored)−nodes). **count** is

0 for the constants TRUE and FALSE.

1 for literals.

1 for a subDAG that has been previously counted.

$count(x)+count(y)+count(z)$ for (**if** x **then** y **else** z), if the triple represents a 2-input AND or OR, that is, if y or z is a constant.

$count(x) + count(y) + count(z) + 1$ for other triples (**if** x **then** y **else** z).

Of these new functions, *count* is the best predictor of area for our benchmarks, with a standard deviation of 13.1%.

Delay estimation may be harder than area estimation. Our best predictor so far is the height of the if-then-else DAG, with a standard deviation of 31.4%. We may be able to estimate delay better by adding a penalty to nodes that are used repeatedly, and by using smaller costs for triples that have a constant **then**- or **else**-parts. An active area of our research is to find good estimators of both area and delay for popular mapper-library pairs.

5 Coverting from BLIF (sum-of-products) format

Most of the standard benchmarks are available in a standard format, the Berkeley Logic Intermediate Format (BLIF) [4], and so we need to convert BLIF files into if-then-else DAGs. In BLIF, combinational logic is described as a directed, acyclic network of gates, and each gate is described in sum-of-products form.

Building an if-then-else DAG from a network of gates is easy if each gate is described as an if-then-else DAG—the only tricky part is converting the sum-of-products descriptions of the gates into if-then-else DAGs. We have several choices:

- Build a canonical DAG for the function expressed by the gate. For gates of the form $ab+cd+ef+gh+\cdots$, the wrong variable ordering can cause an exponential blowup in size.
- Preserve the original and-or structure of the sum-of-products expression. This is guaranteed not to be too big, but offers few advantages over simply using sum-of-product representations.
- Build a partially factored expression for the gate.

  We use a recursive function to get an if-then-else DAG E for a set of terms T. The terms are sorted, grouping together those that don't use the first input variable (T_d), those that use v_1' (T_0), and those that use v_1 (T_1). We then strip the first variable off the terms in each group, and apply the routine recursively to get expressions E_d, E_0, and E_1. We build the expression E as (**if** E_d **then** TRUE **else** (**if** v **then** E_1 **else** E_0)). This idea can be improved by sorting the variables with the most frequently used ones first.

  This algorithm is essentially the same as the popular method of factoring out one-literal cubes, and produces expressions that are often significantly smaller than either the canonical form or the straight sum-of-products form.

After building an expression for a gate, we can try factoring the gate with the `Printform` or `LocalFactor` transformations, or we can try