

Anti-Symmetry and Logic Simulation

Peter M. Maurer
Dept. of Computer Science
Baylor University
Waco, Texas 76798-7356

Abstract – Like ordinary symmetries, anti-symmetries are defined by relations between function cofactors. For ordinary symmetries, two cofactors must be equal, for anti-symmetries two cofactors must be complements of one another. This paper shows that anti-symmetries can be used to improve simulation performance in the same manner as ordinary symmetries. Detailed detection, clustering and simulation algorithms are given along with a set of experimental results to demonstrate the effectiveness of the algorithms. These results show that anti-symmetries can be just as effective as ordinary symmetries in enhancing simulation performance. In fact, in some cases, anti-symmetries give better performance than ordinary symmetries.

1 Introduction

Detecting function symmetries has proven to be useful in many areas of electronic design automation [1-7]. Many symmetry detection algorithms have been created, and permutations are the basis of virtually all of these algorithms [8-10].

The most important types of symmetry are *total symmetry* and *partial symmetry*. The inputs of a totally symmetric function can be rearranged arbitrarily, while a partially symmetric function has subsets of inputs that can be rearranged arbitrarily. Examples of totally symmetric and partially symmetric functions are $abcd$ and $abc + d$ (where $+$ represents OR and multiplication represents AND). There are many other types of permutation-based symmetries (see [9, 11]) which are often lumped together and called *weak symmetries*. But in this paper we will be concerned only with total and partial symmetry.

Total and partial symmetries are common in the circuits we encounter in practice, and can be detected by examining pairs of variables. Two variables constitute a *symmetric variable pair* if they can be exchanged without altering the output of the function. Symmetric variable pairs are transitive. Thus, if (a,b) and (b,c) are symmetric variable pairs, then so is (a,c) . A function is totally symmetric if and only if every pair of input variables is a symmetric variable pair. A function is partially symmetric in the variables x_1, \dots, x_k if (x_i, x_j) is a symmetric variable pair for all i and j , $1 \leq i < j \leq k$.

The symmetric variable pairs of a function f can be detected using the cofactors of f . If f is an n -input Boolean function with input variables x_1, \dots, x_n . The cofactors of f with respect to x_1 are the functions f_{0x_1} and f_{1x_1} , which are computed by setting the variable x_1 to 0 and then to 1. The exact procedure for computing a cofactor depends on the representation of the function. Cofactors can be computed with respect to a single variable or with respect to a set of variables. When there is no opportunity for confusion, we omit the x 's and simply place the 1's and 0's in the subscripts. In symmetry detection, it is common to compute cofactors with respect to pairs of variables. There are four such cofactors f_{00} , f_{01} , f_{10} , and f_{11} .

Different relations between these cofactors can be used to define different types of symmetry [12, 13].

One of the latest developments in symmetry detection is matrix-based symmetry [7]. All permutations can be specified as non-singular matrices, but not all non-singular matrices can be specified as permutations. Thus matrix-based symmetry is an extension of permutation-based symmetry. Conjugate symmetry is one form of matrix-based symmetry, some forms of which can be detected using cofactor-relations.

The algorithms discussed in this paper are based on cofactor relations, the two most important of which are the classical relations and the anti-relations [14].

2 Classical Symmetry

The classical relations are defined in terms of the two-variable cofactors of a function, f_{00} , f_{01} , f_{10} , and f_{11} . There are six possible relations, each of which represents a certain type of symmetry. These relations are given in Figure 1, along with their respective symmetry types.

In a sense, the only relation that truly represents symmetry is $f_{01} = f_{10}$, ordinary symmetry. The other relations represent variable pairs that are *not* symmetric, but can be "corrected" to become symmetric. Our algorithm tests for all six relations to detect symmetric variable pairs. When a symmetric variable pair is detected, it is "corrected," if necessary, and combined into a single clustered variable. In fact, we test only for ordinary symmetry. The other five relations are detected by transforming the state space of the function, and then testing for ordinary symmetry. (See Section 4 for details on the state space transformations.) Our classical symmetry algorithms are described in [7], but for completeness, we repeat some of the details here.

Relation	Symmetry Type
$f_{01} = f_{10}$	Ordinary
$f_{00} = f_{11}$	Multi-Phase
$f_{01} = f_{11}$	Single-Variable A
$f_{10} = f_{11}$	Single-Variable B
$f_{10} = f_{00}$	Multi-Phase Single-Variable A
$f_{01} = f_{00}$	Multi-Phase Single-Variable B

Figure 1. Cofactor Relations.

The multi-phase relation, $f_{00} = f_{11}$, indicates that the function is symmetric, but one variable is inverted with respect to the other. It is possible to treat the multi-phase relation as ordinary symmetry after performing a state-space transformation and adding a NOT gate to one of the inputs.

The single-variable relations represent two types of conjugate symmetry. Not all conjugate symmetries manifest themselves as single-variable symmetries, but the mechanisms used to detect single-variable symmetries can be extended to detect most conjugate symmetries. As with multi-phase symmetries, it is possible to treat

conjugate symmetries as if they were ordinary symmetries using a state-space transformation and a collection of XOR gates on the function inputs. The XOR gates compute a matrix transformation of the inputs prior to passing the inputs into the function. The input transformation is similar to the first layer of logic in some forms of three-level minimization [16, 17].

The multi-phase single-variable relations represent the combination of conjugate symmetry and multi-phase symmetry. These types of symmetry can be handled by combining the techniques for multi-phase and conjugate symmetry.

The result of symmetry detection is a multi-dimensional state machine which represents the state of a Boolean function. Each dimension of the state machine represents a cluster of symmetric variables. It is convenient to think of the multi-dimensional machine as an extended type of hypercube with several states along each dimension. For a simple, non-clustered input variable, the dimension will have two states representing input values of zero and one. For a cluster of n variables, the dimension will have $n+1$ states with the state representing the number of one-inputs in the cluster of n variables. Figure 2 illustrates a function with a simple variable A, and a clustered variable containing three simple variables, B, C, and D.

Another way to view the n -dimensions of a gate state-machine is as a collection of n input-ports. For ordinary and multi-phase symmetries, there is a one-to-one mapping between input ports and clusters of function inputs. For conjugate symmetry, an event on an input can generate events on many different ports. This technique is used to compute the XOR functions on the inputs. We will explain this further in Section 4.

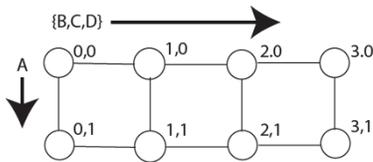


Figure 2. A Gate State-Machine.

3 Anti-Symmetry

Anti-symmetry is known by several other names, including *skew symmetry* and *negative symmetry*. Anti-symmetry is based on the observation that relations of the form $f_{01} = f_{10}$ can be written $f_{01} \oplus f_{10} = \bar{0}$, where \oplus represents the XOR operation and $\bar{0}$ represents the constant zero function. If we reformulate our six relations and replace the constant zero with the constant one function, $\bar{1}$, we obtain the six anti-symmetry relations given in Figure 3. Note that if $f_{01} \oplus f_{10} = \bar{1}$, then f_{01} and f_{10} are inverses of one another.

As with multi-phase and conjugate symmetries, anti-symmetries can be transformed into ordinary symmetries using state-space transformations. These transformations are easier to visualize if we place the four cofactors f_{00} , f_{01} , f_{10} and f_{11} into a hyper-linear structure as shown in Figure 4. To convert the anti-symmetry into an ordinary symmetry, we invert one of the grayed cofactors.

There are several different state-space transformations that will transform an anti-symmetry into an ordinary symmetry, each one of which requires a different corrective action in the final function. The naïve transformation shown in Figure 4 results in a complex corrective action. Let us suppose that an ordinary anti-symmetry is found between the variables are a and b and that f_{01} has been complemented to transform the symmetry into a classical symmetry. This means that when the transformed function is evaluated, the output will be inverted whenever $a=0$ and $b=1$. The corrective

function shown in Figure 5 is applied during the simulation process to produce the correct function output.

Relation	Anti-Symmetry Type
$f_{01} \oplus f_{10} = \bar{1}$	Ordinary
$f_{00} \oplus f_{11} = \bar{1}$	Multi-Phase
$f_{01} \oplus f_{11} = \bar{1}$	Single-Variable A
$f_{10} \oplus f_{11} = \bar{1}$	Single-Variable B
$f_{10} \oplus f_{00} = \bar{1}$	Multi-Phase Single-Variable A
$f_{01} \oplus f_{00} = \bar{1}$	Multi-Phase Single-Variable B

Figure 3. The Anti-Symmetry Relations.

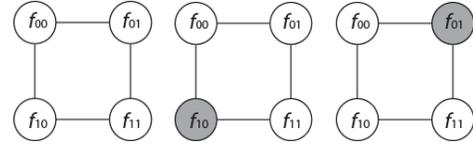


Figure 4. Naïve Corrective Actions.

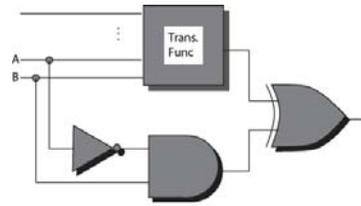


Figure 5. A Naïve Corrective Function.

To simplify the corrective procedure, we use an over-kill method when transforming the function. Instead of just complementing f_{01} (or f_{10}) we also complement f_{11} , as shown in Figure 6. This means that the output of the transformed function is inverted whenever $b=1$ (or $a=1$). This eliminates the AND and NOT gates, as shown in Figure 7. Regardless of how many anti-symmetric variable pairs are detected for a function, only a single XOR gate is required on the output. This XOR gate must have one input for each detected anti-symmetric pair. What is more, in Section 5 we will show how to get the XOR function for free.

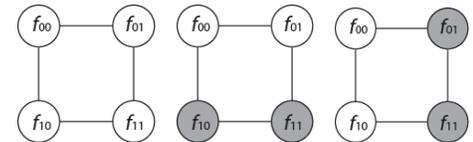


Figure 6. Sophisticated Corrective Actions.

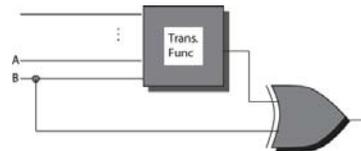


Figure 7. A Simple Corrective Function.

4 Symmetry Detection

There are several problems in determining the inputs of the corrective XOR gate. When combining anti-symmetry with conjugate symmetry, adding inputs to the correcting XOR becomes more complicated. We need to determine what should happen when a variable is added twice to the correcting XOR and we need to

determine how to detect anti-symmetry with respect to clustered variables.

In hyperlinear structures ordinary symmetries can be detected by examining cofactors along the anti-diagonals. For there to be an ordinary symmetry in Figure 2, node (1,0) must equal (0,1), (2,0) must equal (1,1) and (3,0) must equal (2,1). If there are more than two dimensions, the diagonal tests must be repeated for each of the planes containing the two variables. There is no required relationship between separate diagonals or between separate planes.

Multi-phase symmetry can be detected by reversing the structure along one dimension and then testing for ordinary symmetry. Conjugate symmetry can be detected by reversing the odd numbered rows or the odd numbered columns, and then testing for ordinary symmetry. Combined multi-phase and conjugate symmetry is detected by reversing the even numbered rows and columns. The reversals can be done without altering the structure by indexing rows, columns, or the entire dimension in reverse order.

Consider the left-most state machine of Figure 8. This state machine represents a 4-input function with two clustered variables. Assume that the inputs to the function are a, b, c, and d. and have been clustered into two pairs (a,b) and (c,d). The horizontal dimension of the state machine represents the state of the pair (a,b), while the vertical dimension represents the state of the pair (c,d).

To detect symmetries between the clustered variables (a,b) and (c,d) it is necessary to examine the reverse diagonals. If the states with the same letter (L, F, and G) contain the same function, then the two clustered variables (a,b) and (c,d) are symmetric with one another.

When an anti-symmetry exists between any two variables in two different clustered pairs, then an anti-symmetry exists between every pair of variables in the two of clustered variables. This implies that a function must alternate with its complement along each back-diagonal. This condition is shown in the middle state-machine of Figure 8. (See [15] for more detail.)

To convert the anti-symmetry into an ordinary symmetry, we invert the functions in the odd-numbered rows or the odd-numbered columns, as shown in the third state machine of Figure 8, in which the center column is inverted. This column is where the XOR of the individual variables in the clustered variable takes the value 1. Thus to correct the inverted column of Figure 8, we must add the variables *a* and *b* to the corrective XOR gate as shown in Figure 9.

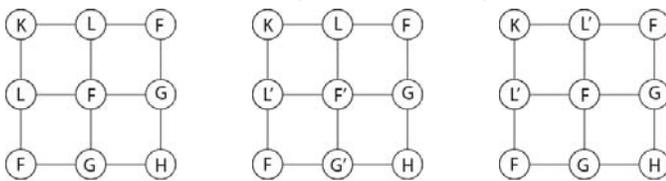


Figure 8. Clustered-Variable State Machine.

When the first anti-symmetry is detected, we add the XOR gate to the output of the function. We also create a list of the input variables that must be added to the inputs of the XOR gate. When new anti-symmetries are detected, new input variables are added to the list.

It is sometimes necessary to add the same input to the list twice. Since the two inputs are identical, the pair will have either the value (0,0) or the value (1,1). In both cases, the XOR of the two values is 0, which will not change the value computed from the other inputs, so both inputs can be removed. Thus, when we add an input to the list, we check to see whether it is currently on the list. If so, then we remove it instead of adding it.

When anti-symmetry is combined with conjugate symmetry we have an additional problem. The function has *n* inputs and the state-machine has *n* input ports, but with conjugate symmetry, the mapping

between inputs and ports is not one-to-one. Several function inputs can be directed into a single port, and a single function input can be directed into several ports. Since symmetry detection is done with respect to ports, and correction is done with respect to function inputs, it is necessary to maintain a mapping between the two. The symmetry detection algorithm maintains a $n \times n$ matrix which shows the input-to-port mapping. When correction must be done with respect to a port, every function input directed into that port is added to (or removed from) the list of XOR inputs.

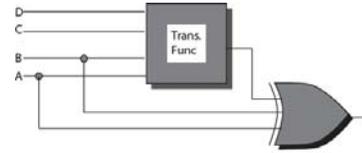


Figure 11. A Clustered Corrective Function.

To avoid conflicts with multi-phase and conjugate symmetry we take note of the reversing operations used during the detection process. To detect multi-phase symmetry it is necessary to reverse either all rows or all columns, but not both. If the rows have been reversed, then we transform the function by inverting the odd rows. If the columns have been reversed then we invert the odd columns. By placing the reversal and the inversion along the same dimension, we invert the same vertices that would have been inverted for an ordinary symmetry.

Conjugate symmetry is handled the same way. If the odd rows are reversed, we invert the odd rows. If the odd columns are reversed, we invert the odd columns. The algorithm for symmetry detection is given in section A1, Figure A1.

The algorithm for detecting symmetric variable pairs is virtually identical for ordinary symmetry and anti-symmetry. State-space transformations are used to combine ordinary and anti-symmetry with conjugate and multiphase symmetry. The basic algorithm is given in Section A1, Figure A2. This algorithm is executed twice, once for anti-symmetry and once for ordinary symmetry.

Checking the back diagonals is done by selecting each diagonal, and obtaining a comparator function from the first element of the diagonal. The comparator function is compared to the function contained in the other vertices along the diagonal. For anti-symmetry, two comparator functions are obtained. The first is taken from the first vertex of the diagonal and the second is obtained by inverting the first function. These functions are compared in alternating fashion along the diagonal. The two algorithms are given in Section A1, Figures A3 and A4.

5 Simulation Code

Detecting the various types of symmetry is beneficial because our simulator can operate faster with symmetric functions. It is possible for the corrective functions to negate the benefit of detecting symmetries, but as we will show in this section, we can either eliminate, or greatly reduce the cost of these functions.

Our simulator is a compiled simulator. Simulation detection is done during the code generation phase so that the cost of symmetry detection can be amortized over many simulations. When we report simulation times, these times do not include the cost of symmetry detection. Symmetry detection is extremely fast and takes less than a second for all but circuit c3540, which required 4 seconds.

At run time, each function is implemented as an *n*-dimensional state machine, with the current state represented as an *n*-element vector. A simple input can either increment or decrement a single index by 1. If $A = \{a_1, \dots, a_k\}$ is a clustered variable, then each of the simple inputs, a_1, \dots, a_k , operates on the same index. A separate index is used for each dimension.

The operations performed with respect to simple variables alternate with one another. If the current action increments the index, then the next action decrements it, and vice versa. It is not necessary to maintain the value of the input as long as we know which operation to perform next. The NOT gates required by multi-phase symmetry can be eliminated as long as the increment/decrement state of the input is initialized properly. Thus we get multi-phase symmetry for free.

Suppose we have two successive events for the same input. The first event will increment (or decrement) an index, and the second event will decrement (or increment) the same index, leaving the state unchanged. In effect we have computed the “exclusive or” of the two events. Since this exclusive-or is inherent in the state machine, we can get the XOR gates required by conjugate symmetry almost for free.

The state of each input port is recorded as 1 or -1 , depending on whether the next operation is an increment or decrement. The data structure representing the event has an array of pointers to port states and port indices. When an event is executed, the port state is added to the port index and is then negated. The indices are used to determine whether the output of the function has changed. If the output changes then an event is queued. If an event is already queued for the output, the queued event is cancelled.

Surprisingly enough, we can get the anti-symmetry correction essentially for free. Since events are processed one at a time, we need only concentrate on the effect of one event. Referring back to Figure 7, suppose an event on input B causes the output of the transformed function to change. This change will propagate to the final XOR gate. The event on B will also propagate to the XOR gate, and the two events will cancel one another. Thus, if the output of the transformed function changes, no output event will be scheduled for the XOR gate.

Now suppose that an event occurs on input B , but the output of the transformed function *does not change*. Because the output does not change, no event propagates into the XOR gate from the function output. However, the event on B still propagates directly into the XOR gate, causing an event on the output of the XOR.

The correcting XOR gate causes the usual effect of an event to be reversed. Events propagate when the output of a function does not change, and no event propagates when the output changes. To take advantage of this, we generate two sets of run-time routines. The first set contains a comparison of the form “if Old = New then propagate” while the second contains a comparison of the form “if Old \neq New then propagate”. The first set is used for those inputs directed into the corrective XOR, the second is used for other inputs. These two routines are virtual functions that are called through function pointers determined at compile time. No run-time test is required to distinguish the two types of inputs. Effectively, the corrective XOR is obtained for free.

Figures A5 and A6 give the run-time code for processing an event. Our simulator does not require gate simulation code, so the algorithms of Figures A5 and A6 represent virtually the entire run-time code of our simulator.

6 Experimental Results

For our experimental results, we used the ISCAS 85 benchmarks [19]. Although these benchmarks are the *de facto* standard for determining simulation performance, they are specified at the gate level rather than at the Boolean function level. In this respect, they are not the most ideal vehicle for determining the effectiveness of our simulator, however they exhibit a wide variety of symmetries of all types.

The main problem with gate-level circuits is that one must attempt to reconstruct the Boolean functions from which the circuit

was created. This is a decidedly non-trivial task. As an approximation to this we first identify the fanout-free networks in the circuit. These networks represent single-output functions, but are only an approximation of the original Boolean functions. About 50% of the gates in each circuit end up as isolated gates. We do not apply our algorithm to isolated gates, because their symmetries are already well-known.

A few circuits have very large fanout-free networks which represent several Boolean functions combined into a single network. To break these giant networks down into something resembling real Boolean functions we limit the number of inputs of a single network. We have experimented with different limits and have found that a limit of eight tends to expose the most symmetries. The limit is only approximate. It is possible for an individual partition to have more than eight inputs.

Our first experiment was to determine the number of anti-symmetries that appear in these circuits. Figure 12 gives the number of anti-symmetries found in each circuit when no other types of symmetries are detected. The numbers are counts of anti-symmetric variable pairs. These are further broken down into ordinary anti-symmetries (Ord.), multi-phase anti-symmetries (M.P.), conjugate anti-symmetries (Conj.), and combined multi-phase/conjugate symmetries (C.M.P). Every one of the ten benchmarks contains some anti-symmetries. The total ranges from a low of 10 for c432, to a high of 944 for c6288. This experiment verifies that anti-symmetries are indeed prevalent enough to be worth pursuing. If we examine the breakdown of sub-types, it is clear that it is necessary to combine anti-symmetry with multi-phase and conjugate symmetry. Note, for example, benchmark c499, which has no ordinary anti-symmetries, but has 104 symmetric pairs of other types.

To compare the prevalence of anti-symmetries with that of classical symmetries we determined the number of classical symmetries in each of the four categories. The results of this experiment are given in Figure 13. For most of the circuits there are significantly more classical symmetries than anti-symmetries, but there are two notable exceptions. Both c1355 and c6288, have more anti-symmetries than classical symmetries. Clearly we should detect anti-symmetries to handle these circuits effectively.

When combining the detection of anti-symmetries with that of classical symmetries, we encounter the phenomenon known as “symmetry masking.” This occurs when the detection of one type of symmetric pair prevents the detection of a different type. This is not necessarily a problem, since the two symmetries are usually with respect to the same pair of variables. Nevertheless, it is not possible to add the results of Figure 12 to those of Figure 13 to determine the total number of symmetric pairs.

To determine the effect of symmetry detection on simulation performance, we compared four different simulations for each benchmark circuit. The four simulations are with no symmetry detection, with anti-symmetry alone, with classical symmetry alone, and with combined classical and anti-symmetry. Each simulation was performed on a dedicated 3.06 Ghz Xeon processor with 2GB of 233 Mhz memory. The results, which are shown in Figure 15, are in seconds of execution time for 500,000 input vectors.

The compile step, which includes all symmetry detection, took less than a second for each circuit, regardless of the types of symmetries detected. The run times given in Figure 15 do not include the time required to detect symmetry.

Ckt	Ord.	M.P.	Conj.	C.M.P.	Total
c432	10	0	0	0	10
c499	0	40	64	0	104
c880	5	12	18	2	37
c1355	104	104	0	0	208
c1908	14	2	99	0	115
c2670	14	11	72	2	99
c3540	173	11	52	9	245
c5315	29	46	139	22	236
c6288	480	464	0	0	944
c7552	53	126	310	44	533

Figure 12. Anti-Symmetries

Ckt	Ord.	M.P.	Conj.	C.M.P.	Total
c432	47	45	9	2	103
c499	110	24	40	0	174
c880	133	4	13	2	152
c1355	14	24	0	104	142
c1908	138	4	4	0	146
c2670	279	10	11	58	358
c3540	198	169	74	23	464
c5315	298	12	120	239	669
c6288	0	0	480	0	480
c7552	710	16	108	119	953

Figure 13. Classical Symmetries.

Ckt	Classical	Anti	Total
c432	103	0	103
c499	174	0	174
c880	152	19	171
c1355	142	104	246
c1908	146	0	146
c2670	357	34	357
c3540	463	5	468
c5315	664	32	696
c6288	480	464	944
c7552	951	99	1050

Figure 14. Classical and Anti-Symmetries

Ckt	None	Anti	Classical	Combined
c432	2.54	2.52	2.35	2.28
c499	3.75	3.43	3.20	3.20
c880	5.54	5.49	4.85	4.78
c1355	5.30	4.91	5.17	5.08
c1908	5.12	5.01	5.04	4.95
c2670	17.77	17.57	17.06	16.97
c3540	12.37	12.20	11.18	11.20
c5315	26.71	26.23	24.11	24.12
c6288	18.81	17.29	19.30	18.47
c7552	31.38	30.73	29.09	29.12

Figure 15. Running Times.

Several conclusions can be drawn from Figure 15. First, using either classical symmetries or anti-symmetries in isolation gives a substantial benefit. Second, combining the two gives improved performance in some cases, and roughly the same performance as Classical symmetries in other cases. When detecting symmetries, each variable pair can be assigned at most one type of symmetry. However for many variable pairs, there is a choice of which type of symmetry to assign. In this respect, the simulator is sensitive to the order in which symmetries are detected. For c6288, anti-symmetries give better performance than either classical symmetries or combined

anti and classical symmetries. We corrected this problem by changing the order in which symmetries were detected. For Figure 15, for each pair of variables, we first detected classical symmetries and then anti-symmetries. We later changed the order to intersperse the detection of classical and anti-symmetries. This caused the anti-ordinary and anti-multiphase symmetries to be detected before the classical single-variable symmetries. This reduced the "combined" time for c6288 to roughly the same as that for anti-symmetries.

7 Conclusion

Detecting and using anti-symmetries can be of benefit for simulation performance. For all of our circuits, simulation detection was essentially instantaneous, even after adding the code for detecting anti-symmetries, so the cost of detecting anti-symmetries is worth the benefit. This is especially true because our simulator is a compiled simulator, and the cost of symmetry detection can be amortized over many hours of simulation. Furthermore, the ability to pick and choose among several different types of symmetry permits us to weigh several different detection options against one another and pick the best. Again, the cost of doing this can be amortized over many simulations. Even in a situation where the circuit is undergoing rapid changes, the cost of detecting symmetry is barely noticeable and will not substantially affect compilation time.

Because of these benefits, anti-symmetry detection is now a permanent part of our simulation engine.

8 References

1. C. E. Shannon, "The synthesis of two-terminal switching circuits," *Bell System Technical Journal*, Vol.28, No.1, pp. 59-98, 1949.
2. C. R. Edwards and S. L. Hurst, "A digital synthesis procedure under function symmetries and mapping methods," *IEEE Transactions on Computers*, Vol.27, No.11, pp. 985-997, 1978.
3. D. Moller, P. Molitor, R. Drechsler and J. W. G. U. Frankfort, "Symmetry based variable ordering for ROBDDs," *IFIP Workshop on Logic and Architecture Synthesis*, pp. 47-53, 1994.
4. C. Scholl, D. Moller, P. Molitor and R. Drechsler, "BDD minimization using symmetries," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol.18, No.2, pp. 81-100, 1999.
5. T. Sasao, "A new expansion of symmetric functions and their application to non-disjoint functional decompositions for LUT type FPGAs," *IEEE International Workshop on Logic Synthesis*, pp. 105-110, 2000.
6. V. N. Kravets and K. A. Sakallah, "Constructive library-aware synthesis using symmetries," *Design Automation and Test in Europe*, pp. 208-213, 2000.
7. P. M. Maurer, "Conjugate Symmetry," *Formal Methods Syst. Des.*, Vol.38, No.3, pp. 263-288, 2011.
8. V. N. Kravets and K. A. Sakallah, "Generalized symmetries in boolean functions," *IEEE International Conference on Computer Aided Design*, pp. 526-532, 2000.
9. J. Mohnke, P. Molitor and S. Malik, "Limits of using signatures for permutation independent Boolean comparison," *Formal Methods Syst. Des.*, Vol.21, No.2, pp. 167-191, 2002.
10. P. M. Maurer, "An application of group theory to the analysis of symmetric gates," Department of Computer Science, Baylor University, Waco, TX 76798, <http://hdl.handle.net/2104/5438>, 2009.
11. V. N. Kravets and K. A. Sakallah, "Generalized symmetries in boolean functions," Advanced Computer Architecture Laboratory Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor, MI 48109, <http://web.eecs.umich.edu/techreports/cse/00/CSE-TR-420-00.pdf>, 2002.

12. C. C. Tsai and M. Marek-Sadowska, "Generalized Reed-Muller forms as a tool to detect symmetries," *IEEE Transactions on Computers*, Vol.45, No.1, pp. 33-40, 1996.
13. M. Chrzanowska-Jeske, A. Mishchenko and J. R. Burch, "Linear Cofactor Relationships in Boolean Functions," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, Vol.25, No.6, pp. 1011-1023, 2006.
14. C. C. Tsai and M. Marek-Sadowska, "Boolean functions classification via fixed polarity Reed-Muller forms," *IEEE Transactions on Computers*, Vol.46, No.2, pp. 173-186, 1997.
15. P. M. Maurer, "Extending symmetric variable-pair transivities using state-space transformations," Department of Computer Science, Baylor University, Waco, Texas 76798, <http://hdl.handle.net/2104/8185>, 2011.
16. F. Luccio and L. Pagli, "On a new Boolean function with applications," *IEEE Transactions on Computers*, Vol.48, No.3, pp. 296-310, 1999.
17. A. Bernasconi, V. Ciriani, F. Luccio and L. Pagli, "Synthesis of Autosymmetric Functions in a New Three-Level Form," *Theory of Computing Systems*, Vol.42, No.4, pp. 450-464, 2008.
18. P. M. Maurer. Efficient event-driven simulation by exploiting the output observability of gate clusters. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on Vol.22, No.11*, pp. 1471-1486.
19. F. Brglez, P. Pownall and R. Hum. Accelerated ATPG and fault grading via testability analysis. Presented at Proceedings of IEEE Int. Symposium on Circuits and Systems.

A1. Appendix 1: Algorithms

Compute f_{00} , f_{01} , f_{10} and f_{11} and place them in a 2-dimensional hypercube structure.

Repeat Until No variables remain

Test the back diagonals for ordinary symmetry

Test the back diagonals for anti-symmetry

If a symmetric variable pair is detected

Collapse the structure by combining vertices along diagonals.

Add any required corrective functions

Endif

If uneliminated variables remain

Compute two new cofactors from each existing cofactor

Double the size of the structure increasing dimensions by 1.

Insert the new cofactors into the new structure

Endif

Figure A1. Symmetry Detection.

Remove all State-Space transformations.

Check Back diagonals, stop if symmetry is detected.

Reverse Hyper Linear structure along one dimension

Check Back diagonals, stop if symmetry is detected.

Restore Hyper Linear Structure

Reverse Odd Rows

Check Back diagonals, stop if symmetry is detected.

Restore Hyper Linear Structure

Reverse Odd Columns

Check Back diagonals, stop if symmetry is detected.

Restore Hyper Linear Structure

Reverse Hyper Linear structure along one dimension

Reverse Odd Rows

Check Back diagonals, stop if symmetry is detected.

Restore Hyper Linear Structure

Reverse Hyper Linear structure along one dimension

Reverse Odd Columns

Check Back diagonals, stop if symmetry is detected.

Figure A2. General Symmetry Pair Detection.

For each plane containing the pair to be tested

For each diagonal

Comparator = HeadVertex.function;

For V = each vertex after the head vertex

If Comparator Not Equal V.function **Then**

Report Failure

Endif

Endfor

Endfor

Report Success

Figure A3. Ordinary Symmetry Diagonal Check.

For each plane containing the pair to be tested

For each diagonal

Comparator = HeadVertex.function;

AntiComparator = Negate(Comparator)

Odd = 1;

For V = each vertex after the head vertex

If Odd = 1 **Then**

If AntiComparator Not Equal V.function **Then**

Report Failure

Endif

Odd = 0;

Else

If Comparator Not Equal V.function **Then**

Report Failure

Endif

Odd = 1;

Endif

Endfor

Endfor

Report Success

Figure A4. Anti-Symmetry Diagonal Check.

GateState[i] = GateState[i] + PortState;

PortState = -PortState;

// note the contents of the Then and Else sections

If Value[GateState] **Not Equal** OldGateState **Then**

If EventQueued **Then**

Dequeue Event

Else

Queue Event

Endif

Else

Do Nothing;

Endif

OldGateState = Value[GateState];

GoTo NextEvent

Figure A5. Ordinary Symmetry Event Processor.

```
GateState[i] = GateState[i] + PortState;
PortState = -PortState;
// note the contents of the Then and Else sections
If Value[GateState] Not Equal OldGateState Then
  Do Nothing;
Else
  If EventQueued Then
    Dequeue Event
```

```
  Else
    Queue Event
  Endif
Endif
OldGateState = Value[GateState];
GoTo NextEvent
Figure A6. Anti-Symmetry Event Processor.
```