

# Fast Heuristic Minimization of Exclusive Sums-of-Products

Alan Mishchenko and Marek Perkowski

Portland State University  
Department of Electrical and Computer Engineering  
Portland, OR 97207, USA  
[alanmi,mperkows]@ee.pdx.edu

## Abstract

Exclusive Sums-Of-Products (ESOPs) play an important role in logic synthesis and design-for-test. This paper presents an improved version of the heuristic ESOP minimization procedure proposed in [1,2]. The improvements concern three aspects of the procedure: (1) computation of the starting ESOP cover; (2) increase of the search space for solutions by applying a larger set of cube transformations; (3) development of specialized data-structures for robust manipulation of ESOP covers. Comparison of the new heuristic ESOP minimizer EXORCISM-4 with other minimizers (EXMIN2 [3], MINT[4], EXORCISM-2[1] and EXORCISM-3[2]) show that, in most cases, EXORCISM-4 produces results of comparable or better quality on average ten times faster.

## Introduction

Two-level Sums-Of-Products (SOP) have been widely used to represent and manipulate Boolean functions. Exact and heuristic SOP minimization has traditionally attracted attention of researchers because in many applications it is important to have as compact SOP representation as possible. The value of Exclusive Sums-Of-Product (ESOP) and ESOP minimization, on the other hand, has been underestimated for the following reasons:

- The EXOR gate, as implemented in the widely used standard cell libraries, is almost two times larger and slower compared to NAND and NOR gates.
- Relatively few algorithms use ESOPs for internal representation of Boolean functions.
- The exact minimization of ESOPs is practical only for arbitrary functions of five input variables and special classes of functions up to ten variables [5], while the heuristic minimization is based on search algorithms and is computationally more expansive than SOP minimization.

However, the EXOR gate and ESOPs play an important role in logic synthesis, design for test, and other areas of computer technology:

- For many Boolean functions, the number of cubes in minimal ESOPs is less than the number of cubes in minimal SOPs [6].
- The EXOR gate has excellent testability properties leading to efficient methods for automatic test pattern generation and design for test [7].
- For many types of practical circuits, the selective use of EXOR gates in logic synthesis yields better implementations in terms of both area and delay [8,9].

Over the years, several ESOP minimization strategies and ESOP minimizers have been created but none of them matched well with our research purposes. Some of the tools did not belong to the public domain; others had limitations on the input data format. Above all, the available tools had performance limitations when applied to large Boolean functions. This motivated our work on a new user-friendly minimizer EXORCISM-4, which follows the tradition of previous ESOP minimizers, EXORCISM-2 [1] and EXORCISM-3 [2], developed at Portland State University.

This paper presents the results of research directed towards development of the new ESOP minimization tool. We have experimented with a number of approaches and algorithms and selected those that proved to be practically efficient. In our presentation below, both the successful and the unsuccessful attempts to improve the performance of ESOP minimization will be mentioned.

The rest of the paper is organized as follows. Section 1 discusses the previous work in ESOP minimization. Section 2 gives the main definitions. Section 3 discusses the general data-flow in our minimization procedure. Section 4 reviews the starting cover computation. Section 5 introduces and illustrates the ExorLink operation. Section 6 presents the look-ahead strategies. Section 7 discusses data structures and implementation issues. Section 8 gives experimental results comparing the new tool with other tools. Section 9 concludes the paper and outlines the future research directions.

## 1 Previous Work

The known approaches to heuristic ESOP minimization are roughly classified as follows:

- Iterative cube transformation [1-4,10,11,12]
- Minimization of subclasses of ESOP expressions [13,14]
- Construction of the resulting ESOP cube by cube [?]
- Simplification based on term rewriting [15]

Most of the ESOP minimizers use the first strategy. They iteratively apply a set of selected cube transformations to single cubes, cube pairs, or larger cube groups. If replacement of the starting cubes by the resulting cubes leads to simplification or some other desirable improvement, the cover is modified without changing the function represented by it.

Our approach also uses cube transformations. The difference is that we use only one very general cube transformation (distance-k ExorLink). When applied in the context of limited look-ahead discussed in Section 7, ExorLink subsumes many other (possibly all) cube transformations known to date.

Also relevant to our approach (in particular, to the starting cover computation) is the second type of methods listed above. These methods attempt heuristic ESOP minimization by developing efficient algorithms to minimize specialized classes of AND-EXOR expressions: Positive Polarity Reed-Muller Forms (PPRMs), Kronecker and Pseudo-Kronecker Forms (KRM and PKRM), Generalized Reed-Muller Forms (GRMs), etc. These classes differ from ESOPs in restrictions applied to polarities of literals or in the type of canonical expansions used to the generic decision diagrams. The reader is referred to [Sasao] for precise definitions.

Because the set of all ESOPs is the superset of expressions belonging to any of the above mentioned classes, an exact or heuristic minimization procedure developed for a class of functions may be also used as an approximation, or a starting point, in ESOP minimization.

Of particular importance to the research in this paper is the algorithm for exact minimization of Pseudo-Kronecker expressions [Drechsler]. This algorithm gives a good trade-off between the quality and the runtime. Pseudo-Kronecker expressions are a relatively comprehensive subclass of ESOPs and therefore finding the exact minimum in this class yields relatively good quality. On the other hand, the Pseudo-Kronecker minimization algorithm exploits the BDD representation and therefore can be applied to a wide range of benchmarks, for which this representation can be constructed. This algorithm is discussed in more detail in Section 4, subsection "Starting Cover Computation".

## 2 Preliminaries

This section gives the basic definitions used in the paper.

The truth tables and Karnaugh maps of OR and Exclusive-OR (EXOR) operations are shown in Fig. 1.




Figure 1. OR and EXOR operations.

A literal is a Boolean variable in the negative or positive polarity. A cube is a product term composed of literals using Boolean AND operation.

Two cubes coincide in variable  $x$  if  $x$  does not appear in the cubes or if  $x$  appears in the cubes in the same polarity. Two cubes differ in variable  $x$  if they do not coincide in variable  $x$ . The distance between two cubes is the number of variables, in which the cubes differ.

A variable appears in the cube in one of the three forms (1) negative polarity; (2) positive polarity; (3) don't-care.

For example, assuming that cube  $\overline{a}b\overline{c}$  belongs to function with input variables  $(a,b,c,d)$ , variable  $a$  appears in positive polarity, variables  $b$  and  $d$  appear in negative polarity, and variable  $c$  appears as a don't-care.

An Exclusive-Sum-Of-Products (ESOP) is an Exclusive-OR of zero or more cubes. An ESOP is reduced if it does not contain identical cubes. An ESOP is minimal if all cube pairs have distance 2 or more. An ESOP is an exact minimum if it contains the minimum number of cubes among all ESOPs representing the given function.

The following propositions can be proved using the fundamental property of EXOR operation.

Proposition 1: Two identical cubes (distance-0 cubes) can be added to any ESOP without changing the function represented by it.

Proposition 2: The EXOR of two cubes that have distance 1 can be represented by a single cube.

The ExorLink operation discussed in Section 5 extends these two propositions to cubes pairs with arbitrary distance between the cubes.

## 3 ESOP Minimization Algorithm

This section outlines the main steps of the heuristic ESOP minimization procedure and data representations used to implement these steps.

### Minimization Algorithm

The pseudo-code of the minimization algorithm is shown in Fig. 2. The minimization procedure takes the multi-output function and the quality parameter, which determines how many minimization loops are performed.

```

esop HeuristicMinimization( func F, quality Q )
{
  esop Cover = GenerateStartingCover( F );
  while ( Q > 0 ) {
    ResetCubePairs( Cover );
    do {
      do {
        Cover= AgressiveMinimization(Cover);
      } while (there is improvement);
      Cover= LastGaspMinimization(Cover);
    } while (there is improvement);
    Q = Q-1;
  }
  Cover = RefinementMinimization(Cover);
  return Cover;
}

```

Figure 2. Pseudo-code of ESOP minimization algorithm.

After computing the starting cover using the BDD-based algorithm [Drechsler], the data flow enters the main minimization loop. Each time the main loop is entered internal data structures storing the candidate pairs for the application of distance-2, distance-3 and distance-4 Exorlink are reset. This is minimization in the internal loops may result in “cube pair leaks”, when some cube pairs fall out of consideration until they are found and returned into the queue by ResetCubePairs().

Inside the main loop, there is a loop that applies aggressive minimization strategy, as described in section “Look-Ahead Strategies”. When aggressive minimization does not lead to improvement, the last gasp is applied, which tries to get the cover out of the local minimum by several additional round of distance-4 ExorLink.

Finally, when the minimization loops terminate, there is a call to the procedure RefinementMinimization() which tries to improve the literal count of the cover without attempting to reduce the number of cubes.

This structure of the minimization procedure was found as a result of extensive experimentation having the goal of finding a sequence of minimization operations that perform relatively well on a wide selection of benchmarks.

#### Data Representation

The input to the ESOP minimization software is the text file in standard PLA or BLIF formats, representing a multi-output Boolean function. In the current version of the minimizer, the don't-cares of the input function are ignored and only the on-set of the function is considered.

The processing starts by creating the shared BDDs for the input function. Dynamic reordering of variables is enabled during the BDD construction, which allows us to work with relatively large multi-output functions.

In general, the minimizer uses mixed explicit and implicit data representations. After computing the starting

cover using shared BDDs, the representation switches to cubes encoded in positional notation and stored explicitly using bit strings. The subsequent cube operations and cover transformations are carried out using the explicit representation, except the verification step at the end. During verification, the resulting cover is again converted into shared BDDs and compared with the BDDs constructed for the input function.

We experimented trying to replace the explicit representation of cubes as bit strings by their implicit representation using Zero-Suppressed Binary Decision Diagrams (ZDDs) [Minato]. This attempt did not succeed because ZDDs speed-up computation only if there is a way to process many cubes in parallel. Meanwhile cube transformations used in our approach, in particular, finding all distance-k cube pairs and performing ExorLink operation, requires processing cubes one at a time.

## 4 Starting Cover Computation

This section summarizes the BDD-based algorithm for minimization of the Pseudo-Kronecker expressions used to compute the starting cover. For a detailed discussion of the algorithm, including the proof that it computes the exact minimum for the given subclass of expressions, the reader is referred to [Drechsler].

The algorithm works in two passes. During the first pass, the shared BDD of the input function is traversed depth-first and in each node the best expansion is found and saved in the lossless cache. The expansions are selected out of the three canonical expansions (Shannon, Positive Davio or Negative Davio) and the number of cubes they add to the solution is used to evaluate their cost. It is possible to count the number of cubes in the resulting Pseudo-Kronecker expressions because they result from “flattening” the Pseudo-Kronecker decision diagrams, which are simulated (not explicitly constructed) by the algorithm presented in Fig. 3. This algorithm also minimizes the number of literals in the minimal Pseudo-Kronecker expression by preferring Positive or Negative Davio to Shannon expansion whenever it does not increase the cube count.

During the second pass, the enumeration of all paths in the diagram is performed (Fig. 4). In each node, the best expansion is retrieved from the lossless cache and, depending on the expansion, one literal is added to the array of the variable values representing the current cube. When the traversing procedure reaches the bottom of the diagram, it has the values of all variables collected in the array, which is now used to generate the cube and add it to the resulting cover. At the end of the traversal, the resulting cover contains all cubes belonging to the exact minimum of the Pseudo-Kronecker expression.

```

(exptype,int) CountCubesInExactPseudoKro( func F )
{
  exptype Exp; int Cost;

  // consider the terminal cases
  if ( F == 0 ) return ( pDavio, 0 );
  if ( F == 1 ) return ( pDavio, 1 );
  if ( ( Exp, Cost ) = CheckCacheForResult( F ) )
    return ( Exp, Cost );

  // determine the cofactors
  ( F0, F1 ) = DecomposeBdd( F, TopVar(F) );

  // recursively solve subproblems
  int N0, N1, N2, Nmax, Cost;
  N0 = CountCubesInExactPseudoKro( F0 );
  N1 = CountCubesInExactPseudoKro( F1 );
  N2 = CountCubesInExactPseudoKro( F0 ⊕ F1 );

  // determine the most costly expansion
  MaxN = max( N0, N1, N2 );

  // choose the least constly expansion
  if ( MaxN == N0 )
    Exp = pDavio; Cost = N1 + N2;
  else if ( MaxN == N1 )
    Exp = nDavio; Cost = N0 + N2;
  else /* if ( MaxN == N2 ) */
    Exp = Shannon; Cost = N0 + N1;

  // cache and return the result
  InsertIntoCache( F, Exp, Cost );
  return ( Exp, Cost );
}

```

Figure 3. Procedure for counting the number of cubes in the exact minimum Pseudo-Kronecker expansion.

## 5 ExorLink Operation

ExorLink operation has been introduced in [N.Song's MStHesis] as a generalization of several simple cube transformation used in various approaches to ESOP minimization. The main idea of ExorLink is to replace two cubes by a set of cubes without changing the function.

The number of cubes to be substituted instead of the initial cube pair is equal to the distance between the starting cubes. If the distance between the cubes is more than 1, there are several possibilities of performing ExorLink, each of them leading to functionally equivalent representation composed of different cubes.

For example, consider the following equivalent transformations of cubes  $ab$  and  $\bar{a}\bar{b}$  constituting the distance-2 ExorLink operation:

$$ab \oplus \bar{a}\bar{b} = a\oplus b = \bar{a} \oplus \bar{b}.$$

In general, distance-k ExorLink operation produces  $k!$  cube groups containing  $k$  cubes each. In all the cube groups there are  $k \cdot 2^{k-1}$  different cubes, some of them appearing is more than one group. The cube set generation is illustrated by distance-3 ExorLink applied to cubes  $abc$  and  $\bar{a}\bar{b}\bar{c}$  in Fig. 5. Zeros are not shown.

```

GenerateExactPseudoKro( bdd F, char * VarValues )
{
  // consider the terminal cases
  if ( F == 0 ) return;
  if ( F == 1 ) {
    cube NewCube = CreateCube( VarValues );
    AddCubeToCover( NewCube );
    return;
  }

  // find the best expansion by a cache lookup
  exptype Exp = CheckCacheForResult( F );

  // determine the top-most variable
  var X = TopVar(F);

  // determine the cofactors
  ( F0, F1 ) = DecomposeBdd( F, X );

  // generate cubes in the left/right branches
  if ( Exp = pDavio ) {
    VarValues[ X ] = VAR_ABSENT;
    GenerateExactPseudoKro( F0, VarValues );
    VarValues[ X ] = VAR_POSITIVE;
    GenerateExactPseudoKro( F0⊕F1, VarValues);
  }
  else if ( Exp = nDavio ) {
    VarValues[ X ] = VAR_ABSENT;
    GenerateExactPseudoKro( F0, VarValues );
    VarValues[ X ] = VAR_NEGATIVE;
    GenerateExactPseudoKro( F0⊕F1, VarValues);
  }
  else /* if ( Exp = Shannon ) */ {
    VarValues[ X ] = VAR_NEGATIVE;
    GenerateExactPseudoKro( F0, VarValues );
    VarValues[ X ] = VAR_POSITIVE;
    GenerateExactPseudoKro( F1, VarValues );
  }
}

```

Figure 4. Procedure for the generation of the exact minimum Pseudo-Kronecker expansion.




Figure 5. Example of distance-3 Exorlink.

Observe that cube groups correspond to different paths in the Boolean space connecting minterms  $abc$  and  $\bar{a}\bar{b}\bar{c}$ . If the distance between the starting cubes is  $k$ , there are  $k$  directions to make the first step,  $k-1$  directions to make the second step, and so on. As a result, there are  $k!$  paths each of them corresponding to one cube group.

## 6 Look-Ahead Strategies

The previous algorithms [Exor-2,Exor-3] used a combination of full backtracking and limited backtracking. Full backtracking means exploring all the branches of the

search tree up to a certain depth, selecting the best way to modify the cubes, and performing the required transformation. The parameter specified on the command line by the switch “-q” stands for exploring a fixed number of branches in the search tree up to a certain depth. The branches to explore may be selected randomly or found heuristically.

For ESOP minimization with ExorLink as a basic cube transformation routine, it was found that “large-depth look-ahead with a small number of backtracks is much faster and usually gives better results than the small-depth look-ahead and a large number of backtracks, although the latter is also sometimes very helpful”. [Song’s PhD Thesis]

We tested these and other look-ahead strategies and finally adopted the one described as “backtrack until the first success”. “Success” is defined differently depending on the minimization strategy and the distance between the cubes, to which ExorLink is applied.

Table 1 lists the criteria, which our algorithm uses to determine whether the transformation is successful and should be applied, or whether backtracking should continue by testing the next available cube group. The next available group may result from the same ExorLink (recall that distance-k ExorLink produces k! cube groups) or from ExorLink applied to the next cube pair in the queue of pairs to be considered.

Table 1. Definition of criteria for accepting cube transformations resulting from the ExorLink operation.

ExorLink Distance	Criteria for Accepting Cube Transformations	
	Aggressive Strategy	Refinement Strategy
2	Leads to reduction in the number of cubes.	Does not increase the number of cubes. Leads to reduction in the number of literals.
3	Does not increase the number of cubes.	Does not increase the number of cubes. Leads to reduction in the number of literals.
4	Does not increase the number of cubes.	Not used

In different modes of operation our algorithm applies a different combination of aggressive and refinement minimization. The general rule is that in each minimization loop, aggressive strategy is applied several times to all cube pairs starting from distance-2 ExorLink-2 up to distance-4 ExorLink.

Depending on the requested minimization quality the algorithm performs different number of minimization loops. This number roughly corresponds to the value of the quality

The trade-off between the quality and the computation time is illustrated in Table 2 using benchmark alu4.pla, which has 14 inputs and 8 outputs. The shared BDD node count is 804. The input file reading and the starting cover computation time is 0.45 sec and 0.22 sec respectively.

Table 2. Quality/runtime trade-off for alu4.pla.

	Quality level (“-qN”), N										
	0	1	2	3	4	5	6	7	8	9	10
Cubes	440	440	435	435	415	415	415	414	411	411	411
Time, c	9.1	10.1	24.2	26.4	65.5	68.9	71.6	87.9	91.7	93.2	97.0

Using this strategy in the context of ExorLink-4 allows us to make the following statement about the quality of minimization achieved by the algorithm.

Theorem. Given enough time to perform relatively complete exploration of the search space, the above minimization procedure always escapes from any local minimum that is one or two cubes deep. (Picture?)

Compared to this, distance-3 ExorLink [] can only guarantee getting out of the local minima one cube deep. This difference explains why in some cases our minimizer has found ESOP covers of smaller cardinality compared to other minimizers (see section “Experimental Results”).

## 7 Implementation Issues

The following general principles of high-performance programming have been used in the implementation of EXORCISM-4.

### Allocate Memory With Consideration

The heuristic minimization algorithm described in this paper generates a large set of intermediate cubes, which only reshape the cover without improving its cardinality. Because calls to system memory allocation are relatively expensive in terms of runtime, EXORCISM-4 performs efficient memory recycling. Memory allocated at the beginning for the starting cover is never reallocated because the size of the cover cannot increase during minimization. To ensure efficient memory recycling, cubes that are not currently in use are kept in a linked list. Adding and removing cubes in the linked list is much more efficient than calling systems memory allocation whenever a new cube is created or deleted.

### Terminate Computation As Early As Possible

In many cases, it is possible to speed up computation by early detection of a situation when the result of computation is useless for whatever reason. The

performance improvement is more or less substantial depending on how often this situation occurs and how large is the part of computation that can be skipped.

EXORCISM-4 uses early termination of computation in the several cases.

- One of the operations repeated very often is computing the distance between two cubes. The algorithm considers only cube pairs with distance four or less. It means that as soon as at least five different variables are detected in the cubes, the rest of computation can be skipped. This simple idea gives up to 20% improvement in runtime for benchmarks with many inputs and outputs.
- As shown in Section 6, distance-k ExorLink operation produces  $k \cdot 2^k$  new cubes grouped into  $k!$  groups in such a way that some of the cubes are encountered in more than one group. Our algorithm considers the groups one by one and as soon as it find the matching cube in a group it is used to reshape the cover. Consequently, ExorLink can be implemented in such a way that instead of generating all cubes at once, it generates them "on demand". This trick is responsible for speeding up distance-4 ExorLink at least 50%, because in practice only 5 to 10 cubes out of 32 are often enough.

#### Reuse Results of Previous Computation

The importance of this principle can hardly be overestimated, as witnessed by the success of caching techniques responsible for speeding-up computation by many orders of magnitude in BDD-based applications.

In case of EXORCISM-4, standard BDD caching techniques are used during the first step of starting cover computation. There are two other cases when caching helps reduce unnecessary computation:

- Each new cube generated by ExorLink is tested for distances with all other cubes present in the cover. Even though at the beginning of testing it is not known whether the cube will be taken into the cover, the testing procedure stores information about those cubes currently in the cover that have distances 2, 3, and 4 with the given cube. As a result, if the cube is eventually accepted into the cover, there is no need to compute candidate pairs for distance 2, 3, or 4 ExorLink in the following iterations.
- Finally, there is a way to reduce computation by saving the results of distance-testing for those cubes that are generated by ExorLink and included into more than one out of  $k!$  groups.

#### Use Sorted Data Structures

Using sorted data structures (ordered lists, heaps, tries, Patricia trees, etc.) often leads to faster search and reduction of the access time.

We tried to store the cubes in an ordered list and in an ordered matrix (sorting them by the number of positive and negative literals). The rationale for doing this would be to facilitate finding small-distance cube pairs, and thereby speeding up cube selection for ExorLink. However, in practice it turned out that sorting cubes does not help reducing the runtime, and eventually we switched back to representing the cover as an unsorted linked list cubes.

## 8 Experimental results

EXORCISM-4 has been implemented in platform-independent C++ using the BDD package CUDD [1] and tested on several Unix and Windows workstations. The experimental results below have been computed on a 966Mhz Pentium III PC under Microsoft Windows 2000. At all times the program used no more than 50Mb of memory. To make a fair runtime comparison with earlier programs, the time measurements have been normalized to reflect the speed of computers used to run these programs.

The correctness of the minimization results has been checked by a BDD-based verifier built into EXORCISM-4.

Table 3 compares the results of minimization by EXORCISM-4 with those of EXMIN2 [2] and MINT [3]. Table 4 compares EXORCISM-4 with EXORCISM-2 [4] and EXORCISM-3 [5]. The comparison is in terms of the number of cubes and literals in the final solution and the CPU time spent in minimization. The last two columns in Tables 3 and 4 give the ratio of the runtime needed by the specified minimizer and the runtime of EXORCISM-4.

It follows from the tables that on small examples, EXORCISM-4 is on average 6 times faster than EXMIN2 and 25 times faster than MINT. On larger examples, for which the experimental results of EXMIN2 and MINT are not available, EXORCISM-4 is 50 times faster than EXORCISM-2 and 1.3 times faster than EXORCISM-3. However, if we look at separate benchmarks, we see that

## 9 Conclusions

Research described in this paper resulted in the development of a new heuristic ESOP minimizer EXORCISM-4. It compares favorably in terms of minimization quality, is faster and allows for processing of larger multi-output Boolean functions. The new tool accepts a wide spectrum of input data formats (PLA, ESOP PLA, BLIF) and unlike EXORCISM-2 and EXORCISM-3, it does not require the input to be a disjoint cover.

The disadvantages of the current implementation are the following: (1) it considers only binary-input data, (2) it does not take don't-cares into account, and (3) on the cover reshaping stage, it uses bit strings to represent cubes. The explicit data structures explain why the performance degrades when increasing the cube cover. However, even functions with covers composed of several thousand cubes can be processed in realistic time (5-10 minutes).

We developed the new efficient ESOP minimizer hoping that it will serve a variety of applications, and in particular that it will open new opportunities for combining the power of SOP and ESOP minimization to produce efficient implementation of large logic functions.

## References

- [1] N. Song, M. Perkowski, "EXORCISM-MV-2: Minimization of Exclusive Sum of Product Expressions for Multiple-Valued Input Incompletely Specified Functions," Proc. IWMVL 1993, pp. 132-137.
- [2] N. Song, M. Perkowski, "Minimization of Exclusive Sum of Products Expressions for Multi-Output Multiple-Valued Input, Incompletely Specified Functions," IEEE Transactions on Computer Aided Design, Vol. 15, No. 4, April 1996, pp. 385-395.
- [3] T. Kozłowski, E. L. Dagless, J. M. Saul. "An enhanced algorithm for the minimization of exclusive-OR sum of products for incompletely specified functions". Proc. of ICCD, 1995, pp. 244-249.
- [4] T. Sasao. "EXMIN2: A Simplified Algorithm for Exclusive-OR-Sum-of Products Expressions for Multiple-Valued-Input Two-Valued-Output Functions". IEEE Trans. CAD. Vol. 12, No. 5, May 1993. pp. 621-632.
- [5] T. Sasao, "An Exact Minimization of AND-EXOR Expressions Using BDDs". Proc. of IFIP WG 10.5 Workshop on Applications of the Reed-Muller Expansion in Circuit Design, 1993, Wilhelm Schickard-Institute fuer Informatik, pp. 91-98.
- [6] N. Koda, T. Sasao, "LP Characteristic Vector for Logic Functions". Proc. of IFIP WG 10.5 Workshop on Applications of the Reed-Muller Expansion in Circuit Design, 1993, Wilhelm Schickard-Institute fuer Informatik, pp. 99-108.
- [7] N. Koda, T. Sasao, "An upper bound on the number of products in minimum ESOPs". Proc. of IFIP WG 10.5 Workshop on Applications of the Reed-Muller Expansion in Circuit Design, 1995, Chiba, Japan, pp. 94-101.
- [8] T. Sasao, ed., Representation of Discrete Functions, Kluwer Academic Publishers, May 1996.
- [9] U. Kalay, M. Perkowski, D. Hall, "A Minimal Universal Test Set for Self Test of EXOR-Sum-Of-Products Circuits". IEEE Trans. CompVol. 49, N3, March 1999, pp.267-276.
- [10] C. Yang, M. Ciesielski, V. Singhal. "BDS: A BDD-based Logic Optimization System". Proc. of DAC 2000, pp. 92-97.
- [11] A. Mishchenko, B. Steinbach, M. Perkowski. "An algorithm for Bi-Decomposition of Logic Functions". Accepted to DAC 2001.
- [12] H. Fleisher, M. Tavel, J. Yeager. "A Computer Algorithm for Minimizing Reed-Muller Canonical Forms". IEEE Trans. Comp.Vol. 36. No. 2, pp. 247-250. February 1987.
- [13] M.Helliwell, M.A. Perkowski. "A Fast Algorithm to Minimize Multi-Output Mixed-Polarity Generalized Reed-Muller Forms".Proc. DAC'88 pp. 427-432.
- [14] J. M. Saul. "An Improved Algorithm for the Minimization of Mixed Polarity Reed-Muller Representation". Proc. ICCD 1990. pp. 372-375.
- [15] T. Sasao. GRM Minimization Paper.
- [16] R. Drechsler, "Pseudo-Kronecker Expressions for Symmetric Functions".IEEE Trans. CompVol. 48. No. 8, Sept. 1999, pp. 987-990.
- [17] D. Brand, T. Sasao. "Minimization of AND-EXOR expressions using rewriting rules".IEEE Trans. CompVol. C42. No. 5, pp. 568-576, May 1993.
- [18] N. Song.Minimization of Exclusive Sum of Product Expressions for Multi-Valued Input Incompletely Specified Functions.M.S. Thesis. EE Dept. Portland State University. Portland, OR, 1992.
- [19] N. Song.A New Design Methodology for Two-Dimensional Logic Cell ArraysPh.D. Thesis. EE Dept. Portland State University. Portland, OR, 1997.
- [20] I. I. Zhegalkin, "O tekhnike vychisleniy predlozheniy v simvolicheskoy logike" (About a Technique of Computation of Expressions in Symbolic Logic), Mat. Sb. Vol. 34, pp. 9-28, 1927.
- [21] I. I. Zhegalkin, "Arifmetizatsiya simvolicheskoy logiki" (Arythmetization of Symbolic Logic), Mat. Sb. Vol. 35, pp. 311-377, 1928.
- [22] S. M. Reed. "A class of multiple-error-correcting codes and their decoding scheme", IRE Trans. Information Theory. Vol. PGIT-4, pp. 38-49. 1954.
- [23] D. E. Muller. "Application of Boolean algebra to switching circuit design and to error detection". IRE Trans. on Electron. Comp. Vol. EC-3, pp. 6-12, 1954.
- [24] F. Somenzi, BDD/ADD/ZDD package "CUDD", v.2.3.0, <http://vlsi.colorado.edu/~fabio/>

Table 3. Comparison of ESOP minimization results with EXMIN2 [] and MINT []

Benchmark			Cubes			Literals			CPU time, c			Comparison	
Name	Ins	Outs	EXM2	MINT	Exor4	EXM2	MINT	Exor4	EXM2	MINT	Exor4	w/EXM2	w/MINT
5xp1	7	10	34	32	31	186	181	175	13	25	3	0.23	0.12
9sym	9	1	53	51	51	433	427	426	25	52	13	0.52	0.25
add6	12	7	127	127	127	872	936	859	430	320	31	0.01	0.10
addm4	9	8	91	90/89	90/89	654	651	624	129	355	39	0.30	0.11
b12	15	9	28	28	28	164	167	166	4	17	1	0.25	0.06
clip	9	5	68	64	64/63	517	492	479	55	140	11	0.20	0.79
ex7	16	5	81	81	81	601	592	584	46	166	13	0.28	0.08
f51m	8	8	32	31	31	161	185	162	10	22	5	0.50	0.23
in7	26	10	35	35	35	333	352	343	12	56	3	0.25	0.05
intb	15	7	307	267	268	3036	2519	2527	1353	2476	134	0.10	0.05
life	9	1	54	52/51	50/48	415	391	370	23	32	9	0.39	0.28
m181	15	9	29	29	29	169	172	171	5	18	1	0.20	0.06
m4	8	16	84	83	77/76	783	897	714	189	178	120	0.63	0.67
max512	9	6	89	88/83	84/82	696	723	672	71	187	64	0.90	0.34
rd53	5	3	15	16/15	14	60	69	57	2	2	1	0.50	0.50
rd73	7	3	42	36	36	221	194	197	20	36	9	0.45	0.25
rd84	8	4	59	55/54	59/58	330	303	333	45	11	17	0.37	1.54
ryy6	16	1	40	40	40	368	368	368	13	18	2	0.15	0.11
sao2	10	4	29	29/27	28	308	311	288	8	12	1	0.13	0.08
seq	41	35	259	249/24	246	5305	5187	5048	2797	15182	378	0.14	0.02
sym10	10	1	84	82	79	751	735	702	154	176	20	0.13	0.11
t3	12	8	25	25	24	209	214	216	5	10	1	0.20	0.10
t481	16	1	13	13	13	53	53	53	677	377	1	0.01	0.01
vg2	25	8	184	184	184	1992	2033	2010	163	1655	42	0.26	0.03
z4	7	4	29	29	29	145	148	133	4	9	3	0.75	0.33
Total			1891	1804	1791	18762	18300	17671	6253	21532	922	0.15	0.04
Gain			0	-87	-100	0	-462	-1085					

Table 4. Comparison of ESOP minimization results with EXORCISM-2 [] and EXORCISM-3 [].

Bench mark			Cubes				Literals			Time			Comparison	
Name	Ins	Outs	Ex2	Ex3	Ex4	Ex4b	Ex2	Ex3	Ex4	Ex2	Ex3	Ex4	w/Ex2	w/Ex3
add6	12	7	127	127	127	127	819	800	832	105	24	5	0.05	0.20
alu4	14	8	447	422	435	411	4816	4430	4520	6828	506	140	0.02	0.28
apex1	45	45	285	286	287	285	3796	3820	3998	4697	109	51	0.01	0.48
apex3	54	50	270	268	260	249	3916	4155	4322	2562	53	250	0.10	4.71
apex5	117	88	400	399	398	398	4038	4027	4065	20156	457	544	0.03	1.19
cps	24	109	135	135	140	135	2462	2625	3546	467	37	24	0.05	0.64
dule2	22	29	79	78	78	78	920	909	935	72	12	3	0.04	0.25
e64	65	65	65	65	65	65	2210	2272	2270	79	12	1	0.01	0.08
ex5	8	63	72	72	72	71	920	904	975	62	10	5	0.08	0.50
misex3	14	14	545	535	510	501	6837	6632	6141	18897	670	418	0.02	0.62
seq	41	35	245	248	247	246	4833	4822	5063	2996	77	52	0.02	0.68
spla	16	46	260	262	267	259	3420	3393	3969	1223	83	83	0.07	1.00
table3	14	14	166	166	166	166	2491	2491	2630	190	24	17	0.09	0.70
table5	17	15	156	156	156	156	2453	2449	2545	150	33	8	0.05	0.24
vg2	25	8	184	184	184	184	1993	1988	2017	225	29	8	0.04	0.28
Total			3436	3413	3392	333	46924	45717	47828	58709	2136	1609	0.02	0.75
Gain			0	-23	-44	-105	0	-1207	+904					