

Chapter 3

Sequential Logic Synthesis Using Symbolic Bi-decomposition

Victor N. Kravets and Alan Mishchenko

Abstract In this chapter we use under-approximation of unreachable states of a design to derive incomplete specification of combinational logic. The resulting incompletely specified functions are decomposed to enhance the quality of technology-dependent synthesis. The decomposition choices are computed implicitly using novel formulation of symbolic bi-decomposition that is applied recursively to decompose logic in terms of simple primitives. The ability of binary decision diagrams to represent compactly certain exponentially large combinatorial sets helps us to implicitly enumerate and explore variety of decomposition choices improving quality of synthesized circuits. Benefits of the symbolic technique are demonstrated in sequential synthesis of publicly available benchmarks as well as on the realistic industrial designs.

3.1 Introduction and Motivation

Due to recent advances in verification technology [2] circuit synthesis of semiconductor designs no longer has to be limited to logic optimization of combinational blocks. Nowadays logic transformations may involve memory elements which change design's state encodings or its reachable state space and still be verified against its original description. In this chapter we focus on a more conservative synthesis approach that changes sequential behavior of a design only in unreachable states, leaving its intended "reachable" behavior unchanged. Unreachable states are used to extract incomplete specification of combinational blocks and are applied as *don't cares* during functional decomposition to improve circuit quality.

To implement combinational logic of a design we rely on a very simple, yet complete, form of functional decomposition commonly referred to as bi-decomposition.

V.N. Kravets (✉)

IBM TJ Watson Research Center, Yorktown, NY

e-mail: kravets@us.ibm.com

Based on Kravets, V.N.; Mishchenko, A.; "Sequential logic synthesis using symbolic bi-decomposition," Design, Automation & Test in Europe Conference & Exhibition, 2009. DATE '09, pp.1458–1463, 20–24 April 2009 © [2009] IEEE.

In general, for a given completely specified Boolean function its bi-decomposition has form

$$f(\mathbf{x}) = h(g_1(\mathbf{x}_1), g_2(\mathbf{x}_2))$$

where h is an arbitrary 2-input Boolean function. This decomposition is not unique and its quality varies depending on selected subsets \mathbf{x}_1 and \mathbf{x}_2 that form possibly overlapping (i.e., non-disjoint) partition of \mathbf{x} . The problem of finding good bi-decomposition has been studied in [1, 10, 18, 19, 21]. The main contribution of the material in this chapter is symbolic formulation of bi-decomposition for incompletely specified functions. The bi-decomposition is used as main computational step in the prototype sequential synthesis tool and is applied recursively to implement logic of combinational blocks whose incomplete specification is extracted from unreachable states of a design. Our symbolic formulation of bi-decomposition finds all feasible solutions and picks the best ones, without explicit enumeration.

Computation of variable partitions in our symbolic formulation of bi-decomposition favors implicit enumeration of decomposition subsets. They are represented compactly with a binary decision diagram (BDD) [4] and are selected based on optimization objective. Unlike previous approaches (e.g., [1, 23]) that rely on BDDs, the decomposition is not checked explicitly for a variable partition and is solved implicitly for all feasible partitions simultaneously utilizing fundamental property of BDDs to share partial computations across subproblems. Thus, no costly enumeration that requires separate and independent decomposability checks is needed. The technique was also used to tune greedy bi-decomposition when handling larger functions.

To overcome limitations of explicit techniques authors in [14] proposed solution that uses a satisfiability solver [11]. Their approach is based on proving that a problem instance is unsatisfied. The unsatisfiable core is then used to greedily select partition of variables that induces bi-decomposition. Authors demonstrate the approach to be efficient in runtime, when determining existence of non-trivial decomposition. The experimental results on a selected benchmark set, however, are primarily focused on the existence of decomposition and do not offer a qualitative synthesis data.

The problem of using unreachable states of a design to improve synthesis and verification quality has been studied before in various contexts. In general, these algorithms either avoid explicit computation of unreachable states or first compute them in pre-optimization stage. Approaches that do not explicitly compute unreachable states are mostly limited to incremental structural changes of a circuit and rely on ATPG environment or induction [5, 8, 12] to justify a change. In contrast, approaches that pre-compute subsets of unreachable states treat them as external don't cares [20] for re-synthesis of combinational logic blocks [6, 15]. In this chapter we adopt the later approach as it offers more flexibility in logic re-implementation through functional decomposition.

This chapter has the following structure. After brief introduction and motivation preliminary constructs are given in Section 3.2. Section 3.3 describes

bi-decomposition existence requirements. They are used in Section 3.4 to formulate implicit computation of decomposition. Implementation details are described in Section 3.5. Experimental results are given in Section 3.6. Section 3.7 gives concluding remarks and possible directions for future work.

3.2 Preliminary Constructs

Basic constructs used by synthesis algorithms of the chapter are introduced in this section.

3.2.1 “Less-Than-or-Equal” Relation

Computational forms constructed in this chapter rely on the partial order relation between Boolean functions. Given functions $f(\mathbf{x})$ and $g(\mathbf{x})$, $f(\mathbf{x}) \leq g(\mathbf{x})$ indicates that $f(\mathbf{x})$ precedes $g(\mathbf{x})$ in the order. This “less-than-or-equal” relation (\leq) between the two functions can be expressed by one of the following three equivalent forms:

$$[f(\mathbf{x}) \Rightarrow g(\mathbf{x})] \equiv [f(\mathbf{x}) \leq g(\mathbf{x})] \equiv [\overline{f(\mathbf{x})} + g(\mathbf{x}) = 1]$$

The relation imposes *consistency* constraint on constructed computational forms. It allows us to represent incompletely specified Boolean functions in terms of *intervals* [3], defined as

$$[l(\mathbf{x}), u(\mathbf{x})] = \{f(\mathbf{x}) | l(\mathbf{x}) \leq f(\mathbf{x}) \leq u(\mathbf{x})\}$$

Here interval represents a set of completely specified functions using its two distinguished members $l(\mathbf{x})$ and $u(\mathbf{x})$, known as upper and lower bounds, respectively. It is non-empty (or consistent) if and only if $l(\mathbf{x}) \leq u(\mathbf{x})$ is satisfied.

Example 3.1 Consider interval $[\bar{x}y, x + y]$ which represents an incompletely specified function. It is composed of four completely specified functions: $\bar{x}y$, y , $x \oplus y$, and $x + y$. Each of them has a don’t care set represented by function x . \square

Application of existential quantification \exists and universal quantification \forall to lower and upper bounds of the interval enables convenient selection of its member functions that are vacuous, i.e., independent in certain variables.

Example 3.2 Consider abstraction of x from the interval in Example 3.1:

$$[\exists x(\bar{x}y), \forall x(x + y)]$$

The abstraction yields non-empty interval that is composed of a unique function that is vacuous in x : $[y, y]$. Abstraction of y , however, results in empty interval since the relation between its lower and upper bounds is not satisfied: $[\bar{x}, x]$ is empty. \square

We will use notation $\nabla_x[l(\mathbf{x}), u(\mathbf{x})]$ to represent abstraction $[\exists x l(\mathbf{x}), \forall x u(\mathbf{x})]$.

3.2.2 Parameterized Abstraction


To determine subsets of variables whose abstraction preserves consistency of a symbolic statement (or a formula) we use *parameterized abstraction* construct. It parameterizes computational form with a set of auxiliary decision variables c that are used to guide variable abstraction decisions. An assignment to c effects consistency of a computational form and thereby determines feasibility of abstracting a corresponding variable subset.

We use the “if-then-else” operator $ITE(c, x, y)$ to encode effect of quantifying variable subsets from a formula. Defined as $cx + \bar{c}y$, the operator selects between variables x and y depending on value of c . As stated, it provides a mechanism to parameterize signal dependencies in a Boolean function. It can be also generalized to the selection between functions. In particular, $ITE(c, f(\mathbf{x}), \exists x f(\mathbf{x}))$ encodes a decision of existential quantification of x from $f(\mathbf{x})$, similarly for the universal quantification.

Example 3.3 We can parameterize abstraction of variable x from interval $[\bar{x}y, x + y]$ using ITE operator and auxiliary variable c_x as

$$[ITE(c_x, \exists x(\bar{x}y), \bar{x}y), ITE(c_x, \forall x(x + y), x + y)]$$

or equivalently $[c_x y + \bar{c}_x(\bar{x}y), c_x y + \bar{c}_x(x + y)]$. Subsequent parameterization of y transforms lower bound $L_x = ITE(c_x, \exists x(\bar{x}y), \bar{x}y)$ into $L_{xy} = ITE(c_y, \exists y L_x, L_x)$. The effect of decisions on c variables then has a form depicted in the tree below:




Parameterization of universal quantification has analogous application to an upper bound of the interval, producing U_{xy} . \square

We rely on consistency of the “less-than-or-equal” relation to find decomposition of a Boolean function. In the example below it is illustrated determining feasible abstraction of the interval variable subsets, and in finding its member functions with smallest support (i.e., with fewest variables).

Example 3.4 For each assignment to c variables the consistency of relation \leq between interval bounds in Example 3.3 determines existence of a function that is independent of the corresponding variables. The tree expansion over decision variables shows that there are only two feasible abstractions, marked with \checkmark :

Of the two, there is only one non-trivial abstraction $\forall x[\bar{x}y, x + y]$ and it yields single member interval $[y, y]$. \square



In a typical computational form, consistent assignment to the decision variables must hold universally, independent of non-decision variables. We can therefore compute *characteristic function* of consistent assignments implicitly, universally abstracting non-decision variables.

Example 3.5 For the parameterized bounds L_{xy} and U_{xy} from Example 3.3 the characteristic function of consistent assignments in $[L_{xy}, U_{xy}]$ is computed implicitly as

$$\forall x, y [L_{xy}, U_{xy}] \equiv \forall x, y [\overline{L_{xy}} + U_{xy} = \mathbf{1}] \equiv \overline{c_x} \overline{c_y} + c_x \overline{c_y} \equiv \overline{c_y}$$

The computed function states that abstraction of variable y yields an empty interval. Therefore, the interval contains no function that is independent of y . □

3.3 Bi-decomposition of Incompletely Specified Functions

This section gives formal statement of bi-decomposition over 2-input decomposition primitives, namely OR and XOR.

3.3.1 OR Decomposition

For a completely specified function $f(\mathbf{x})$, the decomposition of this type is described in terms of equation below:

$$f(\mathbf{x}) = g_1(\mathbf{x}_1) + g_2(\mathbf{x}_2) \tag{3.1}$$

When function is incompletely specified with interval $[l(\mathbf{x}), u(\mathbf{x})]$ we need to make sure that *OR* composition $g_1 + g_2$ is a member function of the interval.

Let $\underline{\mathbf{x}}_1$ and $\underline{\mathbf{x}}_2$ be signal subsets in which decomposition functions g_1 and g_2 are, respectively, vacuous, i.e., are functionally independent. (The underline in $\underline{\mathbf{x}}_i$ indicates that the computed g_i is independent in these variables.) Vacuous in $\underline{\mathbf{x}}_1$ function g_1 must not exceed largest member $u(\mathbf{x})$ in all its minterm points, independent of $\underline{\mathbf{x}}_1$, i.e., relation $g_1(\mathbf{x}_1) \leq \forall \underline{\mathbf{x}}_1 u(\mathbf{x})$ must be satisfied. Otherwise g_1 is either not contained in the interval or not independent of $\underline{\mathbf{x}}_1$. Similarly, $g_2(\mathbf{x}_2) \leq \forall \underline{\mathbf{x}}_2 u(\mathbf{x})$

must hold. Thus, $\forall \underline{\mathbf{x}}_1 u(\mathbf{x})$ and $\forall \underline{\mathbf{x}}_2 u(\mathbf{x})$ give upper bounds on $g_1(\mathbf{x}_1)$ and $g_2(\mathbf{x}_2)$. To ensure that the selection of g_1 and g_2 is “large enough” the following must hold:

$$l(u) \leq \forall \underline{\mathbf{x}}_1 u(\mathbf{x}) + \underline{\mathbf{x}}_2 u(\mathbf{x}) \quad (3.2)$$

The OR composition does not exceed u universally due to the “reducing” effect of \forall on u . Thus, we can determine existence of the bi-decomposition limiting check to relation in (3.2). This check provides necessary and sufficient condition for the existence of OR decomposition and is a re-statement of the result from [17].

AND Decomposition. As indicated in [17], AND decomposition of f can be obtained from OR decomposition utilizing dual property of the two gates. For an incompletely specified function $[l, u]$ we can find complemented g_1 and g_2 by establishing OR decomposability of the interval complement, derived as $\overline{[l, u]} = [\overline{u}, \overline{l}]$.

3.3.2 XOR Decomposition

We first describe XOR decomposability condition for a completely specified function $f(\mathbf{x})$. To derive an existence condition for the XOR decomposition

$$f(\mathbf{x}) = g_1(\mathbf{x}_1) \oplus g_2(\mathbf{x}_2) \quad (3.3)$$

requires partitioning of \mathbf{x}_1 and \mathbf{x}_2 into finer subsets. Let $\underline{\mathbf{x}}_1$ and $\underline{\mathbf{x}}_2$ be subsets of variables in which g_1 and g_2 are, respectively, vacuous, and let \mathbf{x}_3 be a set of variables on which both decomposition functions depend. We can then state necessary and sufficient condition for the existence of XOR decomposition as follows:

Proposition 3.1 *XOR bi-decomposition*

$$f(\mathbf{x}) = g_1(\underline{\mathbf{x}}_2, \mathbf{x}_3) \oplus g_2(\underline{\mathbf{x}}_1, \mathbf{x}_3) \quad (3.4)$$

exists if and only if

$$f(\underline{\mathbf{x}}_1, \underline{\mathbf{x}}_2, \mathbf{x}_3) \neq f(\underline{\mathbf{y}}_1, \underline{\mathbf{x}}_2, \mathbf{x}_3) \quad (3.5)$$

↓

$$\forall \underline{\mathbf{y}}_2 [f(\underline{\mathbf{x}}_1, \underline{\mathbf{y}}_2, \mathbf{x}_3) \neq f(\underline{\mathbf{y}}_1, \underline{\mathbf{y}}_2, \mathbf{x}_3)] \quad (3.6)$$

We derived conditions in the above proposition when analyzing library requirements for an advanced technology [13]. In [14] authors recently and independently stated analogous proposition in terms of the unsatisfiability problem. We therefore show correctness of the above proposition giving only an information-theoretical argument: For (3.4) to hold, it must be that all onset/offset minterms in f that cannot be distinguished by g_1 (3.5) must be distinguished by g_2 (3.6).

For an incompletely specified function $[l(\mathbf{x}), u(\mathbf{x})]$ the consistency constrain (3.5) \Rightarrow (3.6) of Proposition 3.1 changes to

$$\begin{aligned} & [l(\underline{\mathbf{x}}_1, \underline{\mathbf{x}}_2, \mathbf{x}_3) \neq l(\underline{\mathbf{y}}_1, \underline{\mathbf{x}}_2, \mathbf{x}_3)] \wedge [h(\underline{\mathbf{x}}_1, \underline{\mathbf{x}}_2, \mathbf{x}_3) \neq h(\underline{\mathbf{y}}_1, \underline{\mathbf{x}}_2, \mathbf{x}_3)] \\ & \quad \downarrow \\ & \forall \underline{\mathbf{y}}_2 [[l(\underline{\mathbf{x}}_1, \underline{\mathbf{y}}_2, \mathbf{x}_3) \neq h(\underline{\mathbf{y}}_1, \underline{\mathbf{y}}_2, \mathbf{x}_3)] \vee [h(\underline{\mathbf{x}}_1, \underline{\mathbf{y}}_2, \mathbf{x}_3) \neq l(\underline{\mathbf{y}}_1, \underline{\mathbf{y}}_2, \mathbf{x}_3)]] \end{aligned}$$

The above statement extends containment relation (3.5) \Rightarrow (3.6) by reducing lower bound (3.5) and increasing upper bound (3.6) as much as possible. The relation provides the condition for XOR bi-decomposition of incompletely specified functions, previously unsolved problem [14].

3.4 Parameterized Decomposition

Section 3.3 decomposition checks assume that the \mathbf{x}_1 and \mathbf{x}_2 subsets are pre-determined. Finding such feasible subsets, however, may not be straightforward and depending on the objectives could potentially require an exponential search if performed explicitly. Our solution to the problem is to perform the search implicitly, formulating the problem symbolically and solving it by leveraging the capability of binary decision diagrams to compactly represent certain combinatorial subsets.

3.4.1 OR Parameterization

We use (3.2) to find feasible OR decompositions implicitly. It is used to construct a computational form that parameterizes the \forall operation applied to variables \mathbf{x} :

```

U ← u;
for each  $x \in \mathbf{x}$  do
    U ← ITE( $c_x$ , U,  $\forall x U$ );
end for

```

Such iterative parameterization gives function $U(\mathbf{c}, \mathbf{x})$ that encodes the effect of abstracting all variable subsets from u , where variable x is abstracted *iff* $c_x = 0$.

The parameterized function $U(\mathbf{c}, \mathbf{x})$ can be used in (3.2) to encode possible supports to g_1 and g_2 in terms of the decision variables \mathbf{c}_1 and \mathbf{c}_2 :

$$l(\mathbf{x}) \leq U_1(\mathbf{x}, \mathbf{c}_1) + U_2(\mathbf{x}, \mathbf{c}_2) \quad (3.7)$$

For any feasible assignment to \mathbf{c}_1 and \mathbf{c}_2 , the above relation must hold universally, irrespective of values on \mathbf{x} . Thus, computational form

$$\text{Bi}(\mathbf{c}_1, \mathbf{c}_2) \equiv \forall \mathbf{x} [\overline{l(\mathbf{x})} + U_1(\mathbf{x}, \mathbf{c}_1) + U_2(\mathbf{x}, \mathbf{c}_2)] \quad (3.8)$$

yields a characteristic function of all feasible supports for g_1 and g_2 : it evaluates to truth *iff* assignments to \mathbf{c}_1 and \mathbf{c}_2 induce feasible supports for g_1 and g_2 .

We illustrate potentially scalable nature of BDDs to handle computation in (3.8) decomposing multiplexer function for its various support sizes:

Max width		Bi computation		Best partition	
Control	Data	BDD size	Time(s)	(x_1 , x_2)	No. of Choices
2	4	23	0.00	(4, 4)	6
3	8	43	0.01	(7, 7)	70
4	16	79	0.09	(12, 12)	12870
5	32	147	1.35	(21, 21)	6E8.0
6	64	279	20.56	(38, 38)	1.8E18

The above table gives results of the computation in terms of multiplexer widths, BDD size and time required to compute Bi , and the best support sizes of g_1 and g_2 . As the table suggests, the amount of resources required in computation grows moderately for smaller problem instances and is tolerable even for a larger function.

We point out that the exhaustive computational form (3.8) could be relaxed to produce solution subsets to (3.7), instead of producing a complete solution. For example, in place of (3.8) a specialized satisfiability procedure could be used to produce solutions with additional optimization constraints. Specialized BDD-based abstraction techniques that monitor resource consumption could be also deployed to produce solution subsets. Another possibility is to rely on a greedy assignment selection to \mathbf{c}_1 and \mathbf{c}_2 targeting disjoint subsets \mathbf{x}_1 and \mathbf{x}_2 . More detailed discussion on selecting best \mathbf{x}_1 and \mathbf{x}_2 is given later, in Section 3.5.

3.4.2 XOR Parameterization

To simplify presentation we compute characteristic function of all feasible support partitions for a completely specified function. As before, encoding of possible supports for g_1 and g_2 is performed using two sets of auxiliary variables \mathbf{c}_1 and \mathbf{c}_2 . Using \mathbf{c}_1 , (3.5) is transformed into $f(\mathbf{x}) \neq F_1(\mathbf{x}, \mathbf{y}, \mathbf{c}_1)$, where F is derived from $f(\mathbf{x})$ replacing each of its variables x_i with $ITE(c_{1i}, x_i, y_i)$. Similarly, part $f(\underline{\mathbf{x}}_1, \underline{\mathbf{y}}_2, \mathbf{x}_3)$ from (3.6) is parameterized with \mathbf{c}_2 to construct $F'_2(\mathbf{x}, \mathbf{y}, \mathbf{c}_2)$. It encodes selection of vacuous variables for g_2 . The last component $f(\underline{\mathbf{y}}_1, \underline{\mathbf{y}}_2, \mathbf{x}_3)$ is transformed into $F''_2(\mathbf{x}, \mathbf{y}, \mathbf{c}_1, \mathbf{c}_2)$, replacing each variable in $f(\mathbf{x})$ with $ITE(c_{1i} \cdot c_{2i}, x_i, y_i)$. Universally abstracting \mathbf{x} and \mathbf{y} variables gives representation of all feasible supports for g_1 and g_2 :

$$Bi(\mathbf{c}_1, \mathbf{c}_2) \equiv \forall \mathbf{x}, \mathbf{y} [(f \neq F_1) \Rightarrow (F'_2 \neq F''_2)] \quad (3.9)$$

We compare implicit computation of decomposition choices to a greedy algorithm for the XOR decomposition, used by authors in [17, 22]. Starting from a seed partition, the algorithm greedily extends support subsets calling XOR decomposabil-

ity check in its inner loop. Although efficient in general, the check has potentially formidable runtime. The profile of its behavior on a 16-bit adder is given in the table below; it is compared against our implicit computation:

Sum bit	Output		Time(s)	
	No. of Inputs	Best part.	Implicit	[17] Check
s2	7	(2,5)	0.01	0.00
s4	11	(2,9)	0.06	0.13
s6	15	(2,13)	0.12	4.44
s8	19	(2,17)	0.13	71.05
⋮			⋮	
s16	33	(2,31)	0.42	Time out

For a subset of sum-bit functions the table lists runtime for both techniques. (The *best part.* column gives data generated by our implicit enumeration of feasible partitions.) Although not typical, it is interesting that where a rather efficient greedy check times out after an hour, an implicit exhaustive computation takes only 0.42 s.

In general, we can use best partition produced by the exhaustive implicit computation to evaluate and tune greedy algorithm or to improve some other approximate technique.

3.5 Implementation Details of Sequential Synthesis

This section describes a sequential synthesis flow that first extracts incompletely specified logic accounting for unreachable states in a design and then uses bi-decomposition to synthesize technology-independent circuit.

3.5.1 Extraction of Incompletely Specified Logic

Unreachable states of a design form don't cares for the combinational logic. Due to the complexity of computing unreachable states even in designs of modest size, incompletely specified combinational logic is extracted with respect to an approximation of unreachable states. Unlike other partitioning approaches that try to produce a good approximation of unreachable states in reasonable time [10, 16], our objective is to compute a good approximation with respect to support of individual functions. A similar approach to approximate unreachable states using induction was proposed in [7].

We perform state-space exploration with forward reachability analysis for overlapping subsets of registers. These subsets are selected using structural dependence of next-state and primary outputs on the design latches. The selection tries to create partitions maximizing accuracy of reachability analysis for present-state signals $supp_{ps}(f)$ output function f . In particular, the partitioning tries to meet the following goals:

- For each function f , present-state inputs $supp_ps(f)$ are represented in at least one partition.
- Each partition selects additional logic to maximize accuracy of reachability analysis.

After completing reachability analysis for a partition, an incomplete specification of signals that depend on the partition latches becomes available in the form of an interval notation. For each signal, its interval pre-processed with the ∇ operation eliminates vacuous variables, selecting a dependence on the least number of variables. The interval is then used for performing bi-decomposition. Figure 3.1 exemplifies OR bi-decomposition applied to function $f = \overline{a}b + \overline{a}c + bc$ of its output signal. The bi-decomposition of $[f \cdot \overline{a}bc, f + \overline{a}bc]$ finds OR decomposition of f in $g_1 = \overline{a}b + bc$ simplifying the circuit.

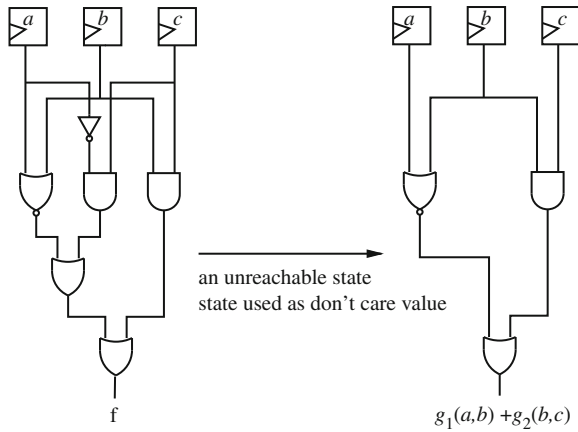


Fig. 3.1 Bi-decomposition with unreachable states. State $\overline{a}bc$ is used as a don't care condition to find OR decomposition that simplifies circuit

3.5.2 Exploring Decomposition Choices

The characteristic function Bi gives all feasible supports for decomposition functions. Since the provided variety of choices could be very large, we restrict them to a subset of desired solutions. The restriction targets minimization and balanced selection of supports in decomposition functions. It is achieved symbolically, as described below.

Let $w_i(\mathbf{c})$ be characteristic function of assignments to \mathbf{c} that have weight i (i.e., have exactly i decision variables set to 1). For a given $n = |\mathbf{c}|$ it represents combinatorial subsets $\binom{n}{i}$. This function has compact representation in terms of BDDs. Given a desired support size $k_1 = |\mathbf{x}_1|$ of g_1 , and of $k_2 = |\mathbf{x}_2|$ of g_2 , existence of the decomposition is determined constraining Bi with its corresponding solution space:

$$Bi(\mathbf{c}_1, \mathbf{c}_2) \cdot w_{k_1}(\mathbf{c}_1) \cdot w_{k_2}(\mathbf{c}_2)$$

If the resulting function is not empty, then the desired decomposition exists. To target balanced decomposition of a function we seek feasible k_1 and k_2 minimizing $\max(k_1, k_2)$. Such simultaneous minimization of k_1 and k_2 balances supports \mathbf{x}_1 and \mathbf{x}_2 , favoring their disjoint selection. (This is in contrast to [14], where different cost measures are used for each of the objectives.)

To avoid “trial-and-error” search for feasible support sizes we select desired k_1 and k_2 from a computed set of feasible pairs for g_1 and g_2 . Suppose function $\kappa_i(\mathbf{e})$ encodes integer k_i in terms of the \mathbf{e} variables. Then function $K(\mathbf{c}, \mathbf{e}) = \sum_{i=0}^n w_i(\mathbf{c}) \cdot \kappa_i(\mathbf{e})$ relates a decision variables assignment to integer encoding of the support size it induces. This function is used within the computational form

$$Bi_\kappa(\mathbf{e}_1, \mathbf{e}_2) \equiv \exists \mathbf{c}_1 \mathbf{c}_2 [Bi(\mathbf{c}_1, \mathbf{c}_2) \cdot K(\mathbf{c}_1, \mathbf{e}_1) \cdot K(\mathbf{c}_2, \mathbf{e}_2)]$$

to generate all feasible size pairs (k_1, k_2) .

The Bi_κ function should be post-processed to purge pairs that are dominated by other, better solutions. For example, pair (3, 5) is dominated by pair (3, 4) since it produces smaller distribution of supports between g_1 and g_2 . Let $\text{gte}(\mathbf{e}, \mathbf{e}')$ describe “greater-than-or-equal” relation between a pair of integers encoded with \mathbf{e} and \mathbf{e}' ; similarly, let equ describe the equality relation. We then define the dominance relation between bi-tuples $\varepsilon \equiv (\mathbf{e}_1, \mathbf{e}_2)$ and $\varepsilon' \equiv (\mathbf{e}'_1, \mathbf{e}'_2)$ as

$$\text{dom}(\varepsilon, \varepsilon') = \text{gte}(\mathbf{e}_1, \mathbf{e}'_1) \cdot \text{gte}(\mathbf{e}_2, \mathbf{e}'_2) \cdot \overline{\text{equ}(\mathbf{e}_1, \mathbf{e}'_1) \cdot \text{equ}(\mathbf{e}_2, \mathbf{e}'_2)}$$

Using this relation subtraction of the dominated solutions from Bi_κ is performed as

$$\forall \varepsilon' [Bi_\kappa(\varepsilon') \leq Bi_\kappa(\varepsilon) \cdot \overline{\text{dom}(\varepsilon, \varepsilon')}]$$

It states that if an assignment to ε' is in Bi_κ (left-hand side of the relation), then its dominated assignments to ε should be subtracted from Bi_κ (right-hand side of the assignment).

To complete decomposition of a function we need to find functions g_1 and g_2 . For the OR decomposition, possible functions g_1 and g_2 can be deduced directly from the corresponding existence condition (3.2), universally quantifying out variables in which g_1 and g_2 are vacuous. To construct XOR decomposition functions we use algorithm from [17].

3.5.3 Synthesis Algorithm

Our logic optimization algorithm selectively re-implements functions of circuit signals relying on bi-decomposition of extracted incompletely specified logic. The pseudocode code in Algorithm 1 captures general flow of the optimization.

Algorithm 1 Logic optimization loop

```

create latch partitions of a design;
selectively collapse logic;
while (more logic to decompose) do
  select a signal and its function  $f(\mathbf{x})$ ;
  retrieve unreachable states  $u(\mathbf{x})$ ;
  abstract vars from interval  $[f \cdot \bar{u}, f + u]$ ;
  apply bi-decomposition to interval;
end while

```

The algorithm first creates overlapping partitions of a design. These partitions are formed according to Section 3.5.2 and are typically limited to 100 latches. Additional connectivity cost measures are used to control size of a partition. For each partition computation of unreachable states is delayed until being requested by a function that depends on its present-state signals. BDDs for computed reachable states are then stored in a separate node space for each partition. When retrieving unreachable states for a given support, their conjunctive approximation is brought together to a common node space.

To re-decompose logic of a design the algorithm first creates functional representation for selected signals in terms of their cone inputs or in terms of other intermediate signals. The decision on whether to select a signal is driven by an assessed impact of bi-decomposition on circuit quality: if it has potential to improve variable partition, logic sharing, or timing over existing circuit structure, then signal is added to a list of re-decomposition candidates.

The logic of candidate signals is processed in topological order until it is fully implemented with simple primitives. This processing constitutes main loop of the algorithm. After a signal and its function $f(\mathbf{x})$ in the loop are selected, a set of unreachable states $u(\mathbf{x})$ are retrieved. This set is derived from reachability information of partitions that $f(\mathbf{x})$ depends on.

Before applying bi-decomposition to the incompletely specified function

$$[f(\mathbf{x}) \cdot \overline{u(\mathbf{x})}, f(\mathbf{x}) + u(\mathbf{x})]$$

the algorithm tries to abstract some of the interval variables while keeping it consistent; this eliminates redundant inputs. The bi-decomposition is then applied targeting potential logic sharing and balanced partition of \mathbf{x} , as described in Section 3.5.3. From a generated set of choices, partition that best improves timing and logic sharing is selected. Figure 3.2 illustrates bi-decomposition that benefits from logic sharing. The transformation re-uses logic of g_1 , which was present in the network but was not in the fanin of f .

3.6 Experimental Evaluation

From a suite of publicly available benchmarks we selected a subset of sequential circuits and assessed effect of unreachable states on bi-decomposition. Three types




Fig. 3.2 Bi-decomposition which benefits from re-using existing logic: node g_1 is shared in the $f = g_1 + g_2$ decomposition

of bi-decomposition were applied to functions of their output and next-state logic: OR, AND, and XOR. They are evaluated in terms of their ability to reduce maximum support of functions g_1 and g_2 . Experiments with and without reachable state-space analysis were performed.

The experimental results are given in Table 3.1. The table first lists circuit name, along with its corresponding number of inputs/outputs and latches. Each circuit was structurally pre-processed to remove cloned, dead, and constant latches. The #dec. column gives number of functions for which non-trivial decomposition was identified. The average ratio between maximum support sizes of g_1 and g_2 and support size of the function being decomposed is given in avg. reduct. column. Note that the reduction of less than 0.5 (as in s713 and s838) indicates that both g_1 and g_2 tend to be vacuous in some of the variables.

Table 3.1 Application of bi-decomposition to functions of next-state and output logic (without and with state analysis)

Name	Original circuit		No states		With states		
	Input/output	Latches	#Dec.	Avg. reduct.	\log_2 states	#Dec.	Avg. reduct
s344	10/11	15	18	0.781	12	18	0.634
s526	3/6	21	21	0.775	14	21	0.556
s713	36/23	19	40	0.652	11	40	0.453
s838	36/2	32	33	0.540	5	33	0.088
s953	17/23	29	29	0.607	13	29	0.565
s1269	18/10	37	39	0.672	31	39	0.671
s5378	36/49	163	145	0.609	125	145	0.603
s9234	36/39	145	97	0.754	141	97	0.774
Average reduction:				0.673			0.54

The results are collected for two experiments: with and without state-space information. The \log_2 of computed reachable states is also listed in the table. Computed average reduction ratios suggest that decomposability of a function improves as the number of unreachable states gets larger. The unreachable states did not contribute

much to s5378 largely because its logic is highly decomposable even in the absence of state-space information. The runtime to compute reachable states for each of the circuits did not exceed 1 min, requiring at most few seconds for circuits with 32 or less latches. Computation of bi-decomposition was limited to 1 min per circuit.

We evaluate our Section 3.5.3 Algorithm 1 synthesizing technology-independent netlists for a set of macro-blocks of a high-performance industrial design. Results of the netlists optimized with bi-decomposition are given in Table 3.2. First four columns list general parameters of each circuit, including number of gates it has in its and/inv expansion. The circuits were first pre-processed using our in-house tool, by optimizing it against publicly available mcnc.genlib library.

Table 3.2 Results of applying bi-decomposition in synthesis of industrial circuits

Name	Original circuit			Pre-processed		Algor. 1	
	Input/Output	Latches	AND	Area	Delay	Area	Delay
seq4	108/202	253	1845	3638	44.8	2921	41.9
seq5	66/12	93	925	1951	47.2	1807	41.6
seq6	183/74	142	811	1578	34.9	1487	36.0
seq7	173/116	423	3173	6435	52.4	5348	48.3
seq8	140/23	201	2922	6183	50.1	5427	48.8
seq9	212/124	353	3896	8250	56.0	6938	45.2
Average reduction:						0.88	0.94

An implementation of the algorithm was then applied to improve each of the circuits. Columns *Pre-processed* and *Algor. 1* compare area (which corresponds to the number of literals) and delay (estimated with a load-dependent model) of mapped netlists before and after running our algorithm. The additional area and timing savings are due to the algorithm, with the average area and delay reductions of 0.88 and 0.94, respectively. We attribute these gains to the algorithm’s ability to implicitly explore reach arsenal of decomposition choices during bi-decomposition. Optimization of each circuit was completed within 4 min of runtime.

3.7 Conclusions and Future Work

Extraction of incompletely specified logic using under-approximation of unreachable states in sequential designs offers valuable opportunity for reducing the circuit complexity. We developed a novel formulation of symbolic bi-decomposition and showed that the extracted logic has better implementation, with substantial area and delay improvements. The introduced symbolic bi-decomposition computes decomposition choices implicitly and enables their efficient subsetting using BDDs. Selecting best decomposition patterns during synthesis, we improved circuit quality of publicly available and realistic industrial design. We are currently working on ways to further maximize logic sharing through bi-decomposition and to apply it in a re-synthesis loop of well-optimized designs.

References

1. Ashenurst, R.L.: The decomposition of switching functions. *Annals of Computation Laboratory, Harvard University* **29**, 74–116 (1959)
2. Baumgartner, J., Mony, H., Paruthi, V., Kanzelman, R., Janssen, G.: Scalable sequential equivalence checking across arbitrary design transformations. In: *Proceedings of ICCD, San Jose, CA*, pp. 259–266. (2006)
3. Brown, F.M.: *Boolean Reasoning*. Kluwer, Boston, MA (1990)
4. Bryant, R.E.: Graph-based algorithms for Boolean function manipulation. *IEEE Transactions on Computers* **C-35**(6), 677–691 (1986)
5. Case, M.L., Kravets, V.N., Mishchenko, A., Brayton, R.K.: Merging nodes under sequential observability. In: *Proceedings of DAC, Anaheim, CA*, pp. 540–545. (2008)
6. Case, M.L., Mishchenko, A., Brayton, R.K.: Inductive finding a reachable state-space over-approximation. In: *IWLS, Vail, CO*, pp. 172–179. (2006)
7. Case, M.L., Mishchenko, A., Brayton, R.K.: Cut-based inductive invariant computation. In: *IWLS, Lake Tahoe, CA*, pp. 172–179. (2008)
8. Cheng, K.T., Entrena, L.A.: Sequential logic optimization by redundancy addition and removal. In: *Proceedings of ICCAD, San Jose, CA*, pp. 310–315. (1993)
9. Cho, J., Hachtel, G., Macii, E., Poncino, M., Somenzi, F.: Automatic state decomposition for approximate FSM traversal based on circuit analysis. *IEEE Transactions on CAD* **15**(12), 1451–1464 (1996)
10. Cortadella, J.: Timing-driven logic bi-decomposition. *IEEE Transactions on CAD* **22**(6), 675–685 (2003)
11. Een, N., Sorensson, N.: An extensible SAT-solver. In: *Proceedings of SAT, Santa Margherita Ligure, Italy*, pp. 502–518. (2003)
12. van Eijk, C.: Sequential equivalence checking based on structural similarities. *IEEE Transactions on CAD* **19**(7), 814–819 (2000)
13. Kravets, V.N. et al.: Automated synthesis of limited-switch dynamic logic (LSDL) circuits. *Prior Art Database (ip.com)* (March 2008)
14. Lee, R.-R., Jiang, J.-H., Hung, W.-L.: Bi-decomposing large Boolean functions via interpolation and satisfiability solving. In: *Proceedings of DAC, Anaheim, CA*, pp. 636–641. (2008)
15. Lin, B., Touati, H., Newton, R.: Don't care minimization of multi-level sequential networks. In: *Proceedings of ICCAD, San Jose, CA*, pp. 414–417. (1990)
16. Mishchenko, A., Case, M.L., Brayton, R.K., Jang, S.: Scalable and scalable-verifiable sequential synthesis. In: *Proceedings of ICCAD, San Jose, CA*, pp. 234–241. (2008)
17. Mishchenko, A., Steinbach, B., Perkowski, M.: An algorithm for bi-decomposition of logic functions. In: *Proceedings of DAC, Las Vegas, NV*, pp. 103–108. (2001)
18. Roth, J.P., Karp, R.: Minimization over Boolean graphs. *IBM Journal of Research and Development* **6**(2), 227–238 (1962)
19. Sasao, T., Butler, J.: On bi-decomposition of logic functions. In: *IWLS, Tahoe City, CA*, (1997)
20. Savoj, H., Brayton, R.K.: The use of observability and external don't cares for the simplification of multi-level networks. In: *Proceedings of DAC, Orlando, FL*, pp. 297–301. (1990)
21. Stanion, T., Sechen, C.: Quasi-algebraic decomposition of switching functions. In: *Proceedings of the 16th Conference on Advance Research in VLSI, Ann Arbor, MI*, pp. 358–367. (1998)
22. Steinbach, B., Wereszczynski, A.: Synthesis of multi-level circuits using EXOR-gates. In: *Proceedings of IFIP WG 10.5 – Workshop on Application of the Reed-Muller Expansion in Circuit Design, Chiba City, Japan*, pp. 161–168. (1995)
23. Yang, C., Cieselski, M., Singhal, V.: BDS: A BDD-based logic optimization system. In: *Proceedings of DAC, Los Angeles, CA*, pp. 92–97. (2000)