

Design Space Exploration as Quantified Satisfaction

Alexander Feldman, Johan de Kleer, Ion Matei

e-mail: a.feldman, dekleer, imatei@parc.com

Palo Alto Research Center Inc.

3333 Coyote Hill Road, Palo Alto, CA 94304, USA

Abstract

We propose novel algorithms for design and design space exploration. The designs computed by these algorithms are compositions of function types specified in component libraries. Our algorithms reduce the design problem to quantified satisfiability and use advanced solvers to find solutions that represent useful systems.

The algorithms we present in this paper are sound and complete and are guaranteed to discover correct designs of optimal size, if they exist. We apply our method to the design of Boolean systems and discover new and more optimal classical and quantum circuits for common arithmetic functions such as addition and multiplication.

The performance of our algorithms is evaluated through extensive experimentation. We have first created a benchmark consisting of specifications of scalable synthetic digital circuits and real-world microchips. We have then generated multiple circuits functionally equivalent to the ones in the benchmark. The quantified satisfiability method shows more than four orders of magnitude speed-up, compared to a generate and test method that enumerates all non-isomorphic circuit topologies.

Our approach generalizes circuit optimization. It uses arbitrary component libraries and has applications to areas such as digital circuit design, diagnostics, abductive reasoning, test vector generation, and combinatorial optimization.

Keywords: design, design space exploration, quantified satisfiability, Boolean circuit design, algorithmics

1. Introduction

Design is a next frontier in artificial intelligence. Providing algorithms and tools for conceiving novel designs benefits many areas such as analog and digital chip design, software development, mechanical design, and systems engineering. Human designers will be assisted in better navigating complex trade-offs such as speed versus number of transistors versus heat dissipation in an Integrated

Circuit (IC). Users will choose from a richer base of trade-offs and this will lead to dramatic improvements in micro-electronics and computing.

10 Computation, representation, and tools have improved tremendously over the last decades so now, one can consider systematic enumeration of the design space. This paper provides a novel encoding scheme for efficient exploration of the design space of digital circuits.

15 The algorithms presented in this paper are more computationally intensive compared to heuristic search (Hansen and Zhou, 2007) and genetic algorithms (Miller et al., 2000) but provide sound and complete enumeration of the design space. Our algorithms exhaustively “prove” that certain designs can or cannot be made with k components where components are drawn from an arbitrary component library.

20 Traditional books on digital design, for example, teach the construction of a full-subtractor with seven components (Maini, 2007) and we found one with only five gates. The five and seven component version of the subtractor will have the same number of transistors but there are other technologies (such as 3-D, or quantum) where the five-component version will have smaller footprint and faster propagation times.

25 As a special case, the circuit generation algorithm presented in this paper, reduces to circuit minimization but its performance should not be compared to other optimization algorithms such as Quine-McCluskey (McCluskey, 1956) or ESPRESSO (Brayton et al., 1984). To illustrate the generality of our approach we have used it to design a reversible quantum circuits of minimal size (Nielsen and Chuang, 2010).

30 Modern satisfiability (SAT) theory (Biere et al., 2009) is widely used in research and in industry. There are SAT solvers that can solve industrial problems with millions of variables (Järvisalo et al., 2012). The algorithms in this paper construct circuit designs by solving Quantified Boolean Formulas (QBFs). QBF satisfiability is a generalization over satisfiability of propositional formulas where universal and existential quantifiers are allowed. The QBFs that our algorithms generate are of interest to designers of quantified satisfiability (QSAT) algorithms as there is always the need of benchmarks with practical applications (Janota et al., 2016a).

40 The algorithms of this paper are validated on an extensive benchmark of combinational circuits with more than seventy successful experiments. We have designed generators of combinational circuits of various size such as adders, multipliers, and multiplexers. These circuits are the basic building blocks of Arithmetic Logic Units and Field Programmable Gate Arrays (FPGAs). In addition to that we consider four digital Integrated Circuits (ICs) from the well-known 74XXX family. We have shown that our QBF-based circuit generation algorithm is multiple orders of magnitude faster compared to a graph-based generate and test algorithm to find minimal circuits.

2. Design Generation and Exploration

50 Technical designs materialize from requirements, specifications, and the designers' experience. The design process is iterative with versions continuously improving and being refined. Incomplete designs often do not meet the requirements and designers "debug" and fix them. The formal underlying problem behind finishing an incomplete design (Gitina et al., 2013) has been studied in
55 the logic and verification communities (Finkbeiner and Tentrup, 2014). In addition to producing an initial design from scratch or continuing an incomplete one, designers often create multiple alternatives for the users and builders to choose from. The later process is called design exploration.

A design is typically specified in some kind of requirements. Depending on
60 the design domain, the requirements can be a mechanical blueprint, an electrical diagram, algorithmic pseudo-code or human readable text. To automate the generation and enumeration of designs, which is the main goal of this paper, we need some formal specification of a function or a design itself.

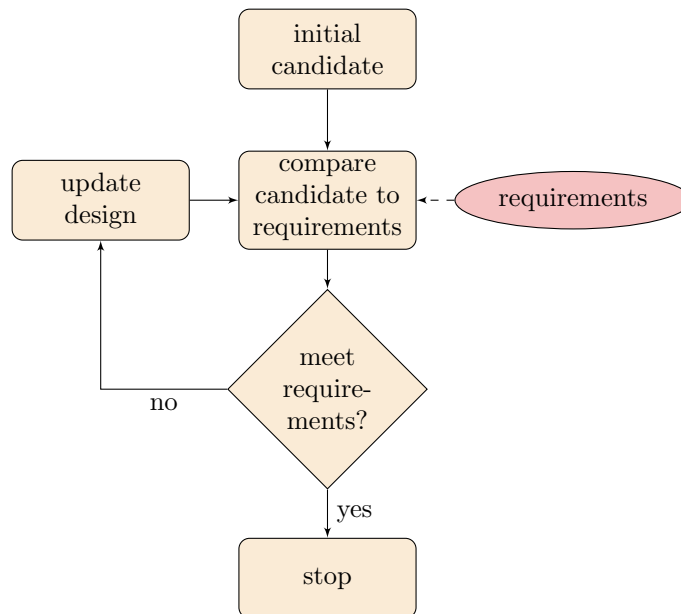


Figure 1: The design process as "generate and test"

Figure 1 illustrates the design generation process. The process is usually supported by Computer Aided Design (CAD) tools, Artificial Intelligence (AI), and combinatorial optimization algorithms. In some cases it is possible to consider
65 the whole design space and completely exhaust the search. Complete algorithms for design and design exploration are the subject of this paper.

The information flow in solving a design problem is shown in Figure 2.
70 The component library (basis) is specified as a set of Boolean functions. An

automated procedure is then used to generate a regular fabric of configurable components and topological interconnections (wires). The configurable fabric is appended to the user requirements which are also specified as a Boolean circuit or a Boolean function. The result is a miter: a formula that checks for Boolean function equivalence. The miter formula is fed to a QBF solver. The QBF solver computes a certificate that contains the configuration of the fabric. The final design is constructed from the certificate of the miter formula.

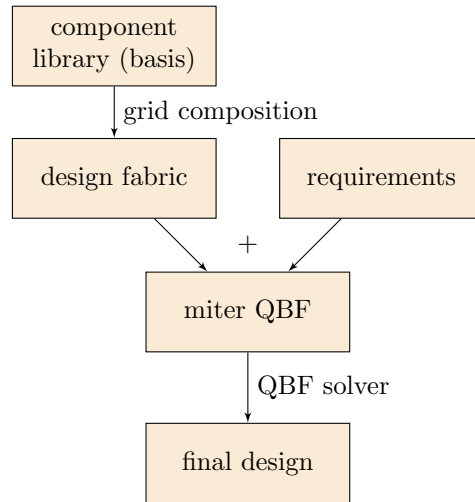


Figure 2: Information flow during the design process

There is only one computationally intensive step in generating a design: solving the QBF miter formula. Finding a satisfiable solution of a QBF is relevant to both satisfiability and game theory and is a prototypical PSPACE-complete problem (Garey and Johnson, 1990).

Consider an arbitrary QBF formula:

$$Q_1x_1Q_2x_2\dots Q_nx_n\varphi(x_1,x_2,\dots,x_n) \quad (1)$$

where Q_1, Q_2, \dots, Q_n are either existential (\exists) or universal (\forall) quantifiers. It can be decided if a formula is true or not by iteratively “unpeeling” the outermost quantifier until no quantifiers remain. If we condition on the value of the first quantifier, we have:

$$A = Q_2x_2\dots Q_nx_n\varphi(0,x_2,\dots,x_n) \quad (2)$$

$$B = Q_2x_2\dots Q_nx_n\varphi(1,x_2,\dots,x_n) \quad (3)$$

The formula is then reduced to $A \wedge B$ if Q_1 is \forall and $A \vee B$ if Q_1 is \exists . This process of recursive formula evaluation resembles a game where alternating the quantifier types forces the solver between making the solver look for primal and dual solution of the formula φ .

The recursive procedure suggested above is inefficient. Modern QBF solvers Janota et al. (2016b) use advanced search methods such as QCDCL (Quantified Conflict-Driven Clause Learning). QBF solvers benefit from knowledge compilation such as OBDD (Coste-Marquis et al., 2005), conflict learning, and even
95 machine learning (Samulowitz and Memisevic, 2007). Some solvers cater to a subclass of QBF formulas such as 2-QBF where there is only one switch between existential and universal quantifiers, others (Janota, 2018) are non-clausal and take directly circuits as their input.

Looking deeper, the QBF solving process resembles the high-level generate
100 and test process of design. Although it is not trivial to reduce design generation and exploration to solving a QBF, in this paper we manage to do that and use the advances in QBF solving to discover novel circuits or circuit topologies.

3. Fundamental Concepts

Definitions 1–3 are directly adopted from Vollmer (2013) and formally
105 introduce the notions of a Boolean function and a Boolean circuit.

Definition 1 (Boolean Function). A multi-output Boolean function is a function $f : \{0, 1\}^m \rightarrow \{0, 1\}^n$ for some $\{m, n\} \in \mathbb{N}$.

Notice that, while in Vollmer (2013) a Boolean function has a single output, we do not have this restriction. Another difference is that we do not use function
110 families, i.e., all our objects are finite.

Some common Boolean functions are negation (\neg), disjunction (\vee), conjunction (\wedge), exclusive or (\oplus), implication¹ (\rightarrow), and equivalence (\leftrightarrow). This paper uses everywhere infix, as opposed to prefix, notation. For example, $p \vee q$ is used instead of $\vee(p, q)$.

We also use equivalence (\leftrightarrow) instead of the equal sign ($=$) to specify Boolean
115 functions. The function output is on the left while the inputs are on the right. For example, the Boolean function $r = p \vee q$ is written as $r \leftrightarrow p \vee q$. When there are multiple outputs, we give a formula for each one of them.

Figure 3 shows the Boolean function $f \leftrightarrow \neg x \wedge y \vee x \wedge \neg y$ as a tree. Notice
120 that only the leaf nodes are variables while all non-leaves are operators.

Definition 2 (Basis). A basis B is defined as a finite set of Boolean functions.

Later in this section we discuss the fine differences between a Boolean circuit and a Boolean function as the two concepts are similar in many ways. One of the most important differences is that circuits use bases while functions do not.
125 A basis B can be thought of as the elementary unit of sharing or as an abstract **component library**. Unlike in the real world, though, each basis function can be used infinitely many times and all functions in a basis have the same cost. Figure 4 shows a basis consisting of typical unary and binary Boolean functions.

¹This paper, similar to many others, shares the same symbol (\rightarrow) for implication and for function mapping. The use is clear from the context.




Figure 3: An example of a Boolean function

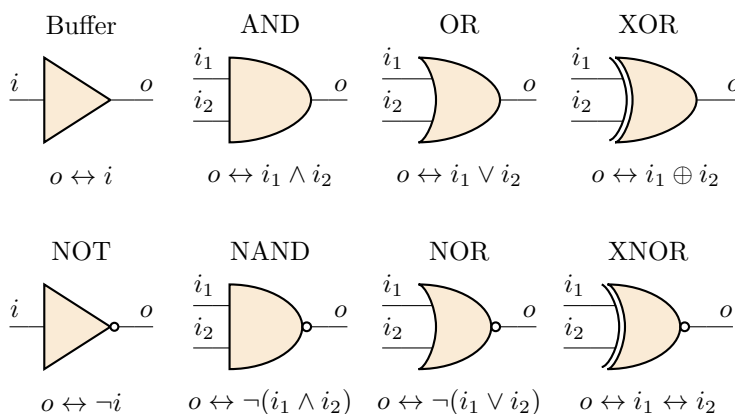


Figure 4: The standard basis

Figure 5 shows bases with multi-input/multi-output components. Figure 5a shows a basis consisting of two multi-output functions. They implement the Fredkin and the Toffoli gates (Fredkin and Toffoli, 1982; Toffoli, 1980). These gates, also known as CSWAP and CCNOT gates, have application in reversible and quantum computing.

Figure 5b shows a basis that contains one component only: a one-bit comparator. Sorting networks are made of chains of comparators. Proving lower-bounds on the number of comparators necessary for the building of a k -input sorting network is an ongoing challenge (Codish et al., 2014). The methods described in this paper provide novel methods for the optimal design and analysis of sorting networks.

It is possible to construct an “if-then-else” basis from the function shown in Figure 6 and the two Boolean constants (\top and \perp). If a circuit uses this base and the output of each gate is connected to exactly one input of another gate, then the problem of synthesizing minimal Binary Decision Diagrams Akers (1978) can be cast as circuit design.

It is possible to work with higher-level components. In the design of an Arithmetic-Logic Unit (ALU), for example, one can consider a basis extending

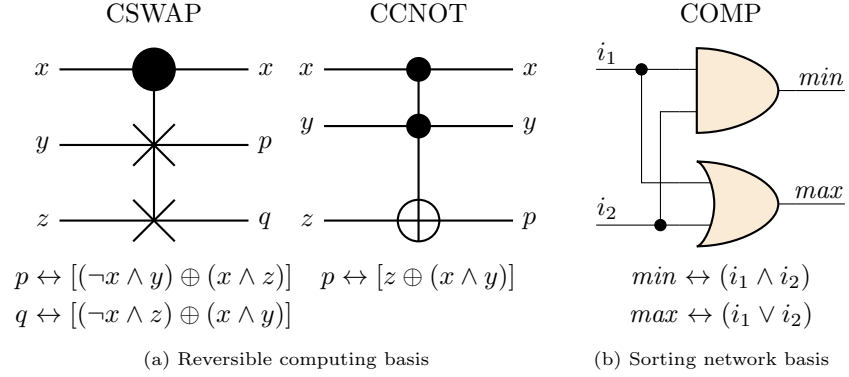


Figure 5: Non-standard bases




Figure 6: The “if-then-else” basis

the standard gates with multi-bit adders, multipliers, barrel shifters, etc.

Definition 3 (Boolean Circuit). Given a basis B , a Boolean circuit C over B is defined as $C = \langle V, E, \alpha, \beta, \chi, \omega \rangle$, where $\langle V, E \rangle$ is a finite directed acyclic graph, $\alpha : E \rightarrow \mathbb{N}$ is an injective function, $\beta : V \rightarrow B \cup \{\star\}$, $\chi : V \rightarrow \{x_1, x_2, \dots, x_n\} \cup \{\star\}$, and $\omega : V \rightarrow \{y_1, y_2, \dots, y_m\} \cup \{\star\}$. The following conditions must hold:

1. If $v \in V$ has an in-degree 0, then $\chi(v) \in \{x_1, x_2, \dots, x_n\}$ or $\beta(v)$ is a 0-ary Boolean function (i.e., a Boolean constant) in B ;
2. If $v \in V$ has an in-degree $k > 0$, then $\beta(v)$ is a k -ary Boolean function from B ;
3. For every $i, 1 \leq i \leq n$, there is exactly one node $v \in V$ such that $\chi(v) = x_i$;
4. For every $i, 1 \leq i \leq m$, there is exactly one node $v \in V$ such that $\omega(v) = y_i$.

The function α determines the ordering of the edges that go into a node when

160 the ordering matters (such as in implication). The function α is not necessary if B consists of symmetric functions only.

The function β determines the type of each node in the circuit: a function in the basis B . The function χ specifies the set of input nodes $\{x_1, x_2, \dots, x_n\}$. The function ω specifies the set of output nodes $\{y_1, y_2, \dots, y_n\}$. A node v is 165 non-output, or computational, if $\chi(v) = \star$ and $\omega(v) = \star$.

Figure 7 shows a simple and frequently used circuit that is used for adding the two binary numbers i_1 and i_2 and a carry input bit c_i . The output is found in the sum bit Σ and in the carry output c_o . Notice that there are two identical subcircuits in Figure 7. These are the two half-adders.




Figure 7: A full-adder

170 Figure 8 shows another circuit that is used for subtracting two binary numbers (i_1 and i_2) and a borrow input bit b_i . The output nodes are the difference d and the borrow output b_o .




Figure 8: A full-subtractor

The circuits shown in Figure 7 and Figure 8 use the standard basis. They are used as running examples for the rest of the paper.

175 Notice that, in a circuit, we use the term gate instead of component. Also, in a circuit the output of each gate is connected to the inputs of one or more other gates, i.e., a gate drives multiple other gates. The number of gates that are connected to a certain output is the gate's *fan-out*.

The size of a circuit is the number of gates.

180 In a Boolean function the result of an operator can be used as an argument of only one another operator. Fan-out does not make sense in a Boolean function. Of course, while it is possible to create an equivalent Boolean function for a circuit with gates with fan-out of more than one, it would require the introduction of new variables and operators. If we measure the size of the Boolean
185 function as the number of operators, then circuits with gates with fan-out of more than one will require fewer wires (variables) and gates. Alternatively, a circuit distinguishes between which variable is a primary input and which not, while in a (single-output) Boolean function a variable is an input.

From a higher-level standpoint, the main difference between Boolean functions and circuits is that **function sharing** is only supported in circuits. It is possible and straightforward to convert a circuit to an equivalent Boolean function but the number of operators in the Boolean function is often larger than the number of gates in the circuit. The full-adder shown in Figure 7, for example, requires at least six operators: $\Sigma \leftrightarrow i_1 \oplus i_2 \oplus c_i$ and $c_o \leftrightarrow i_1 \wedge i_2 \vee (i_1 \oplus i_2) \wedge c_i$.
190 The XOR gate that adds i_1 and i_2 is used both in calculating the sum Σ and the carry-output bit c_o . In some pathological cases, the blow-up can be exponential. The other direction is trivial: all Boolean functions are also circuits with all gates having a fan-out of exactly one.

Sometimes we would like to talk about how the nodes in a circuit are connected, without concerning ourselves with the exact function of each node. This
200 is referred to as the **topology** of a circuit.

Definition 4 (Topology). Given a circuit $C = \langle V, E, \alpha, \beta, \chi, \omega \rangle$, the topology of C is defined by the C sub-tuple $G = \langle V, E, \chi, \omega \rangle$.

The graph in Figure 9 shows the topology of the full-adder circuits shown in
205 Figure 7. There are three types of nodes: the input nodes i_1 and i_2 , the internal nodes that correspond to gates, and the output nodes Σ and c_o .




Figure 9: Full adder topology

The main purpose of this paper is to present an algorithm for synthesizing

circuits of minimal size.

4. Component Selection Problems and the Universal Component Cell

210 Suppose we are given a basis B , a topology $G = \langle V, E, \chi, \omega \rangle$, and a requirements circuit ψ . The purpose of our first algorithm is, given B , G , and ψ to create a circuit φ , such that $\varphi \equiv \psi$.

215 Consider the full-adder from Figure 7 as the requirements circuit ψ . Obtaining the topology G from ψ is trivial as the circuit topology is a sub-tuple of the circuit (see Definition 4). Let B be the standard basis shown in Figure 4. Given that the requirements circuit, itself, uses B , there exists at least one full-adder that uses the standard basis: that is the requirements ψ , itself. It is the trivial solution. We will see that there also exist multiple non-trivial solutions.

220 Figure 10 shows an alternative, non-trivial, implementation φ of the full-adder ψ with gates different from the ones in Figure 7. Instead of using two AND-gates, two XOR-gates, and an OR-gate, the alternative implementation makes two identical subsystems, each one containing an OR-gate and an XNOR-gate. The final carry output bit is computed by an AND-gate.




Figure 10: An alternative implementation of a full-adder

225 We can think of the circuit shown in Figure 10 as a symmetrical equivalence of the circuit shown in Figure 7. In what follows, we present an algorithm that computes and counts these symmetric circuit alternatives. This algorithm, based on QBF, is surprisingly efficient. We will see in the empirical results of Section 8 that circuits implementing common arithmetic and logical operations have many “deep” symmetries.

230 **Problem 1** (Component Selection Problem). Given a basis B , topology $G = \langle V, E, \chi, \omega \rangle$, and requirements ψ , construct a circuit $\varphi = \langle V, E, \alpha, \beta, \chi, \omega \rangle$, such that $\varphi \equiv \psi$.

Problem 1 is concerned with finding the type of each component in φ , or automatically specifying the functions α , and β . In some papers (Haaswijk

235 et al., 2018), Problem 1 is referred to as “labeling” because one can think of the type of a gate as a label in a graph-like topology.

A design exploration problem is to count all possible circuit implementations. Counting has little practical application on its own but the count is an important factor that characterizes the performance of circuit synthesis.

240 **Problem 2** (Counting Component Selection Configurations). Given a basis B , topology $G = \langle V, E, \chi, \omega \rangle$, and requirements ψ , count the number of distinct circuits $\varphi_i = \langle V, E, \alpha_i, \beta_i, \chi, \omega \rangle$, $1 \leq i \leq n$, such that $\varphi_i \equiv \psi$.

A naïve approach to solving Problems 1 and 2 is to consider all possible combinations of component types. There is, of course, the need to perform
245 an equivalence check for each combination of components and there are exponentially many combinations. Equivalence checking is a coNP-hard problem but it is often easy in practice (Matsunaga, 1996). The problem of equivalence checking has been largely solved either by using compilation to Ordered Binary Decision Diagrams (OBDDs) as proposed by Bryant (1986) or through resolution methods (Marques-Silva and Glass, 1999).
250 Despite the practical ease of equivalence checking, solving any instance of Problem 1 would still require an exponential number of coNP-hard calls.

The main idea behind our approach for solving Problems 1 and 2 is the universal component cell: a component that introduces extra selector inputs
255 allowing the choosing of which basis operation to perform. Connecting multiple cells according to the user-specified topology allows the extraction of one solution of Problem 1 from the return value of a single QBF solver call.

4.1. The Universal Component Cell

The universal component cell is a Boolean circuit that can be configured to
260 perform as any of the functions in a basis B . It is shown in Figure 11.

In Figure 11 there is a component c_1, c_2, \dots, c_k for each component of the basis. Suppose that each component of the basis has m inputs and n outputs. All outputs go to a set of n multiplexers m_1, m_2, \dots, m_n .

The configuration of the universal cell is a binary value assigned to a vector
265 of selector lines S . The number of selector inputs is $|S| = \lceil \log_2 n \rceil$ where n is the number of distinct component types in the basis. The actual routing is done by variable-size multiplexer circuits similar to the ones shown in Figure 12.

Figure 12 shows a multiplexer of variable size. Suppose there are n alternative gates and $|S| = \lceil \log_2 n \rceil$ selector lines. The multiplexer needs n multi-input
270 AND-gates and $|S|$ inverters. All AND-gates have $|S| + 1$ inputs. The multiplexer also uses an OR-gate with n inputs. The space complexity of the circuits is $O(|S| \times n)$ when multi-input gates are realized with ladders of two-input ones.

When constructing the cells, we take special care if the components in B have different numbers of inputs and outputs and if $|B| \neq 2^k$, $k \in \mathbb{N}$. The
275 special care is that we augment the miter circuit with gates that “disable” these hanging wires.

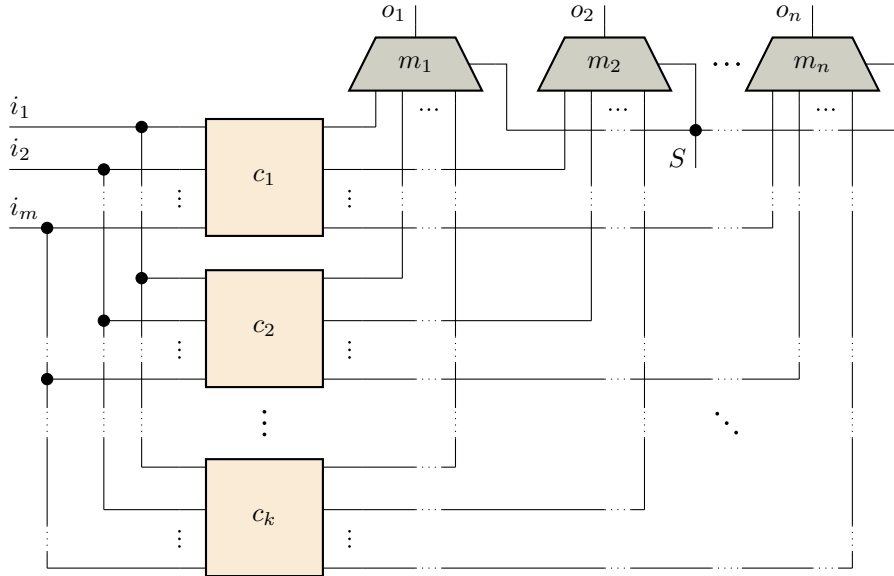


Figure 11: The universal component cell

4.2. An Efficient QBF-Based Algorithm

In what follows we reduce Problem 1 to finding a satisfiable solution of a QBF problem. Most QBF solvers, in addition to determining if a given QBF is satisfiable or not, also compute a partial certificate, or a witness: an assignment to the variables in the outermost quantifier that satisfies or invalidates the formula. We use this assignment for constructing the solution of our problem. The circuit whose partial certificate is a solution of Problem 1 is shown in Figure 13.

The two subcircuits shown in Figure 13 illustrate the concept of a miter (Brand, 1993). The miter is constructed from the requirements circuit ψ and a circuit φ which uses the topology of ψ and instead of gates has universal cells. The corresponding pairs of primary inputs of φ and ψ are joined together and the primary outputs are connected to XNOR gates whose outputs are connected to the constant \top .

The miter is used for equivalence checking. The basic idea of a miter is to pairwise tie all inputs and outputs of the two circuits together and to verify satisfiability. The resulting inputs are $X = \{x_1, x_2, \dots, x_n\}$ and the outputs are $Y = \{y_1, y_2, \dots, y_n\}$.

The subcircuit on the left side of Figure 13 has universal component cells only. The selector lines of all universal component cells make the variable set S . The solution of Problem 1 is an assignment to all S -variables. All internal variables of the requirements circuit ψ and all internal variables of the universal component cell go in the variable set Z .

The circuit that contains the universal cells and the requirements is constructed by the CREATEMITER subroutine of Algorithm 1. The function copies

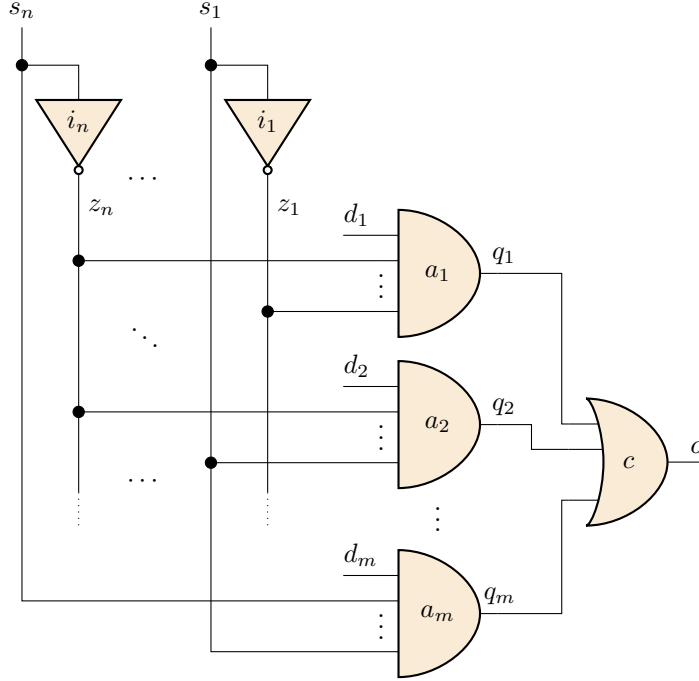


Figure 12: Variable size multiplexer circuit

the requirements circuit ψ under a new name φ ties together each pair of corresponding primary inputs and outputs and replaces all components in φ with universal cells. Each universal cell switches between components in B .

305 The method SOLVEQBF is the actual invocation of the QBF solver. In the case of non-clausal solvers (Lonsing and Egly, 2017; Janota et al., 2016b), one can directly feed the miter as an input. If the solver is clausal, a conversion to quantified Conjunctive Normal Form is needed (CNF). This conversion typically introduces a new set of variables that can affect the performance of the solvers. Clausal QBF solvers benefit from preprocessing the input formula with
 310 approaches such as in BLOQQER (Biere et al., 2011). During preprocessing one should take care that no selector variables are simplified.

315 The counting algorithm works by blocking solutions. This is done by negating a solution and adding a corresponding circuit gates (inverters, AND-gates, and an OR-gate) to the original miter. The size of the miter grows linearly with the number of solutions.

320 The typical miter approach uses XOR gates to compare outputs. The two functions are different if and only if the miter is satisfiable. This is dual to using XNOR gates and checking for validity. Notice, that due to the fact that the XNOR gates are connected to a constant, there is a some constant-folding that simplifies the job of the QBF solver.




Figure 13: Miter

Algorithm 1: LABELCOUNTER(B, ψ)

Input : B , set of Boolean functions, basis
 $\psi = \langle V, E, \alpha, \beta, \chi, \omega \rangle$, Boolean circuit, requirements
Output: $count$, integer, number of configurations

$X \leftarrow \{\chi(v) : v \in V\} \setminus \{\star\}$
 $Y \leftarrow \{\omega(v) : v \in V\} \setminus \{\star\}$
 $count \leftarrow 0$
 $miter, S, Z \leftarrow \text{CREATEMITER}(B, V, E, X, Y)$
while $witness \leftarrow \text{SOLVEQBF}(\exists S \forall X miter)$ **do**
 $miter \leftarrow miter \wedge \neg witness$
 $count \leftarrow count + 1$
end
return $count$

5. Brute-Force Circuit Counting

One can think of Boolean circuit synthesis as having two aspects: (i) coming-up with a topology G and (ii) determining the type of each node in G . Algorithm 1 solves only (ii). Arguably, (i) is the more difficult part, and in general
325 both (i) and (ii) must be solved simultaneously. In this section we combine Algorithm 1 and an exhaustive search over all possible topologies of a certain size.

Circuit design is an optimization problem: the objective is to minimize some property such as primary input to output propagation time or power (if the circuit is implemented electrically).
330 The optimization criterion depends on the use-case. The main goal of this paper is to minimize the complexity of the circuit, i.e., the number of components.

Problem 3 (Optimal Circuit Design). Given a basis B and a requirements circuit ψ , compute a circuit $\varphi = \langle V, E, \alpha, \beta, \chi, \omega \rangle$, such that $\varphi \equiv \psi$ and no other
335 circuit $\varphi' = \langle V', E', \alpha', \beta', \chi', \omega' \rangle$ exists such that $\varphi' \equiv \psi$ and $|V'| < |V|$.

A circuit topology, itself, has two aspects: (i) how components are connected with each other and (ii) how components interface with the outside world in terms of primary inputs and outputs (χ and ω , respectively). This gives rise to a class of graphs that have three types of nodes: (i) primary inputs X ,
 340 (ii) primary outputs Y , and (iii) internal nodes Z . It is assumed that each primary input node $x \in X$ is connected to one or more internal nodes $Z' \subseteq Z$. Each primary output $y \in Y$ is connected to a distinct internal node $r \in Z$.




Figure 14: The fully connected topology $K_{3,5,2}$

Our first approach to solving Problem 3 is to exhaustively enumerate all possible topologies up to a certain size. The topology that has the fully-connected graph is denoted as K . A fully-connected topology where the primary inputs,
 345 outputs, and internal nodes are partitioned is denoted as $K_{|X|,|Y|,|Z|}$, where $|X|$ is the number of primary inputs, $|Y|$ is the number of primary outputs and $|Z|$ is the number of internal variables. A circuit topology of size $|V| = |X| + |Y| + |Z|$ is always a subgraph of $K_{|X|,|Y|,|Z|}$. We can skip circuit topologies where two
 350 primary outputs are tied together.

The number of circuit topologies of a certain size grows rapidly. The number of directed edges in $K_{m,n,k}$ is $|E| = mk + nk + k(k - 1) = k(m + n + k - 1)$. This results in a total of $2^{|E|}$ subsets. Consider the topology of the full-adder with three primary inputs and two primary outputs. The first six elements of the series $|2^{T_{3,k,2}}|$ are 2^5 , 2^{12} , 2^{21} , 2^{32} , 2^{45} , and 2^{60} .
 355

Algorithm 2 is the simplest method for circuit counting. It is guaranteed to terminate as there are upper-bounds for the number of components and for the number of subgraphs of K . Algorithm 2 is also guaranteed to generate a design if all components in ψ correspond to Boolean functions in the basis B .

360 Algorithm 2 computes designs of minimal size. The reason for that is that first all topologies with one internal node are tried, then all topologies with two nodes, etc.

Algorithm 2 solves Problem 1 for each candidate topology $\langle V', E', \chi, \omega \rangle$. The number of invocations of the QBF solver can be significantly reduced if we consider non-isomorphic graphs only. There is no analytic approach to enumerating
 365 all non-isomorphic graphs of size k , the latter is a problem on its own. The world leaders in graph counting are McKay and Piperno (2014).

Algorithm 2: EXHAUSTIVESHARCH(ψ)

Input : B , set of Boolean functions, basis
 $\psi = \langle V, E, \alpha, \beta, \chi, \omega \rangle$, Boolean circuit, requirements
Output: $count$, integer, number of circuits

$X \leftarrow \{\chi(v) : v \in V\} \setminus \{\star\}$
 $Y \leftarrow \{\omega(v) : v \in V\} \setminus \{\star\}$

$count \leftarrow 0$
 $n \leftarrow 1$
while $count = 0 \wedge n \leq |V|$ **do**
 forall $\langle V', E' \rangle \subseteq K_{|X|, |Y|, n}$ **do**
 $miter, S, Z \leftarrow \text{CREATEMITER}(B, V', E', X, Y)$
 $\gamma \leftarrow \text{CIRCUITTOCNF}(miter)$
 while $witness \leftarrow \text{SOLVEQBF}(\exists S \forall X \gamma)$ **do**
 $miter \leftarrow miter \wedge \neg witness$
 $count \leftarrow count + 1$
 end
 end
 $n \leftarrow n + 1$
end

6. Computing Both a Topology and the Component Types from the Solution of a Single 2-QBF Problem

370 The brute-force algorithm of Sec. 5 is too slow. It is possible to encode the whole circuit synthesis, both component selection and topology generation, as a single QBF satisfiability problem. The difficulty of generating a circuit is then left entirely to the QBF solver. The approach is shown in Algorithm 3.

375 Similar to Algorithm 2, Algorithm 3 first tries candidate circuits with one component, then with two, and so on, until an equivalent circuit is discovered. The UNIVERSALCELL subroutine in line 3 adds k universal component cells (see Sec. 4). All inputs of the k universal components are accumulated in X_u and all outputs are accumulated in Y_u . The selector variables for the types of components are accumulated in S_u . The circuit of the universal cells is denoted
380 as φ_u .

6.1. A Configurable Interconnection Fabric

The step after adding the universal cells is to construct the circuit φ_t that represent the interconnection fabric (the wires connecting the gates). Denote the elements of $X \cup Y_u$ as $\{y_1, y_2, \dots, y_m\}$. Consider a single input x of a universal
385 cell. The formula

$$x \leftrightarrow \bigvee_{i=1}^{|X \cup Y_u|} s_i \wedge y_i \quad (4)$$

Algorithm 3: SYNTHESIZECIRCUIT(B, ψ, n)

Input : B , set of Boolean functions, basis
 ψ , requirements (inputs X , outputs Y)
 n , integer, maximum number of components

Output : Φ , a set of circuits

Local Variables : X_u , set of variables, the inputs to all universal cells
 Y_u , set of variables, the outputs of all universal cells
 S_u , set of variables, the inputs to all universal cells
 $S_t = |s_{i,j}|$, matrix of variables, the interconnection

fabric selectors

$\Phi \leftarrow \emptyset$

for $k \in \{1, 2, \dots, n\}$ **do**

$\langle S_u, X_u, Y_u, \varphi_u \rangle \leftarrow \text{UNIVERSALCELLS}(B, k)$

$\langle S_t, \varphi_t \rangle \leftarrow \text{INTERCONNECTIONFABRIC}(X_u \cup Y, X \cup Y_u)$

for $j \in \{1, 2, \dots, |X \cup Y_u|\}$ **do**

$\varphi_t \leftarrow \varphi \wedge$

$\text{COLUMNCARDINALITYCONSTRAINT}(\{s_{1,j}, s_{2,j}, \dots, s_{|X_u \cup Y|,j}\})$

end

for $i \in \{1, 2, \dots, |X_u \cup Y|\}$ **do**

$\varphi_t \leftarrow$

$\varphi \wedge \text{ROWCARDINALITYCONSTRAINT}(\{s_{i,1}, s_{i,2}, \dots, s_{i,|X \cup Y_u|}\})$

end

$\varphi \leftarrow \varphi_u \cup \varphi_t$

if $witness \leftarrow \text{SOLVEQBF}(\exists S_u \cup S_t \forall X \varphi)$ **then**

$\varphi \leftarrow \varphi \wedge \neg witness$

$\Phi \leftarrow \Phi \cup \text{RECONSTRUCTCIRCUIT}(witness)$

end

end

return Φ

“connects” x to all possible outputs of universal cells or primary inputs depending on the values of the selector variables s_i .

Eq. (4) has to be repeated for every possible input of a universal cell or primary output. Let us denote those outputs as $X_u \cup Y = \{x_1, x_2, \dots, x_n\}$.

390

This leads to the formula for the configurable interconnection fabric:

$$\bigwedge_{j=1}^{|X_u \cup Y|} \bigvee_{i=1}^{|X \cup Y_u|} x_j \leftrightarrow s_{i,j} \wedge y_i \quad (5)$$

Most constraints of Eq. (5) are implemented as an array of two input AND-gates and are shown in Figure 15. One of the inputs of each AND-gate from this array is connected to a selector variable $s_{i,j}$. Each selector variable $s_{i,j}$ enables or disables the connection of an output to an input. There is also a multi-input OR-gate for each row of AND-gates.

395

Figure 16 illustrates the high-level structure of the interconnection topology.

6.2. Cardinality and Other Constraints

An unconstrained interconnection fabric would result in malformed circuits: loops, floating wires, wires that are not connected to any components, etc. To
400 avoid these malformed topologies, Algorithm 3 imposes a number of constraints (implemented as sub-circuits) on the encoding. Below is a description of each of these constraint types.

Cycle Breaking (CB): The topology selector variables in the upper right triangle of Fig. 16 are all disabled (assigned \perp). These constraints impose
405 a strict ordering on the components and ensures that the outputs of each universal cell are connected to the inputs of a successor universal cell only. Instead of assigning \perp to the variables there, we simply make the multiplexers of different size and save the computational time for constant-folding.

Row Cardinality Constraints (RCCs): An “exactly-one” constraint is added
410 to each row of the connectivity matrix. These constraints can be implemented either with a sorting network, with a multi-operand adder (see Appendix Appendix A), or with a combination of two-input AND-gates and multi-input OR-gates. The choice of the implementation does not affect
415 the performance because the RCCs constitute a relatively small part of the encoding.

Column Cardinality Constraints (CCCs): These constraints can be either
420 “at-least-one”, “exactly-one”, or a combination of the two. The choice determines the type of the synthesized topology. The options that are of practical significance are:

Circuit Topology: All CCCs are of type “at-least-one”. Notice that an
“at-least-one” constraint is simply a single multi-input OR-gate.

Boolean Function Topology: The first $|X|$ CCCs are of type “at-least-
425 one”, and the remaining $|Y_u|$ columns are of type “one”. This combination of CCCs results in a circuit where the fan-out of primary inputs to gate inputs is unrestricted while the fan-out of each gate is restricted to one. For these circuits there are corresponding Boolean functions whose number of variables is the same as the number of primary inputs in the synthesized circuit. More colloquially: the
430 synthesized circuit is a Boolean function.

Network Topology: All CCCs are of type “exactly one”, same as the
RCCs. This topology is suitable for synthesizing sorting networks and reversible circuits.

Unbalanced Universal Cell Ports (UUCP): The universal cell does not
435 necessarily combine gates with the same number of inputs and outputs,

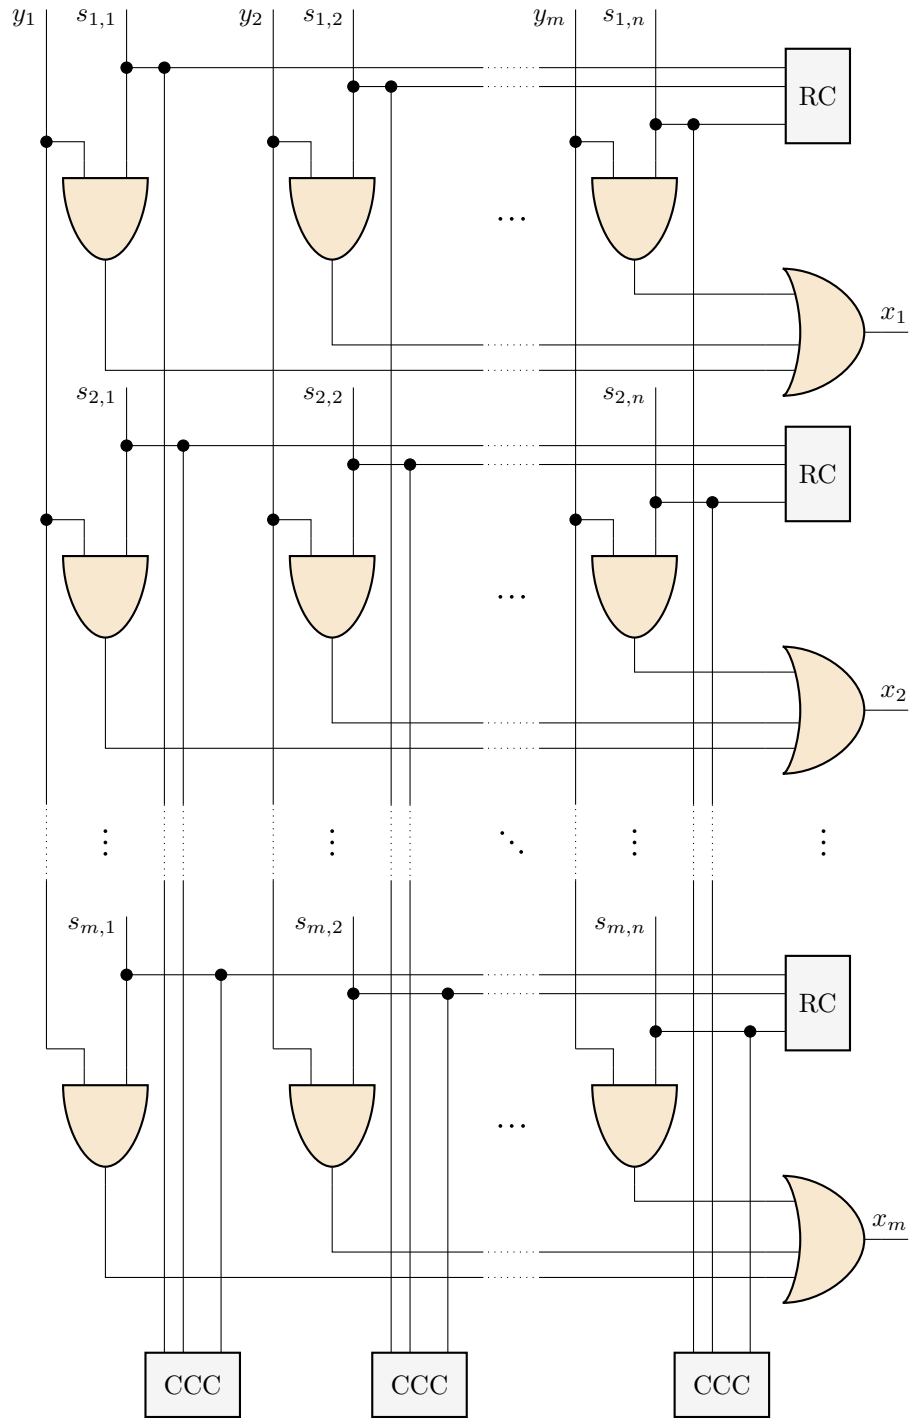


Figure 15: Configurable interconnection fabric and cardinality constraints




Figure 16: Topology constraints

leading to hanging wires. The T_3 constraints prevent components being connected to them. They are implemented as small binary multiplexers that choose which extra topology wires go to a pre-selected input/output of the universal cell.

440 Notice that Algorithm 3 needs an upper-bound for the number of components n . If the basis of the requirements ψ is the same as the basis of the synthesis, then the number of components $|\psi|$ can be used as an upper-bound of n as an existence of a circuit for $n = |\psi|$ is guaranteed. Otherwise, one can use the size of a canonical form. For example in the standard basis any formula corresponding
 445 to a circuit can be converted to Disjunctive Normal Form (DNF). It is possible to use the size of a circuit implementing this DNF as a value for n , although n would be too large. In the case of sorting networks one can take existing upper bounds, for example, the size of the bitonic sorting network corresponding to the desired number of inputs.

450 Having all this in place, we are ready to synthesize some circuits for better understanding of Algorithm 3.

6.3. Examples of Synthesis

Figure 17 shows the result of running Algorithm 3 with the standard basis and the full-subtractor shown in Figure 8 as requirements. The generated circuit
 455 has five components only while the one in the requirements has seven. This is a substantial saving.

Another circuit designed by Algorithm 3 is the reversible adder/subtractor shown in Figure 18. This circuit, using one CSWAP and three CCNOT gates,




Figure 17: An alternative full-subtractor

460 has two constant inputs and two garbage outputs (u_1 and u_2). Synthesizing reversible circuits given bases containing reversible gates has application to the standard model of quantum computing as reversible circuits do not lead to physical increase of entropy.




Figure 18: A reversible full-adder/subtractor

The five-input sorting network shown in Figure 19 is computed by Algorithm 3, configured with a basis containing a comparator only. The requirements circuit ψ is a bitonic sorting network. Proving the size of the optimal sorting network for a certain number of inputs is an open problem Codish et al. (2019). Notice that we can use Algorithm 3 for formally proving minimality of a circuit. For that, we need to prove soundness, correctness, and optimality of Algorithm 3 and also to save and check all resolution-based proofs Heule et al. (2013) of non-existence of circuits of size smaller to the one synthesized.

Figure 20 shows a full-adder implemented with NAND-gates only. The design is the classical one where the two identical half-adder subsystems are visible. A full-adder can also be implemented with NOR-gates only. It has the same topology as the one shown in Figure 20.

475 6.4. Symmetry Breaking for Components Whose Inputs Commute

Component libraries often have components whose inputs commute. For example, all inputs in all components in the standard basis commute. We auto-




Figure 19: An optimal five-input sorting network

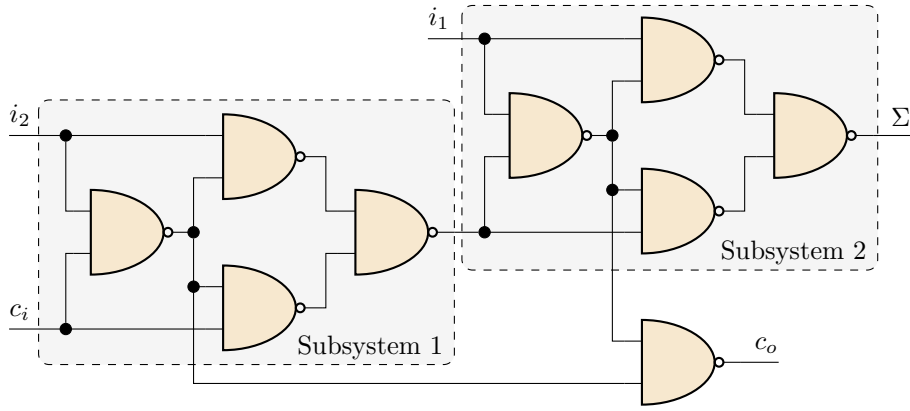


Figure 20: Classical design of a NAND-based full-adder

matically compute the set of all commuting component input pairs by building small miters like the one shown in Figure 13.

480 Consider a pair of commuting component inputs x_1 and x_2 and the set $Y = \{y_1, y_2, \dots, y_n\}$ of all possible component outputs and primary outputs that can be connected to x_1 and x_2 . There are $2n$ selector variables responsible for connecting x_1 and x_2 to Y : $s_{1,1}, s_{1,2}, \dots, s_{1,n}, s_{2,1}, s_{2,2}, \dots, s_{2,n}$. We exclude symmetric topologies by adding the following constraints:

$$\begin{aligned}
 s_{2,1} &\rightarrow \perp \\
 s_{2,2} &\rightarrow s_{1,1} \\
 s_{2,3} &\rightarrow s_{1,1} \vee s_{1,2} \\
 &\vdots \\
 s_{2,n} &\rightarrow s_{1,1} \vee s_{1,2} \vee \dots \vee s_{1,n-1}
 \end{aligned} \tag{6}$$

485 As one can see, Formula 6 essentially orders the outputs of all possible components when they are connected to a pair of commuting inputs. Analogous technique works for sets of commuting inputs of arbitrary size.

6.5. Algorithm Properties

Algorithm 3 generates a candidate 2-QBF circuit representing a topology of
 490 k components and solves it. Analyzing this generated circuit answers question
 about properties such as soundness and completeness. Due to the applied nature
 of this paper we only provide sketches instead of full proofs.

Property 1 (Soundness). *Given a requirements circuit ψ , for any circuit φ
 produced by Algorithm 3, it holds that $\psi \equiv \varphi$.*

495 *Proof (Sketch).* Proving this property can be done directly by analyzing the
 miter formula $\exists S \forall X \varphi \equiv \psi$. The formula is expanded for every possible assign-
 ment to variables in X :

$$\bigwedge_{x \in P^X} \varphi(x) \leftrightarrow \exists S \psi(x), \quad (7)$$

where P^X denotes the set of all possible assignments to variables in X . We
 also denote as $\varphi(x)$ and $\psi(x)$ the values of all primary outputs of the circuits φ
 500 and ψ , respectively, given an assignment to all their primary inputs. We need
 each conjunct in Eq. 7 to be true. This means that we need an assignment
 to the S -variables that makes ψ in every conjunct true if φ is true and false
 otherwise. \square

Completeness means that if there exists a circuit that can be synthesized
 505 from the given basis B , Algorithm 3 is guaranteed to find it.

Property 2 (Completeness). *Given a requirements circuit ψ , Algorithm 3 is
 guaranteed to produce a circuit $\varphi \equiv \psi$ if such a circuit exists.*

Proof (Sketch). This can be shown with the help of a direct proof by analyzing
 the meta-circuit φ generated by Algorithm 3.

510 The idea is to show that each satisfying assignment of φ corresponds to a
 circuit and that for any valid circuit and a fixed k , there exists a corresponding
 satisfiable assignment.

Analyzing the 2-QBF circuit φ generated by Algorithm 3 is too complex, so
 the first step is to split formula in two: 1. topology and topological constraints,
 515 and 2. the universal component cells.

Let us consider a SAT-based algorithm and a circuit φ_t that generates valid
 topologies only (see Sec. 5). This algorithm does not need the requirements
 circuit ψ nor the universal component cells. All variables in φ_t are existentially
 qualified, i.e., varpihi_t is a 1-QBF or a regular circuit. Next, we can show
 520 that each satisfying assignment corresponds to an unique well-formed topology
 graph. In the other direction, we have to show that each good topology can be
 the solution of φ_t .

Each topological constraint type must be analyzed separately and all topo-
 logical constraints must be analyzed together to shown that they do not allow
 525 invalid topologies and that they do not exclude valid ones. An invalid topology

is, for example, a topology in which two component outputs are connected to the same input.

Representing the topology as an adjacency matrix helps with the argument.

The topology result can be next combined with correctness results of Algorithm 1 which should lead to the final conclusion that Algorithm 3 is complete. □

An easy property to show is that of optimality. In this paper, the optimization criterion is the number of components in the circuit. In the application domain of digital design, for example, this corresponds to power consumption. It is possible to introduce other costs and even cost functions in which case Algorithm 3 may lose its optimality or completeness.

Property 3 (Optimality). *Given a requirements circuit ψ , Algorithm 3 is guaranteed to produce a circuit $\varphi \equiv \psi$ with φ of minimal size if such a circuit exists.*

Proof (Sketch). Algorithm 3 attempts to synthesize a circuit for an increasing number of components k , starting from $k = 0$. If the synthesis is sound and complete, then the minimality follows directly from the iteration strategy for k . □

Another property of Algorithm 3 is related to the notion of universality. There are bases for which Algorithm 3 is guaranteed to synthesize a circuit equivalent to any requirements circuit ψ . One such basis is a basis that contains the NAND-gate only.

7. Computational Complexity

For a certain input, Problem 3 becomes the same as the Minimum Equivalent Expression (MEE) problem, classified by Buchfuhrer and Umans (2011). What follows is reformulation of the complexity results of Buchfuhrer and Umans (2011) in the terminology of this paper.

Theorem 1 (Complexity of Circuit Synthesis with Fixed Basis, Single Output and Gate Fan-Out Restricted to One). *Single-output circuit generation with basis $B = \{\neg, \wedge, \vee\}$ and gate fan-out restricted to one is Σ_2^P -complete.*

Proof (Sketch). For a Boolean formula φ with n literals, there exists an $O(n)$ reduction from the Minimum Equivalent Expression (MEE) problem over signature $\{\vee, \wedge, \neg\}$. The MEE problem is classified as L22 in the polynomial-time hierarchy compendium (Schaefer and Umans, 2002) and is shown to be in Σ_2^P by Buchfuhrer and Umans (2011).

The MEE problem asks if, given a Boolean formula φ and a constant k there exists a formula ψ for which $\psi \equiv \varphi$, and $|\psi| < k$. The number of literals in ψ is denoted as $|\psi|$. The circuit generation problem concerns generation of circuits with a minimal number of components. For a basis $B = \{\neg, \wedge, \vee\}$, the number of literals in the Boolean formula equivalent to the generated circuit is equal to the number of literals. □

The complexity of the general circuit synthesis problem can now be shown constructively. The idea is that we have partial input that makes the problem Σ_2^P -hard. On the other hand, we have a constructive proof (Algorithm 3) that can solve Problem 3 by solving a 2-QBF. Of course we also need soundness and completeness of the algorithm.

Theorem 2 (Complexity of Circuit Generation). *Problem 3 is Σ_2^P -hard.*

Proof (Sketch). The lower bound on the worst-time complexity comes from Theorem 1. The upper bound comes, constructively, from Algorithm 3 as it reduces the problem to an $\exists\forall$ QBF. \square

Notice that having a basis with a NAND-gate only is equivalent to DNF minimization which is also in Σ_2^P Umans (2001).

8. Experiments

What follows is an empirical analysis of the encodings and methods introduced in the preceding sections. The high-level miter construction is implemented in PYTHON while the QBF solving is in C/C++. We have compared three award-winning (Janota et al., 2016a) QBF solvers: QFUN (Janota, 2018), RAREQS (Janota et al., 2016b), and DEPQBF (Lonsing and Egly, 2017). The QFUN QBF solver is non-clausal. The QCNF input to RAREQS and DEPQBF has been preprocessed with BLOQQER (Biere et al., 2011) where we had to take special precaution not to eliminate selector variables. The preprocessing step works by eliminating unnecessary clauses and variables. It performs several other optimizations as well. This gives significant speed-up.

In addition to the above three QBF solvers, we have implemented a full expansion of the innermost universal quantifier, resulting in a SAT problem. The SAT problem is then solved with KISSAT (Biere et al., 2020). This approach is suitable for problems with a smaller number of primary inputs. The resulting expansion based 2-QBF solver is called PLQ. PLQ performs constant folding after the expansion and before converting the input to CNF.

To validate the implementation of the algorithms presented in this paper we compare with the help of a miter and a SAT solver the equivalence of each synthesized circuit to the requirements.

All experiments were performed on a 2-CPU (4 cores per CPU) Intel Xeon 3.3 GHz Linux computer with 1.5 TiB of RAM.

8.1. Requirement Circuit Benchmarks

We experiment on three types of circuits. The first type are arithmetic and logic circuit families of variable size. The second type are netlists from real-world ICs. Last, we take as requirements several sets of Boolean functions from exact synthesis Haaswijk et al. (2018).

8.1.1. Arithmetic and Logic Unit Circuits

605 Table 1 shows a scalable synthetic set of combinational arithmetic circuits. The size of each of the eight synthetic circuits, described in Table 1, can be varied by setting a parameter n . Each variable-size circuit shares the same topology. The carry and borrow mechanisms of adders and subtractors, for example, have bus-like topology, while the adder networks of the multipliers
610 resemble two-dimensional meshes.

Name	Description	Role of the independent parameter n
n -mux	Multiplexer	Number of input bits to be multiplexed, selectors are not counted
n -demux	Demultiplexer	Number of output bits
n -add	Full-adder	Number of inputs in one of the addends, carry input is not counted
n -sub	Full-subtractor	Number of inputs in the subtrahend, borrow input is not counted
n -cmp	Comparator	Number of bits in one of the terms
n -shift	Barrel-shifter	Number of input bits to be shifted, selectors are not counted
n -moa	Multi-operand adder	Number of input bits to be added
n -mul	Multiplier	Number of input bits in the multiplicand

Table 1: Role of the n parameter in the ALU- n families

Appendix A provides a detailed description of the circuits in this benchmark.

We have generated two sets of circuit families for $1 \leq n \leq 4$ and $1 \leq n \leq 32$. These two benchmark sets are called ALU-4 and ALU-32, respectively. The former is used to benchmark synthesis while the latter is used for evaluating the
615 performance of gate selection.

8.1.2. 74XXX Integrated Circuits

Table 2 shows the second set of benchmark circuits. These are reverse-engineered real-world ICs (Hansen et al., 1999). The 74XXX circuits can be chained together into larger Arithmetic Logic Units (ALUs).

Name	Description	PIs	POs	Gates
74182	4-bit CLA	9	5	19
74L85	4-bit comparator	11	3	33
74283	4-bit adder	9	5	36
74181	4-bit ALU	14	8	65

Table 2: 74XXX digital circuits

620 The number of gates in the 74XXX circuits are still beyond the ability of the synthesis algorithm. The 74XXX circuits are used for measuring the performance of Algorithm 1 only.

8.1.3. Boolean Functions from Exact Synthesis

625 The performance of Algorithm 3 is compared to the algorithms devised by Haaswijk et al. (2018). The authors of this work have provided function sets from their study on exact synthesis and classification (Haaswijk et al., 2019). Table 3 provides an overview of the benchmark.

Name	Description	PIs	Functions
NPN4	Negation-Permutation-Negation classes	4	222
FDSD6	Fully-DSD decomposable functions	6	1000
PDSD6	Partially-DSD decomposable functions	6	1000
FDSD8	Fully-DSD decomposable functions	8	100
PDSD8	Partially-DSD decomposable functions	8	100

Table 3: Boolean function sets from exact synthesis

We have randomly selected 2422 functions. The sizes of the subsets are the same as in Haaswijk et al. (2018) but the functions, with the exceptions of the one in the NPN4 class, are different.

8.2. Gate Selection

This section empirically evaluates the performance of Algorithm 1. The main results, shown in Figure 21, summarize the QBF performance for ALU-32 and different QBF solvers. The horizontal axis shows the number of variables in the problem and the vertical axis is the time-to-solution. The performance depends on the topology of the requirements circuit and, to some extent, on the choice of the QBF solver.

640 The plots in Figure 21 have logarithmic vertical axes to accommodate the exponential time-to-solution. Contrary to our intuition, the multiplier circuit is not the most difficult one and the full-adder is not the easiest. The performance is best for the demultiplexer, no matter which QBF solver has been used. In general, the QBF solver performance is better for large fan-outs. This can be explained with less back-tracking when there are more outputs.

645 Table 4 characterizes the performance of Algorithm 1 on the 74XXX circuits. The table shows the number of solutions found by each solver in 1 h. Interestingly, the only solver that found solutions for 74181 is the clausal RAREQS, despite solving a 3-QBF problem.

650 The QFUN solver showed better performance than PLQ due to the fact that PLQ spent a lot of time expanding the circuit. Of course, when counting, there is no need to expand the circuit every time after a solution is blocked.

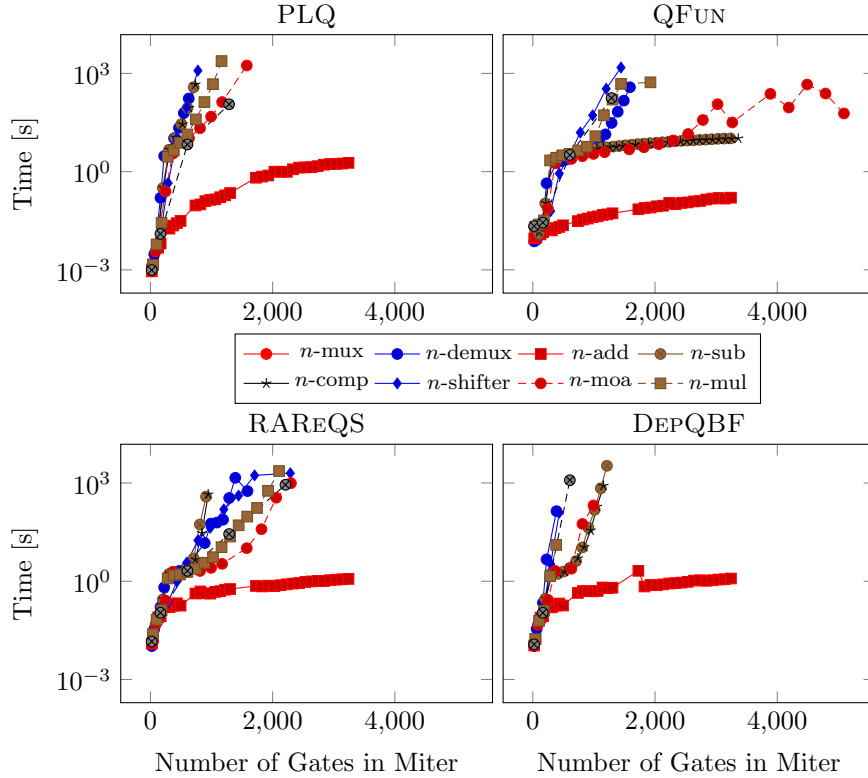


Figure 21: Component selection time-to-solution for ALU-32 circuits

8.3. Circuit Synthesis

The bulk of our experiments is concerned with evaluation the performance of Algorithm 3.

8.3.1. Arithmetic and Logic Unit Circuits

655 Table 5 shows the most important data in this paper. It summarizes the performance of the PLQ and QFUN QBF solvers with and without symmetry breaking. The bounds are on the number of components in a circuit. An upper bound value means that Algorithm 3 has generated a circuit with a certain number of gates. The lower bound values show that Algorithm 3 has proven
 660 non-existence of a circuit of a given size.

Higher values for lower bounds and lower values for upper bounds indicate better result. The best numbers for every circuit are shown in bold. For example, the row for the 1-adder circuit show that all four solver/symmetry breaking configurations prove the non-existence of a 4-component circuit and find a 5-
 665 component one.

In some cases Algorithm 3 could fully solve a circuit. This means that Algorithm 3 found a satisfiable solution for k components and showed non-

Name	PLQ	QFUN	RAREQS	DEPQBF
74182	154	512	512	—
74L85	10	675	134	—
74283	31	691	406	—
74181	—	7	21	—

Table 4: Performance of non-clausal QBF solvers in enumerating component selection for 74XXX circuits

satisfiability for m components, for $0 \leq m \leq k-1$. The names of the fully-solved circuits are also shown in bold in the leftmost column of Table 5.

670 Figure 22 shows the times-to-solution of the QBF solvers for ALU-4. There is significantly more UNSAT data because the search is from small to large number of components. The time-to-solution increases exponentially. The most difficult calls are just one component below the smallest circuit size. Once a circuit has been found it becomes easier for a while and then, when increasing
675 the number of components the QBF solver starts timing out again.




Figure 22: Time-to-solution for various candidate circuit sizes in the ALU-4 benchmark

8.3.2. Reversible Circuits

Table 6 summarizes the results for synthesizing the ALU-4 reversible circuit from the reversible basis shown in Fig. 5a. The data in the table should be interpreted similarly to the data in Table 5, except that there is no symmetry breaking and that in addition to number of gates, there is also number of
680 ancillary inputs.

Coincidentally, what Algorithm 3 could synthesize with the reversible basis, is close to what Algorithm 3 could synthesize with the standard basis. For example all eight circuits in the multiplexer and demultiplexer families could be
685 synthesized and proven minimal.

Name	PIs	Gates	Upper Bound				Lower Bound			
			PLQ		QFUN		PLQ		QFUN	
			SB	No SB	SB	No SB	SB	No SB	SB	No SB
1-mux	2	2	2	2	2	2	1	1	1	1
2-mux	3	4	3	3	3	3	2	2	2	2
3-mux	5	6	–	–	–	–	6	6	6	5
4-mux	6	7	–	–	–	–	6	6	6	5
1-demux	2	2	2	2	2	2	1	1	1	1
2-demux	2	3	2	2	2	2	1	1	1	1
3-demux	3	5	5	5	5	5	4	4	4	4
4-demux	3	6	6	6	6	6	5	5	5	5
1-add	3	5	5	5	5	5	4	4	4	4
2-add	5	10	10	10	10	–	8	7	7	6
3-add	7	15	–	–	–	–	8	8	7	7
4-add	9	20	–	–	–	–	7	7	7	7
1-sub	3	7	5	5	5	5	4	4	4	4
2-sub	5	14	10	10	13	10	7	7	7	7
3-sub	7	21	15	–	–	–	8	8	6	6
4-sub	9	28	–	–	–	–	8	7	7	6
1-comp	2	5	3	3	3	3	2	2	2	2
2-comp	4	10	8	8	8	8	7	7	7	7
3-comp	6	13	13	13	–	–	8	7	7	7
4-comp	8	16	–	–	–	–	8	7	7	6
1-shifter	2	2	2	2	2	2	1	1	1	1
2-shifter	3	5	5	5	5	5	4	4	4	4
3-shifter	5	14	11	13	–	14	7	7	7	6
4-shifter	6	20	–	–	–	–	8	8	7	7
1-moa	1	1	0	0	0	0	1	–	1	–
2-moa	2	2	2	2	2	2	1	1	1	1
3-moa	3	5	5	5	5	5	4	4	4	4
4-moa	4	9	9	9	9	9	7	7	7	6
1-mul	2	1	1	1	1	1	0	0	0	0
2-mul	4	8	7	7	7	7	6	6	6	6
3-mul	6	30	–	–	–	–	9	9	8	8
4-mul	8	64	–	–	–	–	11	10	9	8

Table 5: Optimization performance for ALU-4 circuits and the standard basis

Name	PIs	Gates	Upper Bound				Lower Bound			
			PLQ		QFUN		PLQ		QFUN	
			Gates	Ancil.	Gates	Ancil.	Gates	Ancil.	Gates	Ancil.
1-mux	2	2	2	1	2	1	2	0	2	0
2-mux	3	4	2	0	2	0	1	2	1	2
3-mux	5	6	3	1	3	1	3	0	3	0
4-mux	6	7	3	0	3	0	2	4	2	4
1-demux	2	2	2	1	2	1	2	0	2	0
2-demux	2	3	2	1	2	1	2	0	2	0
3-demux	3	5	3	3	3	3	3	2	3	2
4-demux	3	6	3	3	3	3	3	2	3	2
1-add	3	5	4	1	4	1	4	0	4	0
2-add	5	10	—	—	—	—	5	5	4	8
3-add	7	15	—	—	—	—	4	4	4	8
4-add	9	20	—	—	—	—	3	4	4	8
1-sub	3	7	4	2	4	2	4	1	4	1
2-sub	5	14	—	—	—	—	6	1	5	0
3-sub	7	21	—	—	—	—	5	3	5	1
4-sub	9	28	—	—	—	—	3	5	7	0
1-comp	2	5	3	3	3	3	3	2	3	2
2-comp	4	10	—	—	—	—	6	1	5	1
3-comp	6	13	—	—	—	—	5	5	8	0
4-comp	8	16	—	—	—	—	4	2	16	0
1-shifter	2	2	2	1	2	1	2	0	2	0
2-shifter	3	5	2	1	2	1	2	0	2	0
3-shifter	5	14	—	—	—	—	5	7	5	0
4-shifter	6	20	—	—	—	—	5	3	4	8
1-moa	1	1	0	0	0	0	—	—	—	—
2-moa	2	2	2	2	2	2	2	1	2	1
3-moa	3	5	4	1	4	1	4	0	4	0
4-moa	4	9	—	—	—	—	6	1	5	0
1-mul	2	1	—	—	—	—	1	2	1	2
2-mul	4	8	5	4	—	—	5	3	5	2
3-mul	6	30	—	—	—	—	5	4	5	1
4-mul	8	64	—	—	—	—	3	5	9	0


Table 6: Optimization performance for ALU-4 circuits and the reversible basis

Notice that 1-moa is simply a wire and 1-mul is a single two-input AND-gate. Representatives of successfully synthesized ALU-4 reversible circuits are shown in Appendix Appendix B.


8.3.3. Boolean Functions from Exact Synthesis

690 We next analyze the performance of Algorithm 3 on the function sets from exact synthesis. Each experiment has been repeated twice, for two different topologies of the synthesized circuit. The first topology is the Boolean function one where the gate fan-out is restricted to one. In the second set of experiments there is no restriction on the fan-out. The difference is illustrated in Figure 23a and Figure 23b. Both figures are equivalent to the NPN4 circuit with truth table 0x12D. Notice that the circuit in Figure 23a has one gate less compared to the circuit shown in Figure 23b.

695



(a) Circuit with unrestricted gate fan-out



(b) Circuit with a maximum gate fan-out of one

Figure 23: Minimal implementations of the NPN4 Boolean function with truth-table 0x12D

Table 7 summarizes the experimental results for the PLQ solver. The T/O column shows the number of experiments in which the QBF solver timed out. The time out for each experiment was been set-up to 5 min. Symmetry-breaking was enabled. For the solved problems, we have the mean time μ in s and the standard deviation σ . Columns 2-5 are for circuits with gate fan-out restricted

700

to one. This is the same fan-out as in the experiments of Haaswijk et al. (2018). Columns 6-9 are for circuits with unrestricted gate fan-out.

Name	Solved	Fan-Out = 1			Fan-Out ≥ 1			
		T/O	μ	σ	Solved	T/O	μ	σ
NPN04	221	1	3.3	6.67	222	0	15.58	49.76
PDSD06	922	78	47.64	44.62	908	92	66.51	68.37
FDSD06	999	1	10.53	16.64	999	1	11.85	24.45
PDSD08	0	100	—	—	0	100	—	—
FDSD08	53	47	190.75	39.62	62	38	167.85	60.48

Table 7: Solved instances by PLQ and time-to-solution for Boolean function sets from exact synthesis

705 Table 8 shows the synthesis results for QFUN. Its layout is the same as Table 8. The performance of QFUN is worse compared to the one of PLQ.

Name	Solved	Fan-Out = 1			Fan-Out ≥ 1			
		T/O	μ	σ	Solved	T/O	μ	σ
NPN04	208	14	9.1	39.9	207	15	10.37	16.27
PDSD06	498	502	46.93	57.16	474	526	56.81	62.4
FDSD06	941	59	10.04	22.19	938	62	11.88	29.09
PDSD08	0	100	—	—	0	100	—	—
FDSD08	41	59	59.38	72.14	43	57	44.72	65.29

Table 8: Solved instances by QFUN and time-to-solution for Boolean function sets from exact synthesis

710 In the whole benchmark, PLQ found 10 instances in which an NPN04 circuit could be synthesized with one less gate due to allowing unrestricted fan-out. The QFUN solver found 4 such cases. In the larger sets, PLQ found 26 cases for PDSD06 and 2 for FDSD06 while QFUN did not find any. This initial evidence shows that for the studied function sets, unrestricted fan-out leads to a small size reduction (one gate) in rare cases.

9. Related Work

715 Circuit design is related to diagnostic reasoning (de Kleer and Williams, 1987). Consider Problem 1 and Algorithm 1. The requirements circuit ψ can be thought of as an observation. Instead of augmenting ψ to create ϕ , as done in Algorithm 1, we can augment the buggy system description. The failure modes are “mistaken gate identity”, for example, the modeler has used an AND-gate in place of an OR-gate. Algorithm 1 then computes minimal changes in the system description that explain the observed circuit.

720

The General Diagnostic Engine (GDE) of de Kleer and Williams (1987) can diagnose wiring errors and generate topology. When the problem is reduced to QBF, however, it is easier to avoid “don’t cares” by universally quantifying the primary inputs. Combined with the “connect to successor components only” (see Sec. 6), our approach is more efficient in avoiding loops and exploring the design space.

Some of the motivation for our work comes from Arthur and Polak (2006). The authors of this work show that the evolutionary design of a multi-bit adder takes significantly less steps than anticipated. This “ease” made us attempt a complete algorithm on a seemingly very difficult problem.

The problem of circuit synthesis has been first introduced by Roth and Karp (1962). The authors use a very early computer, an IBM 7090, to solve decomposition problems of four variables in approximately ten minutes. For larger problems they propose a heuristics that would sacrifice the algorithm completeness. Our QBF algorithm, on the other hand, could solve problems of more than 30 variables. This was, of course, done on computers that are orders of magnitude faster but we expect that the difficulty of the synthesis/decomposition problems is at least in the second level of the polynomial hierarchy (Stockmeyer, 1977). Another distinct advantage of our algorithm is that the synthesis/decomposition is in terms of multi-output Boolean functions while the paper of Roth and Karp (1962) supports single output functions only.

The use of the $\exists\forall\exists$ -quantified miter has been proposed for FPGA synthesis (Ling et al., 2005). This paper, however, addresses the component placement problem only and does not consider wiring, routing, and topology. Our paper demonstrates that the combined placement/routing problem can also be solved with a single QBF call and, thus, we have provided a fully automatic solution to the circuit synthesis problem.

There is a large body of work on logic synthesis related to model checking Jr et al. (2018). Typically this type of synthesis is concerned with reasoning about temporal logic. Bloem et al. (2014), for example, uses SAT and QBF for circuit synthesis with emphasis on safety properties.

Problem 1 is closely related to logic synthesis for Filed Programmable Analog Arrays (FPGAs). FPGAs typically consist of array of LookUp Tables (LUTs) and an interconnection network. Programming an FPGA consists of synthesizing the logic elements and configuring the interconnection network. There are multiple methods for doing that (Cong and Ding, 1996) but due to the sizes of the problem, most methods are sub-optimal (Cong and Minkovich, 2007).

10. Discussion

Modern digital designs such as the Pentium CPUs have millions of components. All algorithms in this paper are far from being able to synthesize and enumerate such designs. Large Integrated Circuits (ICs), however, are far from being optimal at the top-level. Companies that make digital circuits integrate subsystems with the designer of each subsystem focusing on the integrity and

optimality of his or her own subsystem. This results in globally suboptimal designs that also have bugs, vulnerabilities and inefficiencies.

The problems we have defined are of industrial interest and create a benchmark that is useful in the QBF competition (Janota et al., 2016a). If accepted the benchmark will help the QBF community to create faster QBF solvers that have practical application. This can be achieved by noticing the structure of the circuit design problems.

We can, at any time, sacrifice completeness and turn the algorithms proposed in this paper into heuristic or stochastic ones. The easiest way to do that is to replace the complete QBF search with stochastic (Gent et al., 2003).

The algorithms in this paper can be adopted to analog designs and design with state. The electronic designs that pose biggest challenge and are of significant practical and theoretical interests are hybrid. It is possible for our synthesis algorithms to work on analogue designs by using QBF modulo theory solvers. These are similar to satisfiability modulo theory solvers (Barrett and Tinelli, 2018) and do not exist at the time of writing of this. The theories can be Ordinary Differential Equations (ODEs) or Differential Algebraic Equations (DAEs). Similarly, the algorithms of this paper, can work for geometric and physical designs with QBF modulo Partial Differential Equations (PDEs).

11. Conclusion

This paper proposes novel and generic solution to the problem of circuit design and exploration. The problem of generating a circuit that is equivalent to a goal is solved similar to how electronic and logic designers solve it: first the component a chosen and placed, and second they are connected with wires. We have given empirical evidence that the complexity of the problem is determined, to a large extent, by the component selection part.

We have proposed a reduction to QBF for solving a difficult problem. We believe that this is the first practical sound and compete algorithm for circuit design and enumeration. The built-in heuristics, compilation and learning in the QBF solvers gives us several orders of magnitude improvement over a baseline graph generation algorithm.

Our method is more generic than anything proposed in literature as it considers arbitrary component libraries, such as ones consisting of reversible gates.

Acknowledgments

We extend our gratitude to Matthew Klenk and John Maxwell from PARC for many discussions and for reviewing this paper. We would also like to thank Florian Lonsing from TU Wien for providing and supporting DEPQBF and for tutoring us on the use of QBF. Thanks to Mikoláš Janota from University of Lisbon for providing RAREQS and useful discussions. Thanks to Martina Seidl from Johannes Kepler University for providing and supporting BLOQQER. Thanks to Marijn Heule from The University of Texas at Austin for useful discussion and reviewing the paper.

References

- Akers, S.B., 1978. Binary decision diagrams. *IEEE Transactions on Computers* 27, 509–516.
- Arthur, W.B., Polak, W., 2006. The evolution of technology within a simple
810 computer model. *Complexity* 11, 23–31.
- Barrett, C., Tinelli, C., 2018. Satisfiability modulo theories, in: *Handbook of Model Checking*. Springer, pp. 305–343.
- Biere, A., Fazekas, K., Fleury, M., Heisinger, M., 2020. CaDiCaL, Kissat, Paracooba, Plingeling and Treengeling entering the SAT Competition 2020,
815 in: Balyo, T., Froylyks, N., Heule, M., Iser, M., Järvisalo, M., Suda, M. (Eds.), *Proc. of SAT Competition 2020 – Solver and Benchmark Descriptions*, University of Helsinki. pp. 51–53.
- Biere, A., Heule, M., van Maaren, H., 2009. *Handbook of Satisfiability*. volume 185. IOS press.
- 820 Biere, A., Lonsing, F., Seidl, M., 2011. Blocked clause elimination for QBF, in: *Proceedings of the Twenty-Third International Conference on Automated Deduction (CADE-2011)*, pp. 101–115.
- Bloem, R., Egly, U., Klampfl, P., Konighofer, R., Lonsing, F., 2014. Sat-based methods for circuit synthesis, in: *2014 Formal Methods in Computer-Aided Design (FMCAD)*, IEEE. pp. 31–34.
825
- Brand, D., 1993. Verification of large synthesized designs, in: *Proceedings of the International Conference on Computer-Aided Design (ICCAD-93)*, IEEE/ACM. pp. 534–537.
- Brayton, R.K., Hachtel, G.D., McMullen, C.T., Sangiovanni-Vincentelli, A.L.,
830 1984. *Logic Minimization Algorithms for VLSI Synthesis*. volume 2. Springer.
- Bryant, R.E., 1986. Graph-based algorithms for Boolean function manipulation. *IEEE Transactions on Computers* 100, 677–691.
- Buchfuhrer, D., Umans, C., 2011. The complexity of Boolean formula minimization. *Journal of Computer and System Sciences* 77, 142–153.
- 835 Codish, M., Cruz-Filipe, L., Ehlers, T., Müller, M., Schneider-Kamp, P., 2019. Sorting networks: to the end and back again. *Journal of Computer and System Sciences* 104, 184–201.
- Codish, M., Cruz-Filipe, L., Frank, M., Schneider-Kamp, P., 2014. Twenty-five comparators is optimal when sorting nine inputs (and twenty-nine for ten),
840 in: *Proceedings of the Twenty-Sixth International Conference on Tools with Artificial Intelligence (ICTAI-2014)*, IEEE. pp. 186–193.

- Cong, J., Ding, Y., 1996. Combinational logic synthesis for LUT based field programmable gate arrays. *ACM Transactions on Design Automation of Electronic Systems (TODAES)* 1, 145–204.
- 845 Cong, J., Minkovich, K., 2007. Optimality study of logic synthesis for lut-based fpgas. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 26, 230–239.
- Coste-Marquis, S., Berre, D.L., Letombe, F., Marquis, P., 2005. Propositional fragments for knowledge compilation and quantified boolean formulae, in: Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI-2005), pp. 288–293.
- 850
- Dadda, L., 1965. Some schemes for parallel multipliers. *Alta Frequenza* 34, 349–356.
- Finkbeiner, B., Tentrup, L., 2014. Fast DQBF refutation, in: Proceedings of the Seventeenth International Conference on Theory and Applications of Satisfiability Testing (SAT-2014), Springer. pp. 243–251.
- 855
- Fredkin, E., Toffoli, T., 1982. Conservative logic. *International Journal of Theoretical Physics* 21, 219–253.
- Fu, Z., Malik, S., 2006. On solving the partial MAX-SAT problem, in: Proceedings of the Ninth International Conference on Theory and Applications of Satisfiability Testing (SAT-2006), pp. 252–265.
- 860
- Garey, M.R., Johnson, D.S., 1990. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co.
- Gent, I.P., Hoos, H.H., Rowley, A.G., Smyth, K., 2003. Using stochastic local search to solve quantified boolean formulae, in: Proceedings of the Ninth International Conference on Principles and Practice of Constraint Programming (CP-2003), Springer. pp. 348–362.
- 865
- Gitina, K., Reimer, S., Sauer, M., Wimmer, R., Scholl, C., Becker, B., 2013. Equivalence checking of partial designs using dependency quantified Boolean formulae, in: Proceedings of the Thirty-First International Conference on Computer Design (ICCD-2013), pp. 396–403.
- 870
- Haaswijk, W., Soeken, M., Mishchenko, A., Micheli, G.D., 2018. SAT based exact synthesis using DAG topology families, in: Proceedings of the Fifty-Fifth IEEE Design Automation Conference (DAC-2018), IEEE. pp. 1–6.
- 875 Haaswijk, W., Soeken, M., Mishchenko, A., Micheli, G.D., 2019. SAT-based exact synthesis: Encodings, topology families, and parallelism. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 39, 871–884.

- 880 Hansen, E.A., Zhou, R., 2007. Anytime heuristic search. *Journal of Artificial Intelligence Research* 28, 267–297.
- Hansen, M., Yalcin, H., Hayes, J., 1999. Unveiling the ISCAS-85 benchmarks: A case study in reverse engineering. *IEEE Design & Test* 16, 72–80.
- Heule, M.J., Hunt, W.A., Wetzler, N., 2013. Verifying refutations with extended resolution, in: *International Conference on Automated Deduction*, Springer. pp. 345–359. 885
- Janota, M., 2018. Towards generalization in QBF solving via machine learning, in: *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Janota, M., Jordan, C., Klieber, W., Lonsing, F., Seidl, M., Gelder, A.V., 2016a. The QBFGallery 2014: The QBF competition at the FLoC olympic games. 890 *Journal on Satisfiability, Boolean Modeling and Computation* 9, 187–206.
- Janota, M., Klieber, W., Marques-Silva, J., Clarke, E., 2016b. Solving QBF with counterexample guided refinement. *Artificial Intelligence* 234, 1–25.
- Järvisalo, M., Berre, D.L., Roussel, O., Simon, L., 2012. The international SAT solver competitions. *AI Magazine* 33, 89–92.
- 895 Jr, E.M.C., Grumberg, O., Kroening, D., Peled, D., Veith, H., 2018. *Model checking*. MIT press.
- de Kleer, J., Williams, B., 1987. Diagnosing multiple faults. *Artificial Intelligence* 32, 97–130.
- Ling, A., Singh, D.P., Brown, S.D., 2005. FPGA logic synthesis using quantified 900 *Boolean satisfiability*, in: *International Conference on Theory and Applications of Satisfiability Testing*, Springer. pp. 444–450.
- Lonsing, F., Egly, U., 2017. DEPQBF 6.0: A search-based QBF solver beyond traditional QCDCL, in: de Moura, L. (Ed.), *Proceedings of the Twenty Sixth Conference on Automated Deduction (CADE-17)*, Springer International Publishing. pp. 371–384. 905
- Maini, A.K., 2007. *Digital Electronics: Principles, Devices and Applications*. John Wiley & Sons.
- Marques-Silva, J., Glass, T., 1999. Combinational equivalence checking using satisfiability and recursive learning, in: *Proceedings of the Conference on Design, Automation and Test in Europe (DATE-99)*, p. 33. 910
- Matsunaga, Y., 1996. An efficient equivalence checker for combinational circuits, in: *Proceedings of the Thirty-Third Annual Design Automation Conference (DAC-1996)*, pp. 629–634.
- McCluskey, Jr, E., 1956. Minimization of Boolean functions. *Bell System Technical Journal* 35, 1417–1444. 915

- McKay, B.D., Piperno, A., 2014. Practical graph isomorphism, II. *Journal of Symbolic Computation* 60, 94–112.
- Miller, J.F., Job, D., Vassilev, V.K., 2000. Principles in the evolutionary design of digital circuits–Part I. *Genetic Programming and Evolvable Machines* 1, 7–35.
- 920 Nielsen, M.A., Chuang, I.L., 2010. *Quantum Computation and Quantum Information*. Cambridge University Press.
- Roth, J.P., Karp, R.M., 1962. Minimization over Boolean graphs. *IBM journal of Research and Development* 6, 227–238.
- 925 Samulowitz, H., Memisevic, R., 2007. Learning to solve qbf, in: *Proceedings of the Twenty-Second Conference on Artificial Intelligence (AAAI-2007)*, pp. 255–260.
- Schaefer, M., Umans, C., 2002. Completeness in the polynomial-time hierarchy: A compendium. *SIGACT News* 33, 32–49.
- 930 Stockmeyer, L., 1977. The polynomial-time hierarchy. *Theoretical Computer Science* 3, 1–22.
- Toffoli, T., 1980. Reversible computing, in: de Bakker, J., van Leeuwen, J. (Eds.), *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP-80)*, pp. 632–644.
- 935 Umans, C., 2001. The minimum equivalent DNF problem and shortest implicants. *Journal of Computer and System Sciences* 63, 597–611.
- Vollmer, H., 2013. *Introduction to Circuit Complexity: A Uniform Approach*. Springer Science & Business Media.
- Wallace, C.S., 1964. A suggestion for a fast multiplier. *IEEE Transactions on Electronic Computers* , 14–17.
- 940

Appendix A. The ALU Circuit Families

Table A.9 gives the number of primary inputs (PIs), primary outputs (POs), and gates as a function of the parameter n . Some of the circuits use a proxy parameter k to avoid the use of logarithms.

945 The multiplexer (see Figure 12) is the same as the one used in the universal component cell. The demultiplexer is similar to the multiplexer and its architecture is shown in Figure A.24. Both can be generated for an arbitrarily sized input/output word.

950 The adder, shown in Figure A.25a and the subtractor, shown in Figure A.25b, are both ripple-carry. Due to the long propagation of carry, they are not used in the design of modern ICs. Used as a requirements circuit and with a sufficiently fast QBF solver Algorithm 3 should be able to enumerate all parallel adders and

Name	Notes	PIs	POs	Gates
n -mux	$n = 2^k, k \geq 1$	$2^k + k$	1	$2^k + k + 1$
n -demux	$n = 2^k, k \geq 1$	$k + 1$	2^k	$2^k + k$
n -add	$n \geq 1$	$2n + 1$	$n + 1$	$5n$
n -sub	$n \geq 1$	$2n + 1$	$n + 1$	$7n$
n -cmp	$n \geq 1$	$2n$	3	$3n + 4$
n -shift	$n \geq 1$	$2^n + n$	2^n	$2^n(3n - 2) + n + 2$
n -moa	$n = 2^k - 1, k \geq 2$	$2^k - 1$	k	$2^{k+1}(k - 2) + 2^k - k + 3$
n -mul	$n \geq 2$	n	$2n$	$6n^2 - 8n$

Table A.9: Size of the circuits in the ALU- n families

subtractors. An example of a real-world four-bit adder with carry look-ahead design is the 74283 IC, which is discussed later.

955 The n -bit comparator, shown in Figure A.25c and Figure A.25d, uses n XNOR gates to check for equality, and inverters and AND-gates to check for “greater than”. The “less than” signal is derived from the other two outputs with the help of an OR-gate and another inverter.

960 Barrel-shifters are used for shifting or rotating the bits in a bit-word and have important application in the design Floating-Point Units (FPUs) and cryptography cores. Figure A.26 shows a variable-size barrel-shifter. It shifts the input word to the right, losing the least-significant bits.

The barrel-shifter shown in Figure A.26 uses a cascade of multiplexers with two inputs and one output. The amount of shifting is specified as a binary number on the selector lines s_1, s_2, \dots, s_n . The total number of multiplexers is $2^n \times n$. There are some multiplexers with an input tied to ground on each column of the array shown in Figure A.26. We have 2^{n-1} such multiplexers per column where n is the column number. Each such multiplexer loses an AND-gate and an OR-gate. This reduces the number of gates as accounted for in 970 Table A.9. All multiplexers of a barrel-shifter reuse the same n inverters. The inverters are not shown in Figure A.26.

The n -bit multi-operand adder circuit, shown in Figure A.27, adds n single-bit numbers. A digital circuit that implements multi-operand addition is useful as a stand-alone circuit and also has application in multipliers Wallace (1964). 975 Multi-operand addition of single-bit numbers is also known as bit-counting or binary vector addition. Applications of satisfiability to optimization use bit-counting for implementing “at-least- k ” or “at-most- k ” constraints (Fu and Malik, 2006).

980 The multi-operand adder is implemented as a chain of multi-operand full-adders (see Figure A.27a). Each full-adder adds one bit to a binary number and consists of k half-adders where k equals the number of bits necessary for representing the binary number (see Figure A.27a). The full-adders can be implemented without a carry-out bit, which saves one AND-gate. The multi-




Figure A.24: Variable size demultiplexer circuit


operand adder uses full-adders of increasing size. The first adder has one input,
 985 the second and third have two inputs, the next four have three inputs, etc.

This particular implementation of a multi-operand adder has no application
 in digital electronics due to the long primary inputs to outputs propagation
 time, but it is useful in constraint programming. The chained multi-operand
 adder can be used as a requirements circuit to allow the automatic discovery of
 990 advanced topologies such as the one in Wallace (1964) or Dadda (1965) trees.


Figure A.28 shows the architecture of a variable size multiplier that imple-
 ments the standard “pen and paper” method. The multiplier consists of
 two subsystems: an array of AND-gates that computes partial products (see
 Figure A.28a and a network of adders that sum the partial products (see Fig-
 995 ure A.28b).

Appendix B. Reversible Circuits from the ALU Families


Figure B.29a shows a 4-to-1 multiplexer. Its functioning can be verified by
 analyzing the circuit. It consists of three CSWAP gates. If the two selector
 lines s_1 and s_2 are both low, then there are no swapped values and the input
 1000 i_1 is copied to the output o . If the left-most CSWAP gate is activated with a
 value of one on s_1 , then i_4 goes to i_2 . If the second selector s_2 is also high, then
 the value of i_4 from i_2 will go to i_1 and o . This is correct because when both




(a) A ripple-carry adder



(b) A ripple-carry subtractor



(c) Equality comparator



(d) "Greater-than" comparator

Figure A.25: Variable size adder, subtractor, and comparator

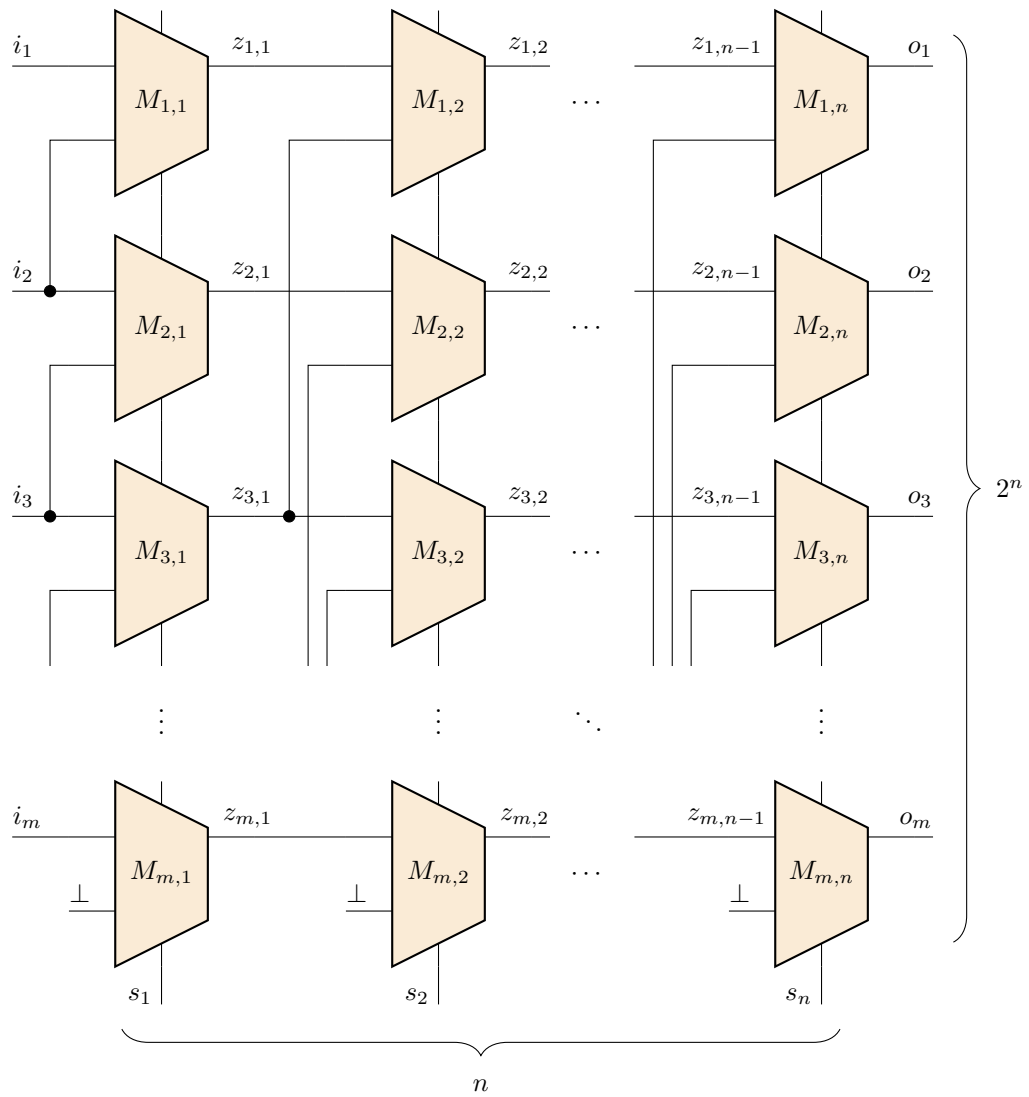

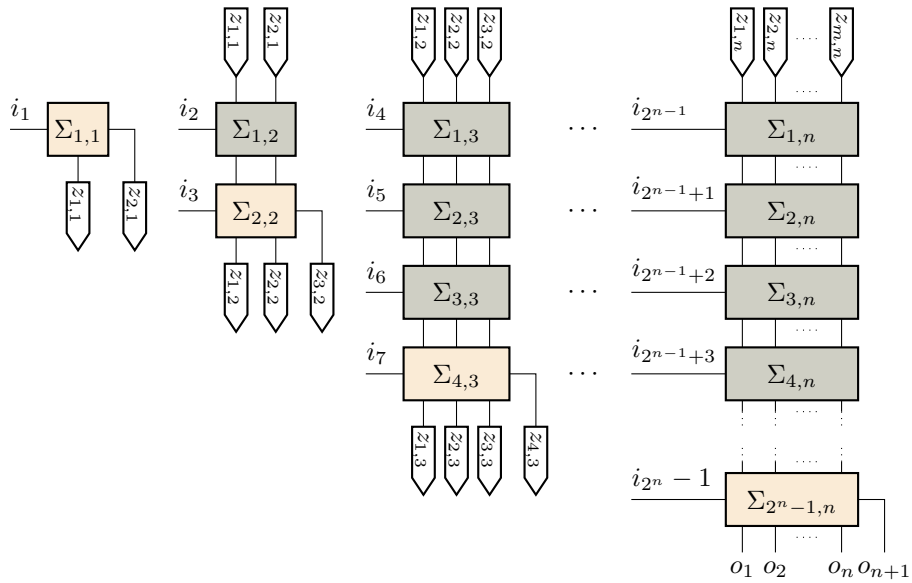


Figure A.26: Variable size barrel-shifter




(a) A half-adder chain with an optional carry out bit

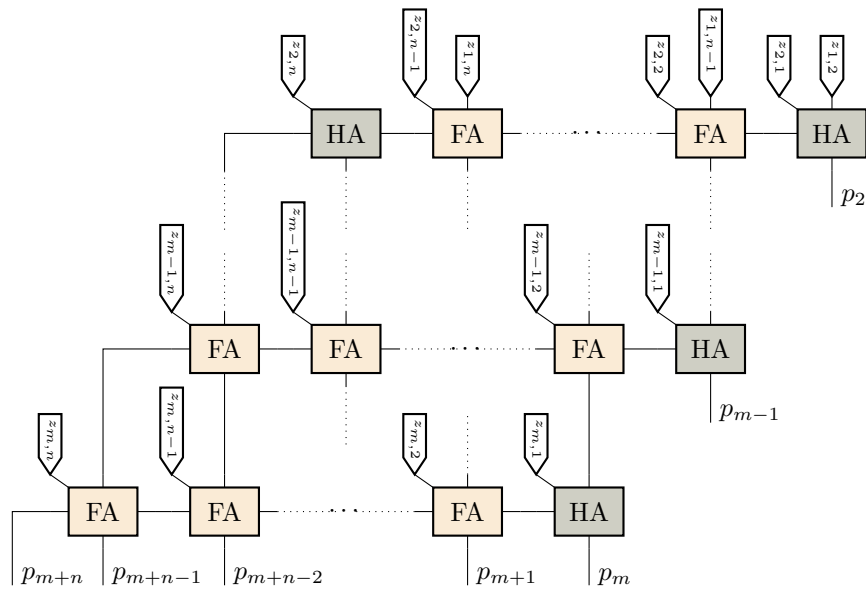


(b) A ladder of half-adder-chains for multi-operand addition

Figure A.27: A binary multi-operand adder of variable size



(a) Partial products



(b) Adder network

Figure A.28: Variable size multiplier

s_1 and s_2 are selected we expect i_4 to send its value to o . The remaining two combinations can be checked in a similar manner.

1005 Figure B.29b shows a 1-to-4 demultiplexer. Similar to the multiplexer from Figure B.29a it is also made of three CSWAP gates. Similar the standard basis multiplexer and demultiplexer, the ones from the reversible basis are very similar.

1010 Figure B.29c and Figure B.29d show a full-adder and a full-subtractor, respectively. They are both made of four gates but the similarities end there. A multi-operand adder with three inputs is equivalent to a regular full-adder, hence Figure B.29c shows both.

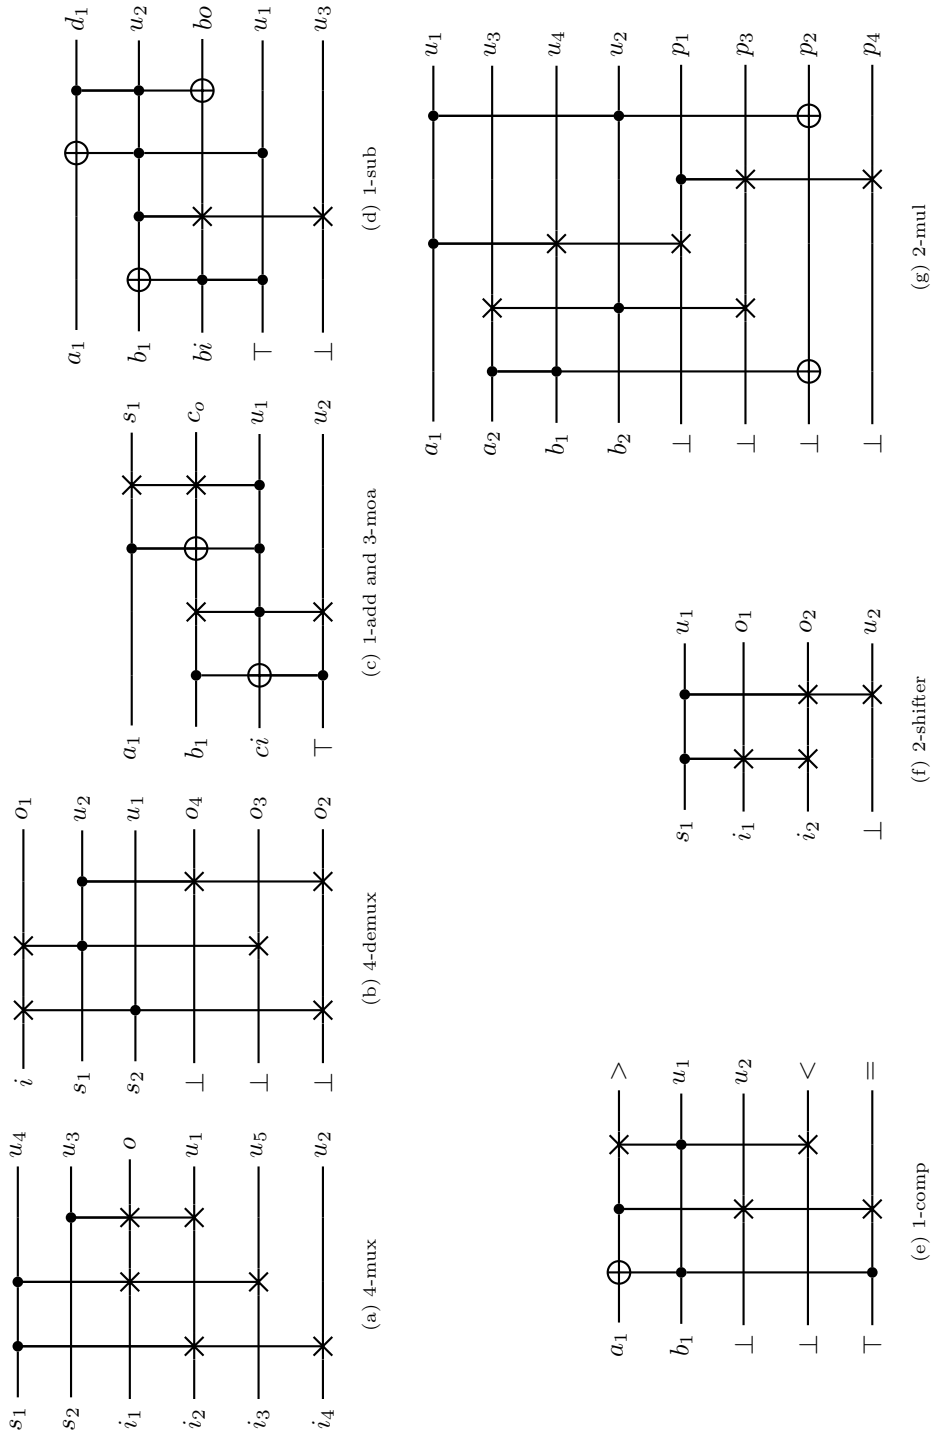


Figure B.29: Representatives of optimal reversible ALU-4 circuits