

PROMISE: Property Mining for Sequential Synthesis

Jiahui Xu*, Jordi Cortadella†, and Lana Josipović*

*ETH Zurich, Department of Information Technology and Electrical Engineering, Zurich, Switzerland

†UPC Barcelona, Department of Computer Science, Barcelona, Spain

Abstract—Modularity—composing a large system using individually designed units—is an essential practice in hardware design. Yet, modularity might compromise quality: when individually designed units are put together, some of their states may become unreachable and, consequently, the logic that implements them is redundant. Sequential synthesis aims to remove redundant circuit logic by leveraging state unreachability. It critically depends on invariants—relations between signals and registers that hold in all reachable states—to prove the validity of redundancies. Yet, existing invariant generation techniques are mostly problem-specific (for a particular circuit or a property) or reliant on localized reasoning. We propose PROMISE, a fast circuit redundancy removal strategy. PROMISE exploits the rich information from simulation traces and uses efficient polynomial-time algorithms to infer global circuit invariants, optimizing the circuit and aiding other sequential synthesis procedures. Experiments show that PROMISE effectively optimizes circuits produced by high-level synthesis tools. PROMISE is open-sourced and available at github.com/ETHZ-DYNAMO/promise.

I. INTRODUCTION

FPGAs offer flexibility, energy efficiency, and high performance; our goal is to make their programming as smooth as traditional software development. The key to devising large designs is modularity: a design is composed of small, manageable pieces that can be easily integrated. Yet, composability comes with a cost. It incurs substantial overhead [1] due to the interface logic required to assemble the modules correctly: when individually designed units are put together, some of their states may become unreachable and, consequently, the logic that implements them is redundant.

Sequential synthesis is a family of logic synthesis techniques that remove redundant circuit logic leveraging state unreachability. Its success critically depends on the available invariants—relations between *flip-flops* (FFs) that hold in all reachable states—to prove the validity of redundancies. However, useful invariants are hard to find: the space of possible properties is enormous, so strategies often rely on localized reasoning (e.g., a subgraph of adjacent circuit gates/FFs) or produce property- or circuit-specific invariants.

We present PROMISE, an invariant generation framework for sequential synthesis. PROMISE leverages the information available in circuit simulation traces and uses efficient polynomial-time algorithms to generate a system of invariants commonly found in circuits produced by *high-level synthesis* (HLS). The invariants characterize the unreachable circuit states, which we use to optimize the circuit's encoding and enhance the effectiveness of existing sequential synthesis approaches. Our result shows that PROMISE-generated invariants bring tangible and justifiable improvements to the circuit quality.

II. BACKGROUND

This section reviews the foundations in formal verification and logic synthesis techniques that PROMISE relies on to optimize redundant circuit logic.

A. Model Checking


In a *finite-state machine* (FSM), such as a sequential circuit, an *invariant* is a property that holds in every reachable state [2]. Model checking [3]–[5] is a formal verification technique that formally proves whether a certain property holds for an FSM. If the property fails, it provides a counterexample. *k-induction* [6], [7] is an important model checking algorithm for invariant properties. *k-induction* verifies if the following two conditions hold: (1) the property holds in any *k* steps starting from the initial state; (2) for any *k* consecutive states where the property holds, the property holds after any transition. In practice, a very large bound *k* is needed for concluding non-trivial properties: when *k* is not big enough, the induction engine will return a counter-example, in which none of the states is reachable. An invariant can act as a constraint during model checking to rule out certain unreachable states (i.e., the model checker ignores the states that violate the invariants), thus it can speed up the verification of the *k-induction* proof of another safety property [8]–[10].

Model checking algorithms like *k-induction* not only apply to verifying the circuit's correctness against a certain specification, as we will see in the next section, but they also have important applications in circuit optimization.

B. Sequential Synthesis

Sequential synthesis refers to circuit transformations that preserve the circuit behavior on the reachable states and allow arbitrary changes in the unreachable states [11], [12]. It leverages state unreachability (i.e., some values never appear in the input of the combinational circuit) to uncover optimizations that are unattainable in combinational synthesis [7], [13]: state reachability is formally verified before being applied to reduce the circuit area. This principle can be generalized into a *suggest-guarantee-optimize* procedure independent of the circuit transformation and the property being verified. This procedure is divided into these steps:


- *Suggest*: Candidate invariants—which indicate state unreachability—are identified using a custom heuristic.
- *Guarantee*: The suggested invariants must be guaranteed by formal verification. The surviving invariants are passed to the *optimize* phase.
- *Optimize*: The circuit is optimized using the unreachable state space generated by the proven invariants.



(a) The unoptimized circuit has a *redundant encoding* (the 5 FFs with shaded background can be reduced to 3 as in Figure 1c).

Cycle	f_1	f_2	f_3	f_4	f_5	f_6
C_0	0	0	0	0	0	0
C_1	1	0	0	0	0	1
C_2	0	1	0	0	0	1
C_3	0	0	1	0	0	1
C_4	0	0	0	1	0	1
C_5	0	0	0	0	1	1
C_6	0	0	0	0	0	0

(b) A simulation trace of the circuit in Figure 1a, which is also the same as the entire reachable set of states.



(c) A circuit diagram (some details omitted) after our encoding optimization (discussed in Section VII-A).

Fig. 1: Circuit with a redundant encoding. We identify optimization opportunities and invariants from simulation traces.


However, current invariant generation methods (the “suggest” phase) do not synergize well with the “optimize” phase: they overlook encoding optimization opportunities and are unable to find useful invariants for effective sequential synthesis, as we will demonstrate next.

III. ARE WE GETTING THE MOST OUT OF OUR INVARIANTS?

This section motivates new approaches for circuit encoding optimizations and invariant generation.

A. Missed Encoding Optimization Opportunities

Figure 1a depicts a sequential circuit with six FFs (f_1 – f_6), all initialized to 0. The input to the start of the FF chain (f_1 – f_5) can be 1 only if the value of f_6 is 0 in the current state; f_6 becomes one after the start of the chain receives a 1, and becomes 0 after the start of the chain outputs a 1. Figure 1b describes a simulation trace of 7 clock cycles (i.e., 7 circuit states) with the values of the FFs. The circuit has simple control logic, and this trace contains all of its reachable states. Notice that, in any state, at most one FF in f_1 – f_5 has an output value equal to 1. In theory, it is possible to replace the chain of 5 FFs with only 3 FFs as in Figure 1c.



(a) The shaded parts can be removed without compromising the circuit’s functionality.

Cycle	pi_1	f_1	f_2	f_3	f_4	f_5	f_6	f_7
C_0	1	0	0	0	0	0	0	0
C_1	1	1	0	0	1	0	0	0
C_2	1	0	1	0	1	1	0	0
C_3	1	1	1	0	1	1	1	0
C_4	0	0	0	1	1	1	1	1
C_5	0	1	1	0	0	1	1	1
C_6	0	0	1	0	0	0	1	1

(b) Simulation trace with 7 states and 7 state variables.

Fig. 2: The invariant extracted from a simulation trace can improve the effectiveness of sequential synthesis (see Section III-B).

This redundancy appears in many circuits and is often *unintentional*. However, state-of-the-art sequential synthesis approaches [7], [13], [14] typically would not optimize away this redundancy, since existing state encoding optimization approaches require complete reachability information [15] (e.g., state-transition graph, or a BDD of the set of reachable states) that is impractical to obtain for large circuits. PROMISE detects these redundancies in the form of a linear inequality¹:

$$f_1 + f_2 + f_3 + f_4 + f_5 \leq 1. \quad (1)$$

When this relation is true for all reachable states, we know that there are only 6 reachable combinations of these 5 FFs, and PROMISE can reencode the 5 FFs into 3 FFs.

B. Localized and Inexpressive Invariants

Consider the circuit in Figure 2a. All FFs are initialized to 0. A chain of FFs f_4 – f_7 receives a 1 only when the 3-bit word w_1 (consisting of f_3 , f_2 , and f_1 ; f_3 is the MSB) has a value less than 5. Counter w_1 counts up when a 1 is loaded into the FF chain f_4 – f_7 , and counts down as f_7 outputs a 1. Since the counter’s value directly corresponds to the number of 1’s in f_4 – f_7 , the “<” gate always evaluates to 1. Theoretically, we can remove or simplify all the shaded wires without altering the functionality. Without the support of invariants, sequential synthesis approaches (e.g., `SCORR` in ABC) require a large induction depth to simplify the circuit (as large as the

¹We use “&”, “|”, and “~” to denote logical *and*, *or*, and *not*. We use “+” for arithmetic sum.

maximum value of the counter). However, existing techniques are unable generate useful invariants for optimizing this circuit, as they are limited to Boolean clauses over a subgraph of adjacent nodes (e.g., a clause $f_1|f_2|f_3$) [16], simple implication or equality between signals [17], [18], are specialized to particular circuits [8], [9], or specific to the invariants relevant for proving one particular property [19].

On the other hand, PROMISE can detect and prove the following invariant, which cannot be obtained using the aforementioned circuit invariant generation techniques:

$$\underbrace{f_1 + 2^1 \cdot f_2 + 2^2 \cdot f_3}_{=w_1} = f_4 + f_5 + f_6 + f_7, \quad (2)$$

which specifies the relation between the 3-bit word w_1 and the FF chain f_4 – f_7 . Once assisted with this invariant, sequential synthesis can easily remove all the shaded logic.

IV. PROMISE: GENERATING CIRCUIT INVARIANTS FROM SIMULATION

We propose PROMISE, an invariant generation framework to enable more effective sequential synthesis. PROMISE leverages polynomial-time algorithms to efficiently *suggest* candidate linear invariants from the circuit’s simulation traces. PROMISE *guarantees* the validity of the invariants using a model checker, and uses the invariants to *optimize* the circuit’s encoding or assist other sequential synthesis approaches [7], [20].

Suppose $c_1, \dots, c_{N+1}$ are coefficients and $f_1, \dots, f_N$ are the circuit’s FFs; PROMISE generates:

- *Inequalities* with coefficients $c_1, \dots, c_N \in \{0, -1, +1\}$ and an integer c_{N+1} :

$$c_1 \cdot f_1 + c_2 \cdot f_2 + \dots + c_N \cdot f_N + c_{N+1} \leq 0; \quad (3)$$

- *Equalities* with integer coefficients:

$$c_1 \cdot f_1 + c_2 \cdot f_2 + \dots + c_N \cdot f_N + c_{N+1} = 0. \quad (4)$$

Such linear relations frequently appear in the control logic of HLS-produced circuits [9], since the control status of a circuit is often realized using linearly evolving elements like counters. PROMISE makes better *suggestions* and therefore, enables better *optimizations*.


The rest of the paper is organized as follows: [Section V](#) describes PROMISE’s *suggest* phase: two mathematical methods for inferring the candidate invariants. [Section VI](#) describes the *guarantee* phase; a standard model checking algorithm to verify the invariants. [Section VII](#) describes the *optimize* phase: how to exploit the properties to optimize the circuit. [Section VIII](#) evaluates the effectiveness of PROMISE.

V. THE SUGGEST PHASE: MATHEMATICAL METHODS FOR PROPERTY MINING

This section describes the mathematical methods for identifying properties in the form of linear equalities and inequalities that we classified above.

Cycle	f_5	f_6	$\sim f_5$	$\sim f_6$
C_0	0	0	1	1
$C_{1\dots5}$	0	1	1	0
C_6	1	1	0	0

(a) f_5 and f_6 from [Figure 1b](#) and their complemented values.



(b) Conflict graph for [Figure 3a](#) colored using 3 colors.

Fig. 3: Coloring a conflict graph using the algorithm in [Section V-A](#).

A. Extracting Mutually Exclusive Sets of Signals by Coloring a Conflict Graph

This subsection devises a systematic strategy for identifying mutually exclusive signals from simulation traces. This enables optimizations such as the one in [Figure 1](#).

FFs $f_1, \dots, f_N$ are mutually exclusive if they are never simultaneously 1, i.e., $f_1 + \dots + f_N \leq 1$ [21]. Mutual exclusiveness translates to these relations: (1) *One-hot*: mutually exclusive FFs are one-hot. (2) *Implication*: if f_1 and f_2 are mutually exclusive, then $\sim(f_1 \& f_2) = 1$ must hold. Now, suppose that f_1 and $\sim f_2$ are mutually exclusive, then $\sim(f_1 \& (\sim f_2)) = (f_1 \rightarrow f_2) = 1$.

PROMISE infers mutually exclusive sets of FFs by *coloring* a *conflict graph* [12], [22]. PROMISE constructs the conflict graph from simulation data as follows:

- For each FF output f_i and its complement $\sim f_i$, add a node to the graph.
- For each simulation cycle, if the corresponding signals of any two nodes are both 1, add an edge between those nodes.
- For each FF f_i , add an edge between f_i and its complement $\sim f_i$ (to avoid a trivial relation like: $\sim f_i + f_i \leq 1$).

Graph coloring assigns different colors to nodes connected by an edge. For scalability, we apply a greedy coloring which runs in linear time [22]; this heuristic leads to excellent results, as we will see in [Section VIII](#). After coloring, each color denotes a set of mutually exclusive FFs. For each color $C := \{f_1, \dots, f_N\}$, we devise the following invariant:

$$\sum_{f_i \in C} f_i \leq 1, \quad (5)$$

which is used by PROMISE to assist other sequential synthesis approaches and to optimize the encoding—[Section VII-A](#) describes how PROMISE carries out the optimization.

For example, [Figure 3](#) describes a conflict graph built from f_5 , f_6 , and their complemented values. After coloring, we get 3 sets: $\{f_5, \sim f_6\}$, $\{f_6\}$, and $\{\sim f_5\}$. The first set corresponds to a non-trivial relation: $f_5 + \sim f_6 \leq 1$.

B. Extracting Equalities Using Gaussian Elimination

This subsection presents a systematic strategy for deriving a system of linear equalities of the signals in the circuit from simulation traces. This improves the effectiveness of sequential synthesis of the circuit, such as the one in [Figure 2](#).

PROMISE infers linear equalities like [Equation 2](#) from simulation traces. Since any valid invariant holds in all reachable

states, it must be at least valid for the observed simulation data. This requirement is equivalent to a system of linear constraints for possible values of coefficients $c_i \in C$ ($f_1^{s_1}$ denotes the value of f_1 in state s_1):

$$\underbrace{\begin{bmatrix} f_1^{s_1} & f_2^{s_1} & \dots & f_N^{s_1} & 1 \\ f_1^{s_2} & f_2^{s_2} & \dots & f_N^{s_2} & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ f_1^{s_M} & f_2^{s_M} & \dots & f_N^{s_M} & 1 \end{bmatrix}}_{\mathbf{A}: M \times (N+1)} \cdot \underbrace{\begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_N \\ c_{N+1} \end{bmatrix}}_{\mathbf{c}: (N+1) \times 1} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}, \quad (6)$$

which is equivalent to equations like Equation 4. A solution to the system ($\mathbf{c} = [c_1 \dots c_{N+1}]^T$) determines the coefficients in Equation 4. This system may have infinitely many solutions; that is, there could be an infinite number of equations to include in the set of invariants. Yet, many equations are redundant—for example, $2 \cdot f_1 + 2 \cdot f_2 = 2 \cdot f_3$ is merely a scaled version of $f_1 + f_2 = f_3$. PROMISE adopts a standard approach [23]–[25] based on Gaussian elimination [26] (with $O(n^3)$ complexity) to determine a minimal set of vectors, ensuring that no vector can be expressed as a linear combination of the others. In this way, PROMISE efficiently infers equations such as Equation 2 and leverages them to aid the sequential synthesis of circuits in Figure 2.

Consider the linear system constructed according to the simulation cycles 0...5 (no cycle 6) in Figure 2b:

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 & 1 & 1 & 1 \end{bmatrix} \cdot \begin{bmatrix} c_1 \\ c_2 \\ \vdots \\ c_7 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}. \quad (7)$$

One possible set of solutions to this system is:

$$[1 \ 0 \ 2 \ -1 \ 1 \ -1 \ 1]^T, \quad (8)$$

$$[0 \ 1 \ 1 \ 0 \ -1 \ 0 \ 0]^T. \quad (9)$$

We plug them separately in Equation 4 (replacing the values of $c_1 \dots c_7$) to obtain two relations:

$$f_1 + 2 \cdot f_3 + f_5 + f_7 = f_4 + f_6, \quad (10)$$

$$f_2 + f_3 = f_5. \quad (11)$$

These relations hold in cycles 0...5, but the second relation failed in cycle 6. In general, the inferred relation might be spurious since a simulation trace is not guaranteed to cover all states. Therefore, the validity of the invariants must be guaranteed by formal verification, as described next.

VI. THE GUARANTEE PHASE: PROVING THE SUGGESTED INVARIANTS

The result of the previous section is a set of *candidate* invariants inferred from simulation. The simulation traces only contain a subset of the reachable states. PROMISE attempts to verify that the conjunction of the invariants is valid in all

reachable states using a model checker. If the proof failed, PROMISE adds the states in the counterexample trace to the set of simulation states to correct the set of invariants. If the invariant holds, we use it to optimize the circuit. For example, the model checker will return a counterexample trace that contains cycle 6 (in Figure 2b) to disprove Equation 11. Having the new state, PROMISE can reexecute the procedure in Section V-B to correctly infer Equation 2.

The verified invariants enable PROMISE to apply the circuit optimization techniques described in the next section.

VII. THE OPTIMIZE PHASE: CAPITALIZING ON THE INVARIANTS

This section describes how PROMISE uses the invariants from Section V to optimize the circuits.

A. Applying Encoding Optimizations to the Circuit

This subsection describes how PROMISE performs encoding optimization using inequalities like Equation 3.

In general, for a set of signals $F := \{s_1, \dots, s_N\}$, a system of inequalities like Equation 3 (that we aim to prove in Section V-A) describes that the sum of the signals is within a set of values $K := \{k_1, \dots, k_M\}$, that is:

$$(s_1 + \dots + s_N) \in \{k_1, \dots, k_M\}. \quad (12)$$

For example, Figure 4a describes a subcircuit with 3 FFs, f_1, f_2 , and f_3 . Assume that the following inequality holds in all reachable states:

$$0 \leq f_1 + f_2 + f_3 \leq 1. \quad (13)$$

Here, $F := \{f_1, f_2, f_3\}$ and $K := \{0, 1\}$. Since f_1, f_2 , and f_3 can only take 4 possible combinations of values, the subcircuit with 3 FFs can be substituted with another subcircuit with 2 FFs as in Figure 4c. The following describes how we construct the substituted circuit.

The substitution must preserve the circuit's functionality. Following a standard pattern as a previous work on encoding optimization [15], PROMISE uses an *encoding circuit* $Enc(\cdot)$ that takes the inputs to the original FFs (e.g., i_1, i_2, i_3 in Figure 4c) and sends the encoded inputs (e_1, e_2) to the encoded FFs (x_1, x_2). PROMISE uses another *decoding circuit* $Dec(\cdot)$ to convert the encoded FFs' outputs back to the original outputs (o_1, o_2, o_3). For any reachable FF value assignment to a set of FFs $f_1, \dots, f_N$, the encoding and decoding functions cancel each other's effect, that is, $f_1, \dots, f_N = Dec(Enc(f_1, \dots, f_N))$, and the decoding circuit preserves the initial state.

PROMISE uses a *state mapping table* to decide on the encoding scheme and the encoding and decoding circuits. It maps a set of reachable FF values—all combinations that satisfy Equation 12—in the unoptimized circuit into the corresponding FF values in the optimized circuit. Figure 4b describes a state mapping table that maps a state in Figure 4a to a state in Figure 4c. For instance, according to the last entry in Figure 4b, when both the original and optimized circuits start from the initial state, if the original circuit (Figure 4a)




Fig. 4: Circuit substitution: 3 FFs to 2 FFs.

reached a state $(f_1, f_2, f_3) = (1, 0, 0)$ after applying certain input sequence, the optimized circuit (Figure 4c) must have $(x_1, x_2) = (1, 1)$. Different mappings potentially have different effects on the cost of the encoding circuit. Although exploring different mappings is beyond the scope of this paper, we will see in Section VIII that the mappings used by PROMISE successfully simplify the circuit.

The encoding and decoding circuits can be built from the state mapping table using a standard logic synthesis technique (e.g., a Karnaugh map). We apply a pre-processing step for FFs with complemented outputs ($\sim f_i$), since they might not exist originally: we insert a pair of inverters at the input and output of each complemented FF ($\sim f_i$), and complement the initial value of that FF [7].

PROMISE performs encoding optimization only for F and K that reduce the number of FFs. The number of FFs in the *substituted subcircuit* $|R_F|$ after applying the encoding optimization is given by

$$|R_F| := \lceil \log_2 \left(\sum_{k \in K} \binom{|F|}{k} \right) \rceil, \quad (14)$$

where $\binom{a}{b}$ denotes the number of b -combinations of a elements. The expression in $\log_2(\cdot)$ describes the number of combinations of FF values that sum up to each value in the set $K \in \{k_1, \dots, k_M\}$. Each combination corresponds to a row in the state mapping table; therefore, the number of FFs needed to represent these states is the $\log_2(\cdot)$ of the number of entries. Encoding optimization is profitable if $|R_F|$ is less than the number of FFs in the original subcircuit $|F|$. In Equation 1, $|F| = 5$ and $K \in \{0, 1\}$, therefore, $|R_F| = \lceil \log_2(\binom{5}{0} + \binom{5}{1}) \rceil = 3$. Since $|R_F| < |F|$, applying the transformation reduces the number of FFs. Consider another case when $K = \{0, 1, 2, 3\}$. Here, $|R_F| = \lceil \log_2(1 + 5 + 10 + 10) \rceil = 5$, and therefore the transformation is not desirable.




Fig. 5: PROMISE's *suggest-guarantee-optimize* circuit optimization procedure.

B. Enhancing Sequential Synthesis Using Invariants

In addition to the encoding optimization above, PROMISE embeds other sequential synthesis procedures in its *optimize* phase and assists them using the invariants in Section V. Without loss of generality, here, we discuss the synergy between PROMISE and one particular sequential synthesis flow [7]; we believe that this also applies to other sequential synthesis approaches (e.g., Marakkalage et al. [14]).

Signal correspondance [7] is a sequential synthesis technique that detects, proves, and merges sequentially equivalent nodes and FFs. This approach relies on k -induction internally to carry out the proof. Without a suitable invariant, the depth k required to prove the property might be prohibitively large. However, in the presence of invariants provided by PROMISE, they can efficiently prove signal equivalence and perform more effective optimizations.

This concludes the optimization methods used by PROMISE. We now evaluate their effectiveness.

VIII. EVALUATION

This section evaluates the effectiveness of PROMISE. Our research artifact is publicly available [27].

A. Methodology

We implemented PROMISE—a *suggest-guarantee-optimize* procedure that uses the techniques we have seen so far. Figure 5 describes our invariant generation and circuit optimization flow. Unless stated otherwise, the parameters discussed apply to all experiments.

Benchmarks. We have a set of benchmarks generated using different circuit compilation tools: *Dynatomic* [28] (an MLIR-based HLS tool that converts an input C code to a dynamically scheduled dataflow circuit) and *XLS* [29] (converts a design specified in its input language into a statically scheduled

TABLE I: Effectiveness of **Suggest** and **Optimize**: The features of PROMISE (EN and IN) consistently improve the cost. Compared with SC only, SC + EN + IN achieved average reductions of 31% 6-LUT and 29% FF.

Benchmark	Enabled features	AIG results			6-LUT results		
		FF	AIG	Depth	FF	6-LUT	Depth
factorial (xls)	SC	37	1566	40	37	214	7
	SC EN	37	1152	40	37	163	8
	SC EN IV	36	1156	40	36	163	8
	SC EN IV	37	1151	40	37	162	8
iterative division (xls)	SC	54	593	20	54	126	5
	SC EN	53	516	17	53	118	4
	SC EN IV	53	561	18	53	127	4
	SC EN IV	54	517	18	54	118	4
iterative sqrt (xls)	SC	43	702	42	43	136	9
	SC EN	43	617	42	43	120	8
	SC EN IV	42	651	42	42	121	8
	SC EN IV	43	622	42	43	122	8
simple loop (xls)	SC	42	626	42	42	121	8
	SC	34	243	17	34	70	4
	SC EN	32	211	13	32	59	3
	SC EN IV	31	225	13	31	59	3
factorial (Dynamatic)	SC	32	211	13	32	58	3
	SC EN	31	220	13	31	60	3
	SC EN IV	31	220	13	31	60	3
	SC EN IV	31	220	13	31	60	3
iterative division (Dynamatic)	SC	497	4817	35	497	986	7
	SC EN	482	4582	35	482	931	7
	SC EN IV	367	4039	35	367	784	7
	SC EN IV	374	2250	28	374	536	5
iterative sqrt (Dynamatic)	SC	328	2411	30	328	523	5
	SC	392	3034	30	392	752	7
	SC EN	310	2369	26	310	545	6
	SC EN IV	135	906	17	135	204	4
simple loop (Dynamatic)	SC	241	874	16	241	302	4
	SC EN	146	924	17	146	217	4
	SC EN IV	146	924	17	146	217	4
	SC EN IV	146	924	17	146	217	4
matvec (Dynamatic)	SC	605	5325	47	605	1189	11
	SC EN	581	4859	41	581	1103	10
	SC EN IV	161	892	17	161	224	4
	SC EN IV	308	877	15	308	355	4
bigc (Dynamatic)	SC	184	910	16	184	251	4
	SC	46	326	20	46	73	5
	SC EN	34	268	16	34	58	4
	SC EN IV	24	156	12	24	45	3
gaussian (Dynamatic)	SC	27	150	12	27	44	3
	SC EN	24	162	16	24	43	3
	SC EN IV	24	162	16	24	43	3
	SC EN IV	24	162	16	24	43	3
gemver (Dynamatic)	SC	262	2435	32	262	586	8
	SC EN	240	2173	31	240	531	7
	SC EN IV	142	1331	27	142	339	5
	SC EN IV	197	1107	27	197	346	5
stencil 2d (Dynamatic)	SC	151	1198	27	151	326	5
	SC	393	3681	32	393	839	8
	SC EN	366	3409	32	366	769	7
	SC EN IV	243	2690	36	243	648	8
2mm (Dynamatic)	SC	278	1815	28	278	527	5
	SC EN	221	1946	28	221	503	5
	SC EN IV	221	1946	28	221	503	5
	SC EN IV	221	1946	28	221	503	5
gemver (Dynamatic)	SC	460	4262	32	460	1002	8
	SC EN	405	3763	31	405	879	7
	SC EN IV	229	2433	29	229	515	6
	SC EN IV	314	1592	23	314	467	5
stencil 2d (Dynamatic)	SC	221	1692	23	221	396	5
	SC	1405	12816	32	1405	2943	8
	SC EN	1285	11686	34	1285	2674	7
	SC EN IV	817	9732	48	817	2279	11
2mm (Dynamatic)	SC	999	6237	31	999	1616	5
	SC EN	758	6655	33	758	1486	6
	SC EN IV	407	3585	33	407	877	8
	SC EN IV	363	3039	30	363	775	6
gemver (Dynamatic)	SC	262	2763	37	262	728	9
	SC EN	280	1414	28	280	457	5
	SC EN IV	215	1614	28	215	442	5
	SC EN IV	215	1614	28	215	442	5
gemver (Dynamatic)	SC	1117	10271	36	1117	2429	8
	SC EN	1031	9338	33	1031	2176	7
	SC EN IV	616	7192	42	616	1586	10
	SC EN IV	801	4639	34	801	1222	5
gemver (Dynamatic)	SC	584	5100	34	584	1121	6
	SC EN	584	5100	34	584	1121	6
	SC EN IV	584	5100	34	584	1121	6
	SC EN IV	584	5100	34	584	1121	6

circuit). They are decent targets for evaluating PROMISE: Dynamatic generates circuits by connecting individually designed dataflow units to ensure performance and flexibility, but suffers from the area overhead due to this modularity. XLS generates coarse-grained, statically scheduled blocks that communicate via handshake interfaces—offering flexibility, but also introducing redundancy. Each benchmark without memory accesses (XLS cannot handle memory accesses) has two functionally identical implementations in two tools. We also include several standard HLS benchmarks (*matvec*, *bigc*,

gemver, *kernel 2mm*, *stencil 2d*) [30]. For each synthesized top-level module, we added a wrapper to simplify the communication between the circuit and its environment. Each module includes a "go" input pin to start the execution and a "done" output pin that signals when the computation is complete. All designs have been functionally verified using a set of representative input vectors.

Simulate. We synthesize the circuit using Yosys [31]. We use Verilator [32] to simulate each circuit with random inputs for 4 rounds, each for 25000 cycles.

Suggest. We use the algorithms described in Section V to infer invariants from the set of simulated states. We only extract relations between the FFs that dictate the control status in the design (e.g., the data validity flag of a buffer).

Guarantee. We use the rIC3 model checker [33] to carry out the formal verification of the conjunction of the invariants inferred in the *suggest* phase. rIC3 uses the PDR model checking algorithm [19], [34]. If rIC3 reports that the property is invalid, the states in the counterexample will be added to the set of simulated states. We iterate between the *suggest* and *guarantee* phases until rIC3 confirms that the conjunction of the invariants holds in the circuit.

Optimize. Our baselines are the unoptimized circuit and the circuit optimized only with the `scorr` command (the signal correspondence optimization) available in ABC. We then complement the `scorr` command with our encoding optimization (Section VII-A) and invariants (Section V).

We need to specify our invariants when running the `scorr` command: `scorr` allows declaring a certain PO as a constraint; during the proving step, the algorithm ignores the states where the constraint fails. We construct a logic cone as a constraint from the conjunction of the invariants. We strip away the logic cone after the optimization.

Our optimization metrics. To report the area results, we use ABC to convert the circuit to an AIG network using the `st` command and map the circuit to an FPGA LUT network using `if -K 6`. As with any sequential synthesis optimization, we do not alter the circuit's latency (i.e., the clock cycle count); therefore, we do not report it, as it remains consistent across all designs. Whenever the complexity permits, we apply sequential equivalence checking to formally verify that our modifications preserve the behaviors of the original circuits.

B. Effectiveness of PROMISE: Suggest and Optimize

Table I reports the ABC synthesis results of PROMISE. The columns grouped with "Enabled features" indicate which optimization techniques are applied: Column EN indicates whether encoding optimization is applied, SC indicates whether signal correspondence [7] (`scorr` in ABC) is applied, and IN indicates whether the invariants are used in `scorr`. The columns AIG results and 6-LUT results report the achieved area and delay. The best synthesis result of each single benchmark is highlighted in green.

Design with a network of modular units. The benchmarks labeled with "(Dynamatic)" are generated by Dynamatic [28].




Fig. 6: Effect of using *different induction depths* and *inclusion or exclusion of invariants* when using `scorr` in ABC.

Dynamic implements handshake modules at the operation level (e.g., a multiplier has its handshake interface) to ensure the best composability and latency [35], but has a very large resource overhead [1]. In these benchmarks, signal correspondence alone (i.e., the rows with only SC) sees improvement over pure combinational synthesis (i.e., the rows with no techniques). On the other hand, we see a substantial logic reduction when signal correspondence is used with our additions—either encoding optimization (EN) or invariants (IN); in each benchmark, the best metric is achieved by applying at least one of our optimizations (i.e., EN or IN).

Design with a single module. The benchmarks labeled “(xls)” are generated by XLS [29]. As a single-module design, the circuit has less redundancy to explore. When adding a wrapper to simplify communication (see Section VIII-A), some interface logic in XLS-produced circuits becomes redundant. Our invariant supports the signal correspondence procedure in removing this redundancy.

Encoding optimization vs. invariant-enhanced signal correspondence. Generally, it is expected to see that encoding optimization reduces the FF count at the cost of a more complex logic. However, circuits with SC + EN consistently have fewer AIG or 6-LUT nodes than those with SC alone. This is because reducing the FFs also reduces the combinational logic’s primary inputs and outputs, which potentially simplifies the logic function. While the encoding optimization may increase the maximum logic depth, this effect is offset by the logic savings enabled by the invariants—the rows with SC + EN + IN always have the best logic depth.

TABLE II: Comparing the effectiveness of PROMISE’s invariants vs. using the reachable set of states as an invariant for benchmark “simple loop (Dynamic)”.

Invariants	AIG results			6-LUT results		
	FF	AIG	Depth	FF	6-LUT	Depth
No invariants	34	268	16	34	58	4
Reachable states	46	152	13	46	65	3
Promise’s invariants	27	150	12	27	44	3

TABLE III: Runtime statistics of **Suggest** and **Guarantee**. Columns **Sim**, **Proof**, **Linear equality**, and **Mutual exclusion** are reported in seconds.

XLS-Produced Circuits					
Benchmark	Sim	Proof	Linear equality	Mutual exclusion	Iter.
factorial	0	0.1	0	0	0
iterative division	0	0.4	0	0	1
iterative sqrt	0	0.6	0	0	2
simple loop	0	0.1	0	0	0
Dynamic-Produced Circuits					
Benchmark	Sim	Proof	Linear equality	Mutual exclusion	Iter.
factorial	0.8	4.7	0	0.2	1
iterative division	0.8	1.7	0	0	0
iterative sqrt	1.7	2.4	0.1	0.1	0
simple loop	0	0.2	0	0	0
matvec	0.4	3.4	0	0	0
bicg	0.8	7.6	0	0	0
gaussian	1.2	16.3	0.1	0.1	0
gemver	7.6	119.2	0.9	0.7	0
stencil 2d	0.8	30.9	0	0.1	0
2mm	5.7	334.4	0.6	0.5	0

Increasing induction depth vs. using invariants. Figure 6 describes the effect of inclusion and exclusion of invariants and varying the induction depth (1 to 250) when using `scorr` for benchmarks `matvec`, `bicg`, and `gaussian`. In each benchmark, increasing the induction depth alone does not improve the area. Surprisingly, the number of 6-LUTs after applying `scorr` *increases* when the induction depth is large. This shows that PROMISE-generated invariants greatly improve the effectiveness of `scorr`.

Effectiveness of PROMISE’s invariants vs. reachability. BDD-based reachability analysis produces a set of reachable states—this can be formulated as an invariant (i.e., given the set of all reachable states $\{s_1, s_2, \dots, s_N\}$, we can format them as a invariant: $(state = s_1) \vee (state = s_2) \vee \dots \vee (state = s_N)$). Table II reports the synthesis result of the benchmark “simple loop (Dynamic)”. The table reports the optimization result after using `scorr`, assisted by no invariants (No invariants), the set of reachable states as an invariant (Reachable states), and PROMISE’s invariants (PROMISE’s invariants). From the result, using the set of all reachable states as an invariant is less effective than our invariants.

We omit the further comparison with encoding optimization that requires a complete set of reachable states (e.g., Sentovich et al. [15]) due to the poor scalability of the reachability analysis [9]. For one of our medium-sized benchmarks, `matvec` (Dynamic), the reachability analysis in ABC could not converge after 48 hours.

C. Scalability of PROMISE: Suggest and Guarantee

Table III reports the runtime statistics of the property mining procedures (Section V) and time needed to prove the properties (Section VI). Column **Iter.** reports the number of failed proof attempts (that trigger re-execution of the **suggest** phase). Columns **Mutual exclusion** and **Linear equality** report the total runtime of the techniques introduced in Section V-A and Section V-B. Column **Sim** reports the total time spent on circuit simulation. Our results indicate that circuit simulation is a scalable source for gathering information for invariant generation—the total simulation time is always under 10 seconds. This can be further improved by parallelization. Column **Proof** reports the total runtime of running the rIC3 model checker. All model checking runs converge in a reasonable time without any abstraction technique applied (often necessary for model checking to converge [1], but not needed in our experiments).

Scalability of property mining. In benchmarks with data-dependent control flow (e.g., the “factorial” benchmarks generated using Dynamatic), few corrections were done (see Section VI) before our property mining procedure could infer a verifiable invariant. Yet, our property mining procedures are scalable: most property generation takes less than 1 second. All the mutual exclusion properties are generated within 0.1 seconds. The largest benchmark “kernel 2mm (Dynamatic)” only requires less than 2 seconds to generate the properties.

IX. RELATED WORK

Dataflow designs produced from high-level languages.

There has been an increasing interest in HLS tools to produce dataflow processing networks [28], [29], [35]–[41]. Dataflow processing systems are composable and deliver high performance due to their dynamic nature [42]. This paradigm comes with a resource overhead (up to 50% of the dataflow circuit logic is redundant, bypassing multiplexers and buffer slots that are never occupied with valid data [1]), and many research works aim at removing it [1], [35], [43]–[46]. Unlike these efforts, PROMISE offers a more general solution, independent of specific circuit generation methods.

Encoding optimization. Sentovich et al. [15] focus on optimizing encoding in circuits generated from ESTEREL. They greedily remove one register at a time and require computation and analysis of the entire reachable state space. Computing the reachable state space is typically impractical without problem-specific abstractions. Empirical evidence suggests that induction in the presence of invariants is more scalable than reachability analysis [9]. We efficiently infer encoding optimization opportunities using simulation. Sentovich et al. [21] also describe an FF optimization approach that leverages the knowledge of their circuits; yet, their technique is specific to ESTEREL-produced circuits.

Redundancy removal in sequential circuits. Sequential synthesis techniques like those of Mishchenko et al. [7] and Marakkalage et al. [14] remove redundancy using induction; yet, they do not take advantage of any invariants. Many research efforts explore redundancy removal approaches in a

limited setting, such as combinational circuits [47], feedback-free circuits, and a particular redundancy structure [1], [48]. These approaches do not aim to improve the circuits’ encoding, and our invariants can be used to improve the effectiveness of their optimization.

Generating invariants for circuits. There exist techniques for automatically deriving inductive invariants for dataflow circuits [8], [9] to improve verification runtime, but they are limited to a set of predefined units. There is a family of model checking algorithms that aim to synthesize an inductive invariant to prove the safety property [19], [34], [49], [50]. These methods are specific to generating an inductive invariant for a single safety property, operate on unreachable states, and the operators used to construct the invariants have either too limited expressivity (i.e., pure Boolean formulas [19]) or are too general and overfit the observed states [49] (i.e., any first-order logic operator). Our method operates on simulation traces, and we aim to infer expressions (e.g., Equation 3 and Equation 4) that commonly appear in the control logic of the circuits generated from hardware compilers.

Property generation from simulation. Using simulation—referred to as *dynamic analysis* in software engineering—to derive loop invariants has been studied in the software verification domain [23]–[25], [51]. Empirical results show that these invariants can support proving the equivalence between the program before and after a certain optimization [51]. Execution traces, from both software and hardware, have also been leveraged to infer temporal specifications [52], [53]; they generate more expressive properties than ours (e.g., LTL formulas instead of invariants). However, these properties are not readily applicable to circuit optimization. Inspired by them, we adopted a similar insight—using simulation to support verification—but specifically for circuit optimization.

X. CONCLUSION

We presented PROMISE, a framework that utilizes simulation data to detect redundancy in the sequential circuit’s state encoding and extract invariants to speed up the circuit verification. PROMISE efficiently detects state encoding optimization opportunities in cases where conventional techniques are prohibitively expensive, and derives linear invariants to make existing sequential synthesis procedures fundamentally more effective. PROMISE synergizes comprehensive simulation-based testing, formal verification, and logic synthesis to uncover new opportunities for more effective and scalable optimization. PROMISE is open-sourced and available at github.com/ETHZ-DYNAMO/promise.

ACKNOWLEDGEMENT

This work has been supported by the Swiss National Science Foundation (grant number 215747) and the ETH Future Computing Laboratory (donation from Huawei Technologies).

REFERENCES

- [1] J. Xu, E. Murphy, J. Cortadella, and L. Josipović, "Eliminating excessive dynamism of dataflow circuits using model checking," in *Proceedings of the 31st ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, Monterey, CA, Feb. 2023, pp. 27–37.
- [2] S. Chaki and A. Gurfinkel, "BDD-based symbolic model checking," in *Handbook of Model Checking*. Springer International Publishing, 2018, pp. 219–245. [Online]. Available: https://doi.org/10.1007/978-3-319-10575-8_8.
- [3] E. Clarke, K. McMillan, S. Campos, and V. Hartonas-Garmhausen, "Symbolic model checking," in *Proceedings of the 8th International Conference on Computer Aided Verification*, New Brunswick, NJ, Jun. 1996, pp. 419–22.
- [4] E. M. Clarke, T. A. Henzinger, and H. Veith, "Introduction to model checking," in *Handbook of Model Checking*. Springer International Publishing, 2018, pp. 1–26. [Online]. Available: https://doi.org/10.1007/978-3-319-10575-8_1.
- [5] C. Kern and M. R. Greenstreet, "Formal verification in hardware design: A survey," *ACM Transactions on Design Automation of Electronic Systems*, vol. 4, no. 2, pp. 123–93, Apr. 1999.
- [6] M. Sheeran, S. Singh, and G. Stålmarck, "Checking safety properties using induction and a SAT-solver," in *Proceedings of the 3rd International Conference on Formal Methods in Computer-Aided Design*, Austin, TX, Nov. 2000, pp. 127–44.
- [7] A. Mishchenko, M. Case, R. Brayton, and S. Jang, "Scalable and scalably-verifiable sequential synthesis," in *Proceedings of the 27th International Conference on Computer-Aided Design*, San Jose, CA, Nov. 2008, pp. 234–41.
- [8] S. Chatterjee and M. Kishinevsky, "Automatic generation of inductive invariants from high-level microarchitectural models of communication fabrics," *Formal Methods in System Design*, vol. 40, no. 2, pp. 147–69, 2012.
- [9] J. Xu and L. Josipović, "Automatic inductive invariant generation for scalable dataflow circuit verification," in *Proceedings of the 42nd International Conference on Computer-Aided Design*, San Francisco, CA, Oct. 2023, pp. 1–9.
- [10] C. A. Furia, B. Meyer, and S. Velder, "Loop invariants: Analysis, classification, and examples," *ACM Computing Surveys*, vol. 46, no. 3, Jan. 2014.
- [11] E. M. Sentovich, K. J. Singh, C. Moon, H. Savoj, R. K. Brayton, and A. Sangiovanni-Vincentelli, "Sequential circuit design using synthesis and optimization," in *Proceedings 1992 IEEE International Conference on Computer Design*, Cambridge, MA, Oct. 1992, pp. 328–33.
- [12] G. De Micheli, *Synthesis and Optimization of Digital Circuits*. New York: McGraw-Hill, 1994.
- [13] R. Brayton and A. Mishchenko, "ABC: An academic industrial-strength verification tool," in *Proceedings of the 22nd International Conference on Computer Aided Verification*, Edinburgh, Jul. 2010, pp. 24–40.
- [14] D. S. Marakkalage, E. Testa, W. L. Neto, A. Mishchenko, G. De Micheli, and L. Amarù, "Scalable sequential optimization under observability don't cares," in *Proceedings of 2024 Design, Automation & Test in Europe Conference & Exhibition*, Valencia, Spain, Mar. 2024, pp. 1–6.
- [15] E. M. Sentovich, H. Toma, and G. Berry, "Latch optimization in circuits generated from high-level descriptions," in *Proceedings of the 15th International Conference on Computer-Aided Design*, San Jose, CA, Nov. 1996, pp. 428–35.
- [16] M. Case, A. Mishchenko, and R. Brayton, "Cut-based inductive invariant computation," in *Proceedings of the 17th International Workshop on Logic Synthesis*, Lake Tahoe, CA, Jun. 2008, pp. 253–58.
- [17] G. Cabodi, S. Nocco, and S. Quer, "Strengthening model checking techniques with inductive invariants," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 28, no. 1, pp. 154–58, Jan. 2009.
- [18] C. van Eijk, "Sequential equivalence checking based on structural similarities," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 19, no. 7, pp. 814–19, Jul. 2000.
- [19] A. R. Bradley, "SAT-based model checking without unrolling," in *Proceedings of the 12th International Workshop on Verification, Model Checking, and Abstract Interpretation*, Austin, TX, Jan. 2011, pp. 70–87.
- [20] ABC: System for sequential logic synthesis and formal verification, Commit: ca78f5e, berkeley-abc. [Online]. Available: <https://github.com/berkeley-abc/abc/tree/ca78f5e65308df420ffc5c709e6d37caf97e40b>.
- [21] E. M. Sentovich, H. Toma, and G. Berry, "Efficient latch optimization using exclusive sets," in *Proceedings of the 34th Annual Design Automation Conference*, Anaheim, CA, Jun. 1997, pp. 8–11.
- [22] Graph coloring. [Online]. Available: https://en.wikipedia.org/w/index.php?title=Graph_coloring&oldid=1279292851.
- [23] M. D. Ernst, J. Cockrell, W. G. Griswold, and D. Notkin, "Dynamically discovering likely program invariants to support program evolution," in *Proceedings of the 21st International Conference on Software Engineering*, Los Angeles, CA, May 1999, pp. 213–24.
- [24] T. Nguyen, D. Kapur, W. Weimer, and S. Forrest, "Using dynamic analysis to discover polynomial and array invariants," in *Proceedings of 34th International Conference on Software Engineering*, Zurich, Switzerland, Jun. 2012, pp. 683–93.
- [25] R. Sharma, S. Gupta, B. Hariharan, A. Aiken, P. Liang, and A. V. Nori, "A data driven approach for algebraic loop invariants," in *Proceedings of 22nd European Symposium on Programming*, Rome, Italy, Mar. 2013, pp. 574–92.
- [26] Kernel (linear algebra). [Online]. Available: [https://en.wikipedia.org/w/index.php?title=Kernel_\(linear_algebra\)&oldid=1261076439#Computation_by_Gaussian_elimination](https://en.wikipedia.org/w/index.php?title=Kernel_(linear_algebra)&oldid=1261076439#Computation_by_Gaussian_elimination).
- [27] J. Xu, *Research artifact of Promise: Property mining for sequential synthesis*, Jul. 2025. [Online]. Available: <https://doi.org/10.5281/zenodo.16333795>.
- [28] Dynamic, Commit: 999dc3c, EPFL-LAP. [Online]. Available: <https://github.com/EPFL-LAP/dynamic/tree/999dc3ce2fb95eac1dd39cad441fbfd6b8389aee>.
- [29] Xls: Accelerated hw synthesis, Google, Inc. [Online]. Available: <https://github.com/google/xls>.
- [30] L.-N. Pouchet, *Polybench: The polyhedral benchmark suite*, 2012. [Online]. Available: <https://sourceforge.net/p/polybench/wiki/Home/>.
- [31] Yosys Open SYnthesis Suite, Commit: 29cf4a9. [Online]. Available: <https://github.com/YosysHQ/yosys/tree/29cf4a919062fe7b6a6f21b946dbec15a3d2114a>.
- [32] W. Snyder, *Verilator: Verilator open-source SystemVerilog simulator and lint system*, 2025. [Online]. Available: <https://www.veripool.org/verilator>.
- [33] Y. Su, Q. Yang, Y. Ci, T. Bu, and Z. Huang, "The r1c3 hardware model checker," *arXiv preprint arXiv:2502.13605*, Feb. 2025.
- [34] N. Een, A. Mishchenko, and R. Brayton, "Efficient implementation of property directed reachability," in *Proceedings of 14th International Conference on Formal Methods in Computer-Aided Design*, Austin, TX, Oct. 2011, pp. 125–34.
- [35] L. Josipović, R. Ghosal, and P. lenne, "Dynamically scheduled high-level synthesis," in *Proceedings of the 26th ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, Monterey, CA, Feb. 2018, pp. 127–36.
- [36] L. Josipović, A. Guerrieri, and P. lenne, "Dynamic: From C/C++ to dynamically scheduled circuits," in *Proceedings of the 28th ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, Monterey, CA, Feb. 2020, pp. 1–10.
- [37] Vivado design suite user guide: High-level synthesis, Xilinx Inc., 2018. [Online]. Available: https://www.xilinx.com/support/documentation/sw_manuals/xilinx2017_4/ug902-vivado-high-level-synthesis.pdf.
- [38] Y. Chi, L. Guo, J. Lau, Y.-k. Choi, J. Wang, and J. Cong, "Extending high-level synthesis for task-parallel programs," in *Proceedings of the 29th IEEE Symposium on Field-Programmable Custom Computing Machines*, Orlando, FL, May 2021, pp. 204–13.
- [39] L. Guo, Y. Chi, J. Wang, et al., "AutoBridge: Coupling coarse-grained floorplanning and pipelining for high-frequency HLS design on multi-die FPGAs," in *Proceedings of the 29th ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, Virtual Event, Mar. 2021, pp. 81–92.
- [40] L. Guo, P. Maidee, Y. Zhou, et al., "RapidStream: Parallel physical implementation of FPGA HLS designs," in *Proceedings of the 30th ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, Virtual Event, Feb. 2022, pp. 1–12.
- [41] A. Elakhras, A. Guerrieri, L. Josipović, and P. lenne, "Unleashing parallelism in elastic circuits with faster token delivery," in *Proceedings of the 32nd International Conference on Field-Programmable Logic and Applications*, Belfast, UK, Aug. 2022, pp. 253–61.
- [42] A. Elakhras, A. Guerrieri, L. Josipović, and P. lenne, "Survival of the fastest: Enabling more out-of-order execution in dataflow circuits," in *Proceedings of the 32nd International Symposium on Field-Programmable Gate Arrays*, Monterey, CA, Mar. 2024, pp. 44–54.
- [43] J. Xu and L. Josipović, "Suppressing spurious dynamism of dataflow circuits via latency and occupancy balancing," in *Proceedings of the 32nd ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, Monterey, CA, Mar. 2024, pp. 188–98.
- [44] R. Nigam, S. Thomas, Z. Li, and A. Sampson, "A compiler infrastructure for accelerator generators," in *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, Virtual, Apr. 2021, pp. 804–17.
- [45] A. Elakhras, J. Xu, M. Erhart, P. lenne, and L. Josipović, "ElasticMiter: Formally verified dataflow circuit rewrites," in *Proceedings of the 30th International Conference on Architectural Support for Programming Languages and Operating Systems*, Rotterdam, The Netherlands, Apr. 2025, pp. 293–308.
- [46] L. Josipović, A. Marmet, A. Guerrieri, and P. lenne, "Resource sharing in dataflow circuits," in *Proceedings of the 30th IEEE Symposium on Field-Programmable Custom Computing Machines*, New York, May 2022, pp. 1–9.
- [47] S.-Y. Lee, H. Rienar, A. Mishchenko, R. K. Brayton, and G. De Micheli, "A simulation-guided paradigm for logic synthesis and verification," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 41, no. 8, pp. 2573–86, May 2022.
- [48] Q. Tan, A. Gupta, and S. Malik, "Usage-based RTL subsetting for hardware accelerators," in *Proceedings of the 41st International Conference on Computer-Aided Design*, San Diego, CA, Dec. 2022, pp. 1–9.
- [49] A. Goel and K. Sakallah, "AVR: Abstractly verifying reachability," in *Proceedings of 26th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, Dublin, Ireland, Apr. 2020, pp. 413–22.
- [50] M. L. Case, A. Mishchenko, and R. K. Brayton, "Automated extraction of inductive invariants to aid model checking," in *Proceedings of the 7th International Conference on Formal Methods in Computer Aided Design*, Nov. 2007, pp. 165–72.
- [51] B. Churchill, O. Padon, R. Sharma, and A. Aiken, "Semantic program alignment for equivalence checking," in *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*, Phoenix, AZ, Jun. 2019, pp. 1027–40.
- [52] M. Gabel and Z. Su, "Symbolic mining of temporal specifications," in *Proceedings of the 30th International Conference on Software Engineering*, Leipzig, Germany, May 2008, pp. 51–60.
- [53] W. Li, A. Forin, and S. A. Seshia, "Scalable specification mining for verification and diagnosis," in *Proceedings of the 47th Design Automation Conference*, Anaheim, CA, Jun. 2010, pp. 755–60.