

Efficiently Synthesizing Lowest Cost Rewrite Rules for Instruction Selection

Ross Daly

Stanford University

Stanford, CA, USA

rdaly525@cs.stanford.edu



Jackson Melchert

Stanford University

Stanford, CA, USA

melchert@stanford.edu



Pat Hanrahan

Stanford University

Stanford, CA, USA

hanrahan@cs.stanford.edu



Caleb Donovick

Stanford University

Stanford, CA, USA

donovick@cs.stanford.edu



Priyanka Raina

Stanford University

Stanford, CA, USA

praina@stanford.edu



Caleb Terrill

Stanford University

Stanford, CA, USA

cdterrill26@gmail.com

Clark Barrett

Stanford University

Stanford, CA, USA

barrett@cs.stanford.edu



Abstract—Compiling programs to an instruction set architecture (ISA) requires a set of rewrite rules that map patterns consisting of compiler instructions to patterns consisting of ISA instructions. We synthesize such rules by constructing SMT queries, whose solutions represent two functionally equivalent programs. These two programs are interpreted as an instruction selection rewrite rule. Existing work is limited to single-instruction ISA patterns, whereas our solution does not have that restriction. Furthermore, we address inefficiencies of existing work by developing two optimized algorithms. The first only generates unique rules by preventing synthesis of duplicate and composite rules. The second only generates lowest-cost rules by preventing synthesis of higher-cost rules. We evaluate our algorithms on multiple ISAs. Without our optimizations, the vast majority of synthesized rewrite rules are either duplicates, composites, or higher cost. Our optimizations result in synthesis speed-ups of up to 768× and 4004× for the two algorithms.

I. INTRODUCTION

As we approach the end of Moore’s law and Dennard scaling, drastically improving computing performance and energy efficiency requires designing domain-specific hardware architectures (DSAs) or adding domain-specific extensions to existing architectures [22]. As a result, many DSAs have been developed in recent years [4], [8], [24], [27], [30], each with its own custom instruction set architecture (ISA) or ISA extension.

Targeting such ISAs from a compiler’s intermediate representation (IR) requires a custom library of instruction selection rewrite rules. A rewrite rule is a mapping of an IR pattern to a functionally equivalent ISA pattern. Manual specification of rewrite rules is error-prone, time-consuming, and often incomplete. It is therefore desirable to automatically generate valid rewrite rules.

When specifying instruction selection rewrite rules, there are two common cases. When ISAs have complex instructions, rewrite rules will often map multi-instruction IR patterns to a single ISA instruction. When ISAs have simple instructions, rewrite rules will often map a single IR instruction to a multi-instruction ISA pattern. A rewrite rule generation tool should be able to create rewrite rules for both cases. We call such rewrite rules *many-to-many* rules.

Generating instruction selectors is not a new idea. Most relevant to this work is Gulwani et al. [21] who use a satisfiability modulo theories (SMT) solver to synthesize a loop-free program that is functionally equivalent to a given specification. Their approach is called component-based program synthesis (CBPS), as each synthesized program must include functional components from a given component library. Buchwald et al. [6] use and extend CBPS to efficiently generate multi-instruction loop-free IR programs equivalent to a single ISA instruction program; that is, they solve the many-to-one rewrite rules synthesis problem. However, multi-instruction ISA programs cannot be synthesized.

Both of these algorithms produce many *duplicate* rules, which are removed during a post-processing step. As we show, this adds significant additional cost. Another issue is that CBPS as currently formulated does not incorporate the notion of optimizing for cost. In practice, we often want only the set of lowest-cost rules, making it unnecessary (and expensive) to generate equivalent higher-cost rules.

This paper presents an algorithm for automatically generating a complete set of many-to-many rewrite rules. We address the above issues by preventing the synthesis of both duplicate and high-cost rules at rule generation time, using exclusion techniques. As a further optimization, we generate

rules in stages and exclude *composite* rules, i.e., rules that can be composed of smaller rules found in previous stages. These ensure we produce a small but complete set of rewrite rules. Compared to previous work, our approach eliminates unnecessary rules and significantly reduces the time required to produce the unique necessary ones.

Our contributions are as follows:

- We define generalized component-based program synthesis (GCBPS) as the task of synthesizing two functionally equivalent programs using two component libraries. We then present an SMT-based synthesis approach inspired by Gulwani et al. to solve it.
- We present an iterative algorithm *genAll* to generate all unique many-to-many rules up to a given size. We identify a set of equivalence relations for patterns encoded as programs and for rules that map IR programs to ISA programs. We use these relations to enumerate and exclude duplicate rules. Furthermore, we directly exclude composite rewrite rules. These result in up to a $768\times$ synthesis speed-up.
- We present an algorithm *genAll_{LC}* which generates only the lowest-cost rules by incorporating a cost metric in addition to excluding duplicate and composite rewrite rules. This results in a synthesis speed-up up to $4004\times$.

The rest of the paper is organized as follows. Section II discusses instruction selection, existing rule generation methods, SMT, and program synthesis. Section III describes a program synthesis query for generating many-to-many rules. Section IV presents an algorithm for generating only unique rewrite rules and defines duplicates and composites. Section V presents an algorithm for synthesizing only the lowest-cost rules. Section VI evaluates both algorithms, and Section VII discusses limitations and further optimizations.

II. BACKGROUND AND RELATED WORK

A. Instruction Selection

Instruction selection is the task of translating code in the compiler’s intermediate representation (IR) to functionally equivalent code for a target ISA. Typically, a library of rewrite rules is used in instruction selection. A rewrite rule is a mapping from an IR pattern consisting of IR instructions to a functionally equivalent ISA pattern consisting of ISA instructions. Such patterns can be expression trees or directed acyclic graphs (DAGs).

Significant work has been devoted to developing rewrite rule tiling algorithms to perform instruction selection [1], [5], [12], [14]–[17], [19], [26], [29]. For each rule in the rule library, a tiling algorithm first finds all fragments from the IR program in which the rule’s IR pattern exactly matches that fragment. Then, the instruction selector finds a tiling of these matches that completely covers the basic block and minimizes the total rule cost according to some cost metric.

Simple instruction selectors only handle tree-based IR patterns, which is inefficient for reused computations. Modern instruction selectors like LLVM use DAG-based matching that

allows for both richer rules and better tiling. Koes et al. [26] describe a similar near-optimal DAG-based instruction selection algorithm [5]. We want to generate rules that can be used with such modern DAG-based instruction selectors.

B. Generating Instruction Selectors

Generating instruction selectors from instruction semantics has been a topic of research interest [6], [7], [9], [10], [23]. Dias and Ramsey [10] introduce an algorithm for generating rewrite rules based on a declarative specification of the ISA. While this solves part of the many-to-many rule task, their work relies on an existing set of algebraic rewrite rules for synthesizing semantically equivalent rules. Our work uses SMT for the instruction and program semantics. However, incorporating certain kinds of algebraic rewrite rules could be an avenue for future optimizations.

Daly et al. [9] propose a way to synthesize instruction selection rewrite rules from the register-transfer level (RTL) specification of a processor. Their algorithm requires a set of pre-specified IR patterns. In contrast, we can efficiently synthesize rules that consider all possible multi-instruction IR patterns up to a given size. Their approach for synthesizing complex instruction constants and handling floating point types could be combined with the approaches in this paper.

The most relevant to this work is the work by Buchwald et al. [6], which leverages component-based program synthesis to generate rules with multi-instruction IR patterns and single-instruction ISA patterns. In contrast, our work synthesizes rules with both multi-instruction IR patterns and multi-instruction ISA patterns. We additionally prevent the synthesis of duplicate, composite, and high-cost rewrite rules, unlike any of the above approaches.

C. Program Synthesis and Equivalence

We use SMT-based program synthesis to enumerate a complete set of instruction selection rewrite rules. In program synthesis enumeration, it is common to remove equivalent solutions [3]. We use the equivalence relation defined in Section IV-A to determine equivalent rewrite rules. In prior work [2], observational equivalence (i.e., programs with the same semantics) has been used for de-duplication [2], however observational equivalence does not take into account the structure of the program, which is essential for rewrite rule pattern matching.

D. Logical Setting and Notation

We work in the context of many-sorted logic (e.g., [13]), where we assume an infinite set of variables of each sort and the usual notions of terms, formulas, assignments, and interpretations. Terms are denoted using non-boldface symbols (e.g., X). Boldface symbols (e.g., \mathbf{X}) are used for sets, tuples, and multisets, whose elements are either terms or other collections of terms. $\mathbf{Y} := (Y_1, \dots, Y_N)$ defines a tuple, where $|\mathbf{Y}| = N$ and Y_i refers to the i -th element. $\mathbf{Z} := \{z^n\}$ defines a multiset, where the multiplicity of element z is $n \in \mathbb{N}$. Both ψ and ϕ are used to denote formulas. $\psi(\mathbf{X})$ is a formula

whose free variables are a subset of \mathbf{X} . We use $\mathcal{M} \models \psi(\mathbf{X})$ to denote the *satisfiability* relation between the interpretation \mathcal{M} and the formula ψ . Assuming \mathbf{X} is a collection of variables, $\mathcal{M}_{\mathbf{X}}$ denotes the *assignment* to those variables induced by \mathcal{M} . For an assignment α , we write $\alpha \models \psi(\mathbf{X})$ if $\mathcal{M} \models \psi(\mathbf{X})$ for every interpretation \mathcal{M} such that $\mathcal{M}_{\mathbf{X}} = \alpha$.

E. Component-based Program Synthesis

CBPS is a program synthesis task introduced by Gulwani et al. The inputs to the task are:

- A *specification* $\mathbf{S} := (\mathbf{I}^S, O^S, \phi_{spec}(\mathbf{I}^S, O^S))$ containing a tuple of input variables \mathbf{I}^S , a single output variable O^S , and a formula $\phi_{spec}(\mathbf{I}^S, O^S)$ relating the inputs and the output.
- A *library of components* (e.g., instructions) \mathbf{K} , where the k -th component $\mathbf{K}_k := (\mathbf{I}_k, O_k, \phi_k(\mathbf{I}_k, O_k))$ consists of a tuple of input variables \mathbf{I}_k , a single output variable O_k , and a formula $\phi_k(\mathbf{I}_k, O_k)$ defining the component's semantics.

An example component for an addition instruction is shown below using the theory of bit-vectors, QF_BV , where $BV_{[n]}$ is an n -bit sort and $+_{[n]}$ is addition modulo 2^n .

$$((I_0 : BV_{[16]}, I_1 : BV_{[16]}), O : BV_{[16]}, I_0 +_{[16]} I_1 = O)$$

The task is to synthesize a valid program functionally equivalent to the specification using each component from \mathbf{K} exactly once.

For notational convenience, we group together the set of all inputs and outputs of the components: $\mathbf{W} := \cup_{(\mathbf{I}_k, O_k, _) \in \mathbf{K}} (O_k \cup (\cup \mathbf{I}_k))$. Gulwani et al. encode the program structure using a *connection constraint*: $\phi_{conn}(\mathbf{L}, \mathbf{I}^S, O^S, \mathbf{W})$. This is a formula representing how the program inputs (\mathbf{I}^S) and program output (O^S) are connected via the components. The connections are specified using *location variables* \mathbf{L} . We do not go into the details of how location variables encode connections (they are in [21]). It is sufficient for our purposes to know that these are integer variables, and an assignment to them uniquely determines a way of connecting the components together into a program. The *program semantics* ϕ_{prog} are defined as the components' semantics conjoined with the connection constraint:

$$\phi_{prog}(\mathbf{L}, \mathbf{I}^S, O^S, \mathbf{W}) := \left(\bigwedge_k \phi_k(\mathbf{I}_k, O_k) \right) \wedge \phi_{conn}(\mathbf{L}, \mathbf{I}^S, O^S, \mathbf{W}). \quad (1)$$

They define a verification constraint that holds if a particular program is both well-formed (specified using a well-formedness constraint ψ_{wfp}) and satisfies the specification ϕ_{spec} :

$$\phi_{verif} := \psi_{wfp}(\mathbf{L}) \wedge \forall \mathbf{I}^S, O^S, \mathbf{W}. \quad (2)$$

$$\phi_{prog}(\mathbf{L}, \mathbf{I}^S, O^S, \mathbf{W}) \implies \phi_{spec}(\mathbf{I}^S, O^S).$$

A *synthesis formula* ϕ_{synth} existentially quantifies \mathbf{L} in (2):

$$\phi_{synth} := \exists \mathbf{L}. \forall \mathbf{I}^S, O^S, \mathbf{W}. \quad (3)$$

$$\psi_{wfp}(\mathbf{L}) \wedge (\phi_{prog}(\mathbf{L}, \mathbf{I}^S, O^S, \mathbf{W}) \implies \phi_{spec}(\mathbf{I}^S, O^S)).$$

This formula can be solved using a technique called counter-example guided inductive synthesis (CEGIS). CEGIS solves such exist-forall formulas by iteratively solving a series of quantifier-free queries and is often more efficient than trying to solve the quantified query directly. More details are in [21]. For our purposes, we assume the existence of a CEGIS implementation, *CEGIS*, which takes an instance of ϕ_{synth} and returns a model \mathcal{M} with the property that $\mathcal{M}_{\mathbf{L}} \models \phi_{verif}$, from which a program that is a solution to CBPS can be constructed.

III. COMPONENT-BASED PROGRAM SYNTHESIS FOR MANY-TO-MANY RULES

Given the IR and ISA instruction sets \mathbf{K}^{IR} and \mathbf{K}^{ISA} , Buchwald et al. [6] use CBPS to synthesize rewrite rules. They use a single ISA instruction $\mathbf{k}^{ISA} \in \mathbf{K}^{ISA}$ for the CBPS specification and a subset of the IR instructions for the CBPS components. A solution to the resulting ϕ_{synth} formula gives a program \mathbf{P}^{IR} . If \mathbf{P}^{ISA} is the single-instruction program consisting of \mathbf{k}^{ISA} , they interpret the pair $(\mathbf{P}^{IR}, \mathbf{P}^{ISA})$ as an instruction selection rewrite rule.

However, Buchwald et al.'s solution is insufficient for generating many-to-many rules, as they cannot synthesize IR and ISA programs that both contain multiple instructions. Instead, two functionally equivalent programs need to be synthesized. We first define an extension to CBPS called generalized component-based program synthesis (GCBPS) to address this problem. Then, we show how to construct a synthesis query whose solutions represent pairs of functionally equivalent programs.

A. Generalized Component-based Program Synthesis

We define the GCBPS task as that of synthesizing two programs, \mathbf{P}^a and \mathbf{P}^b , represented using location variables \mathbf{L}^a and \mathbf{L}^b , given two sets of components \mathbf{K}^a and \mathbf{K}^b , two sets of inputs $\mathbf{I}^a, \mathbf{I}^b$ where $|\mathbf{I}^a| = |\mathbf{I}^b|$, and two outputs O^a, O^b where the following conditions hold true:

- 1) \mathbf{P}^a uses each component in \mathbf{K}^a exactly once.
- 2) \mathbf{P}^b uses each component in \mathbf{K}^b exactly once.
- 3) \mathbf{P}^a is functionally equivalent to \mathbf{P}^b .

B. Solving GCBPS

We start with the CBPS verification constraint from (2) using components \mathbf{K}^a (and a corresponding set of inputs and outputs \mathbf{W}^a), but modify it slightly by introducing variables (\mathbf{I}^a, O^a) that are fresh copies of (\mathbf{I}^S, O^S) :

$$\psi_{wfp}(\mathbf{L}^a) \wedge \forall \mathbf{I}^a, O^a, \mathbf{W}^a, \mathbf{I}^S, O^S. \quad (4)$$

$$(\phi_{prog}^a(\mathbf{L}^a, \mathbf{I}^a, O^a, \mathbf{W}^a) \wedge \phi_{spec}(\mathbf{I}^S, O^S)) \implies ((\bigwedge_i I_i^a = I_i^S) \implies O^a = O^S).$$

Assuming the formulas for both the program and the specification, if their inputs are the same, their outputs must also be the same.

We next replace the specification program with a different component-based program using components \mathbf{K}^b and quantify over that program's inputs \mathbf{I}^b , output O^b , and component variables \mathbf{W}^b :

$$\phi_{\text{verif}} := \psi_{\text{wfp}}(\mathbf{L}^a) \wedge \psi_{\text{wfp}}(\mathbf{L}^b) \wedge \forall \mathbf{I}^a, \mathbf{I}^b, O^a, O^b, \mathbf{W}^a, \mathbf{W}^b. \quad (5)$$

$$\begin{aligned} & (\phi_{\text{prog}}^a(\mathbf{L}^a, \mathbf{I}^a, O^a, \mathbf{W}^a) \wedge \phi_{\text{prog}}^b(\mathbf{L}^b, \mathbf{I}^b, O^b, \mathbf{W}^b)) \implies \\ & ((\wedge_i I_i^a = I_i^b) \implies O^a = O^b). \end{aligned}$$

This is our generalized verification constraint stating the correctness criteria for when two component-based programs are semantically equivalent.

To synthesize such a pair of programs, a synthesis formula ϕ_{synth} is defined by existentially quantifying \mathbf{L}^a and \mathbf{L}^b in the verification formula (5):

$$\phi_{\text{synth}} := \exists \mathbf{L}^a, \mathbf{L}^b. \forall \mathbf{I}^a, \mathbf{I}^b, O^a, O^b, \mathbf{W}^a, \mathbf{W}^b. \quad (6)$$

$$\begin{aligned} & \psi_{\text{wfp}}(\mathbf{L}^a) \wedge \psi_{\text{wfp}}(\mathbf{L}^b) \wedge \\ & \left((\phi_{\text{prog}}^a(\mathbf{L}^a, \mathbf{I}^a, O^a, \mathbf{W}^a) \wedge \phi_{\text{prog}}^b(\mathbf{L}^b, \mathbf{I}^b, O^b, \mathbf{W}^b)) \implies \right. \\ & \left. ((\wedge_i I_i^a = I_i^b) \implies O^a = O^b) \right). \end{aligned}$$

As above, we assume that calling *CEGIS* on ϕ_{synth} returns a model \mathcal{M} such that $\mathcal{M}_{\mathbf{L}^a \cup \mathbf{L}^b} \models \phi_{\text{verif}}$. This can be converted into a pair of programs $(\mathbf{P}^a, \mathbf{P}^b)$ representing a rewrite rule that is a solution for the GCBPS task. We write $\text{rewriteRule}(\mathbf{K}^a, \mathbf{K}^b, \mathcal{M}_{\mathbf{L}^a}, \mathcal{M}_{\mathbf{L}^b})$ for the rewrite rule constructed from a specific model \mathcal{M} using the component sets \mathbf{K}^a and \mathbf{K}^b .

IV. GENERATING ALL MANY-TO-MANY REWRITE RULES

Buchwald et al. [6] describe an iterative algorithm, *IterativeCEGIS*, to synthesize rewrite rules using CBPS. This algorithm iterates over all multisets of IR instructions up to a given size and only runs synthesis on each such multiset. Compared to running synthesis using all the IR instructions at once, this iterative algorithm works better in practice.

However, *IterativeCEGIS* cannot synthesize rewrite rules with both multi-instruction IR programs and multi-instruction ISA programs. Furthermore, it produces duplicate rewrite rules which are then filtered out in a post-synthesis filtering step. Although the results are correct, this approach is highly inefficient because each call to *CEGIS* is expensive, and a *CEGIS* call is made, not just for some duplicate rules, but for every duplicate rule. In our approach, we make the requirement that a solution is not a duplicate part of the *CEGIS* query itself, ensuring that each successful *CEGIS* query finds a new, non-redundant rewrite rule.

Our iterative algorithm, *genAll*, is shown in Figure 1. It takes as parameters the IR and ISA component sets, \mathbf{K}^{IR}

```

1  genAll( $\mathbf{K}^{IR}, \mathbf{K}^{ISA}, N^{IR}, N^{ISA}$ ):
2   $\mathbb{S}_R \leftarrow \{\}$ 
3  for  $n_1, n_2 \in [1, N^{IR}] \times [1, N^{ISA}]$ :
4  for  $\mathbf{m}^{IR} \in \text{multicomb}(\mathbf{K}^{IR}, n_1)$ :
5  for  $\mathbf{m}^{ISA} \in \text{multicomb}(\mathbf{K}^{ISA}, n_2)$ :
6  for  $\mathbf{I}^{IR}, \mathbf{I}^{ISA} \in \text{allInputs}(\mathbf{m}^{IR}, \mathbf{m}^{ISA})$ :
7   $\phi, \mathbf{L}^{IR}, \mathbf{L}^{ISA} \leftarrow$ 
   GCBPS( $\mathbf{m}^{IR}, \mathbf{m}^{ISA}, \mathbf{I}^{IR}, \mathbf{I}^{ISA}$ )
8   $\phi \leftarrow \phi \wedge \neg \text{AllComposites}(\mathbb{S}_R, \dots)$ 
9   $\mathbb{S}_R \leftarrow \mathbb{S}_R \cup$ 
   CEGISAll( $\phi, \mathbf{m}^{IR}, \mathbf{m}^{ISA}, \mathbf{L}^{IR}, \mathbf{L}^{ISA}$ )
10 return  $\mathbb{S}_R$ 

```

Fig. 1: Iterative algorithm to generate all unique rewrite rules up to a given size.

```

1  CEGISAll( $\phi, \mathbf{m}^{IR}, \mathbf{m}^{ISA}, \mathbf{L}^{IR}, \mathbf{L}^{ISA}$ ):
2   $\mathbb{S}_R = \{\}$ 
3  while True:
4   $\mathcal{M} \leftarrow \text{CEGIS}(\phi)$ 
5  if  $\mathcal{M} = \perp$ : return  $\mathbb{S}_R$ 
6   $\mathbf{R} \leftarrow \text{rewriteRule}(\mathbf{m}^{IR}, \mathbf{m}^{ISA}, \mathcal{M}_{\mathbf{L}^{IR}}, \mathcal{M}_{\mathbf{L}^{ISA}})$ 
7   $\mathbb{S}_R \leftarrow \mathbb{S}_R \cup \{\mathbf{R}\}$ 
8   $\phi \leftarrow \phi \wedge \neg \psi_{\text{dup}}(\mathbf{R}, (\mathbf{L}^{IR}, \mathbf{L}^{ISA}))$ 

```

Fig. 2: AllSAT algorithm to synthesize all unique rules. Line 8 excludes all rules that are duplicates of the current synthesized rewrite rule.

and \mathbf{K}^{ISA} respectively, as well as a maximum number of components of each kind to use in rewrite rules, N^{IR} and N^{ISA} , and iteratively builds up a set \mathbb{S}_R of rewrite rules, which it returns at the end. Line 3 shows that n_1 and n_2 iterate up to these maximum sizes. Line 4 iterates over all multisets of elements from \mathbf{K}^{IR} of size n_1 using a standard multicomposition algorithm *multicomb* [25] (not shown). Line 5 is similar but for multisets from \mathbf{K}^{ISA} of size n_2 . Next, for a given choice of multisets, line 6 enumerates all possible ways of selecting input vectors from those multisets that could create well-formed programs by constructing two fresh sets of input variables. Line 7 constructs fresh sets of location variables \mathbf{L}^{IR} and \mathbf{L}^{ISA} and returns them along with the instantiated GCBPS synthesis formula (using Equation (6)).¹ Line 8 excludes all *composite rules* from the synthesis search space. Composite rules are rules that can be constructed using the current set of rules \mathbb{S}_R and are thus unnecessary for instruction selection. We discuss this in more detail in Section IV-B. Finally, on line 9, the current set of rules \mathbb{S}_R is updated with the result of calling *CEGISAll*, which we describe next.

Figure 2 shows the *CEGISAll* algorithm that performs the AllSAT [20], [31] task. Its parameters are the synthesis formula ϕ , the multisets \mathbf{m}^{IR} and \mathbf{m}^{ISA} , and the location variables \mathbf{L}^{IR} and \mathbf{L}^{ISA} . It returns a set \mathbb{S}_R of rewrite rules. Initially, this set is empty. The algorithm iteratively calls

¹We augment the well-formed program constraint in (6) to prevent synthesizing programs containing dead code and unused inputs. This can be accomplished by enforcing that each input and intermediate value is used in at least one location.

a standard *CEGIS* algorithm to solve the synthesis query, constructing a new rewrite rule \mathbf{R} , which is added to the set \mathbb{S}_R of rewrite rules, when the call to *CEGIS* is successful. The iteration repeats until the *CEGIS* query returns \perp , indicating that there are no more rewrite rules to be found. Note that after each iteration, the ϕ_{synth} formula is refined by adding the negation of a formula capturing the notion of duplicates for this rule. We describe how this is done next.

A. Excluding Duplicate Rules

Consider the two distinct rules below. As a syntactical convention, infix operators are used for IR patterns and function calls for ISA patterns.

$$\begin{aligned} I_1 + (I_2 \cdot I_3) &\rightarrow add(I_1, mul(I_2, I_3)) \\ (I_1 \cdot I_3) + I_2 &\rightarrow add(I_2, mul(I_1, I_3)) \end{aligned}$$

The two IR patterns represent the same operation despite the fact that the variable names and the order of the commutative arguments to addition are both different. Both rules would match the same program fragments in an instruction selector and would result in the same rewrite rule application. Thus, we consider such rules to be equivalent and would like to ensure that only one is generated by our algorithm.

We first define a rewrite rule equivalence relation, \sim_{rule} . Informally, two rules are equivalent if replacing either one by the other has no discernible effect on the execution of an instruction selection algorithm. We make this more formal by considering various attributes of standard instruction selection algorithms.

Commutative Instructions Modern pattern matching algorithms used for instruction selection try all argument orderings for commutative instructions [5]. We define the commutative equivalence relation \sim_{CIR} as $\mathbf{P}_1^{IR} \sim_{CIR} \mathbf{P}_2^{IR}$ iff \mathbf{P}_2^{IR} is a remapping of \mathbf{P}_1^{IR} 's commutative instruction's arguments.

Same-kind Instructions Programs \mathbf{P} generated by GCPBS have a unique identifier, the program line number, for each instruction. This means that if two instructions of the same kind appear in a program, interchanging their line numbers results in a different program, even though it makes no difference to the instruction selection algorithm. We define the same-kind equivalence relation \sim_{KIR} as $\mathbf{P}_1^{IR} \sim_{KIR} \mathbf{P}_2^{IR}$ iff \mathbf{P}_2^{IR} is the result of remapping the line numbers for same-kind instructions in \mathbf{P}_1^{IR} .

Data Dependency Modern instruction selection algorithms perform pattern matching, not based on a total order of instructions, but on a partial order determined by data dependencies. Many different sequences may thus lead to the same partial order. We define \sim_{DIR} as $\mathbf{P}_1^{IR} \sim_{DIR} \mathbf{P}_2^{IR}$ iff \mathbf{P}_1^{IR} and \mathbf{P}_2^{IR} have the same data dependency graph.

Rule Input Renaming For a given rewrite rule, the input variables used for the IR program must match the input variables used for the ISA program, but the specific variable identifiers used do not matter. We define the equivalence relation \sim_{Irule} on rules (i.e., pairs of programs) as $\mathbf{R}_1 \sim_{Irule} \mathbf{R}_2$ iff \mathbf{R}_2 is the result of remapping variable identifiers in \mathbf{R}_1 .

Rule Equivalence The first three equivalence relations defined above are for IR programs, but the analogous relations (\sim_{CISA} , \sim_{KISA} , \sim_{DISA}) for ISA instructions are also useful.

Putting everything together, we define rule equivalence \sim_{rule} as follows.

$$\sim_{IR} := \uplus\{\sim_{CIR}, \sim_{KIR}, \sim_{DIR}\} \quad (7)$$

$$\sim_{ISA} := \uplus\{\sim_{CISA}, \sim_{KISA}, \sim_{DISA}\} \quad (8)$$

$$\sim_{rule} := \uplus\{(\sim_{IR} \otimes \sim_{ISA}), \sim_{Irule}\} \quad (9)$$

Overall IR equivalence is defined as the transitive closure of the union (notated with \uplus) of the three individual IR relations. ISA equivalence is defined similarly. Overall rewrite rule equivalence is then defined using the \otimes operator, where $\sim_{\otimes} = \sim_a \otimes \sim_b$ is defined as: $(a_1, b_2) \sim_{\otimes} (a_2, b_2)$ iff $a_1 \sim_a a_2$ and $b_1 \sim_b b_2$. Specifically, rule equivalence is obtained by combining IR equivalence in this way with ISA equivalence, and then combining the result with \sim_{Irule} using \uplus .

The set of all duplicates of rule \mathbf{R} is the rule equivalence class $[\mathbf{R}]_{rule}$, where $\mathbf{R}' \in [\mathbf{R}]_{rule} \iff \mathbf{R} \sim_{rule} \mathbf{R}'$. ψ_{dup} can be constructed as the disjunction of all elements of the equivalence class $[\mathbf{R}]_{rule}$.

B. Excluding Composite Rules

We also exclude any rule whose effect can already be achieved using the current set of generated rules (line 8 of Figure 1). We elucidate this using a simple example. Assume the algorithm just constructed a new query for the multisets \mathbf{m}^{IR} , \mathbf{m}^{ISA} , and the input \mathbf{I}^{IR} (line 7 of Figure 1), and assume that the rule library \mathbb{S}_R currently contains rules for addition ($I_1 + I_2 \rightarrow add(I_1, I_2)$) and multiplication ($I_1 \cdot I_2 \rightarrow mul(I_1, I_2)$). Consider the following cases.

- 1) If $\mathbf{I}^{IR} = (I_1)$, $\mathbf{m}^{IR} = \{+\}$, and $\mathbf{m}^{ISA} = \{add\}$, then the rule $I_1 + I_1 \rightarrow add(I_1, I_1)$ will be synthesized by *CEGISAll*. But this rule is a *specialization* of the existing rule for addition. Any use of this specialized rule could instead be replaced by the more general rule, and this rule can thus be excluded. Note that we order the inputs on line 6 of Figure 1 to guarantee that the most general version of a rule is found first.
- 2) If $\mathbf{I}^{IR} = (I_1, I_2, I_3)$, $\mathbf{m}^{IR} = \{+, \cdot\}$, and $\mathbf{m}^{ISA} = \{add, mul\}$, then the composite rule $(I_1 + (I_2 \cdot I_3)) \rightarrow add(I_1, mul(I_2, I_3))$ will be synthesized by *CEGISAll*. Using similar logic, any use of this composite rule could instead use the simpler and more general rules for addition and multiplication, and this rule can thus be excluded. The multiset ordering used in lines 4 and 5 of Figure 1 ensures that subsets are visited before supersets, guaranteeing that smaller rules are found first. A specialized rule can be interpreted as a composite rule composed of the general rule with fewer inputs.

Only composite rules that would have been synthesized for a particular query need to be excluded. In general, for a specific query based on \mathbf{m}^{IR} , \mathbf{m}^{ISA} , and \mathbf{I}^{IR} , we enumerate and exclude composite rules $\mathbf{R} := (\mathbf{P}^{IR}, \mathbf{P}^{ISA})$ that meet the following criteria:

```

1  genAllLC( $\mathbf{K}^{IR}, \mathbf{K}^{ISA}, N^{IR}, N^{ISA}, cost$ ):
2     $\mathbf{K}_{sorted} \leftarrow sortByCost(\mathbf{K}^{ISA}, N^{ISA}, cost)$ 
3     $\mathbb{S}_R \leftarrow \{\}$ 
4    for  $n \in [1, N^{IR}]$ :
5      for  $\mathbf{m}^{IR} \in multicomcomb(\mathbf{K}^{IR}, n)$ :
6        for  $\mathbf{m}^{ISA} \in \mathbf{K}_{sorted}$ :
7           $c_{cur} \leftarrow cost(\mathbf{m}^{ISA})$ 
8          for  $\mathbf{I}^{IR}, \mathbf{I}^{ISA} \in allInputs(\mathbf{m}^{IR}, \mathbf{m}^{ISA})$ :
9             $\phi, \mathbf{L}^{IR}, \mathbf{L}^{ISA} \leftarrow$ 
10              $GCBPS(\mathbf{m}^{IR}, \mathbf{m}^{ISA}, \mathbf{I}^{IR}, \mathbf{I}^{ISA})$ 
11             $\phi \leftarrow \phi \wedge \neg AllComposites_{LC}(\mathbb{S}_R, c_{cur}, \dots)$ 
12             $\mathbb{S}_R \leftarrow \mathbb{S}_R \cup$ 
13              $CEGISAll_{LC}(\phi, \mathbf{m}^{IR}, \mathbf{m}^{ISA}, \mathbf{L}^{IR}, \mathbf{L}^{ISA})$ 
14  return  $\mathbb{S}_R$ 

```

Fig. 3: Iterative algorithm to generate all lowest-cost rules. ISA multisets are ordered by cost. *CEGISAll* is modified to exclude rules with duplicate IR programs.

- \mathbf{R} has exactly $|\mathbf{I}^{IR}|$ inputs.
- \mathbf{P}^{IR} has the same components as \mathbf{m}^{IR} .
- \mathbf{P}^{ISA} has the same components as \mathbf{m}^{ISA} .
- \mathbf{P}^{IR} is built from the IR programs of already-found rules in \mathbb{S}_R .
- \mathbf{P}^{ISA} is the result of applying the rewrite rules used to build \mathbf{P}^{IR} .

This enumeration is encapsulated by the call to *AllComposites* on line 8 of Figure 1.

V. GENERATING ALL LOWEST-COST RULES

Because all duplicates are excluded, the *genAll* algorithm generates only unique rewrite rules. However, two unique rules can share the same IR pattern. For a particular IR pattern, only the lowest-cost rule is needed for some cost metric. Knowing the instruction selection cost metric at rule-generation time presents another time-saving opportunity because we can also prevent the synthesis of high-cost rules.

We make a few assumptions about such a cost metric.

- The cost for an instruction selection tiling is equal to the sum of the costs of each tiling rule’s ISA program.
- The cost of an ISA program \mathbf{P}^{ISA} only depends on the instruction *contents*, not the program structure. This cost is the sum of the cost of each instruction in the program.

While these assumptions are a restriction on the space of possible cost metrics, they are sufficient to represent common ones like code size and energy. If the compiler’s cost metric violates these assumptions, the *genAll* algorithm can be used instead. This restricted space of cost metrics has the important property that the cost of any rule that would be synthesized using the components \mathbf{m}^{ISA} can be determined up front as the sum of the cost of each component.

Figure 3 shows our synthesis algorithm updated to only synthesize the lowest-cost rules for each unique IR pattern. The first change is to sort all possible multisets of ISA instructions up to size N^{ISA} by cost (lower cost first) (line 2). This ordering ensures that the first rule synthesized for a

particular IR program will be the lowest-cost version of that rule. Therefore, after synthesizing a new rule, all rules with a duplicate IR program can be excluded. The second change excludes rules with duplicate IR programs. A duplicate IR program is defined using the IR equivalence relation:

$$\sim_{IR_{LC}} := \uplus \{ \sim_{C^{IR}}, \sim_{K^{IR}}, \sim_{D^{IR}}, \sim_{I^{IR}} \} \quad (10)$$

This is the same definition as (7), but with an additional relation $\sim_{I^{IR}}$ defined as $\mathbf{P}_1^{IR} \sim_{I^{IR}} \mathbf{P}_2^{IR}$ iff \mathbf{P}_2^{IR} is the result of remapping variable identifiers in \mathbf{P}_1^{IR} . The *CEGISAll*_{LC} function called on line 11 is the same as *CEGISAll*, except that it uses $\sim_{IR_{LC}}$ instead of \sim_{IR} when constructing ψ_{dup} .

The third change modifies *AllComposites* to use the known up-front cost $cost(\mathbf{m}^{ISA})$. To see how this works, we consider again the example from Section IV-B. As before, we assume \mathbb{S}_R currently contains two rules: one for addition ($I_1 + I_2 \rightarrow add(I_1, I_2)$) and one for multiplication ($I_1 \cdot I_2 \rightarrow mul(I_1, I_2)$). We assume the target (ISA) expressions for these rules have cost 5 and 10, respectively. Consider the following situation:

- Suppose $\mathbf{I}^{IR} = (I_1, I_2, I_3)$, and $\mathbf{m}^{IR} = \{+, \cdot\}$. It might be possible to synthesize a rule that has IR pattern $(I_1 + (I_2 \cdot I_3))$. We know that the composite rule $(I_1 + (I_2 \cdot I_3)) \rightarrow add(I_1, mul(I_2, I_3))$ would have a cost of 15 since rule costs are additive. Therefore, we can exclude any rule that matches this IR pattern and has $cost(\mathbf{m}^{ISA}) \geq 15$.

To implement this, only one adjustment needs to be made to the conditions in Section IV-B. Instead of requiring \mathbf{P}^{ISA} to have the same components as \mathbf{m}^{ISA} , we simply require $cost(\mathbf{P}^{ISA}) \geq cost(\mathbf{m}^{ISA})$, i.e., for rules matching the other conditions, if the ISA program has a cost equal to or greater than cost of the ISA program in the current rule, it is excluded. These conditions are encapsulated by the call to *AllComposites*_{LC} (line 10).

VI. EVALUATION

Our evaluation strategy is threefold. We first show that our algorithm is capable of producing a variety of many-to-many rules. A good set of rewrite rules involves both many-to-one and one-to-many rules. We also show that by removing duplicate, composite, and high-cost rules, we produce a much smaller set of rewrite rules. Second, we analyze the effect on performance of the optimizations described above. We show that they all significantly reduce the time spent in synthesis. Finally, we show that by using different cost metrics, we can generate different sets of lowest-cost rewrite rules.

A. Implementation

All instructions are formally specified using the hwtypes Python library [11], which leverages pySMT [18] to construct (quantifier-free) SMT queries in the theory of bit-vectors. We also use annotations indicating which instructions are commutative. We use Boolector [28] as the SMT solver and set a timeout of 12 seconds for each CEGIS invocation. Every synthesized rewrite rule is independently verified to be valid.

B. Instruction Specifications

To evaluate our algorithms, we selected small but non-trivial sets of IR and ISA instructions operating on 4-bit bit-vectors.

IR We define the IR instruction set to be constants (0, 1), bitwise operations (*not*, *and*, *or*, *xor*), arithmetic operations (*neg*, *add*, *sub*), multiplication (*mul*), unsigned comparison operations (*ult*, *ule*, *ugt*, *uge*), equality (*eq*), and dis-equality (*neq*).

ISA 1 This is a minimal RISC-like ISA containing only 6 instructions: *nand*, *sub*, three comparison instructions (*cmpZ*, *cmpN*, *cmpC*) which compute the zero (Z), sign (N), and carry (C) flags respectively for a subtraction, and a flag inverting instruction (*inv*).

ISA 2 This is an ISA specialized for linear algebra. It supports the 5 instructions: *neg*, *add*, *add3* (addition of 3 values), *mul*, and *mac* (multiply-accumulate).

C. Rewrite Rule Synthesis

For each ISA we run three experiments. The first experiment (All Rules) is the baseline that generates all many-to-many rules, including duplicate, composite, and high cost rules. This is an implementation of Buchwald et al.’s *IterativeCEGIS* algorithm extended to use GCBPS for many-to-many rules (notated as *IterativeCEGIS_{GCBPS}*). The second (Only Unique) generates only unique rules by excluding all duplicates and composites using the *genAll* algorithm. The third (Only Lowest-Cost) generates only the lowest-cost rules using the *genAll_{LC}* algorithm in Figure 3. A code-size cost metric is used, i.e., $cost(\mathbf{K})$ is just the number of components in \mathbf{K} .

For ISA 1, we split the rule generation into two parts. The first part (ISA 1a) synthesizes rules composed of bitwise and arithmetic IR instructions using the ISA’s *nand* and *sub* instructions. The second part (ISA 1b) synthesizes rules composed of constants and comparison instructions using the four instructions *cmpZ*, *cmpN*, *cmpC*, and *inv*.

For 1a and 1b, we synthesize rewrite rules up to an IR program size of 2 and an ISA program size of 3 (written 2-to-3). For (Only Lowest-Cost), we increase the ISA program size to 5 and 4 respectively. For ISA 2, we synthesize all rewrite rules composed of constant, and arithmetic (including *mul*) IR instructions up to size 3-to-2.

The number of rewrite rules produced for each configuration for ISA 1a, 1b, and 2 is shown in Tables I, II, and III, respectively. Each table entry is the number of rewrite rules synthesized for a particular IR and ISA program size. For all ISAs, the extra synthesized rules in (All Rules) were compared against the duplicate and composite rules excluded by (Only Unique). Entries in (All Rules) marked with a ‘(-n)’ represent ‘n’ rules that (Only Unique) synthesized, but (All Rules) missed due to CEGIS timeouts. The (All Rules) experiment for the entry marked with an asterisk could not complete in 70 hours, so the number calculated from (Only Unique) is shown.

For both ISAs we were able to synthesize 1-to-many and many-to-1 rules for both IR and ISA instructions.

genAll produced a more complete set of rules than *IterativeCEGIS_{GCBPS}*.

Table IV shows the percentage of rules that are duplicates or composites in the first column, and the percentage of rules that are high cost in the second column. Most rules in (All Rules) are duplicates, composites, or high cost. Out of the 349179 rules up to size 3-to-2 for ISA 2 (i.e., the sum of the (All Rules)), 99.5% are duplicates or composites. Similarly, most rules are high cost. In ISA 1a, 59672 out of 59822 rules (99.7%) up to size 2-to-3 are high cost.

D. Synthesis Time Improvement with *genAll*

In this section we showcase the synthesis time improvements of *genAll*. The first experiment is the baseline *IterativeCEGIS_{GCBPS}*. The second excludes duplicate rules (i.e., with line 8 of Figure 2). The third, *genAll*, excludes both duplicates and composites (i.e. with line 8 of both Figure 2 and Figure 1).

For each *GCBPS* query, we note the time required (t_{sat}) to run *CEGISAll*. Next, we measure the number of unique rules (N_{unique}) found by *CEGISAll*. We then add the pair (N_{unique}, t_{sat}) to our dataset. We plot the cumulative synthesis time versus the number of unique rules found by doing the following. Each data point is sorted by its slope (t_{sat}/N_{unique}). Then, the increase in both t_{sat} and N_{unique} is plotted for each sorted point. Some data points have $N_{unique} = 0$ indicating that every synthesized rule was redundant and is shown using a vertical slope.

The synthesis time plot for unique rewrite rules for ISA 1b up to size 2-to-3 is shown in Figure 4a. Excluding all duplicates shows a $5.3\times$ speedup. Excluding both duplicates and composites shows a $6.2\times$ speedup. Both optimizations find an additional 5 unique rules.

E. Synthesis Time Improvement with *genAll_{LC}*

We also showcase the synthesis time improvements of *genAll_{LC}* using a similar setup. The first experiment is the baseline *IterativeCEGIS_{GCBPS}*. The second excludes IR duplicate rules. The third, *genAll_{LC}*, excludes both IR duplicates and IR composites.

We use the same experimental setup as before, except when computing N_{unique} , all higher-cost rules are filtered instead. The synthesis time plot for lowest-cost rewrite rules for ISA 1b up to size 2-to-3 is shown in Figure 4b.

Excluding rules with duplicate IR programs provides a $41\times$ speed-up. Also excluding high-cost composites provides a $1254\times$ speed-up over the baseline (All Rules) configuration.

F. Total Speed-up

We summarize the speed-ups of *genAll* and *genAll_{LC}* compared to the *IterativeCEGIS_{GCBPS}* baseline for all configurations in Table V. We compare the synthesis time in the ‘‘Synth’’ column. We compare the total algorithm runtime in the ‘‘Total’’ column (including time for iterating, solving, rule filtering, etc.). The last row’s baseline did not complete in 70 hours, so we provide lower bounds for speed-up.

		ISA Program Size										
		All Rules			Only Unique			Only Lowest-Cost				
		1	2	3	1	2	3	1	2	3	4	5
IR Prog Size	1	5	32	1096	3	10	96	3	4	2	1	0
	2	76	1719	56894	40	189	1940	40	67	34	12	6


TABLE I: Number of synthesized rewrite rules for ISA 1a.

		ISA Program Size									
		All Rules			Only Unique			Only Lowest-Cost			
		1	2	3	1	2	3	1	2	3	4
IR Program Size	1	17	71	3662	9	51	873	7	3	0	0
	2	89	3942 (-5)	199572	78	717	21511	52	64	9	0

TABLE II: Number of synthesized rewrite rules for ISA 1b.

		IR Program Size								
		All Rules			Only Unique			Only Lowest-Cost		
		1	2	3	1	2	3	1	2	3
ISA Program Size	1	11	287	3998	3	14	315	3	14	315
	2	10	3115	341758*	3	69	1337	1	32	760

TABLE III: Number of synthesized rewrite rules for ISA 2.



(a) $genAll$

(b) $genAll_{LC}$

Fig. 4: Cumulative synthesis time comparison for ISA 1b up to size 2-to-3.

ISA	Rule Size up to (IR, ISA)	% Duplicate or Composite	% High-cost
1a	(2, 3)	96.2%	99.7%
1b	(2, 3)	88.8%	99.9%
2	(3, 2)	99.5%	99.7%

TABLE IV: Percent of rewrite rules up to (IR, ISA) size that are a duplicate or a composite, and percent that are high-cost.

ISA	Rule Size up to (IR, ISA)	$genAll$ Speed-up		$genAll_{LC}$ Speed-up	
		Synth	Total	Synth	Total
1a	(2, 2)	3.5 \times	1.3 \times	11 \times	2.8 \times
1b	(2, 2)	3.1 \times	1.7 \times	26 \times	2.8 \times
2	(2, 2)	11 \times	2 \times	53 \times	2.5 \times
1a	(2, 3)	12 \times	6.8 \times	601 \times	57 \times
1b	(2, 3)	6.2 \times	2.7 \times	1254 \times	63 \times
2	(3, 2)	> 768 \times	> 81 \times	> 4004 \times	> 171 \times

TABLE V: Speed-ups compared to *IterativeCEGIS_{GC}BPS*.

The speed-ups depend on many parameters including the maximum size of the rewrite rules, the number of possible instructions, the commutativity of the instructions, and the semantics of the instructions. The optimizations discussed produce several orders of magnitude speed-ups. Further optimizing the non-solver portions (e.g., re-coding in C) would drastically increase the “Total” speed-ups to be closer to the “Synth” ones. Clearly, the combination of all optimizations discussed in this paper can produce speed-ups of several orders of magnitude.

G. Cost Metric Comparisons

Our final experiment explores how the choice of cost metric influences the rules. We have implemented two cost metrics: a code size metric (CS) and an estimated energy metric (E).

ISA	Rule Size up to (IR, ISA)	Unique (CS)	Unique (E)	Common
1a	(2, 5)	121	161	48
1b	(2, 4)	99	198	36
2	(3, 2)	134	137	991

TABLE VI: Number of unique and common rewrite rules synthesized for code size (CS) and energy (E) cost metrics.

The energy metric was created to correspond to real hardware energy data. For example the cost ratio for *mul* and *add* is 1 : 1 for code size, but is 2.5 : 1 for energy. The number of common and unique lowest-cost rewrite rules for each ISA is shown in Table VI.

While there is some overlap in common rules, each cost metric produces a differing set of unique lowest-cost rules.

VII. CONCLUSION AND FUTURE WORK

We showed that many-to-many instruction selection rewrite rules can be synthesized for various ISAs using program synthesis. This supports two major trends in computer architecture. The first is the trend towards simple or reduced instruction architectures where multiple instructions are needed for simple operations. It also supports the trend to introduce more complex domain-specific instructions for energy efficiency. In this case, a single instruction can implement complex operations.

We showed that our algorithms are efficient. Removing duplicates, composites, and higher-cost rules results in multiple orders of magnitude speed-ups. Synthesizing many-to-many rewrite rules for modern IRs and ISAs may require further optimizations. Many of our synthesized rules contain program fragments that a compiler would optimize during IR optimization or peephole optimization. A modified version of GCBPS could be used to directly synthesize and exclude such program fragments.

Buchwald et al. [6] presented generalizations for multi-sorted instructions, multiple outputs, preconditions, and internal attributes, enabling the modeling of memory and control flow instructions. Our synthesis query and algorithms are orthogonal and could incorporate these features, allowing for a broader range of possible instruction sets.

As is the case in prior work, we limit synthesis to loop free patterns. Relaxing this constraint and using other instruction selection algorithms would be an interesting research avenue.

Another promising research direction involves exploring the trade-offs between synthesis time, compile time, and code quality. This could be done by varying the maximum size of rewrite rules, changing the instruction selection algorithms, relaxing the completeness guarantee, or incorporating IR or peephole optimizations.

We believe this research area is fertile ground and hope our work inspires and enables future research endeavors towards the goal of automatically generating compilers for emerging domain-specific architectures.

REFERENCES

- [1] Alfred V. Aho, Mahadevan Ganapathi, and Steven W. K. Tjiang. Code generation using tree matching and dynamic programming. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 11(4):491–516, 1989.
- [2] Aws Albarghouthi, Sumit Gulwani, and Zachary Kincaid. Recursive program synthesis. In *Computer Aided Verification: 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13–19, 2013. Proceedings 25*, pages 934–950. Springer, 2013.
- [3] Rajeev Alur, Arjun Radhakrishna, and Abhishek Udupa. Scaling enumerative program synthesis via divide and conquer. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 319–336. Springer, 2017.
- [4] Rick Bahr, Clark Barrett, Nikhil Bhagdikar, Alex Carsello, Ross Daly, Caleb Donovick, David Durst, Kayvon Fatahalian, Kathleen Feng, Pat Hanrahan, et al. Creating an agile hardware design flow. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2020.
- [5] Eli Bendersky. *A deeper look into the LLVM code generator, Part 1*, Feb 2013.
- [6] Sebastian Buchwald, Andreas Friedl, and Sebastian Hack. Synthesizing an instruction selection rule library from semantic specifications. In *Proceedings of the 2018 International Symposium on Code Generation and Optimization*, pages 300–313, 2018.
- [7] R. G. Cattell. Automatic derivation of code generators from machine descriptions. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 2(2):173–190, 1980.
- [8] Yu-Hsin Chen, Joel Emer, and Vivienne Sze. Eyeriss: A spatial architecture for energy-efficient dataflow for convolutional neural networks. *ACM SIGARCH Computer Architecture News*, 44(3):367–379, 2016.
- [9] Ross Daly, Caleb Donovick, Jackson Melchert, Rajsekhar Setaluri, Nestan Tsiskaridze Bullock, Priyanka Raina, Clark Barrett, and Pat Hanrahan. Synthesizing instruction selection rewrite rules from RTL using SMT. In *Conference on Formal Methods in Computer-Aided Design (FMCAD)*, page 139, 2022.
- [10] Joao Dias and Norman Ramsey. Automatically generating instruction selectors using declarative machine descriptions. *ACM Sigplan Notices*, 45(1):403–416, 2010.
- [11] Caleb Donovick, Ross Daly, Jackson Melchert, Lenny Truong, Priyanka Raina, Pat Hanrahan, and Clark Barrett. Peak: A single source of truth for hardware design and verification. *arXiv preprint arXiv:2308.13106*, 2023.
- [12] Helmut Emmelmann, F.-W. Schröer, and Rudolf Landwehr. BEG: A generator for efficient back ends. *ACM Sigplan Notices*, 24(7):227–237, 1989.
- [13] Herbert Enderton and Herbert B. Enderton. *A mathematical introduction to logic*. Elsevier, 2001.
- [14] Christopher W. Fraser and David R. Hanson. *A retargetable C compiler: Design and implementation*. Addison-Wesley Longman Publishing Co., Inc., 1995.
- [15] Christopher W. Fraser, David R. Hanson, and Todd A. Proebsting. Engineering a simple, efficient code-generator generator. *ACM Letters on Programming Languages and Systems (LOPLAS)*, 1(3):213–226, 1992.
- [16] Mahadevan Ganapathi. *Retargetable code generation and optimization using attribute grammars*. PhD thesis, 1980. AAI8107834.
- [17] Mahadevan Ganapathi and Charles N. Fischer. Description-driven code generation using attribute grammars. In *Proceedings of the 9th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL ’82, page 108–119, New York, NY, USA, 1982. Association for Computing Machinery.
- [18] Marco Gario and Andrea Micheli. PySMT: A solver-agnostic library for fast prototyping of SMT-based algorithms. In *SMT Workshop 2015*, 2015.
- [19] R. Steven Glanville and Susan L. Graham. A new method for compiler code generation. In *Proceedings of the 5th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, POPL ’78, page 231–254, New York, NY, USA, 1978. Association for Computing Machinery.
- [20] Orna Grumberg, Assaf Schuster, and Avi Yadgar. Memory efficient all-solutions SAT solver and its application for reachability analysis. In *Formal Methods in Computer-Aided Design: 5th International Conference, FMCAD 2004, Austin, Texas, USA, November 15–17, 2004. Proceedings 5*, pages 275–289. Springer, 2004.
- [21] Sumit Gulwani, Susmit Jha, Ashish Tiwari, and Ramarathnam Venkatesan. Synthesis of loop-free programs. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2011.
- [22] John L. Hennessy and David A. Patterson. A new golden age for computer architecture. *Commun. ACM*, 62(2):48–60, January 2019.
- [23] Roger Hoover and Kenneth Zadeck. Generating machine specific optimizing compilers. In *Proceedings of the 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 219–229, 1996.
- [24] Norman P. Jouppi, Cliff Young, Nishant Patil, and David Patterson. A domain-specific architecture for deep neural networks. *Communications of the ACM*, 61(9):50–59, 2018.
- [25] Donald E. Knuth. *The Art of Computer Programming, Volume 4, Fascicle 3: Generating All Combinations and Partitions*. Addison-Wesley Professional, 2005.
- [26] David Ryan Koes and Seth Copen Goldstein. Near-optimal instruction selection on DAGs. In *Proceedings of the 6th Annual IEEE/ACM*

International Symposium on Code Generation and Optimization, pages 45–54, 2008.

- [27] Jackson Melchert, Kathleen Feng, Caleb Donovan, Ross Daly, Ritvik Sharma, Clark Barrett, Mark A Horowitz, Pat Hanrahan, and Priyanka Raina. APEX: A framework for automated processing element design space exploration using frequent subgraph analysis. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*, pages 33–45, 2023.
- [28] Aina Niemetz, Mathias Preiner, and Armin Biere. Boolector 2.0. *J. Satisf. Boolean Model. Comput.*, 9(1):53–58, 2014.
- [29] Eduardo Pelegri-Llopert and Susan L. Graham. Optimal code generation for expression trees: An application BURS theory. In *Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 294–308, 1988.
- [30] Raghu Prabhakar, Yaqi Zhang, David Koeplinger, Matt Feldman, Tian Zhao, Stefan Hadjis, Ardavan Pedram, Christos Kozyrakis, and Kunle Olukotun. Plasticine: A reconfigurable architecture for parallel patterns. *ACM SIGARCH Computer Architecture News*, 45(2):389–402, 2017.
- [31] Takahisa Toda and Takehide Soh. Implementing efficient all solutions sat solvers. *Journal of Experimental Algorithmics (JEA)*, 21:1–44, 2016.