# Syntax-Guided Synthesis with Quantitative Syntactic Objectives

Qinheping Hu$^{(\boxtimes)}$, Loris D'Antoni

University of Wisconsin-Madison, Madison, USA
{qhu28,loris}@cs.wisc.edu

**Abstract.** Automatic program synthesis promises to increase the productivity of programmers and end-users of computing devices by automating tedious and error-prone tasks. Despite the practical successes of program synthesis, we still do not have systematic frameworks to synthesize programs that are "good" according to certain metrics—e.g., produce programs of reasonable sizes or with good runtime—and to understand when synthesis can result in such good programs. In this paper, we propose QSYGuS, a unifying framework for describing syntax-guided synthesis problems with quantitative objectives over the syntax of the synthesized programs. QSYGuS builds on weighted (tree) grammars, a clean and foundational formalism that provides flexible support for different quantitative objectives, useful closure properties, and practical decision procedures. We then present an algorithm for solving QSYGuS. Our algorithm leverages closure properties of weighted grammars to generate intermediate problems that can be solved using non-quantitative SYGuS solvers. Finally, we implement our algorithm in a tool, QuaSi, and evaluate it on 26 quantitative extensions of existing SYGuS benchmarks. QuaSi can synthesize optimal solutions in 15/26 benchmarks with times comparable to those needed to find an arbitrary solution.

## 1 Introduction

The goal of program synthesis is to find a program in some search space that meets a specification—e.g., a set of examples or a logical formula. Recently, a large family of synthesis problems has been unified into a framework called syntax-guided synthesis (SYGuS). A SYGuS problem is specified by a context-free grammar describing the search space of programs, and a logical formula describing the specification. Many synthesizers now support this format [2] and annually compete in synthesis competitions [4]. Thanks to these competitions, these solvers are now quite mature and are finding wide application [14].

While the logical specification mechanism provided by SYGuS is powerful, it can only capture the functional requirements of the synthesis problem—e.g., the program should perform correctly on a given set of input/output examples. When multiple possible programs can satisfy the specification, SYGuS *does not* provide a way to prefer one to the other—e.g., one cannot ask a solver to return the program with the fewest if-statements. As a consequence, existing synthesis

tools do not provide guarantees about what solution is returned if multiple ones exist. While a few synthesizers have attempted to include some form of specification to express this kind of quantitative intents [7,19,16,15], these approaches are domain-specific, do not apply to SyGuS problems, and do not provide a simple and flexible specification mechanism. The lack of a formal treatment of quantitative requirements stands in the way of designing synthesizers that can take advantage quantitative of objectives to perform more efficient forms of synthesis.

In this paper, we propose QSyGuS, a unifying framework for describing syntax-guided synthesis problems with quantitative objectives over the syntax of the synthesized programs—e.g., find the most likely program with respect to a given probability distribution—and present an algorithm for solving synthesis problems expressed in this framework. We focus on syntactic objectives because they are the most common ones in practical applications of program synthesis. For example, in programming by examples it is desirable to produce small programs with fewer constants because these programs are more likely to generalize to examples outside of the specification [13]. QSyGuS extends SyGuS in two ways. First, in QSyGuS the search space is represented using weighted grammars, which augment context-free grammars with the ability to assign weights to programs. Second, QSyGuS allows the user to specify constraints over the weight of the program, including optimization objectives—e.g., find the program with the fewest if-statements and with the lowest depth.

QSyGuS is a natural, general, and flexible formalism and is grounded in the well-studied theory of weighted grammars. We leverage this theory and design an algorithm for solving QSyGuS problems using closure properties of weighted grammars. Given a QSyGuS problem, our algorithm generates a SyGuS problem that can be delegated to existing SyGuS solvers. The algorithm then iteratively refines the solution returned by the SyGuS solver to find an optimal one by further generating new SyGuS instances using weighted grammar operations. We implement our algorithm in a tool, QuaSi, and evaluate it on 26 quantitative extensions of existing SyGuS benchmarks. QuaSi can synthesize optimal solutions in 15/26 benchmarks with times comparable to those needed to find a solution that does not need to satisfy any quantitative objective.

**Contributions** In summary, our contributions are:

– QSyGuS, a formal framework grounded in the theory of weighted grammars that can describe syntax-guided synthesis problems with quantitative objectives over the syntax of the synthesized programs. (§ 3)
– An algorithm for solving QSyGuS problems that leverages closure properties of weighted grammars and existing SyGuS solvers. (§ 4)
– QuaSi, a tool for specifying and solving QSyGuS problems that interfaces with existing SyGuS solvers and a comprehensive evaluation of QuaSi, which shows that QuaSi can efficiently solve QSyGuS problems over different types of weights, including additive weights, probabilities, and combinations of multiple weights. (§ 5)

$$\text{Start} ::= \text{Start} + \text{Start}/(\mathbf{0},\mathbf{1}) \qquad\qquad \text{BExpr} ::= \text{Start} > \text{Start}$$
$$| \ \ \text{if(BExpr) then Start else Start}/(\mathbf{1},\mathbf{0}) \qquad\qquad | \ \ \neg\text{BExpr}$$
$$| \ \ x \ | \ y \ | \ 0 \ | \ 1 \qquad\qquad\qquad\qquad\qquad\qquad | \ \ \text{BExpr} \wedge \text{BExpr}$$

Fig. 1: Weighted grammar that assigns weight $(w_1, w_2) \in \mathsf{Nat} \times \mathsf{Nat}$ to a program where $w_1$ is the number of if-statements and $w_2$ is the number of plus-statements.

## 2  Illustrative Example

In this section, we illustrate the main components of our framework using an example. We start with a Syntax-Guided Synthesis (SyGuS) problem in which no quantitative objective is provided. We recall that the goal of a SyGuS problem is to synthesize a function $f$ of a given type that is accepted by a context-free grammar $G$, and such that $\forall x.\phi(f, x)$ holds (for a given Boolean constraint $\phi$).

The following SyGuS problem asks to synthesize a function that is accepted by the following grammar and that computes the max of two numbers.

$$\text{Start} ::= \text{Start} + \text{Start} \ | \ \text{if(BExpr) then Start else Start} \ | \ x \ | \ y \ | \ 0 \ | \ 1$$
$$\text{BExpr} ::= \text{Start} > \text{Start} \ | \ \neg\text{BExpr} \ | \ \text{BExpr} \wedge \text{BExpr}$$

The semantic constraint is given by the following formula.
$$\psi(f) \stackrel{\text{def}}{=} \forall x, y. f(x, y) \geq x \wedge f(x, y) \geq y \wedge (f(x, y) = x \vee f(x, y) = y)$$
The following three programs are semantically equivalent, but syntactically different solutions.
$$max_1(x, y) = \text{if}(x > y) \text{ then } x \text{ else } y$$
$$max_2(x, y) = \text{if}(x > y) \text{ then } (x + 0) \text{ else } (y + 0)$$
$$max_3(x, y) = \text{if}(x > y) \text{ then } x \text{ else } (\text{if}(y > x) \text{ then } y \text{ else } x)$$

All solutions are correct, but the user might, for example, prefer the smallest one. However, SyGuS does not provide ways to specify this quantitative intent.

*Adding weights* In our formalism, QSyGuS, we augment context-free grammars to assign weights to programs in the search space. Concretely, we adopt weighted grammars [10], a well-studied formalism with many desirable properties. In a weighted grammar, each production is assigned a weight. For example, the weighted grammar shown in Figure 1 extends the one from the previous SyGuS example to assign to each program $p$ a pair of weights $(w_1, w_2)$ where $w_1$ is the number of if-statements and $w_2$ is the number of plus operators in $p$. In this case, the weights are pairs of integers and the weight of a grammar derivation is the pairwise sum of all the weights of the productions involved in the derivation—e.g., the sum of $(w_1, w_2)$ and $(w'_1, w'_2)$ is $(w_1 + w'_1, w_2 + w'_2)$. In the figure, we write $/(w_1, w_2)$ to assign weight $(w'_1, w'_2)$ to a production. We omit the weight for productions with cost $(0, 0)$. The functions $max_1$, $max_2$ and $max_3$ have weights $(1, 0)$, $(1, 2)$, and $(2, 0)$ respectively.

*Adding and solving quantitative objectives* Once we have a way to assign weights to programs, QSyGuS allows the user to specify quantitative objectives over

the weights of the productions—e.g., only allow solutions with fewer than 4 if-statements. In our example, we could require the solution to be minimal with respect to the number of if-statements, i.e., minimize the first component of the paired weight. With these constraints both $max_1$ and $max_2$ would be considered optimal solutions because there exists no solution with 0 if-statements. If we require the solution to also be minimal with respect to the second component of the paired weight, $max_1$ will be *a possible* optimal solution.

Our tool QuaSi can automatically discover solutions in both these cases. Let's consider the last minimization objective. In this case, QuaSi first uses existing SyGuS solvers to synthesize an initial solution using the non-weighted version of the grammar. Let's say that the returned solution is, for example, $max_3$ of weight $(2, 0)$. QuaSi uses this solution to build a new SyGuS instance that only accepts programs with at most one if-statement. Solving this SyGuS problem can, for example, result in the program $max_2$ of weight $(1, 2)$, which will require our solver to build yet another SyGuS instance. This approach is repeated and if it terminates, an optimal program is found.

## 3  SyGuS with Quantitative Objectives

In this section, we introduce our framework for defining syntax-guided synthesis problems with quantitative objectives over the syntax of the synthesized programs. We first provide preliminary definitions for notions such as semirings (Section 3.1) and weighted tree grammars (Section 3.2), and then use these notions to augment SyGuS problems with quantitative objectives (Section 3.3).

### 3.1  Weights over Semirings

We now define the universe of weights we will assign to programs. In general, weights are defined using monoids—i.e., sets equipped with an addition operator—but when a grammar is nondeterministic—i.e., it can produce the same term using multiple derivations—the same term might be assigned multiple weights. Hence, we choose to use semirings. Since we also care about optimization objectives, we assume all our semirings are equipped with a partial order.

**Definition 1 (Semiring).** *A (ordered) semiring is a pair $(\boldsymbol{S}, \preceq)$ where (*i*) $\boldsymbol{S} = (S, \oplus, \otimes, 0, 1)$ is an algebra consisting of a commutative monoid $(S, \oplus, 0)$ and a monoid $(S, \otimes, 1)$ such that $\otimes$ distributes over $\oplus$, $0 \neq 1$, and, for every $x \in S$, $x \otimes 0 = 0$, (*ii*) $\preceq \; \subset S \times S$ is a partial order over $S$.*

We often use the word semiring to refer to just the algebra $\mathbf{S}$.

*Example 1.* In this paper, we focus on semirings with the following algebras.

**Boolean** $\mathsf{Bool} = (\mathbb{B}, \vee, \wedge, 0, 1)$. This semiring only contains the values *true* and *false* and is used to represent non-quantitative problems.

**Tropical** $\mathsf{Trop} = (\mathbb{Z} \cup \{\infty\}, \min, +, \infty, 0)$. This semiring is the most common one and is used to assign additive weights—e.g., term sizes and term depth. In this case, we typically consider the order $\preceq \overset{\text{def}}{=} \leq$.

**Probabilistic** $\mathsf{Prob} = ([0,1], +, \cdot, 0, 1)$. This semiring is used to assign probabilities to terms in a grammar. □

In our framework, we allow synthesis problems to have multiple objectives. Hence, we define a product operation to compose semirings. Intuitively, the following operation composes algebras of semirings to create a pair and applies the operation of each algebra to the corresponding projections of the pair. Similarly, two orders can be composed to create an order over pairs of elements. We propose two such compositions, one which assigns equal weights to the two orders (Pareto) one one which prefers one order over the other (Sorted).

**Definition 2 (Products).** *Given two algebras $\boldsymbol{S}_1 = (S_1, \oplus_1, \otimes_1, 0_1, 1_1)$ and $\boldsymbol{S}_2 = (S_2, \oplus_2, \otimes_2, 0_2, 1_2)$, the* product algebra *is the tuple $\boldsymbol{S}_1 \times_{\boldsymbol{S}} \boldsymbol{S}_2 = (S_1 \times S_2, \oplus, \otimes, (0_1, 0_2), (1_1, 1_2))$ such that for every $x_1, x_2 \in S_1$ and $y_1, y_2 \in S_2$, we have $(x_1, y_1) \oplus (x_2, y_2) \overset{\text{def}}{=} (x_1 \oplus_1 x_2, y_1 \oplus_2 y_2)$ and $(x_1, y_1) \otimes (x_2, y_2) \overset{\text{def}}{=} (x_1 \otimes_1 x_2, y_1 \otimes_2 y_2)$.*

*Given two partial orders $\preceq_1 \subset S_1 \times S_1$ and $\preceq_2 \subset S_2 \times S_2$, the* Pareto product *of the two orders is defined as the partial order $\preceq_p = \text{PAR}(\preceq_1, \preceq_2) \subseteq (S_1 \times S_2) \times (S_1 \times S_2)$ such that, for every $x_1, x_2 \in S_1$ and $y_1, y_2 \in S_2$, we have $(x_1, y_1) \preceq_p (x_2, y_2)$ iff $x_1 \preceq_1 x_2$ and $y_1 \preceq_2 y_2$.*

*Given two partial orders $\preceq_1 \subset S_1 \times S_1$ and $\preceq_2 \subset S_2 \times S_2$, the* Sorted product *of the two orders is defined as the partial order $\preceq_s = \text{SORT}(\preceq_1, \preceq_2) \subseteq (S_1 \times S_2) \times (S_1 \times S_2)$ such that, for every $x_1, x_2 \in S_1$ and $y_1, y_2 \in S_2$, we have $(x_1, y_1) \preceq_s (x_2, y_2)$ iff $x_1 \preceq_1 x_2$ or $(x_1 = x_2$ and $y_1 \preceq_2 y_2)$.*

*Example 2.* The weights in the grammar in Figure 1 are from the product semiring $\mathsf{Trop} \times_{\mathbf{S}} \mathsf{Trop}$. When using the Pareto partial orders, we have, for example, $(1,0) \preceq (2,0)$ and $(1,0) \preceq (1,2)$, but $(1,2)$ is incomparable to $(2,0)$. When using the Sorted product, we have, for example, $(1,0) \preceq (1,2) \preceq (2,0)$. □

### 3.2 Weighted Tree Grammars

Since SYGuS defines search spaces using context-free grammars, we propose to extend this formalism with weights to assign costs to terms in the grammar. We focus our attention on a restricted class of context-free grammars called regular tree grammars—i.e., grammars generating regular tree languages—because, to our knowledge, the benchmarks appearing in the SYGuS competition [3] and in practical applications of SYGuS operate over tree grammars. Moreover, it was recently shown that SYGuS problems that are undecidable for context-free grammars become decidable with weighted tree grammars [8].

*Trees* A *ranked alphabet* is a tuple $(\Sigma, rk_\Sigma)$ where $\Sigma$ is a finite set of symbol and $rk_\Sigma : \Sigma \to \mathbb{N}$ associates a rank to each symbol. For every $m \geq 0$, the set of all symbols in $\Sigma$ with rank $m$ is denoted by $\Sigma^{(m)}$. In our examples, a ranked

alphabet is specified by showing the set $\Sigma$ and attaching the respective rank to every symbol as superscript—e.g., $\Sigma = \{+^{(2)}, c^{(0)}\}$. We use $T_\Sigma$ to denote the set of all (ranked) trees over $\Sigma$—i.e., $T_\Sigma$ is the smallest set such that $(i)$ $\Sigma^{(0)} \subseteq T_\Sigma$, $(ii)$ if $\sigma \in \Sigma^{(k)}$ and $t_1, \ldots, t_k \in T_\Sigma$, then $\sigma(t_1, \cdots, t_k) \in T_\Sigma$. In the following we assume a fixed ranked alphabet $(\Sigma, rk_\Sigma)$.

*Weighted Tree Grammars* Tree grammars are similar to word grammars but they generate ranked trees instead of words. Weighted tree grammars augment tree grammars by assigning weights from a semiring to trees. They do so by associating weights to productions in the grammar. Weighted grammars can, for example, compute the height of a tree, the number of occurrences of some node in the tree, or the probability of a tree with respect to some distribution In the following, we assume a fixed semiring $(\mathbf{S}, \preceq)$ where $\mathbf{S} = (S, \oplus, \otimes, 0, 1)$.

**Definition 3 (Weighted Tree Grammar).** *A weighted tree grammar (WTG) is a tuple $G = (N, Z, P, \mu)$, where $N$ is a set of non-terminal symbols with arity 0, $Z$ is an axiom with $Z \in N$, $P$ is a set of production rules of the form $A \rightarrow \beta$ where $A \in N$ is a non-terminal and $\beta$ is a tree of $T(\Sigma \cup N)$, and $\mu : P \rightarrow S$ is a function assigning to each production a weight from the semiring.*

We can now define the semantics of a WTG as a function $\mathrm{w}_G : T_\Sigma \mapsto S$, which assigns weights to trees. Intuitively, the weight of a tree is $\oplus$-sum of the weight of every possible derivation of that tree in a grammar and the weight of a derivation is the $\otimes$-product of the weights of the productions appearing in the derivation. We use $MS(\beta) = \langle X_1, \ldots, X_k \rangle$ to denote the multi-set of all nonterminals appearing in $\beta$ and $\beta[t_1/X_1, \ldots, t_k/X_k]$ to denote the result of simultaneously substituting each $X_i$ with $t_i$ in $\beta$. Given a derivation $p = A \rightarrow \beta$ such that $MS(\beta) = \langle X_1, \ldots, X_k \rangle$, we assume that $p$ is a symbol of arity $k$. A derivation $d$ starting at non-terminal $X$ is a tree of productions $d \in T(P)$ representing one possible way to derive a tree starting from $X$. The derivation has to be such that: $(i)$ the root of $d$ is a production of the form $X \rightarrow \beta$, $(ii)$ for every node $p = A \rightarrow \beta$ in $d$, if $MS(\beta) = \langle X_1, \ldots, X_k \rangle$, then, for every $1 \leq i \leq k$, the $i$-th child of $p$ is a production $X_i \rightarrow \beta_i$. Given a derivation $d$ with root $p = X \rightarrow \beta$, such that $MS(\beta) = \langle X_1, \ldots, X_k \rangle$ and $p$ has children subtrees $d_1, \ldots, d_k$, the tree generated by $d$ is recursively defined as $tree(d) = \beta[tree(d_1)/X_1, \ldots, tree(d_k)/X_k]$. We use $\mathrm{DER}(X, t)$ to denote the set of all derivations $d$ starting at $X$, such that $tree(d) = t$. The weight $\mathrm{DW}(d)$ of a derivation $d$ is the $\otimes$-product of the weights of the productions appearing in the derivation. Finally, the weight of a tree $t$ is the $\oplus$-sum of the weights of all the derivations of $t$ from the initial nonterminal $\mathrm{w}_G(t) = \bigoplus_{d \in \mathrm{DER}(Z,t)} \mathrm{DW}(d)$. A weighted tree grammar is *unambiguous* iff, for every $t \in T_\Sigma$, there exists at most one derivation—i.e., $|\mathrm{DER}(Z, t)| \leq 1$.

Weighted tree grammars generalize weighted tree automata. In particular, a *weighted tree automaton* (WTA) is a WTG in which every production is of the form $A \rightarrow \sigma(T_1, \ldots, T_n)$, where $A \in N$, each $T_i \in N$, and $\sigma \in \Sigma^{(n)}$. Finally, a *tree automaton* (TA) is a WTA over the Boolean semiring—i.e., the TA accepts all trees with some derivations yielding *true*. Similarly, a *tree grammar* (TG) is

a WTG over the Boolean semiring. Given a TA (resp. TG) $G$, we use $L(G)$ to denote the set of trees accepted by $G$—i.e., $L(G) = \{t \mid \mathrm{w}_G(t) = true\}$.

*Example 3.* The weighted grammar in Fig. 1 operates over the semiring $\mathsf{Trop} \times \mathsf{Trop}$, $N = \{\mathrm{Start}, \mathrm{BExpr}\}$, $Z = \mathrm{Start}$, $P$ contains 9 productions, and $\mu$ assigns non-zero weights to two of them. □

Aside from being a natural formalism for assigning weights to trees, TGs and WTGs enjoy properties that make them a good choice for our model. First, WTGs (resp. TGs) are equi-expressive to WTAs (resp. TAs) and have logic characterizations [9,11,10]. Due to this reason, tree grammars are closed under Boolean operations and enjoy decidable equivalence [9]. Second, WTGs enjoy many closure and decidability properties—e.g., given two WTGs $G_1$ and $G_2$, we can compute the grammars $G_1 \oplus G_2$ and $G_1 \otimes G_2$ such that, for every $f$, $\mathrm{w}_{G_1 \oplus G_2}(f) = \mathrm{w}_{G_1}(f) \oplus \mathrm{w}_{G_2}(f)$ and $\mathrm{w}_{G_1 \otimes G_2}(f) = \mathrm{w}_{G_1}(f) \otimes \mathrm{w}_{G_2}(f)$. This operation is convenient for building grammars over product semirings.

### 3.3 QSYGUS

In this section, we formally define QSYGUS, which extends SYGUS with quantitative objectives. In SYGUS a problem is specified with respect to a background theory $T$—e.g., linear arithmetic—and the goal is to synthesize a function $f$ that satisfies two constraints provided by the user. The first constraint describes a *functional semantic property* that $f$ should satisfy and is given as a predicate $\psi(f) \stackrel{\text{def}}{=} \forall x.\phi(f,x)$. The second constraint limits the *search space* $S$ of $f$ and is given as a set of expressions specified by a context-free grammar $G$ defining a subset of all the terms in $T$. A solution to the SYGUS problem is an expression $e$ in $S$ such that the formula $\psi(e)$ is valid.

We augment such a framework in two ways. First, we replace context free grammars with WTGs, which we use to assign weights (from a given semiring) to terms. Second, we augment the problem formulation with constraints over the weight of the synthesized program—i.e., only consider programs of weight greater than 2—and optimization objectives over the same weight—i.e., find the solution of minimal weight. Weight constraints range over the grammar

$$WC := WC \wedge WC \mid WC \vee WC \mid \neg WC \mid w \preceq s \mid s \preceq w \mid w \prec s \mid s \prec w,$$

where $w$ is a special variable and $s$ is an element of the semiring under consideration. Given a constraint $\omega \in WC$, we write $\omega(t)$ to denote the term obtained by replacing $w$ with $t$ in $\omega$.

**Definition 4 (QSYGUS).** *A QSYGUS problem is a tuple $(T, (\boldsymbol{S}, \preceq), \psi(f), G, \omega, \mathrm{OPT})$ where:*

- *$T$ is a background theory.*
- *$(\boldsymbol{S}, \preceq)$ is an ordered semiring defining the set of weights and their operations.*
- *$G$ is a weighted tree grammar with weights over the semiring $\boldsymbol{S}$ and that only contains terms in $T$—i.e., $L(G) \subseteq T$.*

---

**Algorithm 1** QSYGUS synthesis algorithm

---

1: **procedure** QSYGUS-SOLVE$(T, \mathbf{S}, \psi, G, \omega, \text{OPT})$
2:     $G' \leftarrow \text{REDUCEGRAMMAR}(G, \omega)$            $\triangleright$ extract grammar satisfying $\omega$
3:     $f^* \leftarrow \text{SYGUS}(T, \psi, G')$            $\triangleright$ solve corresponding SYGUS problem
4:     **if** $\text{OPT} = false$ **then return** $f^*$
5:     **while** true **do**
6:        $G' \leftarrow \text{REDUCEGRAMMAR}(G', w \prec \text{W}_G(f^*))$
7:        $f \leftarrow \text{SYGUS}(T, \psi, G')$            $\triangleright$ Try to find better solution
8:        **if** $f = \bot$ **then return** $f^*$          $\triangleright$ Return the optimal solution
9:        $f^* \leftarrow f$

---

- $\psi(f) \stackrel{\text{def}}{=} \forall x.\phi(f, x)$ *is a Boolean formula constraining the semantic behavior of the synthesized program* $f$.
- $\omega \in WC$ *is a set of constraints over the weight* $w$ *of the synthesized program.*
- OPT *is a Boolean denoting whether the solution has to have minimal weight with respect to* $\preceq$.

*A solution to the* QSYGUS *problem is a term* $e$ *such that* $e \in L(G)$, $\psi(e)$ *is true, and* $\omega(\text{W}_G(e))$ *is true. If* OPT *is true, we also require that there is no* $g$ *that satisfies the previous conditions and such that* $\omega(\text{W}_G(g)) \prec \omega(\text{W}_G(e))$.

A SYGUS problem is a QSYGUS problem without weight constraints—i.e., $\omega \equiv true$ and $\text{OPT} = false$. We denote such problems just as triples $(T, \psi(f), G)$.

*Example 4.* Consider the QSYGUS problem described in Section. 2. We already described all the components but $\omega$ and OPT in the rest of this section. In this example, $\omega = true$ and $\text{OPT} = true$ because we want to synthesize the solution with minimal weight.

## 4   Solving QSYGUS Problems via Grammar Reduction

In this section, we present an algorithm for solving QSYGUS problems (Algorithm 1), which works as follows. First, given a QSYGUS problem, we construct (under certain assumptions) a SYGUS problem for which the solution is guaranteed to satisfy the weight constraints $\omega$ (line 2) and use existing SYGUS solvers to find a solution to such a problem (line 3). If the QSYGUS problem requires minimization, our algorithm produces a new SYGUS instance to search for a solution that is better than the previously found one and tries to solve it (lines 6-7). This procedure is repeated until an optimal solution is found (line 8).

### 4.1   From QSYGUS to SYGUS

The first step of our algorithm is to construct a SYGUS problem characterizing exactly all the solutions of the QSYGUS problem that satisfy the weight

constraints. Given a QSYGUS problem $P = (T, (\mathbf{S}, \preceq), \psi(f), G, \omega, \text{OPT})$, we construct a SYGUS problem $P' = (T, \psi(f), G')$ such that a function $g$ is a solution to the SYGUS problem $P'$ iff $g$ is a solution of $P = (T, (\mathbf{S}, \preceq), \psi(f), G, \omega, false)$, where the optimization constraint has been dropped. We denote the grammar reduction operation as $G' \leftarrow \text{REDUCEGRAMMAR}(G, \omega)$.

*Base case* First we show how to solve the problem when $\omega$ is an atomic formula—i.e. of the form $w \preceq s$, $s \preceq w$, $w \prec s$, or $s \prec w$. We start by showing how to solve the problem for $w \preceq s$ as the construction is identical for the other constraints.

Concretely, we are given a WTG $G = (N, Z, P, \mu)$ and we want to construct a TG $G_{\preceq s} = (N', Z', P')$ such that $t \in L(G_{\preceq s})$ iff $\text{w}_G(t) \preceq s$. In general, it is not possible to perform this construction for arbitrary semirings and grammars. We first present our algorithm and then describe sufficient conditions under which we can ensure termination and correctness.

The idea behind our construction is to introduce new nonterminals in the grammar $G_{\preceq s}$ to keep track of the weight of the trees that can be produced from those nonterminals. For example, a nonterminal pair $(X, s')$ will derive all trees derivable from $X$ using a single derivation of weight $s'$. Therefore, the set of nonterminals $N'$ is a subset of $N \times S$ (plus an initial nonterminal $Z'$), where $S$ is the universe of the WTG's semiring. We construct our set of nonterminals $N'$ starting from the leaf productions of $G$ and then recursively explore other productions. At the same time we generate the set of productions $P'$. Formally, $N'$ and $P'$ are the smallest sets such that the following conditions hold.

1. $Z' \in N'$ (the initial nonterminal).
2. For every production $p \in P$ such that $p = (A \rightarrow \beta)$ and $\beta \in T_\Sigma$—i.e., $p$ is a leaf—and $\mu(p) \preceq s$, then $(A, \mu(p)) \in N'$ and $((A, \mu(p)) \rightarrow \beta) \in P'$. If $A = Z$, then $Z' \rightarrow (A, \mu(p)) \in P'$.
3. For every production $p \in P$ such that $p = (A \rightarrow \beta)$, $MS(\beta) = \langle X_1, \ldots, X_k \rangle$, $(X_1, s_1), \ldots, (X_k, s_k) \in N'$ (for some values $s_i \in S$), and $\mu(p) \otimes s_1 \otimes \ldots \otimes s_k = s'$, $s' \preceq s$, then $(A, s') \in N'$, and $((A, s') \rightarrow \beta[(X_1, s_1)/X_1, \ldots, (X_k, s_k)/X_k]) \in P'$. If $A = Z$, then $Z' \rightarrow (A, s') \in P'$.

*Example 5.* We illustrate our construction using the grammar in Figure 1 . Assume the weight constraint is $w \preceq (1, 0)$ and the partial order is built using a Pareto product—i.e., we accept terms with 1 or less if-statements and no plus-statements. Our construction yields the following grammar.

```
        Z' ::=  (Start,1,0) | (Start,0,0)
  (Start,1,0) ::=  if((BExpr,0,0)) then (Start,0,0) else (Start,0,0) | x | y | 0 | 1
  (Start,0,0) ::=  x | y | 0 | 1
  (BExpr,0,0) ::=  (Start,0,0) > (Start,0,0) | ¬(BExpr,0,0) | (BExpr,0,0) ∧ (BExpr,0,0)
```

$\square$

The construction of $G_{\preceq s}$ only terminates for certain semirings and grammars, and only guarantees that individual derivations yield the correct cost—i.e., it does not account for the $\oplus$-sum of multiple derivations.

*Example 6.* The following WTG over Prob is ambiguous and, if we apply the grammar reduction algorithm for $\omega := w \preceq 0.6$, the resulting grammar will be empty. However, the tree $1 + 1$ has weight $0.9 \preceq 0.6$ $(0.9 \geq 0.6)$.

$$
\begin{array}{ll}
\text{Start} ::= & \text{Start} + \text{Start}/0.5 \qquad \text{Expr} ::= \text{Expr} + \text{Expr}/0.4 \\
\mid & x \mid 0 \mid 1 \mid \text{Expr} \qquad\qquad\qquad \mid x \mid 0 \mid 1
\end{array}
$$

□

We now identify sufficient conditions under which the construction of $G_{\preceq s}$ terminates and is sound. In particular, we start by restricting our attention to unambiguous WTGs, which are the common ones in practice. We use $\text{WEIGHTS}(G) = \{s \mid p \in P \wedge \mu(p) = s\}$ to denote the set of weights used by $G$ and $M_{\mathbf{S},G} = (S', \otimes, 1)$ to denote the submonoid of $\mathbf{S}$ generated by $\text{WEIGHTS}(G)$—i.e., the set of all weights we can generate using $\otimes$ and $\text{WEIGHTS}(G)$.

**Theorem 1.** *Given an unambiguous WTG $G$ over a semiring $\mathbf{S}$ such that $M_{\mathbf{S},G} = (S', \otimes, 1)$, and a weight $s \in S$, the construction of $G_{\preceq s}$ terminates if the set $\{s' \mid s' \preceq s \wedge w \in S'\}$ is finite. Moreover, if the set of weights $\text{WEIGHTS}(G)$ is monotonically increasing with respect to $\preceq$—i.e. for every $s \in S$ and $s' \in \text{WEIGHTS}(G)$, $s \preceq s \otimes s'$—then $L(G_{\preceq s})$ contains exactly every tree $t$ such that $\text{W}_G(t) \preceq s$.*

The theorem above also holds for other atomic constraints $w \prec s$, $s \preceq w$, or $s \prec w$ (for these last two, the direction of the monotonicity is reversed). Moreover, in certain cases, even if the construction may not terminate for, let's say $s \preceq w$, it might terminate for the negated constraint $w \prec s$. In such a case, we can use the closure properties of regular tree grammars/automata to construct the reduced grammar for $s \preceq w$ as $G_{\preceq w} = \text{INTERSECT}(G, \text{COMPLEMENT}(G_{\succ w}))$. The same idea can be applied to all atomic constraints.

In practice, the restriction of Theorem 1 holds for grammars that operate over the Boolean and probabilistic semirings, and the tropical semiring only with positive weights. Theorem 1 never holds when $\mathbf{S}$ is the tropical semiring and the grammar contains negative weights. In general, one cannot construct the constrained grammar in this case. However, it is easy to modify our algorithm to work with grammars that do not contain loops—i.e., derivations from a nonterminal to a tree containing the same nonterminal—with negative weights.

Intuitively, when the grammar contains no negative loops, we can find a constant $SH$ such that any intermediate derivation with weight greater than $s + SH$ will never result in tree with weight smaller than $s$. We use this idea to modify the construction of $G_{\leq s}^{\text{Trop}}$—i.e., $G_{\leq s}$ for Trop—as follows. First, this constant is bounded by $ck^{n+1}$ where $c$ is the absolute value of the smallest negative weight in the grammar, $k$ is the largest number of nonterminals appearing in one grammar production, and $n = |N|$ is the number of nonterminals. Second, in steps 2 and 3 of the construction, a new nonterminal and the corresponding productions are produced if $\mu(p) \leq s + |SH|$ (previously $\mu(p) \leq s$). However, if $A = Z$ in steps 2 and 3, we add a new production $Z' \rightarrow (A, s')$ only if $s' \preceq s$.

We now show when this construction terminates and return correct values. Since the tropical semiring combines multiple runs using the min operator, we can *drop* the requirement that the grammar has to be unambiguous.

**Theorem 2.** *Given a WTG $G$ over $\mathsf{Trop}$ and a weight $s \in \mathbb{Z}$, the construction of $G^{\mathsf{Trop}}_{\leq s}$ terminates if $G$ contains no loop with cumulative negative weight. Moreover, $G^{\mathsf{Trop}}_{\leq s}$ contains exactly every tree $t$ such that $\mathrm{W}_G(t) \leq s$.*

*Composing semirings* We next discuss how Theorem 1 relates to product semirings. Given a grammar $G = (N, Z, P, \mu)$ over a semiring $\mathbf{S}_1 \times_{\mathbf{S}} \mathbf{S}_2$, we use $G^{\mathbf{S}_i}$ to denote the grammar $(N, Z, P, \mu_i)$ in which the weight function outputs the corresponding projected weight—i.e., if $\mu(p) = (s_1, s_2)$, then $\mu_i(p) = s_i$.

Let's first consider the case where the product semiring uses a Pareto partial order. In this case, if Theorem 1 holds for each grammar $G^{\mathbf{S}_i}$ and $w_i \preceq_i s_i$, then it holds for $G$ and $(w_1, w_2) \preceq_p (s_1, s_2)$. However, the other direction is not true. Theorem 3 proves this intuition and states that, in some sense, solving Pareto partial orders is easier than solving the individual partial orders.

**Theorem 3** *Given an unambiguous WTG $G$ over the semiring $\boldsymbol{S} = \boldsymbol{S}_1 \times_{\boldsymbol{S}} \boldsymbol{S}_2$ with Pareto partial order $\preceq_p = \mathrm{PAR}(\preceq_1, \preceq_2)$ and a weight $s = (s_1, s_2) \in S$, if the constructions $G^{\boldsymbol{S}_1}_{\preceq_1 s_1}$ and $G^{\boldsymbol{S}_2}_{\preceq_2 s_2}$ terminate, then the construction of $G_{\preceq s}$ terminates.*

When we move to Sorted partial order we cannot get an analogous theorem: if Theorem 1 holds for each grammar $G^{\mathbf{S}_i}$ and $w_i \preceq_i s_i$, then it does not necessary hold for $G$ and $(w_1, w_2) \preceq_s (s_1, s_2)$. In particular, if the semiring $\mathbf{S}_2$ is infinite and there exists an $s'_1 \prec s_1$, there will be infinitely many elements $(s'_1, \_) \prec (s_1, s_2)$. Using this observation, we devise a modified algorithm for reducing grammars with sorted objectives. First, we compute the grammars $G^{\mathbf{S}_1}_{\prec_1 s_1}$, $G^{\mathbf{S}_1}_{=s_1}$, and $G^{\mathbf{S}_2}_{\prec_2 s_2}$. Second, we use WTG closure properties to compute $G_{\preceq_s}(s_1, s_2)$ as the union of $G^{\mathbf{S}_1}_{\prec_1 s_1}$ and $\mathrm{INTERSECT}(G^{\mathbf{S}_1}_{=s_1}, G^{\mathbf{S}_2}_{\prec_2 s_2})$.

*General formulas* We can now inductively construct the grammar accepting only terms satisfying all constraints in $\omega$. We again use the fact that tree grammars are closed under Boolean operations to compute intersections and unions and correctly characterize all conjunctions and unions appearing in the formulas.

## 4.2 Finding an Optimal Solution

If our QSYGUS problem does not require minimization—i.e., OPT $= false$—the technique presented in Section 4.1 can be used to generate an equivalent SYGUS problem $P' = (T, \psi(f), G')$, which can be solved using off-the-shelf SY-GUS solvers. In this section, we show how to extend this technique to handle minimization objectives. Our idea is to use SYGUS solvers to find a non-optimal solution for $P'$ and then iteratively refine our grammar $G'$ to search for a better solution. This loop is illustrated in Algorithm 1 (lines 5-9). Given the initial solution $f^*$ to $P'$ such that $\mathrm{W}_G(f^*) = s$, we can construct a new grammar $G_{\prec s}$ and look for a solution with lower weight. If the SYGUS solver we use is sound—it can find a solution if it exists—and complete—it can detect if a solution does not exist—Algorithm 1 terminates with an optimal solution.

In general, the above conditions are too strict and in practice this implies that the algorithm will often not terminate. However, if the SyGuS solver is sound, the Algorithm 1 will eventually find the optimal solution, but it will not be able to prove that no smaller one exists. In our experiments, we will show that this approach can yield better solutions than those given by vanilla SyGuS solvers even when Algorithm 1 does not terminate.

## 5 Implementation and Evaluation

First, We extended the SyGuS format with new syntax for expressing QSy-GuS problems. Our format supports all semirings presented in Section 3.1 as well as additional ones. The format also allows creating tuples of semirings using the product operation described in Section 3.1. We augment the original SyGuS syntax to support weights on grammar productions. Weight constraints are added using an SMT-like syntax.

Second, we implemented Algorithm 1 in a tool called QuaSi. QuaSi already interfaces with three SyGuS solvers: CVC4 [6], ESolver [4], and EUSolver [5]. QuaSi supports all the semirings allowed in our format and implements a library for tree automata/grammars and weighted tree automata/grammars operations, as well as several optimizations we did not discuss in the paper. In particular, QuaSi often uses simple grammar reduction techniques to simplify the generated grammars, remove unnecessary productions, and consolidate equivalent ones.

We evaluate QuaSi through the following questions (experiments performed on an Intel Core i7 4.00GHz CPU with 32GB/RAM).

**Q1** Can QuaSi solve quantitative variants of real SyGuS benchmarks? (§ 5.1)
**Q2** What is the overhead of synthesizing optimal solutions? (§ 5.2)
**Q3** How do multiple iterations of Alg. 1 affect the solution's weight? (§ 5.3)
**Q4** Can QuaSi solve QSyGuS problems with multiple objectives? (§ 5.4)

*Benchmarks* We perform our evaluation on 26 quantitative extensions of existing SyGuS competition benchmarks taken from 4 SyGuS benchmark tracks [4]: Hackers Delight, Integers, ICFP and Bitvector. 18 of our benchmarks only use a minimization objective over a single semiring (Table 1), while 8 use a minimization objective (Pareto or Sorted) over a product semiring (Table 2). We select SyGuS benchmarks using the following criteria: (*i*) the benchmark can be solved by either CVC4 [6] or ESolver [4], and (*ii*) the solution is not optimal according to some reasonable metric—e.g., size or number of if statements.

### 5.1 Effectiveness of QSyGuS Solver

We evaluate the effectiveness of QuaSi on the 18 single-minimization-objective benchmarks. For each benchmark, we run QuaSi using either CVC4 or ESolver as the backend SyGuS solver (we also evaluated QuaSi using EUSolver [5], but, due to its poor performance, we do not report the results). The results are shown in Table 1. The timeout for each iteration of Alg. 1 is 10 minutes.

Table 1: Performance of QUASI. **Time** shows the sequence of times taken to solve individual iterations of Alg. 1. **Largest** is the size of the largest SYGUS sub-problem. **Grammar Size** is the number of rules in the original grammar.

| Problem | CVC4 | | ESolver | | Grammar |
|---|---|---|---|---|---|
| | **Time**[sec] | **Largest** | **Time**[sec] | **Largest** | **Size** |
| max_ite(2,3) | 0.1+0.1 | 42 | 0.1 | 42 | 13 |
| max_ite(2,15) | 0.1+0.1 | 239 | 0.3 | 239 | 13 |
| max_ite(3,15) | 0.1+0.1+0.1 | 238 | OOM | 238 | 13 |
| max_ite(10,15) | 0.5+0.5+0.9 | 226 | OOM | 226 | 13 |
| parity_not | 0.1+TO | 301 | 26.9+TO | 43 | 6 |
| max3_ite | 0.1+TO | 31 | OOM | - | 14 |
| array_search_3 | 0.1+TO | 135 | TO | - | 15 |
| array_search_5 | 0.1+TO | 108 | TO | - | 16 |
| hackers_5 | 0.1+0.1 | 27 | 0.1+0.1+0.1 | 35 | 13 |
| hackers_7 | 0.1+0.3 | 35 | 0.1+0.1+0.2 | 41 | 13 |
| hackers_17 | 0.1+0.7 | 41 | 2.8+3.0+1.0 | 62 | 13 |
| hackers_19 | 0.2+TO | 174 | TO | - | 13 |
| icfp_7 | 0.2+TO | 146 | TO | - | 11 |
| LinExpr_eq1ex | 0.7+TO | 1717 | TO | - | 14 |
| hackers_2_prob | 0.6+4.1+0.1 | 95 | 0.8+0.1+0.2 | 154 | 13 |
| hackers_5_prob | 0.1+0.9+0.1 | 96 | 0.1+0.2+0.1 | 154 | 13 |
| hackers_7_prob | 0.1+TO | 162 | 0.1+0.1+0.2 | 212 | 13 |
| hackers_17_prob | 0.1+TO | 187 | 3.4+6.5+OOM | 291 | 13 |

With CVC4, QUASI terminates with an optimal solution in 9/18 benchmarks, taking less than 5 seconds (avg: 0.7s) to solve each sub-problem. In 3 of these cases, the initial solution is already optimal and the second iteration is used to prove optimality. With ESolver, QUASI terminates with an optimal solution in 8/18 benchmarks, taking less than 7 seconds (avg: 0.9s) to solve each sub-problem. In 2 cases, it can find a better solution than the original one, but it cannot prove that the solution is optimal. Overall, by combining solvers, QUASI can find a better solution than the original SYGUS solution given by one of the two solvers in 9/18 benchmarks. QUASI cannot improve the initial solution of the linear integer arithmetic benchmarks (array_search and LinExpr_eq1ex).

Both solvers timeout on large grammars. The grammars in Table 1 are 1 to 2 order of magnitude larger than those in existing SYGUS benchmarks (avg: 224 vs 13 rules) and existing solvers have not yet been optimized for this parameter. In some cases, the solver times out for intermediate grammars that do not contain a solution, but that generate infinitely many terms. In general, existing SYGUS solvers cannot prove unsatisfiability for these types of problems. To answer **Q1**, QUASI can **solve quantitative variants of 10/18 real SyGuS benchmarks**.

## 5.2 Solving Time for Different Iterations

In this section, we evaluate the time required by each iteration of Alg. 1. Figure 2 shows the ratio of time taken by each iteration with respect to the initial non-

quantitative SYGUS solving time. Some of the iterations shown in Figure 1 do not appear in Figure 2 since they resulted in no solution—i.e., the initial solution was minimal. CVC4 is typically slower in subsequent iterations and can take up to 10 times the original solving time, while ESolver has comparable runtime to the initial run and is often faster. These numbers are largely due to how the two solvers work: CVC4 is optimized to solve problems where the grammar imposes no restrictions on the structure of the solution, while ESolver performs enumerative search and takes advantage of more restrictive grammars.

One interesting point is the `parity_not` benchmark. ESolver takes 26.9s to find an initial solution. But, with a weight constraint $w < 11$, an solution can be found in 2.2s. CVC4 can find the initial solution with weight 11 in 0.1s but cannot solve the next iteration. We tried using different solvers in different iterations of our algorithm and, in fact, found that, if we use CVC4 to find an initial solution and then ESolver in subsequent iterations with restricted grammars we can fully solve this benchmark in a total of 2.3s which is much better than the time taken by a single solver. To answer **Q2**, with appropriate choices of solvers **the overhead of synthesizing optimal solutions is minimal**.



Fig. 2: Solving time across iterations

## 5.3 Solution Weight across Iterations

In this section, we present how the weight of the synthesized solutions change across each iteration of Alg. 1. Figure 3 shows the percentage of weight of solutions synthesized at each iteration with respect to the weight of the initial SYGUS solution. The result shows that we can improve the solutions of CVC4 by 15-25% in one iteration, and the solutions of ESolver by 20-50% when taking one iteration and 50-60% when taking two. The Prob benchmarks, which require two iterations, can be improved more when using ESolver because ESolver tends to synthesize small terms whose probability may also be small. To answer **Q3**, QUASI can **improve the weights of SYGUS solutions by 20-60%**.



Fig. 3: Solution weight across iterations.

## 5.4 Multi-Objective Optimization

In this section, we evaluate the effectiveness of QUASI on the 8 benchmarks involving two minimization objectives. The benchmarks consists of two families,

Table 2: Performance of QuaSi on multi-objective benchmarks. **Weight** denotes the sequence of weights explored during minimization.

| | Problem | Time[sec] | Weight | Largest | Size |
|---|---|---|---|---|---|
| | `array_search_sorted` | TO | - | - | 15 |
| | `hackers_5_sorted` | 0.1+0.1+01 | $(0,3) \to (0,2)$ | 31 | 13 |
| | `hackers_7_sorted` | 0.1+0.3+0.1 | $(1,4) \to (0,5) \to (0,3)$ | 72 | 13 |
| Trop × Trop | `hackers_17_sorted` | 0.1+156.1+TO | $(2,5) \to (1,4) \to (0,6)$ | 97 | 13 |
| | `array_search_pareto` | TO | - | - | 15 |
| | `hackers_5_pareto` | 0.1+0.1+01 | $(0,3) \to (0,2)$ | 31 | 13 |
| | `hackers_7_pareto` | 0.1+0.3+0.1 | $(1,4) \to (1,3) \to (0,3)$ | 74 | 13 |
| | `hackers_17_pareto` | 0.1+9.1+0.1 | $(2,5) \to (2,4) \to (1,4)$ | 54 | 13 |

4 for sorted optimization and 4 for Pareto optimization. The sorted optimization benchmarks ask to minimize first the number of occurrences of specified operator (`bvand` in `hacks` and `ite` in `array_search`) and then the size of the solution. The Pareto optimization benchmarks have the same objectives as sorted optimization but here we are synthesizing a Pareto optimal solution instead of sorted optimal one. The results are shown in Table 2. We do not present the results using CVC4 because it cannot solve any of the benchmarks.

The `array_search` times out since it is already hard on a single objective. For the `hackers_5` benchmarks, the initial solution is already optimized for the first objective, so the problem degenerates to the single-objective optimization problem. For the `hackers_7` and `hackers_17`, we present the weights of the intermediate solutions we can see that Pareto and Sorted optimizations yield different solutions. To answer **Q4**, QuaSi can **solve problems with multiple objectives** when the same problems are feasible with a single objective.

## 6   Related Work

*Qualitative Synthesis* Existing program synthesizers fall in three categories: (*i*) enumeration solvers, which typically output the smallest program [1], (*ii*) symbolic solvers, which reduce the synthesis problem to a constraint solving problem and output whatever program is produced by the constraint solver [21], (*iii*) probabilistic synthesizers, which randomly search the space for a solution and are typically unpredictable [18]. Since the introduction of the SyGuS format [2], these techniques have been used to build several SyGuS solvers that have competed in SyGuS competitions [4]. The most effective ones, which are used in this paper are ESolver and EUSolver [1] (enumeration), and CVC4 [6] (symbolic).

*Quantitative synthesis* Domain-specific synthesizers typically employ hard-coded ranking functions that guide the search towards a "preferable" program [17], but these functions are typically hard to write and are decoupled from the functional specification. Unlike QSyGuS, these synthesizers allow arbitrary ranking functions to be expressed in general purpose languages, but typically only support limited grammars for synthesis. Moreover, in many practical applications the

ranking functions are very simple. For example, the popular spreadsheet formula synthesizer FlashFill [12] uses a ranking function to prefer small programs with few constants. This type of objective is expressible in our framework.

The Sketch synthesizer supports optimization objectives over variables in sketched programs [20]. This work differs from ours in that sketches are a different specification mechanism from SyGuS. In Sketch the search space is encoded as a program with holes to facilitate synthesis by constraint solving. Translating SyGuS problems into sketches is non-trivial and results in poor performance.

The work closest to ours is Synapse [7], which combines sketching with an approach similar to ours. For the same reasons as for Sketch, Synapse differs from our work because it proposes a different search space mechanisms. However, there are a few analogies between our work and Synapse that are worth explaining in detail. Synapse supports syntactic cost functions that are defined using a decidable theory, and separately from the sketch search space. Synthesis is done using an iterative search where sketches—i.e., set of partial programs with holes—of increasing sizes are given to the synthesizer. At the high level, the intermediate sketches are related to our notion of reduced grammars—i.e., they accept solution of weight less than a given constant. However, while our algorithm generates reduced grammars automatically for a well-defined family of semirings, Synapse requires the user to provide a function for generating the intermediate sketches. Moreover, since Synapse requires cost functions that are defined using a decidable theory, it would not support certain families of costs QSyGuS supports—e.g., the probabilistic semiring.

Koukoutos et al. [15] have proposed the use of probabilistic tree grammars to guide the search of enumerative synthesizers on applications outside of SyGuS. Their algorithm enumerates all terms accepted by the grammar in decreasing probability using a variant of the search algorithm $A^*$ and requires the grammar to not contain transitions of weight 1 to avoid getting stuck. Probabilistic tree grammars are a special case of QSyGuS and our algorithm does not impose limitations of what weights can appear in the grammar. Moreover, our algorithm does not require implementing a new solver when changing the cost semiring.

## 7 Conclusion

We presented QSyGuS, a general framework for defining and solving SyGuS problems in the presence of quantitative objectives over the syntax of the programs. QSyGuS is (*i*) *natural*: requires minimal modification to the SyGuS format, (*ii*) *general*: it supports complex but practical types of weights, (*iii*) *formal*: it is grounded in the theory of weighted tree grammars, (*iv*) *effective*: our tool QuaSi can solve quantitative variations of existing SyGuS benchmarks with little overhead. In the future, we plan to extend our framework to handle probabilistic objectives and quantitative objectives over the semantics of the program—e.g., synthesize programs that satisfy most of the specification.

# References

1. *ESolver*. `https://github.com/abhishekudupa/sygus-comp14`.
2. R. Alur, R. Bodik, G. Juniwal, M. M. Martin, M. Raghothaman, S. A. Seshia, R. Singh, A. Solar-Lezama, E. Torlak, and A. Udupa. Syntax-guided synthesis. In *Formal Methods in Computer-Aided Design (FMCAD), 2013*, pages 1–8. IEEE, 2013.
3. R. Alur, D. Fisman, R. Singh, and A. Solar-Lezama. Results and analysis of SyGuS-comp'15. *arXiv preprint arXiv:1602.01170*, 2016.
4. R. Alur, D. Fisman, R. Singh, and A. Solar-Lezama. Sygus-comp 2016: results and analysis. *arXiv preprint arXiv:1611.07627*, 2016.
5. R. Alur, A. Radhakrishna, and A. 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.
6. C. Barrett, C. L. Conway, M. Deters, L. Hadarean, D. Jovanović, T. King, A. Reynolds, and C. Tinelli. Cvc4. In *Proceedings of the 23rd International Conference on Computer Aided Verification*, CAV'11, pages 171–177, Berlin, Heidelberg, 2011. Springer-Verlag.
7. J. Bornholt, E. Torlak, D. Grossman, and L. Ceze. Optimizing synthesis with metasketches. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '16, pages 775–788, New York, NY, USA, 2016. ACM.
8. B. Caulfield, M. N. Rabe, S. A. Seshia, and S. Tripakis. What's decidable about syntax-guided synthesis? *CoRR*, abs/1510.08393, 2015.
9. H. Comon, M. Dauchet, R. Gilleron, C. Löding, F. Jacquemard, D. Lugiez, S. Tison, and M. Tommasi. Tree automata techniques and applications. Available on: `http://www.grappa.univ-lille3.fr/tata`, 2007. release October, 12th 2007.
10. M. Droste, W. Kuich, and H. Vogler. *Handbook of Weighted Automata*. Springer Publishing Company, Incorporated, 1st edition, 2009.
11. M. Droste and H. Vogler. Weighted tree automata and weighted logics. *Theoretical Computer Science*, 366(3):228 – 247, 2006. Automata and Formal Languages.
12. S. Gulwani. Automating string processing in spreadsheets using input-output examples. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, pages 317–330, 2011.
13. S. Gulwani. Programming by examples: Applications, algorithms, and ambiguity resolution. In *Automated Reasoning - 8th International Joint Conference, IJCAR 2016, Coimbra, Portugal, June 27 - July 2, 2016, Proceedings*, pages 9–14, 2016.
14. Q. Hu and L. D'Antoni. Automatic program inversion using symbolic transducers. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2017, Barcelona, Spain, June 18-23, 2017*, pages 376–389, 2017.
15. M. Koukoutos, M. Raghothaman, E. Kneuss, and V. Kuncak. On repair with probabilistic attribute grammars. *CoRR*, abs/1707.04148, 2017.
16. V. C. Ngo, M. Dehesa-Azuara, M. Fredrikson, and J. Hoffmann. Verifying and synthesizing constant-resource implementations with types. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 710–728, May 2017.
17. O. Polozov and S. Gulwani. Flashmeta: a framework for inductive program synthesis. In *Proceedings of the 2015 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications, OOPSLA*

18

*2015, part of SPLASH 2015, Pittsburgh, PA, USA, October 25-30, 2015*, pages 107–126, 2015.

18. E. Schkufza, R. Sharma, and A. Aiken. Stochastic program optimization. *Commun. ACM*, 59(2):114–122, 2016.

19. R. Singh and S. Gulwani. Predicting a correct program in programming by example. In *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I*, pages 398–414, 2015.

20. R. Singh, S. Gulwani, and A. Solar-Lezama. Automated feedback generation for introductory programming assignments. In *Proceedings of PLDI'13*, pages 15–26, New York, NY, USA, 2013. ACM.

21. A. Solar-Lezama. Program sketching. *International Journal on Software Tools for Technology Transfer*, 15(5):475–495, Oct 2013.

## A    Format Examples

The following are three examples of QSYGUS file formats.

### A.1    max_ite(2,3)

```
(set-weight (w1 TROP))
(set-logic LIA)

(declare-var x Int)
(declare-var y Int)
(synth-fun max2 ((x Int) (y Int)) Int
        ((Start Int (x y 0 1
                (+ Start Start)
                (- Start Start)
                ((ite StartBool Start Start) : 1)))
        (StartBool Bool ((and StartBool StartBool)
                (or StartBool StartBool)
                (not StartBool)
                (<= Start Start)
                (= Start Start)
                (>= Start Start)))))

(constraint (>= (max2 x y) x))
(constraint (>= (max2 x y) y))
(constraint (or (= x (max2 x y)) (= y (max2 x y))))

(weight-constraint (and (<= w1 3) (<= 2 w1)))
(optimize w1)

(check-synth)
```

### A.2    hackers_17_prob

```
(set-weight (w1 PROB))
(set-logic BV)

(define-fun hd17 ((x (BitVec 32))) (BitVec 32)
        (bvand (bvadd (bvor x (bvsub x #x00000001)) #x00000001) x))

(synth-fun f ((x (BitVec 32))) (BitVec 32)
        ((Start (BitVec 32) (((bvnot Start):0.125)
                ((bvxor Start Start):0.125)
```

```
                          ((bvand Start Start):0.125)
                          ((bvor Start Start) : 0.125)
                          ((bvneg Start):0.125)
                          ((bvadd Start Start):0.25)
                          ((bvmul Start Start):0.25)
                          ((bvudiv Start Start):0.25)
                          ((bvurem Start Start):0.25)
                          ((bvlshr Start Start):0.25)
                          ((bvashr Start Start):0.25)
                          ((bvshl Start Start):0.25)
                          ((bvsdiv Start Start):0.25)
                          ((bvsrem Start Start):0.25)
                          ((bvsub Start Start):0.25)
                          x
                          #x0000001F
                          #x00000001
                          #x00000000
                          #xFFFFFFFF))))

(declare-var x (BitVec 32))
(constraint (= (hd17 x) (f x)))
(optimize w1)
(check-synth)
```

## A.3   hackers_17_pareto

```
(set-weight (w1 TROP) (w2 TROP))
(set-logic BV)

(define-fun hd17 ((x (BitVec 32))) (BitVec 32)
        (bvand (bvadd (bvor x (bvsub x #x00000001)) #x00000001) x))

(synth-fun f ((x (BitVec 32))) (BitVec 32)
        ((Start (BitVec 32) (((bvnot Start): 0 1)
                  ((bvxor Start Start): 0 1)
                  ((bvand Start Start): 1 1)
                  ((bvor Start Start) : 0 1)
                  ((bvneg Start):0 1)
                  ((bvadd Start Start):0 1)
                  ((bvmul Start Start):0 1)
                  ((bvudiv Start Start):0 1)
                  ((bvurem Start Start):0 1)
                  ((bvlshr Start Start):0 1)
                  ((bvashr Start Start):0 1)
                  ((bvshl Start Start):0 1)
                  ((bvsdiv Start Start):0 1)
                  ((bvsrem Start Start):0 1)
                  ((bvsub Start Start):0 1)
                  x
                  #x0000001F
                  #x00000001
                  #x00000000
                  #xFFFFFFFF))))

(declare-var x (BitVec 32))
(constraint (= (hd17 x) (f x)))
(optimize (PARETO w1 w2))
(check-synth)
```

## B   Benchmarks

The following are the full description of benchmarks presented in Table 1.

max_ite(a,b) These benchmarks extend the SYGUS benchmark max2 by re-
     stricting the number of if statements between $a$ and $b$ and then require to

minimize the total size of the solution. These benchmarks operate over the semiring Trop $\times_{\mathbf{S}}$ Trop, but only impose one minimization objective.

parity_not minimizes number of not in corresponding SyGuS benchmark.

max3_ite minimizes number of ite in corresponding SyGuS benchmark.

Rest of Trop minimizes size of the solution in corresponding SyGuS benchmark.

hackers_a_prob Probabilistic extensions of corresponding SyGuS benchmarks (Prob semiring). The probability scheme we use assigns probability $\frac{1}{8}$ to shift operators, probability $\frac{1}{4}$ to arithmetic operators and $\frac{1}{2}$ to logical operators. The goal is to find the most probable solution.

## C   Proofs of Theorems

**Theorem 1** *Given an unambiguous WTG $G$ over a semiring $\mathbf{S}$ such that $M_{\mathbf{S},G} = (S', \otimes, 1)$, and a weight $s \in S$, the construction of $G_{\preceq s}$ terminates if the set $\{s' \mid s' \preceq s \,\wedge\, w \in S'\}$ is finite. Moreover, if the set of weights WEIGHTS$(G)$ is monotonically increasing with respect to $\preceq$—i.e. for every $s \in S$ and $s' \in$ WEIGHTS$(G)$, $s \preceq s \otimes s'$—then $L(G_{\preceq s})$ contains exactly every tree $t$ such that $w_G(t) \preceq s$.*

*Proof.* We first show that each step of the algorithm terminates. Steps 1 and 2 terminate since the grammar $G$ is finite. Step 3 only produces nonterminals that belongs to $N \times \{s' \mid s' \preceq s \wedge w \in S'\}$ which is also finite.

We now prove soundness. It is straightforward to prove the following claim by induction: for every nonterminal $(A, s) \in N'$, tree $t \in T_\Sigma$, we have that $d \in$ DER$((A, s'), t)$ iff there exists a derivation $d' \in$ DER$(A, t)$ such that DW$(d') = s'$ and $s' \preceq s$. Because $G$ is unambiguous, every tree has at most one derivation. Therefore DW$(d') = w_G(t)$ and $w_G(t) \preceq s$. □

**Theorem 2** *Given a WTG $G$ over Trop and a weight $s \in \mathbb{Z}$, the construction of $G_{\leq s}^{\text{Trop}}$ terminates if $G$ contains no loop with with cumulative negative weight. Moreover, $G_{\leq s}^{\text{Trop}}$ contains exactly every tree $t$ such that $w_G(t) \leq s$.*

*Proof.* First, we show that any tree with weight $\leq s$ must be accepted by $G_{\leq s}^{\text{Trop}}$. We do so by showing that if a tree $t$ is not accepted by $G_{\leq s}^{\text{Trop}}$—i.e., $t$ has some subtree $\beta$ with weight greater than $s + SH$—the weight of $t$ must be greater than $s$. Note that the modified algorithm can track weights $\leq s + |SH|$ in the intermediate nonterminals but still accept only trees with weight $\leq s$. According to the definition, the weight of $t$ is the sum of all rules used to derived $t$, that is, $w_G(t) = w_G(\beta) + w_G(t[B/\beta])$ where $t[B/\beta] \in T_{\Sigma \cup \{B\}}$ is the result of substituting the node corresponding to $\beta$ with $B$. Then if $t[B/\beta]$ contains loops, we can eliminate all loops from it to get a tree $t'[B/\beta]$ such that $w_G(t[B/\beta]) \geq w_G(t'[B/\beta])$ because loops have non-negative weights. If $t[B/\beta]$ contains no loop, its weight $w_G(t[B/\beta]) \geq -ck^{n+1} = -SH$ since the size of $t[B/\beta]$ is no more than $k^{n+1}$ (the height of $t[B/\beta]$ is no more than $n$ since there is no loop in it) and each production used to derive $t[B/\beta]$ has weight greater than $-c$. Therefore, using the fact that $w_G(t[B/\beta]) \geq -SH$, we have that the weight of $t$ is $w_G(\beta) + w_G(t[B/\beta]) > s + SH + w_G(t[B/\beta]) \geq s$.

Now, we show that the algorithm terminates. We observe that there is only a finite number of trees without loops so the set of their weights is also finite, namely the minimum weight of any tree without loops is $w^*$. On the other hand, for any tree $t$ containing loops, the weight $w$ of $t$ must be greater or equal to the weight of some tree without loops—i.e., $w > w^*$. This is because we can eliminate loops, whose weights are non-negative, from $t$ and the obtained tree has greater or equal weight to $t$. So the weights of nonterminal produced by our constructions are fall in the range $[w^*, s + SH]$ which is finite. Finally, the construction only needs to produce finite number of nonterminals and will always terminate. □

**Theorem 3** *Given an unambiguous WTG $G$ over the semiring $\boldsymbol{S} = \boldsymbol{S}_1 \times_{\boldsymbol{S}} \boldsymbol{S}_2$ with Pareto partial order $\preceq_p = \mathrm{PAR}(\preceq_1, \preceq_2)$ and a weight $s = (s_1, s_2) \in S$, if the constructions $G^{\boldsymbol{S}_1}_{\preceq_1 s_1}$ and $G^{\boldsymbol{S}_2}_{\preceq_2 s_2}$ terminate, then the construction of $G_{\preceq s}$ terminates.*

*Proof.* We first show by induction that if a nonterminal $(X, w_1, w_2)$ is produced in the construction of $G^{\mathbf{S}}_{\preceq s}$, the nonterminals $(X, w_1)$ and $(X, w_2)$ must be produced in the construction of $G^{\mathbf{S}_1}_{\preceq s_1}$ and $G^{\mathbf{S}_1}_{\preceq s_1}$ respectively. For the base case, we consider the nonterminals $(\bar{X}, w_1, w_2)$ produced in step 2 with production $p$. The condition $\mu(p) \preceq s$ in the step 2 implies that $\mu_1(p) \preceq s_1$ and $\mu_2(p) \preceq s_2$ which means that $(X, w_1)$ and $(X, w_2)$ are also produced in the construction of the corresponding grammar. Then, for every nonterminal $(X, w_1, w_2)$ produced in the step 3 with rule $p$ and $\{(X_i, w_i^{(1)}, w_i^{(2)})\}_i \subseteq N'$ where $(w_1, w_2) := \mu(p) \otimes \bigotimes_i (w_i^{(1)}, w_i^{(2)}) \preceq (s_1, s_2)$, according to the induction hypothesis, nonterminals in $\{(X_i, w_1^i)\}_i$ and $\{(X_i, w_2^i)\}_i$ are already produced in the grammars $G^{\mathbf{S}_1}_{\preceq s_1}$ and $G^{\mathbf{S}_2}_{\preceq s_2}$. Therefore, we can apply the step 3 with $p$ and nonterminals $G^{\mathbf{S}_1}_{\preceq s_1}$ (or $G^{\bar{\mathbf{S}}_2}_{\preceq s_2}$) to produce a new nonterminal $(X, w_1)$ (or $(X, w_2)$). Note that $w_1 = \mu(p) \otimes \bigotimes_i w_i^{(1)} \preceq s_1$, and $w_2 = \mu(p) \otimes \bigotimes_i w_i^{(2)} \preceq s_2$,.

Since both of the constructions of $G^{\mathbf{S}_1}_{\preceq_1 s_1}$ and $G^{\mathbf{S}_2}_{\preceq_2 s_2}$ terminate, the number of nonterminals they produce, namely $n_1$ and $n_2$, must be finite. We have shown that the number of nonterminals produced in $G^{\mathbf{S}}_{\preceq s}$ is less than $n_1 \times n_2$, which is also finite. At last, the construction of $G^{\mathbf{S}}_{\preceq s}$ terminates. □