Extending Parikh's Theorem to Weighted and Probabilistic Context-Free Grammars

Vijay Bhattiprolu[1][*], Spencer Gordon[2], and Mahesh Viswanathan[2][**]

[1] Carnegie Mellon University, Pittsburgh PA 15213, USA,
vpb@cs.cmu.edu
[2] University of Illinois at Urbana-Champaign, Urbana 61801, USA,
{slgordo2,vmahesh}@illinois.edu

Abstract. We prove an analog of Parikh's theorem for weighted context-free grammars over commutative, idempotent semirings, and exhibit a stochastic context-free grammar with behavior that cannot be realized by any stochastic right-linear context-free grammar. Finally, we show that every unary stochastic context-free grammar with polynomially-bounded ambiguity has an equivalent stochastic right-linear context-free grammar.

1 Introduction

Two words u, v over an alphabet Σ are said to be Parikh equivalent, if for each $a \in \Sigma$, the number of occurrences of a in u and v are the same. The Parikh image of a language L, is the set of Parikh equivalence classes of words in L. One of the most celebrated results in automata theory, Parikh's theorem [26], states that for any context-free language L, there is a regular language L' such that the Parikh images of L and L' are the same. For example, the context-free language $\{a^n b^n \mid n \geq 0\}$ has the same Parikh image as the regular language $(ab)^*$; both the Parikh images only consist of those equivalence classes where the numbers of as is equal to the number of bs. An important and immediate consequence of this result is that every context-free language over the unary alphabet is in fact regular. Parikh's theorem has found many applications — in automata theory to prove non-context-freeness of languages [12], decision problems for membership, universality and inclusions involving context-free languages and semi-linear sets [16,17,18,10]; in verification of subclasses and extensions of counter machines [19,14,36,10,7,32,13,34,22]; automata and logics over unranked trees with counting [2,33]; PAC-learning [22].

Weighted automata [31,9] are a generalization of classical automata (finite or otherwise) in which each transition has an associated weight from a semiring. Recall that a semiring is an algebra with two operations $\oplus$ and $\otimes$ such that $\oplus$ is a commutative monoid operation, $\otimes$ is a monoid operation, $\otimes$ distributes over $\oplus$, and the identity of $\oplus$ is an anhilator for $\otimes$. Unlike classical automata that compute Boolean-valued functions over words, weighted automata compute more general functions over words — the weight of an accepting computation is the product of

---

[*] This work was started while this author was at the University of Illinois, Urbana-Champaign.
[**] Partially suported by NSF CNS 1314485.

the weights of its transitions, and the weight of a word is the sum of the weights of all its accepting computations. Since the seminal paper of Schützerberger [31], weighted automata have inspired a wealth of extensions and further research (see [9] for a recent handbook compilation). Weighted automata have found applications in verification [8,4,5,23], reasoning about competitive ratio of online algorithms [1], digital image compression [6,15,20,21], in speech-to-text processing [24,25,3], and data flow analysis [30,29]. A special case of weighted automata are probabilistic automata [28,27] that model randomized algorithms and stochastic uncertainties in the system environment.

In this paper, we investigate whether Parikh's theorem can be generalized to the weighted case. In particular we investigate if for any weighted context-free grammar G there is a weighted right-linear grammar G′ such that for any Parikh equivalence class C, the sum of the weights of words in C under G and G′ is the same. It is easy to see that if the weight domain is not commutative (i.e., ⊗ is not a commutative operation) then Parikh's theorem does not hold. Thus we focus our attention on commutative weight domains.

Our first result concerns weight domains that are additionally idempotent, which means that ⊕ is an idempotent operation. A classical example of such a semiring is the min-plus or tropical semiring over the natural numbers where min is the "addition" operation, and + is the "product" operation. We show that Parikh's theorem does indeed hold for weighted automata over commutative, idempotent semirings.

Next, we show that our assumption about idempotence of semirings is necessary. In particular, we give an example of a stochastic context-free grammar G over the unary alphabet such that the function computed by G cannot be realized by any stochastic right linear grammar.

Our last result concerns unary grammars that are polynomially ambiguous. Recall that a grammar is polynomially ambiguous if there is a polynomial p such that on any word of length n in the language, the number of derivation trees for the word is bounded by $p(n)$. We prove that Parikh's theorem extends for such grammars. Specifically, we show that, over the unary alphabet, any probability function realized by a stochastic context-free grammar can also be realized by a right-linear grammar. Though we present this result in the context of stochastic grammars, the proof applies to any polynomially ambiguous weighted context-free grammar over a semiring that is commutative, but not necessarily idempotent.

The rest of the paper is organized as follows. We introduce the basic models and notation in Section 2. The Parikh's theorem for weighted automata over commutative, idempotent semirings is presented in Section 3. In Section 4, we present an example unary stochastic context-free grammar, and show that there is no stochastic right-linear grammar that is equivalent to it. Section 5 contains our proof for Parikh's theorem for polynomially ambiguous grammars. Eventhough this proof is presented in the context of stochastic grammars, it is easy to see that it extends to any weighted context free grammar over a commutative (but not necessarily idempotent) semiring. Finally, we present our conclusions and directions for future work in Section 6.

2   Preliminaries

Strings. Let us fix a finite string/word $w \in \Sigma^*$ over Σ. For a subset $\Gamma \subseteq \Sigma$, $w{\restriction}_\Gamma$ will denote the string over Γ obtained by removing the symbols not in Γ from w. The **Parikh map**, or Parikh image, of $w \in \Sigma^*$, denoted by $\mathrm{Pk}(w)$, is a mapping from Σ to $\mathbb{N}$, such that for $a \in \Sigma$, $\mathrm{Pk}(w)(a)$ is the number of occurrences of a in w. The **Parikh equivalence class** of w, $[w]_{\mathrm{Pk}} = \{w' \mid \mathrm{Pk}(w') = \mathrm{Pk}(w)\}$, is the set of all words with the same Parikh image as w. We can extend the Parikh map to languages $L \subseteq \Sigma^*$, defining $\mathrm{Pk}(L) \triangleq \{\mathrm{Pk}(w) \mid w \in L\}$.

Context Free Grammars. We will consider context free grammars in **Greibach Normal Form**. Formally (in this paper) a **context-free grammar** is $G = (V, \Sigma, P, S)$, where V and Σ are disjoint sets of variables (or non-terminals) and terminals, respectively; $S \in V$ is the start symbol; and $P \subseteq V \times \Sigma V^*$ is a finite set of productions where each production is of the form $A \rightarrow \mathtt{a}\beta$ with $\mathtt{a} \in \Sigma$ and $\beta \in V^*$. Without loss of generality, we assume that every production in the grammar is used in some derivation from S to a string in Σ^*. A **sentence** is a string in $(\Sigma \cup V)^*$. A **right-linear grammar** is a context-free grammar where the productions have at most one non-terminal on the right-hand side, i.e., $P \subseteq V \times (\Sigma(\{\epsilon\} \cup V))$. It is well known that a language is generated by a right-linear grammar if and only if it is regular.

We will find it convenient to partition the variables of a grammar into those that have exactly one derivation tree and those that have more than one. Formally, the set of **single-derivation variables** $X \subseteq V$ is the smallest set containing all variables A with exactly one production of the form $A \rightarrow \mathtt{a}$ (with $\mathtt{a} \in \Sigma$) and having the property that if a variable A has exactly one production of the form $A \rightarrow \mathtt{a}\alpha$ where $\mathtt{a} \in \Sigma$ and $\alpha \in X^*$ then $A \in X$. The remaining variables, i.e. $Y = V \setminus X$, are **multiple-derivation variables**.

Prioritized leftmost derivations. In this paper we will consider special derivation sequences of a context-free grammar that expand the leftmost variable while giving priority to single-derivation variables. We call these **prioritized leftmost (PLM) derivations**, and we define them precisely next.

Definition 1. *Consider a context-free grammar* $G = (V, \Sigma, P, S)$*, where the non-terminals* V *have been partitioned into the set of single-derivation variables* X *and multiple-derivation variables* Y*. We say that* $\alpha A \beta$ *rewrites in a* single prioritized leftmost derivation step *to* $\alpha\gamma\beta$ *(denoted as* $\alpha A \beta \Rightarrow_{\mathrm{plm}} \alpha\gamma\beta$*) iff* $\exists \pi \in P, \pi = (A \rightarrow \gamma)$ *such that either*

1. *$A \in X$, $\alpha \in (\Sigma \cup Y)^*$, and $\beta \in V^*$, or*
2. *$A \in Y$, $\alpha \in \Sigma^*$, and $\beta \in (\Sigma \cup Y)^*$*

In other words, either A is the leftmost single-derivation variable in $\alpha A \beta$, or A is the leftmost multiple-derivation variable and $\alpha A \beta$ has no single-derivation variables. If $\alpha \Rightarrow_{\mathrm{plm}} \beta$ by application of π, we'll write $\alpha \overset{\pi}{\Rightarrow}_{\mathrm{plm}} \beta$. Note that if $\alpha \Rightarrow_{\mathrm{plm}} \beta$ there is always a unique π such that $\alpha \overset{\pi}{\Rightarrow}_{\mathrm{plm}} \beta$.

A prioritized leftmost (PLM) derivation *is a sequence $\psi = \alpha_1, \ldots, \alpha_n$ such that $\alpha_1 \Rightarrow_{\mathrm{plm}} \alpha_2 \Rightarrow_{\mathrm{plm}} \cdots \Rightarrow_{\mathrm{plm}} \alpha_n$. The set of all PLM derivations is denoted $\mathrm{Der}_{\mathrm{plm}}(G)$.*

The **language generated by** G is $L(G) \triangleq \left\{ \alpha \in \Sigma^* \mid S \Rightarrow^*_{\text{plm}} \alpha \right\}$ where $\Rightarrow^*_{\text{plm}}$ is the reflexive and transitive closure of $\Rightarrow_{\text{plm}}$. Finally, the **parse** of a word $w \in (\Sigma \cup V)^*$, denoted $\text{parse}_G(w)$, is the set of all PLM derivations yielding w:

$$\text{parse}_G(w) \triangleq \left\{ \alpha_1, \ldots, \alpha_n \in \text{Der}_{\text{plm}}(G) \mid \alpha_1 = S \text{ and } \alpha_n = w \right\}.$$

Example 1. We present a simple example to illustrate the definitions. Consider the grammar $G = (\{S, B\}, \{a, b\}, P, S)$ where P consists of the following productions: $\pi_1 = S \to aSB$, $\pi_2 = S \to aB$, and $\pi_3 = B \to b$. The set of single-derivation variables is $\{B\}$ and the set of multiple-derivation variables is $\{S\}$. An example of a prioritized leftmost derivation is

$$S \xoverset{\pi_1}{\Rightarrow}_{\text{plm}} aSB \xoverset{\pi_3}{\Rightarrow}_{\text{plm}} aSb \xoverset{\pi_2}{\Rightarrow}_{\text{plm}} aaBb \xoverset{\pi_3}{\Rightarrow}_{\text{plm}} aabb$$

The language generated by this grammar is $\{a^n b^n \mid n \geq 1\}$.

Derivation trees. The set of all derivation trees for G will be denoted as Δ_G. For a derivation tree τ, a node n in τ, and a path p from the root in τ, $\ell(\tau)$, $\ell(n)$ and $\ell(p)$ will denote the label of the root, the node n, and the node reached by path p in τ, respectively. For any node n in a tree τ and path p from the root, we denote the subtree rooted at n by $\tau(n)$, and the subtree rooted at the node reached by path p by $\tau(p)$. The **frontier of a tree** τ, denoted $\text{Fr}(\tau)$ is the sentence $\ell(n_1)\ell(n_2)\ldots\ell(n_k)$ where $n_1, \ldots, n_k$ are the leaves of τ in left-to-right order.

For any variable $A \in V$, $\Delta_G(A) \triangleq \{\tau \in \Delta_G \mid \ell(\tau) = A\}$ is the subset of derivation trees rooted at A. A tree τ for which $\text{Fr}(\tau) \in \Sigma^*$ is called a **complete derivation tree**, and the set of all complete derivation trees rooted at A is $\Delta_G^\Sigma(A) \triangleq \{\tau \in \Delta_G(A) \mid \text{Fr}(\tau) \in \Sigma^*\}$. The set of all complete derivation trees is $\Delta_G^\Sigma \triangleq \{\tau \in \Delta_G \mid \text{Fr}(\tau) \in \Sigma^*\}$. A tree $\tau \in \Delta_G(A)$ is said to be an *A-pumping tree* if $\text{Fr}(\tau)\!\restriction_V = A$. The set of *A*-pumping trees is $\Delta_G^p(A) \triangleq \{\tau \in \Delta_G(A) \mid \text{Fr}(\tau)\!\restriction_V = A\}$. The set of all pumping trees is given by $\Delta_G^p = \{\tau \in \Delta_G \mid \text{Fr}(\tau)\!\restriction_V \in V\}$.

Remark 1. In a context-free grammar (where all productions are "useful"), every single-derivation variable is the root of exactly one complete derivation tree, and every multiple-derivation variable is the root of at least two complete derivation trees.

Tree Notation. We will use the following notation on derivation trees. Let $\tau \in \Delta_G$, n be a node in τ, and p be a path from the root in τ. The leftmost child of the node reached by path p, will be the one reached by the path $p \cdot 0$ with the other children corresponding to the paths $p \cdot 1$, $p \cdot 2$, etc. For $\tau_1 \in \Delta_G(\ell(n))$ ($\tau_1 \in \Delta_G(\ell(p))$), $\tau[n \mapsto \tau_1]$ ($\tau[p \mapsto \tau_1]$) denotes the derivation tree obtained by replacing $\tau(n)$ ($\tau(p)$) by the tree τ_1. We denote by $\text{rem}_p(\tau)$ the tree obtained by replacing $\tau(p)$ by the root of $\tau(p)$, i.e., by "removing" all the children of p. Finally, for a rule $A \to a\alpha$ with $\alpha = A_1 A_2 \cdots A_k$, and trees $\tau_i \in \Delta_G(A_i)$, $A_{a\alpha}(\tau_1, \tau_2, \ldots \tau_k)$ denotes the tree with root labeled A and children $a, \tau_1, \ldots \tau_k$ from left-to-right. Thus, A_a denotes the tree with root labeled A and one child labeled a.

Cuts. Observe that, for any string $\alpha \in (V \cup \Sigma)^*$, there is a bijection between derivation trees τ with $\mathrm{Fr}(\tau) = \alpha$ and PLM derivations in $\mathrm{parse}_G(\alpha)$. A set of nodes separating the root of τ from all of the leaves in τ is a **cut** of τ. Now consider the unique PLM derivation Ψ corresponding to τ. Every sentence in Ψ corresponds to a cut $\mathcal{C}$ in τ. We call any such $\mathcal{C}$ a **prioritized leftmost (PLM) cut** of τ. For a set of trees T and a variable $A \in V$, the **Parikh supremum** of variable A in T, denoted by $\sup_{\mathrm{Pk}}(A, T)$, is the maximum number of occurrences of A in any PLM cut of any tree $\tau \in T$. Observe that any PLM derivation sequence corresponding to a tree τ in T can have at most $\sup_{\mathrm{Pk}}(A, T)$ occurrences of the variable A in any sentence.

Ambiguity. We will say that a set of trees Γ is **ambiguous** if there are two distinct trees τ_1, τ_2 such that $\mathrm{Fr}(\tau_1) = \mathrm{Fr}(\tau_2)$; if Γ is not ambiguous, we say it is **unambiguous**. The **ambiguity function** $\mu_G : \mathbb{N} \to \mathbb{N}$ for a grammar G is a function mapping every natural number n to the maximal number of PLM derivations which a word of length n may have. Formally, $\mu_G(n) = \max_{w \in L(G), |w|=n} |\mathrm{parse}_G(w)|$. A grammar is said to have **exponential ambiguity** if its ambiguity function is in $2^{\Theta(n)}$, and it is said to have **polynomially-bounded ambiguity**, or to be polynomially ambiguous, if its ambiguity function is in $O(n^d)$ for some $d \in \mathbb{N}_0$. Any grammar G has either exponential ambiguity or polynomially-bounded ambiguity [35]. The following characterization of polynomial ambiguity was proved in [35].

Theorem 1 ([35]). *A context-free grammar G has polynomially-bounded ambiguity if and only if Δ_G^p is unambiguous.*

We conclude the preliminaries by recalling a classical result due to Parikh [26].

Theorem 2 (Parikh's Theorem [26]). *For every context-free grammar G, there is a right-linear context-free grammar G′ such that $\mathrm{Pk}(L(G)) = \mathrm{Pk}(L(G'))$.*

2.1 Weighted and Stochastic Context-Free Grammars

Weighted context-free grammars define a function that associates a value in a semiring with each string. Stochastic context-free grammars are special weighted context-free grammars that associate probabilities with strings. We recall these classical definitions in this section. We begin by defining a semiring.

Semiring. A **semiring** is a structure $\mathbb{D} = (D, \oplus, \otimes, 0_D, 1_D)$ where $(D, \oplus, 0_D)$ is a *commutative monoid* with identity 0_D, $(D \setminus \{0_D\}, \otimes, 1_D)$ is a *monoid* with identity 1_D, $\otimes$ distributes over $\oplus$ (i.e., $(a \oplus b) \otimes c = (a \otimes c) \oplus (b \otimes c)$ and $a \otimes (b \oplus c) = (a \otimes b) \oplus (a \otimes c)$, for every $a, b, c \in D$), and 0_D is an *annhilator* for $\otimes$ (i.e., $a \otimes 0_D = 0_D \otimes a = 0$ for every $a \in D$). We abuse notation and use $\mathbb{D}$ to denote the semiring and the underlying set where the meaning is clear from context. We define $\mathbb{D}_0 = D \setminus \{0_D\}$. When considering an abstract semiring $\mathbb{D}$, we'll write $0_{\mathbb{D}}$ and $1_{\mathbb{D}}$ for 0_D and 1_D respectively. An **idempotent semiring** satisfies the additional requirement that for all $a \in \mathbb{D}$, $a \oplus a = a$. A **commutative semiring** is one where $\otimes$ is commutative, i.e., $(D \setminus \{0_D\}, \otimes, 1_D)$ is a commutative monoid as well.

Example 2. Classical examples of a semiring are the tropical semiring and the probability semiring. The tropical or min-plus semiring is $(\mathbb{N} \cup \{\infty\}, \min, +, \infty, 0)$, where ∞ is taken to be larger than everything in $\mathbb{N}$. It is commutative and idempotent as $\min(a, a) = a$ for any a. The probability semiring is $([0, 1], +, \times, 0, 1)$, where $[0, 1]$ is the set of reals between 0 and 1. It is commutative as $\times$ is commutative. However, since the addition of two numbers is not idempotent, the probability semiring is not idempotent.

Weighted context-free grammars. A **weighted context-free grammar** is a pair (G, W) where $G = (V, \Sigma, P, S)$ is a context-free grammar, and $W : P \to \mathbb{D}$ assigns a weight from $\mathbb{D}$ to each production in P, for some semiring $\mathbb{D}$. (Note that W may assign $0_{\mathbb{D}}$ to some productions in P.) The **weight** of a PLM derivation $\psi = \alpha_1 \overset{\pi_1}{\Rightarrow}_{\text{plm}} \alpha_2 \overset{\pi_2}{\Rightarrow}_{\text{plm}} \cdots \overset{\pi_{n-1}}{\Longrightarrow}_{\text{plm}} \alpha_n$ of G, is given by $W(\psi) \triangleq \otimes_{i=1}^{n-1} W(\pi_i)$. For $w \in \Sigma^*$, $W(w) \triangleq \oplus_{\psi \in \text{parse}_G(w)} W(\psi)$; we assume that if $\text{parse}_G(w) = \emptyset$ (i.e., $w \notin L(G)$) then $W(w) = 0_{\mathbb{D}}$. The **semantics** of a weighted grammar (G, W), denoted $[\![G]\!]_W : \Sigma^* \to \mathbb{D}$, is the function mapping each word to its weight in G, i.e., $[\![G]\!]_W(w) \triangleq W(w)$.

Example 3. Let G be the grammar described in Example 1. Consider a weight function W that assigns weights from the tropical semiring, with the weight of every production $\pi \in P$ being equal to 1. Then the semantics of (G, W) is given as $[\![G]\!]_W(a^n b^n) = 2n$ and $[\![G]\!]_W(w) = 0$, when $w \notin L(G)$.

Definition 2. *The* Parikh image *of a weighted context-free grammar* (G, W)*, written as* $\text{Pk}[\![G]\!]_W$ *is function defined as*

$$\text{Pk}[\![G]\!]_W(w) \triangleq \bigoplus_{w' \in [w]_{\text{Pk}}} [\![G]\!]_W(w')$$

Stochastic Context-free Grammars. A **stochastic context-free grammar** is a weighted context-free grammar $(G = (V, \Sigma, P, S), W)$ where the weight domain is the probability semiring $([0, 1], +, \times, 0, 1)$, and for any $A \in V$ and $a \in \Sigma$, we have

$$\sum_{\alpha \in V^*:(A \to a\alpha) \in P} W(A \to a\alpha) \in [0, 1].$$

3 A Parikh's Theorem for Weighted CFGs

The main result of this section is that for any weighted context-free grammar over an idempotent, commutative semiring (like the tropical semiring), there is a Parikh equivalent weighted right-linear context-free grammar. Thus, this observation extends the classical result to weighted CFGs over idempotent semirings.

Theorem 3 (Weighted Parikh's Theorem). *For every weighted context-free grammar* (G, W) *over an* idempotent, commutative *semiring, there exists a Parikh-equivalent weighted right-linear grammar* (G_*, W_*)*, that is, we have*

$$\text{Pk}[\![G]\!]_W = \text{Pk}[\![G_*]\!]_{W_*}.$$

Proof. The full proof for this result is presented in Section A in the Appendix. Here we present the broad ideas.

Let $G = (V, \Sigma, P, S)$ be a context-free grammar and let $W : P \to \mathbb{D}$ be a weight function over a commutative, idempotent weight domain $\mathbb{D}$. Consider the following homomorphism $h : P^* \to \Sigma^*$ defined as $h(\pi) = \mathsf{a}$, where $\pi = A \to \mathsf{a}\alpha \in P$.

We begin by first constructing a weighted context-free grammar (G_1, W_1) over the alphabet P, whose image under h gives us G. Formally, $G_1 = (V, P, P_1, S)$ has as productions $P_1 = \{A \to \pi\alpha \mid \exists \mathsf{a} \in \Sigma. \ \pi = A \to \mathsf{a}\alpha \in P\}$. In addition, take W_1 to be $W_1(A \to \pi\alpha) = W(\pi)$. It is easy to see that $h(L(G_1)) = L(G)$ by construction. Moreover, given W_1 and W, we can conclude $[\![G]\!]_W(w) = \bigoplus_{\omega \in P^* : h(\omega) = w} [\![G_1]\!]_{W_1}(\omega)$.

By Parikh's theorem (Theorem 2), there is a right-linear grammar $G_2 = (V_2, P, P_2, S_2)$ such that $\mathrm{Pk}(L(G_2)) = \mathrm{Pk}(L(G_1))$. Define the weight function W_2 as $W_2(A \to \pi B) = W(\pi)$ to give us the weighted CFG (G_2, W_2). Using the fact that $\otimes$ is commutative, and $\oplus$ is idempotent, we can prove that $\mathrm{Pk}[\![G_1]\!]_{W_1} = \mathrm{Pk}[\![G_2]\!]_{W_2}$; the full proof is in Section A.

Finally, we consider the context free grammar G_3 obtained by "applying the homomorphism h" to G_2. Formally, $G_3 = (V_2, \Sigma, P_3, S_2)$, where $P_3 = \{A \to h(\pi)B \mid A \to \pi B \in P_2\}$. The weight function W_3 is defined in such a way that weight of $A \to \mathsf{a}B$ is the sum of the weights of all productions $A \to \pi B$, where $h(\pi) = \mathsf{a}$, i.e.,

$$W_3(A \to \mathsf{a}B) = \bigoplus_{\pi : h(\pi) = \mathsf{a}} W_2(A \to \pi B).$$

(G_3, W_3) and (G_2, W_2) share the same relationship as (G, W) and (G_1, W_1). That is, we have $h(L(G_2)) = L(G_3)$ and $[\![G_3]\!]_{W_3}(w) = \bigoplus_{\omega \in P^* : h(\omega) = w} [\![G_2]\!]_{W_2}(\omega)$.

The weighted CFG (G_3, W_3) is the desired weighted grammar, i.e., $\mathrm{Pk}[\![G_3]\!]_{W_3} = \mathrm{Pk}[\![G]\!]_W$; the full proof is in Section A. $\qquad\square$

Corollary 1. *If* (G, W) *is a weighted context-free grammar over an* idempotent, commutative *weight domain and a unary alphabet, then there exists a weighted right-linear context-free grammar* (G', W') *such that* $[\![G']\!]_{W'} = [\![G]\!]_W$.

Example 4. Starting with the weighted grammar (G, W) from Example 3, the construction used in the proof of Theorem 3 would have P_1 containing the following productions: $S \to \pi_1 SB$, $S \to \pi_2 B$, $B \to \pi_3$. The language of this grammar is $L(G_1) = \{\pi_1^n \pi_2 \pi_3^{n+1} \mid n \geq 0\}$.

One candidate for G_2 would have as productions $S \to \pi_1 J$, $S \to \pi_2 K$, $J \to \pi_3 S$, and $K \to \pi_3$. The language is $L(G_2) = \{(\pi_1 \pi_3)^n \pi_2 \pi_3 \mid n \geq 0\}$.

In that case (G_3, W_3) would have productions and weights as follows:

$$W_3(S \to \mathsf{a}J) = 1 \qquad\qquad W_3(S \to \mathsf{a}K) = 1$$
$$W_3(J \to \mathsf{b}S) = 1 \qquad\qquad W_3(K \to \mathsf{b}) = 1$$

The language of the underlying grammar would be $L(G_3) = \{(\mathsf{ab})^n \mathsf{ab} \mid n \geq 0\}$, and

$$[\![G_3]\!]_{W_3}(w) = \begin{cases} 2k & \text{if } w = (\mathsf{ab})^k \text{ for some } k \geq 1 \\ 0 & \text{otherwise} \end{cases}$$

4 A Counterexample to Parikh's Theorem for Stochastic Grammars

Theorem 3 crucially relies on the semiring being idempotent. In this section, we show that Theorem 3 fails to generalize if we drop the requirement of idempotence. We give an example of a stochastic context-free grammar over the unary alphabet that is not equivalent to any stochastic right-linear grammar. Before presenting the example stochastic context-free grammar and proving the inexpressivity result, we recall some classical observations about unary stochastic right linear grammars.

4.1 Properties of Unary Stochastic Right-linear Grammars

Stochastic right-linear grammars satisfy pumping lemma type properties. Here we recall such an observation for unary stochastic right-linear grammars.

Theorem 4 (Pumping Lemma). *Let* $(G = (V, \{a\}, P, S), W)$ *be a stochastic right-linear grammar over the unary alphabet. There is a number k, and real numbers $c_0, c_1, \ldots c_k$ with $c_0 + c_1 + \cdots + c_k = 1$ such that for every $\ell \in \mathbb{N}$*

$$\llbracket G \rrbracket_W (a^{\ell+k+1}) = \sum_{i=0}^{k} c_i \llbracket G \rrbracket_W (a^{\ell+i})$$

Proof. The result is a consequence of the Cayley-Hamilton theorem and the fact that 1 is an eigen value of stochastic matrices. We skip the proof as it is a specialization of Theorem 2.8 in Chapter II.C in [27]. □

Let (G, W) be a unary weighted context-free grammar. The generating function of such a grammar is $P(x) = \sum_{k=0}^{\infty} \llbracket G \rrbracket_W (a^k) x^k$. We conclude this section by observing that if G is right-linear, then its generating function must be a rational function, i.e., $P(x)$ is an algebraic fraction where both the numerator and denominator are polynomials.

Theorem 5. *Let* (G, W) *be a stochastic right-linear grammar over the unary alphabet. Then the generating function $P(x) = \sum_{k=0}^{\infty} \llbracket G \rrbracket_W (a^k) x^k$ is a rational function.*

Proof. Observe that Theorem 4 says that the sequence $\langle \llbracket G \rrbracket_W (a^n) \rangle_{n \in \mathbb{N}}$ satisfies a linear homogeneous recurrence with constant coefficients. Thus, its generating function must be rational. □

4.2 The Counterexample

We now present a unary weighted CFG and show that there is no weighted right-linear CFG that is equivalent to it. Consider the grammar $G_* = (\{S\}, \{a\}, \{(S \to a), (S \to aSS)\}, S)$. Let p be some number in $(0, 1)$. The weight function W_* is defined as follows: $W_*(S \to a) = 1-p$, and $W_*(S \to aSS) = p$. Taking c_n to be the nth Catalan number, we can see that $\llbracket G_* \rrbracket_{W_*} (a^{2k+1}) = c_k p^k (1-p)^{k+1}$;

this is because the probability of any PLM derivation for a^{2k+1} is $p^k(1-p)^{k+1}$ and there are c_k elements in $\text{parse}_{G_*}(a^{2k+1})$. Taking $b_k = [\![G_*]\!]_{W_*}(a^k)$, we have

$$b_k = \begin{cases} c_{(k-1)/2}p^{(k-1)/2}(1-p)^{(k-1)/2+1} & \text{if } k \text{ is odd} \\ 0 & \text{otherwise} \end{cases}$$

Recall that the generating function for the Catalan numbers, $C(z) = \sum_{k\geq 0} c_k z^k$, is given by $C(z) = \frac{1-\sqrt{1-4z}}{2z}$. Based on the above observations, the generating function for the grammar (G_*, W_*), $P(z) = \sum_{k\geq 0} b_k z^k$ can be written as follows.

$$\begin{aligned} P(z) &= \sum_{k\geq 0} b_k z^k = \sum_{k\geq 0} b_{2k+1} z^{2k+1} \\ &= z\sum_{k\geq 0} b_{2k+1}\left(z^2\right)^k = z\sum_{k\geq 0} c_k p^k (1-p)^{k+1}\left(z^2\right)^k \\ &= z(1-p)\sum_{k\geq 0} c_k \left(z^2 p(1-p)\right)^k \\ &= z(1-p)C(z^2 p(1-p)) \\ &= z(1-p)\frac{1-\sqrt{1-4z^2 p(1-p)}}{2z^2 p(1-p)} \\ &= \frac{1}{2zp}\left(1-\sqrt{1-4z^2 p(1-p)}\right) \end{aligned}$$

Having identified an expression for the generating function for (G_*, W_*), we are ready to prove that there is no Parikh equivalent right-linear grammar for (G_*, W_*). First notice that if a weighted grammar (G, W) is over the unary alphabet, then $[\![G]\!]_W = \text{Pk}[\![G]\!]_W$. Therefore, to establish the result of this section, it suffices to prove the statement that there is no right-linear grammar that is equivalent to (G_*, W_*); this is the content of the next theorem.

Theorem 6. *There is no stochastic right-linear grammar* (G, W) *such that* $[\![G]\!]_W = [\![G_*]\!]_{W_*}$.

Proof. Given Theorem 5, it suffices to prove that the generating function $P(z)$ for (G_*, W_*) is not rational. Taking $Q(z) = \sqrt{1-4z^2 p(1-p)}$, we see that $P(z) = \frac{1}{2zp}(1-Q(z))$. Given this relationship, we can conclude that if $P(z)$ is rational function, then so is $Q(z)$. Thus, our goal will be to prove that $Q(z)$ is not a rational function.

Assume for contradiction that $Q(z)$ is rational. Then $Q(z) = f(z)/g(z)$ where f and g are both polynomials with greatest common divisor having degree 0. Then $Q^2(z) = f^2(z)/g^2(z)$ and $Q^2(z)g^2(z) = \left(1-4z^2 p(1-p)\right)g^2(z) = f^2(z)$. Thus, $Q^2(z)$ must divide $f^2(z)$. We observe that $Q^2(z) = (1-2z\sqrt{p(1-p)})(1+2z\sqrt{p(1-p)})$ is square-free so $Q^2(z)$ must divide $f(z)$, and $f(z) = Q^2(z)h(z)$ for some polynomial h. Substituting for $f(z)$ in $Q^2(z)g^2(z) = f^2(z)$ and rearranging we obtain $Q^2(z)g^2(z) = (Q^2(z))^2 h^2(z) \implies g^2(z) = Q^2(z)h^2(z)$, and by the same argument as above, $Q^2(z)$

divides $g(z)$. Thus, $Q^2(z)$ divides both $f(z)$ and $g(z)$. Since $Q^2(z)$ is not a degree 0 polynomial, we contradict the assumption that the greatest common divisor of f and g has degree 0. □

5 Parikh's Theorem for Unary Polynomially Ambiguous Stochastic Grammars

The weighted stochastic context-free grammar (G_*, W_*) in Section 4.2 is exponentially ambiguous; the ambiguity function μ_{G_*} is bounded by the Catalan numbers. Exponential ambiguity turns out to be critical to construct such counterexamples. In this section, we prove that any unary stochastic context-free grammar with polynomial ambiguity is equivalent to a unary stochastic right-linear grammar. The proof of this result relies on an observation that in any PLM cut in a complete derivation tree of a unary polynomially ambiguous grammar, the number of occurences of any variable is bounded by a constant dependent on the grammar. The unary alphabet assumption is crucial in obtaining such a bound (Lemma 3). In the next two subsections we present a proof of this observation by first bounding the number of occurrences in cuts of pumping trees and then using it to bound it in complete derivation trees. In Section 5.3, we then present the construction of the right-linear grammar. Though we present this result in the context of stochastic grammars, it applies to any weighted CFG over a commutative (but not necessarily idempotent) semiring.

In the rest of this section, let us fix a unary, polynomially ambiguous, context-free grammar $G = (V, \{a\}, P, S)$ and a stochastic weight function Pr (for probability). We assume that the set of variables V is partitioned into single-derivation variables X and multiple-derivation variables Y. As we have done throughout this paper, we assume that every production in G is "useful", that is, is used in some complete derivation tree whose root is labeled S. Finally we will assume that m is the maximum length of the right-hand side of any production in P.

5.1 Parikh Suprema in Pumping trees

In this section we will bound the number of times a variable can appear in any PLM cut of a pumping tree in G. We begin by observing some simple properties about single-derivation variables X and multiple-derivation variables Y. Since every production in the grammar is useful, we can conclude that there is a unique complete derivation tree with root A if $A \in X$, and that there are at least two complete derivation trees with root A if $A \in Y$, i.e., $|\Delta_G^\Sigma(A)| = 1$ if $A \in X$, and $|\Delta_G^\Sigma(A)| > 1$ if $A \in Y$. Next, for $A \in X$, the unique $\tau \in \Delta_G^\Sigma(A)$ has the following properties: (a) no node is labeled by a variable in Y; (b) each variable in X labels at most one node along any path in τ. Property (b) holds because if $A \in X$ has a derivation $A \Rightarrow_{plm}^* \alpha A \beta$, then A cannot have any complete derivation tree and it would be useless. This also means that any pumping tree $\tau \in \Delta_G^p$ must have $\ell(\tau) \in Y$. These properties allow us to bound the size of the unique complete derivation for variable $A \in X$.

Lemma 1. *For any $A \in X$, the unique tree $\tau \in \Delta_G^\Sigma(A)$ has size at most $m^{|X|}$.*

Proof. Since only variables in X can appear as labels in τ and no variable appears more than once in any path, the height of τ is $\leq |X|$. Finally, since any node has at most m children, we get the bound on the size of τ. $\qquad\square$

Next we prove that Lemma 1 allows one to bound the number of times any single-derivation variable appears in any PLM cut of a pumping tree.

Lemma 2. *For any $A \in Y$ and $B \in X$, $\sup_{\text{Pk}}(B, \Delta_G^p(A)) \leq m^{|X|+1}$.*

Proof. Let $\tau \in \Delta_G^p(A)$ be an arbitrary A-pumping tree, where $A \in Y$. Let $\mathcal{C}$ be an arbitrary PLM cut of τ. We will prove a slightly stronger statement; we will show that the total number of single-derivation variables in $\mathcal{C}$ is $\leq m^{|X|+1}$. This will bound the Parikh supremum for any single-derivation variable.

Without loss of generality, assume that $\mathcal{C}$ has at least one node with label in X. Amongst all nodes in $\mathcal{C}$ that are labeled by a variable in X, let n be the node that is closest to the root, and if there are multiple such nodes, take n to be the leftmost one. From the definition of PLM cuts, the following property holds for $\mathcal{C}$ and n: (a) any node to right of n in $\mathcal{C}$ that is labeled by a variable in X must be a sibling of n, and (b) all nodes to the left of n in $\mathcal{C}$ labeled by variables of X must be descendents of some left sibling (say n_1) of n that is also labeled by a variable in X. Thus, the number of nodes to the right of n (including n) in $\mathcal{C}$ labeled by X is at most m, and, by Lemma 1, the number of nodes to the left of n in $\mathcal{C}$ labeled by X is at most $m^{|X|}$. Putting these together, the total number of nodes in $\mathcal{C}$ labeled by some variable in X is at most $m + m^{|X|} \leq m^{|X|+1}$. $\qquad\square$

Lemma 3. *For any $A \in V$, and $B \in Y$, $\sup_{\text{Pk}}(B, \Delta_G^p(A)) \leq 2$.*

Proof. Let τ be a A-pumping tree, for some variable A. Note that A must be a multiple-derivation variable because of property (b) before Lemma 1. Let $\mathcal{C}$ be any PLM cut of τ. Since τ is an A-pumping tree it must contain A in its frontier. Then there must be some node n in $\mathcal{C}$ such that the subtree $\tau(n)$ contains A in its frontier. Let $C = \ell(n)$. Observe that C is a multiple-derivation variable because a node labeled $A \in Y$ is a descendent. Thus, there are two complete derivation trees τ_1^C, τ_2^C with roots labeled C (Remark 1).

We'll first show that there cannot be more than two occurrences of nodes labeled C in $\mathcal{C}$. Assume towards the contrary that there are at least three nodes n_1, n_2, n_3 in $\mathcal{C}$ with $\ell(n_1) = \ell(n_2) = \ell(n_3) = C$. Without loss of generality, assume n_1, n_2, and n_3 are in left-to-right order in τ and $n \in \{n_1, n_2, n_3\}$. Since n_1, n_2, n_3 belong to a cut, they are not related by the ancestor/descendent relationship.

Let τ_1 be the tree $\tau[n_1 \mapsto \tau_1^C, n_2 \mapsto \tau_2^C, n_3 \mapsto \tau(n)]$, and let τ_2 be the tree $\tau[n_1 \mapsto \tau_2^C, n_2 \mapsto \tau_1^C, n_3 \mapsto \tau(n)]$. By construction, τ_1 and τ_2 are both A-pumping trees with $\text{Fr}(\tau_1) = \text{Fr}(\tau_2)$ and $\tau_1 \neq \tau_2$. However, since G is polynomially ambiguous, by Theorem 1, the set of pumping trees is unambiguous, giving us the desired contradiction.

Next, we show that there cannot be more than two nodes labeled $B \in Y$ in $\mathcal{C}$, where $B \neq C$. Assume that there are at least three nodes n_1, n_2, n_3 in $\mathcal{C}$ with $\ell(n_1) = \ell(n_2) = \ell(n_3) = B$. Again assume n_1, n_2, and n_3 are in left-to-right order in

τ. Further, since $B \in Y$, there are two complete derivation trees τ_1^B and τ_2^B with root labeled B (Remark 1).

Observe that at least two nodes of $\{n_1, n_2, n_3\}$ must lie to one side of n in τ. Without loss of generality we may assume that n_1 and n_2 are those nodes. Let τ_1 be the tree $\tau[n_1 \mapsto \tau_1^B, n_2 \mapsto \tau_2^B]$, and let τ_2 be the tree $\tau[n_1 \mapsto \tau_2^B, n_2 \mapsto \tau_1^B]$. Clearly, τ_1, τ_2 are A-pumping trees with $Fr(\tau_1) = Fr(\tau_2)$, and $\tau_1 \neq \tau_2$. □

5.2 Parikh Suprema in Complete Derivation Trees

We will now use the results in Section 5.1 to bound the Parikh supremum of any variable in a complete derivation tree of G. The key property we will exploit is the fact that any complete derivation tree can be written as the "composition" of a small number of pumping trees (see Figure 1) such that any PLM cut is the union of cuts in each of these pumping trees. The bounds on Parikh suprema will then follow from the observations in Section 5.1.

We begin with some convenient notation. For a $\tau \in \Delta_G$, let longestpath(τ) denote the longest path from the root of τ to a node labeled $\ell(\tau)$. If there are multiple such paths, longestpath(τ) is the lexicographically-first path among them. Note that longestpath(τ) can be ε if the root is the only node with label $\ell(\tau)$ in τ. Let depth(τ) denote the length of the longest path from root to leaf in τ.

We now describe two procedures `compress` and `decompress`. Let us fix a complete derivation tree τ. The procedure `compress` returns a data structure of pumping trees. These pumping trees are small in number and τ is the "composition" of these pumping trees. Let n be the lowest node in τ that has the same label as the root. `compress` identifies the pumping tree obtained by removing the children of n, and recursively compresses the subtrees rooted at the children of n. Note that if n is the same as the root, then the pumping tree identified by `compress` will just be the tree with one node.

`compress`(τ):

   If $\tau = A_a$ for some $A \in V$, **return** τ
   $p \leftarrow$ longestpath(τ)
   Let A, a and α be such that $\tau(p) = A_{a\alpha}(\tau(p \cdot 1), \ldots, \tau(p \cdot k))$
   $\tau_{pump} \leftarrow rem_p(\tau)$
   Return $[\tau_{pump}, A \to a\alpha, \text{compress}(\tau(p \cdot 1)), \ldots, \text{compress}(\tau(p \cdot k))]$

The tree τ is the "composition" of pumping trees in the data structure returned by `compress`. We describe this "composition operation" itself by an algorithm `decompress`.

`decompress`(τ_c):

   If $\tau_c = A_a$ for some $A \in V$, **return** τ_c
   Let τ_c be of the form $[\tau_{pump}, A \to a\alpha, \tau_c^1, \ldots, \tau_c^k]$
   $\tau' \leftarrow A_{a\alpha}(\text{decompress}(\tau_c^1), \ldots, \text{decompress}(\tau_c^k))$
   Return $\tau_{pump}[\text{longestpath}(\tau_{pump}) \mapsto \tau']$

The following lemma characterizing the relationship between `compress` and `decompress` is easy to see.

Fig. 1. A complete derivation tree for the grammar in Example 5 on the left and the compressed tree data structure with removed pumping trees on the right.

Lemma 4. *For any complete derivation tree τ, $\tau = \texttt{decompress}(\texttt{compress}(\tau))$.*

Example 5. Consider a grammar $(\{S, B, C\}, \{\texttt{a}\}, \text{P}, S)$ with productions

$$S \to \texttt{a}SB \,|\, \texttt{a}BB \,|\, \texttt{a}B, \quad B \to \texttt{a}BC \,|\, \texttt{a}, \quad C \to \texttt{a}.$$

Consider the complete derivation tree shown on the left in Figure 1. The output of `compress` will be

$$[\tau_1, S \to \texttt{a}BB, [\tau_2, B \to \texttt{a}], B_{\texttt{a}}]$$

We will now show that the data structure returned by `compress` has a constant number of pumping trees. Consider a call of $\texttt{compress}(\tau)$, where p is the $\texttt{longestpath}(\tau)$. The key property that we exploit is the fact that the label $\ell(\tau)$ does not appear in the subtrees rooted at the children of p.

Lemma 5. *For any complete derivation tree τ, the number of trees in the data structure returned by $\texttt{compress}(\tau)$ is at most $m^{|V|}$.*

Proof. Let $p = \texttt{longestpath}(\tau)$, and let $\tau(p \cdot 0), \tau(p \cdot 1), \ldots, \tau(p \cdot k)$ be the children of $\tau(p)$. As observed before, the label $\ell(\tau)$ does not appear in the subtrees $\tau(p \cdot i)$. Thus, the depth of the recursion in `compress` is bounded by $|V|$. Finally, observing that $k \le m$, we get the desired bound. $\qquad\square$

We are now ready to prove the main result of this section — bound on the Parikh supremum of any variable in a complete derivation tree.

Lemma 6. *For any variable A, $\sup_{\text{Pk}}(A, \Delta_G^\Sigma) \le m^{|X|+|V|+1}$*

Proof. By Lemma 5, we know that the number of trees in $\texttt{compress}(\tau)$ is at most $m^{|V|}$. Consider any PLM cut $\mathcal{C}$ of τ. Any node in $\mathcal{C}$ belongs to at most one tree in $\texttt{compress}(\tau)$. Further for any $\tau_1 \in \texttt{compress}(\tau)$, $\mathcal{C}$ restricted to τ_1 is a PLM cut of τ_1. Thus, $\mathcal{C}$ can be seen as the union of at most $m^{|V|}$ PLM cuts in pumping trees. By Lemma 2 and Lemma 3, the Parikh supremum of any variable in any of these pumping trees is at most $m^{|X|+1}$. Putting these observations together, gives us the desired bound. $\square$

5.3 Right-Linear SCFG for Polynomially Ambiguous SCFGs

For this section, let us fix $k = m^{|V|+|X|+1}$. By Lemma 6, in any PLM derivation of G, any variable appears at most k times at any step. Since k is a constant, the right-linear grammar can simulate every PLM derivation of G by explicitly keeping track of only k copies of any variable. This idea is very similar to the one used in [11]. We now give the formal definition of the right-linear grammar based on this intuition.

For a sentence $\alpha \in (\Sigma \cup V)^*$, we define $\boldsymbol{\ell f}(\alpha) \triangleq \alpha{\upharpoonright}_\Sigma \alpha{\upharpoonright}_X \alpha{\upharpoonright}_Y$. The stochastic right-linear grammar $(G_1 = (V_1, \{\texttt{a}\}, P_1, S_1), \Pr_1)$ is formally defined as follows.

1. $V_1 = \{\langle \alpha \rangle \mid \alpha \in X^*Y^*$ such that for each $A \in V$. $\text{Pk}(\alpha)(A) \le k\}$. Thus, the variables of G_1 are sequences of single-derivation variables followed by multiple-derivation variables from G in which each variable appears at most k times.
2. $S_1 = \langle S \rangle$
3. For any production $\pi = (A \to \texttt{a}\beta) \in P$ and sentence $\alpha \in V^*$ we define a production

$$\pi^\alpha = (\langle A\alpha \rangle \to \texttt{a} \langle \boldsymbol{\ell f}(\beta\alpha) \rangle)$$

   corresponding to applying the production π as a PLM step from the sentence $A\alpha$. The set P_1 is defined as

$$P_1 = \{\pi^\alpha \mid \pi = (A \to \texttt{a}\beta) \in P \wedge \langle A\alpha \rangle, \langle \beta\alpha \rangle \in V_1\}$$

4. Finally $\Pr_1$ is defined as $\Pr_1(\pi^\alpha) = \Pr(\pi)$ for all $\pi^\alpha \in P_1$.

We first observe that $(G_1, \Pr_1)$ is a stochastic CFG. The proof is in Section B.

Proposition 1. $(G_1, \Pr_1)$ *is a stochastic CFG.*

$(G_1, \Pr_1)$ is equivalent to $(G, \Pr)$. Its proof is in Section C.

Theorem 7. *For any unary, stochastic grammar* $(G, \Pr)$ *of polynomial ambiguity, there is a stochastic right-linear grammar* $(G_1, \Pr_1)$ *such that* $[\![G]\!]_{\Pr} = [\![G_1]\!]_{\Pr_1}$.

6 Conclusions

In this paper we investigated whether Parikh's theorem generalizes to weighted automata. We proved that it does indeed when the weighted context-free grammar is over a commutative, idempotent semiring. We showed that idempotence of the weight domain is necessary by demonstrating that Parikh's theorem does not extend

to unary, stochastic grammars. However, we proved that if the context-free grammar is polynomially ambiguous, then idempotence of the weight domain is not required for Parikh's theorem to hold.

Our proof for Parikh's theorem for commutative and idempotent semirings extends (as is) to pushdown automata (as opposed to context-free grammars). However, the same does not apply to our result for unary, polynomially ambiguous grammars over non-idempotent rings. Our current proof subtly relies on the "one state" property of context-free grammars. It would be interesting to see how to generalize these ideas to the case of pushdown automata. Finally, stochastic context-free grammars have a (weaker) semantics as language acceptors — the grammar accepts a word if its weight is $> \frac{1}{2}$. Our results imply that every unary language accepted by a polynomially ambiguous, stochastic context-free grammar is also accepted by a probabilistic automata (with probability $> \frac{1}{2}$). It is open if this also holds when the grammar is exponentially ambiguous; our counterexample in this paper only shows that there is no probabilistic automaton that satisfies the stronger requirement that words are accepted with the *same* probability.

References

1. Aminof, B., Kupferman, O., Lampert, R.: Reasoning about online algorithms with weighted automata. ACM Transactions on Algorithms 6(2) (2010)
2. Beeri, C., Milo, T.: Schemas for integration and translation of structured and semi-structured data. In: Proceedings of the International Conference on Database Theory. pp. 296–313 (1999)
3. Buchsbaum, A., Giancarlo, R., Westbrook, J.: On the determinization of weighted finite automata. SIAM Journal on Computing 30(5), 1502–1531 (2000)
4. Chatterjee, K., Doyen, L., Henzinger, T.: Quantitative languages. In: Proceedings of the Annual Conference of the European Association for Computer Science Logic. pp. 385–400 (2008)
5. Chatterjee, K., Doyen, L., Henzinger, T.: Alternating weighted automata. In: Proceedings of the International Symposium on Fundamentals of Computation Theory. pp. 3–13 (2009)
6. Culik, K., Kari, J.: Image compression using weighted finite automata. Computer and Graphics 17, 305–313 (1993)
7. Dang, Z., Ibarra, O., Bultan, T., Kremmerer, R., Su, J.: Binary reachability analysis of discrete pushdown timed automata. In: Proceedings of the International Conference on Computer Aided Verification. pp. 69–84 (2000)
8. Droste, M., Gastin, P.: Weighted automata and weighted logics. In: Proceedings of the International Colloquium on Automata, Languages and Programming. pp. 513–525 (2005)
9. Droste, M., Kuich, W., Vogler, H. (eds.): Handbook of Weighted Automata. Springer (2009)
10. Esparza, J.: Petri nets, Commutative context-free grammars, and Basic Parallel Processes. Fundamenta Informaticae 31(1), 13–25 (1997)
11. Esparza, J., Ganty, P., Kiefer, S., Luttenberger, M.: Parikh's theorem: A simple and direct automaton construction. Inf. Process. Lett. 111(12), 614–619 (2011)
12. Goldstine, J.: A simplified proof of Parikh's theorem. Discrete Mathematics 19, 235–239 (1977)
13. Göller, S., Mayr, R., To, A.: On the computational complexity of verifying one-counter processes. In: Proceedings of the IEEE Symposium on Logic in Computer Science. pp. 235–244 (2009)

14. Gurari, E., Ibarra, O.: The complexity of decision problems for finite-turn multi-counter machines. Journal of Computer and System Sciences 22, 220–229 (1981)
15. Hafner, U.: Low bit-rate image and video coding with weighted finite automata. Ph.D. thesis, Universität Würzburg (1999)
16. Huynh, D.: The complexity of semilinear sets. In: Proceedings of the International Colloquium on Automata, Languages and Programming. pp. 324–337 (1980)
17. Huynh, D.: Deciding the ineuivalence of context-free grammars with 1-letter terminal alphabet is σ_2^p-complete. Theoretical Computer Science 33, 305–326 (1984)
18. Huynh, D.: Complexity of equivalence problems for commutative grammars. Information and Control 66(1), 103–121 (1985)
19. Ibarra, O.: Reversal-bounded multi-counter machines and their decision problems. Journal of the ACM 25, 116–133 (1978)
20. Jiang, Z., Litow, B., de Vel, O.: Similarity enrichment in image compression through weighted finite automata. In: Proceedings of the International Computing and Combinatorics Conference. pp. 447–456 (2000)
21. Katritzke, F.: Refinements of data compression using weighted finite automata. Ph.D. thesis, Universität Siegen (2001)
22. Kopczyński, E., To, A.: Parikh images of grammars: Complexity and Applications. In: Proceedings of the IEEE Symposium on Logic in Computer Science. pp. 80–89 (2010)
23. Kuperberg, D.: Linear temporal logic for regular cost functions. In: Proceedings of the Symposium on Theoretical Aspects of Computer Science. pp. 627–636 (2011)
24. Mohri, M.: Finite-state transducers in language and speech processing. Computational Linguistics 23, 269–311 (1997)
25. Mohri, M., Pereira, F., Riley, M.: The design principles of weighted finite-state transducer library. Theoretical Computer Science 231, 17–32 (1997)
26. Parikh, R.J.: On context-free languages. Journal of the ACM 13(4), 570–581 (1966)
27. Paz, A.: Introduction to Probabilistic Automata. Academic Press (1971)
28. Rabin, M.: Probabilistic automata. Information and Control 6(3), 230–245 (1963)
29. Reps, T., Lal, A., Kidd, N.: Program analysis using weighted pushdown systems. In: Proceedings of the International Conference on Foundations of Software Technology and Theoretical Computer Science. pp. 23–51 (2007)
30. Reps, T., Schwoon, S., Jha, S., Melski, D.: Weighted pushdown systems and their application to interprocedural dataflow analysis. Science of Computer Programming 58(1-2), 206–263 (2005)
31. Schützenberger, M.: On the definition of a family of automata. Information and Control 4, 245–270 (1961)
32. Sen, K., Viswanathan, M.: Model checking mutli-threaded programs with asynchronous atomic methods. In: Proceedings of the International Conference on Computer Aided Verification. pp. 300–314 (2006)
33. Siedl, H., Schwentick, T., Muscholl, A.: Counting in trees. Texts in Logic and Games 2, 575–612 (2007)
34. To, A., Libkin, L.: Algorithmic meta-theorems for decidable LTL model checking over infinite state systems. In: Proceedings of the International Conference on the Foundations of Software Science and Computation Structures. pp. 221–236 (2010)
35. Wich, K.: Exponential ambiguity of context-free grammars. In: Developments in Language Theory, Foundations, Applications, and Perspectives, Aachen, Germany, 6-9 July 1999. pp. 125–138 (1999)
36. Yen, H.: On reachability equivalence for BPP-nets. Theoretical Computer Science 179, 301–317 (1996)

A Proof of Theorem 3

Let $G = (V, \Sigma, P, S)$ be a context-free grammar and let $W : P \to \mathbb{D}$ be a weight function over a commutative, idempotent weight domain $\mathbb{D}$. Consider the following homomorphism $h : P^* \to \Sigma^*$ defined as $h(\pi) = \mathtt{a}$, where $\pi = A \to \mathtt{a}\alpha \in P$.

We begin by first constructing a weighted context-free grammar (G_1, W_1) over the alphabet P, whose image under h gives us G. Formally, $G_1 = (V, P, P_1, S)$ has as productions $P_1 = \{A \to \pi\alpha \mid \exists \mathtt{a} \in \Sigma.\ \pi = A \to \mathtt{a}\alpha \in P\}$. In addition, take W_1 to be $W_1(A \to \pi\alpha) = W(\pi)$. It is easy to see that $h(L(G_1)) = L(G)$ by construction. Moreover, given W_1 and W, we can conclude

$$\llbracket G \rrbracket_W (w) = \bigoplus_{\omega \in P^* : h(\omega) = w} \llbracket G_1 \rrbracket_{W_1} (\omega) \tag{1}$$

Next, observe that the weight of a derivation in G_1 is determined by the terminal symbols. Further, since $\oplus$ is idempotent, we have for any $\omega = \pi_1 \pi_2 \cdots \pi_n \in L(G_1)$,

$$W_1(\omega) = W(\pi_1) \otimes W(\pi_2) \otimes \cdots \otimes W(\pi_n) \tag{2}$$

By Parikh's theorem (Theorem 2), there is a right-linear grammar $G_2 = (V_2, P, P_2, S_2)$ such that $Pk(L(G_2)) = Pk(L(G_1))$. Define the weight function W_2 as $W_2(A \to \pi B) = W(\pi)$ to give us the weighted CFG (G_2, W_2). Like G_1, the weight of a derivation in G_2 is determined by the terminal symbols. Together with idempotence of $\oplus$, we get, for $\omega = \pi_1 \pi_2 \cdots \pi_n \in L(G_2)$,

$$W_2(\omega) = W(\pi_1) \otimes W(\pi_2) \otimes \cdots \otimes W(\pi_n) \tag{3}$$

Putting together Equation 2, Equation 3, commutativity of $\otimes$, and the fact that $Pk(L(G_2)) = Pk(L(G_1))$, we get

$$Pk \llbracket G_1 \rrbracket_{W_1} = Pk \llbracket G_2 \rrbracket_{W_2}. \tag{4}$$

Finally, we consider the context free grammar G_3 obtained by "applying the homomorphism h" to G_2. Formally, $G_3 = (V_2, \Sigma, P_3, S_2)$, where $P_3 = \{A \to h(\pi)B \mid A \to \pi B \in P_2\}$. The weight function W_3 is defined in such a way that weight of $A \to \mathtt{a}B$ is the sum of the weights of all productions $A \to \pi B$, where $h(\pi) = \mathtt{a}$, i.e.,

$$W_3(A \to \mathtt{a}B) = \bigoplus_{\pi : h(\pi) = \mathtt{a}} W_2(A \to \pi B).$$

(G_3, W_3) and (G_2, W_2) share the same relationship as (G, W) and (G_1, W_1). That is, we have $h(L(G_2)) = L(G_3)$ and

$$\llbracket G_3 \rrbracket_{W_3} (w) = \bigoplus_{\omega \in P^* : h(\omega) = w} \llbracket G_2 \rrbracket_{W_2} (\omega) \tag{5}$$

The weighted CFG (G_3, W_3) is the desired weighted grammar. We complete the proof by showing $Pk[\![G_3]\!]_{W_3} = Pk[\![G]\!]_W$ through the following reasoning.

$$Pk[\![G_3]\!]_{W_3}(w) = \bigoplus_{w' \in [w]_{Pk}} [\![G_3]\!]_{W_3}(w')$$

$$= \bigoplus_{w' \in [w]_{Pk}} \left[\bigoplus_{\omega \in P^*: h(\omega)=w'} [\![G_2]\!]_{W_2}(\omega) \right] \qquad \text{by Equation 5}$$

$$= \bigoplus_{\omega \in P^*: h(\omega)=w} \left[\bigoplus_{\omega' \in [\omega]_{Pk}} [\![G_2]\!]_{W_2}(\omega') \right] \qquad \text{commutativity of } \oplus$$

$$= \bigoplus_{\omega \in P^*: h(\omega)=w} Pk[\![G_2]\!]_{W_2}(\omega)$$

$$= \bigoplus_{\omega \in P^*: h(\omega)=w} Pk[\![G_1]\!]_{W_1}(\omega) \qquad \text{by Equation 4}$$

$$= \bigoplus_{\omega \in P^*: h(\omega)=w} \left[\bigoplus_{\omega' \in [\omega]_{Pk}} [\![G_1]\!]_{W_1}(\omega') \right]$$

$$= \bigoplus_{w' \in [w]_{Pk}} \left[\bigoplus_{\omega \in P^*: h(\omega)=w'} [\![G_1]\!]_{W_1}(\omega) \right] \qquad \text{commutativity of } \oplus$$

$$= \bigoplus_{w' \in [w]_{Pk}} [\![G]\!]_W(w) \qquad \text{by Equation 1}$$

$$= Pk[\![G]\!]_W(w).$$

B Proof of Proposition 1

For any $\langle A\alpha \rangle \in V_1$ and $a \in \Sigma$ we have

$$\sum_{\beta \in V^*:(\langle A\alpha \rangle \to a\langle \beta \rangle) \in P_1} Pr_1(\langle A\alpha \rangle \to a\langle \beta \rangle) = \sum_{\gamma \in V^*:(\langle A\alpha \rangle \to a\langle \ell f(\gamma\alpha) \rangle) \in P_1} Pr_1(\langle A\alpha \rangle \to a\langle \ell f(\gamma\alpha) \rangle)$$

$$= \sum_{\gamma \in V^*(A \to a\gamma) \in P} Pr_1((A \to a\gamma)^\alpha)$$

$$= \sum_{\gamma \in V^*(A \to a\gamma) \in P} Pr(A \to a\gamma)$$

$$\in [0, 1].$$

C Proof of Theorem 7

Before presenting the proof, we highlight the relationship between G and G_1.

Lemma 7. *For any PLM derivation* $\alpha_0 \overset{\pi_1}{\Rightarrow}_{plm} \alpha_1 \overset{\pi_2}{\Rightarrow}_{plm} \cdots \overset{\pi_n}{\Rightarrow}_{plm} \alpha_n$ *in* G *and for all* $i \in \{0, 1, \ldots, n\}$,

$$Pr(\alpha_0 \overset{\pi_1}{\Rightarrow}_{plm} \alpha_1 \overset{\pi_2}{\Rightarrow}_{plm} \cdots \overset{\pi_i}{\Rightarrow}_{plm} \alpha_i) = Pr_1(\ell f(\alpha_0) \overset{\pi_1^{\ell f(\alpha_0) \upharpoonright_V}}{=\!=\!=\!\Longrightarrow}_{plm} \ell f(\alpha_1) \overset{\pi_2^{\ell f(\alpha_1) \upharpoonright_V}}{=\!=\!=\!\Longrightarrow}_{plm} \cdots \overset{\pi_i^{\ell f(\alpha_{i-1}) \upharpoonright_V}}{=\!=\!=\!\Longrightarrow}_{plm} \ell f(\alpha_i))$$

Lemma 7 can be proved by induction on the length of the PLM derivation. Using this lemma, we can finally prove that $[\![G]\!]_{\mathrm{Pr}} = [\![G_1]\!]_{\mathrm{Pr}_1}$.

Construct the stochastic right-linear grammar (G_1, Pr_1) as described in Section 5.3.

For a derivation $\Psi = \alpha_0 \overset{\pi_1}{\Rightarrow}_{\mathrm{plm}} \alpha_1 \overset{\pi_2}{\Rightarrow}_{\mathrm{plm}} \cdots \overset{\pi_n}{\Rightarrow}_{\mathrm{plm}} \alpha_n$ of G, let $f(\Psi) = \boldsymbol{\ell f}(\alpha_0) \overset{\pi_1^{\ell f(\alpha_0)\restriction_V}}{=\!=\!=\!\Longrightarrow}_{\mathrm{plm}}$
$\boldsymbol{\ell f}(\alpha_1) \overset{\pi_2^{\ell f(\alpha_1)\restriction_V}}{=\!=\!=\!\Longrightarrow}_{\mathrm{plm}} \cdots \overset{\pi_n^{\ell f(\alpha_{n-1})\restriction_V}}{=\!=\!=\!\Longrightarrow}_{\mathrm{plm}} \boldsymbol{\ell f}(\alpha_n)$ as in Lemma 7. Now, for any $w \in \{\mathrm{a}\}^*$, we have

$$
\begin{aligned}
[\![G]\!]_{\mathrm{Pr}}(w) &= \sum_{\Psi \in \mathrm{parse}_G(w)} \mathrm{Pr}(\Psi) \\
&= \sum_{\Psi \in \mathrm{parse}_G(w)} \mathrm{Pr}_1(f(\Psi)) \\
&= \sum_{\psi \in \mathrm{parse}_{G_1}(w)} \mathrm{Pr}_1(\psi) \\
&= [\![G_1]\!]_{\mathrm{Pr}_1}(w).
\end{aligned}
$$