

This is a repository copy of *Random Graph Generation in Context-Free Graph Languages*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/211740/>

Version: Accepted Version

Proceedings Paper:

Vastarini, Federico and Plump, Detlef orcid.org/0000-0002-1148-822X (2024) Random Graph Generation in Context-Free Graph Languages. In: Proc. 13th International Workshop on Developments in Computational Models (DCM 2023). Electronic Proceedings in Theoretical Computer Science . Open Publishing Association , pp. 107-120.

<https://doi.org/10.4204/EPTCS.408.7>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Random Graph Generation in Context-Free Graph Languages

Federico Vastarini

University of York
York, UK

federico.vastarini@york.ac.uk

Detlef Plump

University of York
York, UK

detlef.plump@york.ac.uk

We present a method for generating random hypergraphs in context-free hypergraph languages. It is obtained by adapting Mairson’s generation algorithm for context-free string grammars to the setting of hyperedge replacement grammars. Our main results are that for non-ambiguous hyperedge replacement grammars, the method generates hypergraphs uniformly at random and in quadratic time. We illustrate our approach by a running example of a hyperedge replacement grammar generating term graphs.

1 Introduction

We present a novel approach to the generation of random hypergraphs in user-specified domains. Our method extends a method of Mairson for generating strings in context-free languages [8] to the setting of context-free hypergraph languages specified by hyperedge replacement grammars. Generating (or “sampling”) graphs and hypergraphs according to a given probability distributions is a problem that finds application in testing algorithms and programs working on graphs. Other potential application areas include molecular biology and post-quantum cryptography.

Our generation algorithm uses as input a hyperedge replacement grammar in Chomsky normal form and a positive integer n . The former specifies the hypergraph language to sample from, the latter the size of the hypergraph to be generated. The algorithm then chooses a hypergraph at random from the slice of the language consisting of all members of size n . We show that if the grammar is non-ambiguous, the generated samples are uniformly distributed. The only requirements for our method are that the properties sought for the generated hypergraphs are representable by a hyperedge replacement language and that, to guarantee a uniform distribution, a non-ambiguous grammar is used as input.

We also show that our method generates a random hypergraph of size n in time $O(n^2)$. This is the same time bound established by Mairson (for his first method) in the setting of random string generation in context-free languages.

2 Hyperedge Replacement Grammars

This section gives a concise overview of the definitions needed to understand the generation process. We also introduce our running example of a language of term graphs specified by hyperedge replacement. For comprehensive treatments of the theory of hyperedge replacement grammars and languages, we refer to Courcelle [2], Drewes et al. [3] and Engelfriet [4].

Let $type: C \rightarrow \mathbb{N}_0$ be a typing function for a fixed set of labels C , then a *hypergraph* over C is a tuple $H = (V_H, E_H, att_H, lab_H, ext_H)$ where V_H is a finite set of *vertices*, E_H is a finite set of *hyperedges*, $att_H: E_H \rightarrow V_H^*$ is a mapping assigning a sequence of *attachment nodes* to each $e \in E_H$, $lab_H: E_H \rightarrow C$

is a function that maps each hyperedge to a *label* such that $\text{type}(\text{lab}_H(e)) = |\text{att}_H(e)|$, $\text{ext}_H \in V_H^*$ is a sequence of pairwise distinct *external nodes* (Fig. 1). The class of all hypergraphs over C is denoted by \mathcal{H}_C .

We write $\text{type}(H)$ for $|\text{ext}_H|$ and call H an n -hypergraph if $\text{type}(H) = n$.

We call the length of the sequence of *attachments* $|\text{att}_H(e)|$ the *type* of a hyperedge and e an m -hyperedge if $\text{type}(\text{lab}(e)) = m$. We also use $\text{type}(e) = m$ if the context is clear.

We call a n -hypergraph H a *handle* if $|V_H| = n$, $E_H = \{e\}$ and $\text{att}_H(e) = \text{ext}_H$ and denote H by A^\bullet , the handle induced by A , if $\text{lab}_H(e) = A$.

We denote by $\text{att}_H(e)_i$ the i th attachment node of $e \in E_H$ and by $\text{ext}_{H,i}$ the i th external node of H .

The set $E_H^X = \{e \in E \mid \text{lab}_H(e) \in X\}$ denotes the subset of E_H with labels in $X \subseteq C$.

We define $|H| = |V_H| + |E_H|$ as the *size* of H and we call H a *size- n* -hypergraph if $|H| = n$.

Figure 1 shows a *term graph*, a form of acyclic hypergraphs that represent functional expressions with possibly shared subexpressions. (See [9] for an introduction to the area of term graph rewriting.) Grey boxes represent hyperedges labelled with the function symbols $*$, $+$ and 1 , while nodes are drawn as black bullets. Lines connect hyperedges with their attachment nodes, whose position in the attachment sequence is given by small numbers.

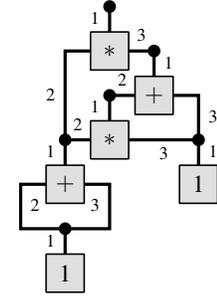


Figure 1: A term graph

Two hypergraphs $H, H' \in \mathcal{H}_C$ are *isomorphic*, denoted $H \cong H'$, if there are bijective mappings $h_V: V_H \rightarrow V_{H'}$ and $h_E: E_H \rightarrow E_{H'}$ such that $h_V^*(\text{att}_H(e)) = \text{att}_{H'}(h_E(e))$ and $\text{lab}_H(e) = \text{lab}_{H'}(h_E(e))$ for each $e \in E_H$ and $h_V^*(\text{ext}_H) = \text{ext}_{H'}$. Two isomorphic hypergraphs are considered to be the same. A hypergraph H is a *subgraph* of H' , denoted as $H \subseteq H'$ if $V_H \subseteq V_{H'}$ and $E_H \subseteq E_{H'}$.

Hypergraphs are generated by replacement operations. Let $H \in \mathcal{H}_C$ and $B \subseteq E_H$ and let $\text{repl}: B \rightarrow \mathcal{H}_C$ be a mapping with $\text{type}(\text{repl}(e)) = \text{type}(e)$ for each $e \in B$. Then the *replacement* of the hyperedges in B with respect to $\text{repl}(e)$ is defined by the operations: remove the subset B of hyperedges from E_H ; for each $e \in B$, disjointly add the vertices and the hyperedges of $\text{repl}(e)$; for each $e \in B$ and $1 \leq i \leq \text{type}(e)$, fuse the i th external node $\text{ext}_{\text{repl}(e),i}$ with the i th attachment node $\text{att}_B(e)_i$.

We denote the resulting hypergraph by $H[e_1/R_1, \dots, e_n/R_n]$, where $B = \{e_1, \dots, e_n\}$ and $\text{repl}(e_i) = R_i$ for $1 \leq i \leq n$, or $H[\text{repl}]$. The replacement preserves the external nodes, thus $\text{ext}_{H[\text{repl}]} = \text{ext}_H$.

Given the subsets $\Sigma, N \subseteq C$ used as terminal and non-terminal labels, with $\Sigma \cap N = \emptyset$, we denote E_H^Σ and E_H^N respectively the subsets of terminal and non-terminal hyperedges of H .

The replacements applied during the generation are defined in productions: $p = (A, R)$ is a *production* over N , where $\text{lhs}(p) = A \in N$ is the label of the replaced hyperedge and $\text{rhs}(p) = R \in \mathcal{H}_C$ is a hypergraph with $\text{type}(R) = \text{type}(A)$. If $|\text{ext}_R| = |V_R|$ and $E_R = \emptyset$, then p is said to be *empty*.

Let $H \in \mathcal{H}_C$ and let $p = (\text{lab}(e), R)$, with $e \in E_H$, then a *direct derivation* $H \Rightarrow_p H'$ is obtained by the replacement $H' = H[e/R]$.

A sequence d of direct derivations $H_0 \Rightarrow_{p_1} \dots \Rightarrow_{p_k} H_k$ of length k with $(p_1, \dots, p_k) \in P$ is denoted as $H \Rightarrow^k H_k$ or $H \Rightarrow_p^* H_k$ if the length is not relevant. We denote it as $H \Rightarrow^* H_k$ if the sequence is clear from the context.

A derivation $H \Rightarrow^* H'$ of length 0 is given if $H \cong H'$.

Given an ordered set $\{\alpha_1, \dots, \alpha_n\}$ where $\alpha_i < \alpha_j$ if $i < j \in \mathbb{N}$ we define a *hyperedge replacement grammar*, or *HRG* as a tuple $G = (N, \Sigma, P, S, (\text{mark}_p)_{p \in P})$ where $N \subseteq C$ is a finite set of non-terminal labels, $\Sigma \subseteq C$ is a finite set of terminal labels with $N \cap \Sigma = \emptyset$, P is a finite set of productions, $S \in N$ is the starting symbol, $(\text{mark}_p)_{p \in P}$ is a family of functions $\text{mark}_p: E_R \rightarrow \{\alpha_1, \dots, \alpha_n\}$ assigning a mark to

each hyperedge in the right-hand side of a production p (Fig. 2). For each pair $e_i, e_j \in E_R$ with $i \neq j$, $\text{mark}(e_i) \neq \text{mark}(e_j)$.

We denote as $P^A \subseteq P$ the subset of productions where $\text{lhs}(p) = A$. We call a production $p = (A, R) \in P_N \subseteq P$ *non-terminal* if $E_R^N \neq \emptyset$ or *terminal* if $p = (A, R) \in P_\Sigma = P \setminus P_N$.

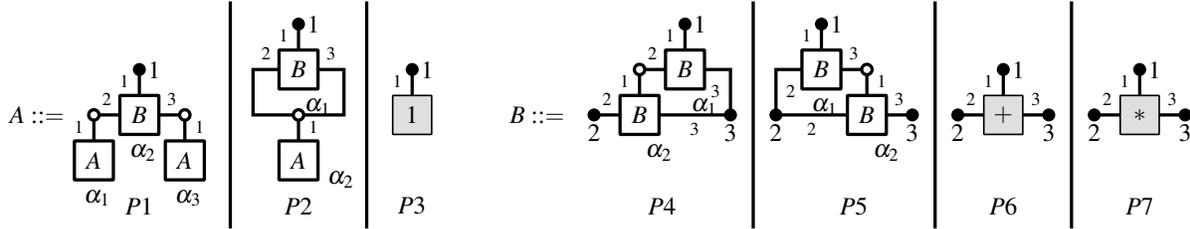


Figure 2: An ambiguous hyperedge replacement grammar for term graphs

The marking of the hyperedges in the *rhs* of each production, represents the order in which the replacements are carried out $(\alpha_1, \alpha_2, \dots, \alpha_{n-1}, \alpha_n)$. It allows for the definitions of ordered derivation tree and leftmost derivation.

Given a set of productions P , we denote by T_P the set of all ordered trees over P which is inductively defined as follows: for each $p \in P$, $p \in T_P$; for $t_1, \dots, t_n \in T_P$ and $p \in P$, $p(t_1, \dots, t_n) \in T_P$.

Then, given a *HRG* G , an *ordered derivation tree* t for e such that $\text{lab}(e) = X \in N$, is a tree $p(t_{\alpha_1}, \dots, t_{\alpha_n})$ in T_P , such that $p = (X, R)$ is a production in P , and $t_{\alpha_1}, \dots, t_{\alpha_n}$ are derivation trees for $e_1 \dots e_n$, such that $X_1 \dots X_n$ are the labels of the non-terminal hyperedges in R marked with $\alpha_1 \dots \alpha_n$, respectively (Fig. 11).

We define the *yield* of t , denoted with $\text{yield}(t)$, as the sequence of replacements: $\text{yield}(p(t_{\alpha_1}, \dots, t_{\alpha_n})) = \text{rhs}(p)[e_1/\text{yield}(t_{\alpha_1}), \dots, e_n/\text{yield}(t_{\alpha_n})]$.

Let t be an ordered derivation tree for a hypergraph H obtained from a derivation $d = S^\bullet \Rightarrow_p^* H$ and $\text{trav}(t)$ its pre-ordered visit. Then d is said to be a *leftmost derivation*, denoted as $\text{lmd}(H)$, if and only if the order of the applied productions of d corresponds to $\text{trav}(t)$.

Since we need a measure to ensure the termination of the proposed algorithm, we define a *HRG* to be *non-contracting* if for each direct derivation $H \Rightarrow_p H'$, $|H| \leq |H'|$. We call a grammar *essentially non-contracting* if there exists $p = (S, R) \in P$ such that p is the empty production.

The *hyperedge replacement language* (*HRL*) generated by a *HRG* is the set $L(G) = \{H \in \mathcal{H}_\Sigma \mid S^\bullet \Rightarrow_p^* H\}$. We define for each $A \in N$, $L^A(G) = \{H \in \mathcal{H}_\Sigma \mid A^\bullet \Rightarrow_p^* H\}$. We also define for $n \in \mathbb{N}$, $L_n^A(G) = \{H \in \mathcal{H}_\Sigma \mid A^\bullet \Rightarrow_p^* H \wedge |H| = n\}$. Clearly $L_n^A(G) \subseteq L^A(G)$. We denote as $|L_n^A(G)|$ the size of the set of all size- n -hypergraphs in L that can be derived from A^\bullet .

For example, the hyperedge replacement grammar in Figure 2 generates the class of all term graphs with function symbols in $\{*, +, 1\}$. Note that hyperedges with non-terminal labels are depicted as white boxes. A derivation with the Chomsky normal form version of this grammar is given in Figure 10.

We define a grammar G to be *ambiguous* if there are ordered derivation trees $t_1, t_2 \in T_P$, such that $t_1 \neq t_2$ and $\text{yield}(t_1) \cong \text{yield}(t_2)$, or equivalently, if there exist $H, H' \in L(G)$ such that $H \cong H'$ and $\text{lmd}(H) \neq \text{lmd}(H')$. If $\text{yield}(t_1), \text{yield}(t_2) \in L_n(G)$ we say that G is *n-ambiguous*. A non-ambiguous version of the term graphs grammar is given in Figure 3.

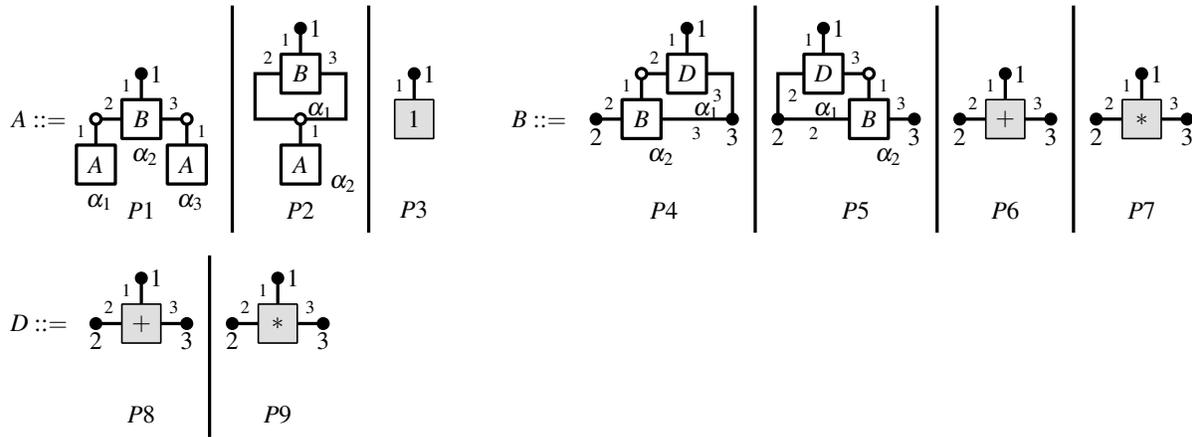


Figure 3: A non-ambiguous hyperedge replacement grammar for term graphs

3 Random hypergraph generation

In 1994, Mairson proposed a pair of methods for the sampling of strings from context-free grammars [8]. His approach requires, as input, a grammar G in Chomsky normal form and the length n of the word to be generated. He proves that, if G is non-ambiguous, such a word is generated uniformly at random. The first method has a time complexity of $O(n^2)$ while requiring $O(n)$ space, and vice versa, the second method runs in linear time using quadratic space. In the following we adapt the first of Mairson's methods to hyperedge replacement grammars. We use our running example of a term graph language to illustrate the generation process.

We define a *Chomsky normal form (CNF)* for hyperedge replacement grammars as a tuple $G_{CNF} = (N, \Sigma, P, S, (mark_p)_{p \in P})$ where:

- $N \subseteq C$ is a finite set of non-terminal labels
- $\Sigma \subseteq C$ is a finite set of terminal labels with $N \cap \Sigma = \emptyset$
- P is a finite set of productions
- $S \in N$ is the starting symbol
- $(mark_p)_{p \in P}$ is a family of functions $mark_p : E_R \rightarrow \{\alpha, \beta\}$ assigning a mark to each hyperedge in the right-hand side of a production p

Each production $p = (A, R) \in P$ satisfies one of the following constraints:

- $E_R = \{e_1, e_2\}$ where $lab(e_1), lab(e_2) \in N$ and $mark(e_1) \neq mark(e_2)$, in which case the replacement is firstly carried out on the hyperedge marked with α , then on the one marked with β
- $E_R = \{e_1\}$ where $lab(e_1) \in \Sigma$ and $mark(e_1) = \alpha$
- $E_R = \emptyset, |V_R| > |ext_R|$
- $A = S, p$ is the empty production and for each $q \in P$, for each $e \in rhs(q)$, $lab(e) \neq S$

Note that in the first two cases, $rhs(p)$ contains either exactly two non-terminal hyperedges or a single terminal hyperedge and may also contain isolated nodes. Productions according to the third case are considered as *terminal* productions. The last case specifies that the empty production is only allowed if there is no other production having the starting symbol in its right-hand side. The grammar in Figure 4 is the Chomsky normal form version of the term graph grammar in Figure 2.

Lemma 3.1. *There exists an algorithm that for every hyperedge replacement grammar G produces a grammar G' in Chomsky normal form such that $L(G) = L(G')$.*

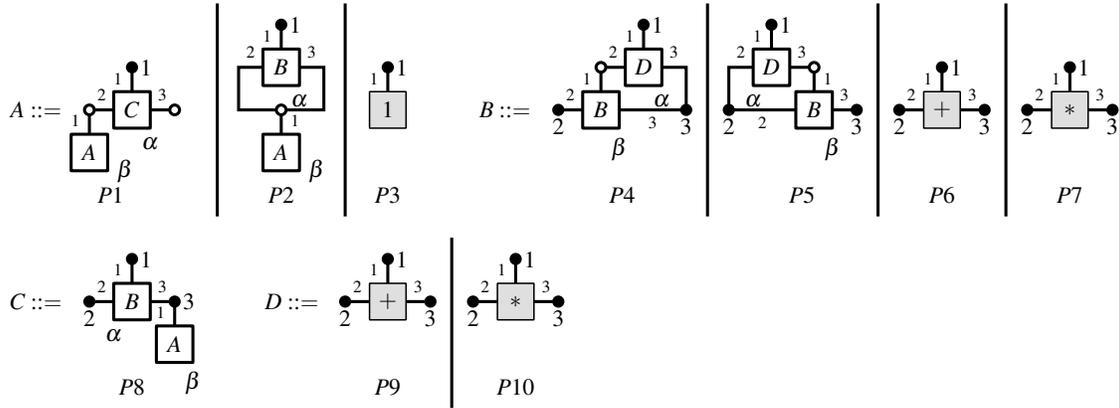


Figure 4: Chomsky normal form version of the grammar in Figure 2

Proof. We present a set of rules to transform any grammar G , into an equivalent grammar G' such that, for each direct derivation $H \Rightarrow_p H'$ with $p \in P_G$, it exists an equivalent derivation $H \Rightarrow_Q^* H'$ with $Q \subseteq P_{G'}$. The proof is provided with a running example showing the application of the rules. The grammar in Figure 5 contains productions that are not in Chomsky normal form: $P1$ has more than 2 hyperedges; $P2$ has a single non-terminal hyperedge; $P3$ is an empty production, but its *lhs* is not S ; $P4$ has 2 hyperedges one of which is terminal.

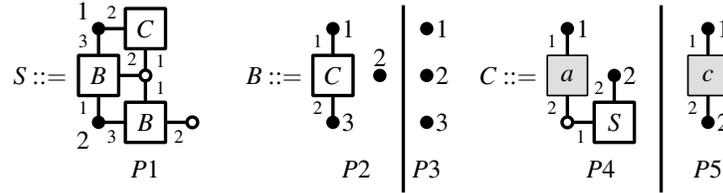


Figure 5: Starting grammar for the proof of *CNF* equivalence.

For a production $p = (A, R) \in P$, that is not already in *CNF*, we consider the following set of rules, applied in this order, to obtain a corresponding equivalent set of productions P' in *CNF*:

1. If p is the empty production, for each production $q = (B, X) \in P$ having $e \in E_X$ with $lab(e) = A$ in its *rhs*, for each production $q' = (A, Y) \in P$ having A in its *lhs* we apply the substitution $R' = X[e, Y]$ and add the productions $p = (B, R')$. We then remove the productions that are no longer needed. The proof of equivalence of the derivations $H \Rightarrow_q H' \Rightarrow_{q'} H''$ and $H \Rightarrow_{p'} H''$ is the following: if e' with $lab(e') = B$ is the hyperedge involved in the derivation $H \Rightarrow_q H'$ then $H'' = H[e'/X[e/Y]] = H[e'/R']$ since $R' = X[e, Y]$.

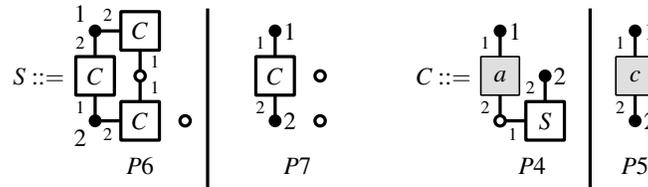


Figure 6: Removal of the empty production $P3$.

order in which the terminal hyperedges are chosen is irrelevant. The proof of equivalence of the derivations $H \Rightarrow_p H''$ and $H \Rightarrow_{p'} H' \Rightarrow_{p''} H''$ is the following: if $e' \in E_R$ with $\text{lab}(e') \in \Sigma$ is the hyperedge involved in the derivation $H \Rightarrow_p H''$ then $H' = H[e/R] = H[e/R'[e'/e'^\bullet]]$ since $R = R'[e'/e'^\bullet]$.

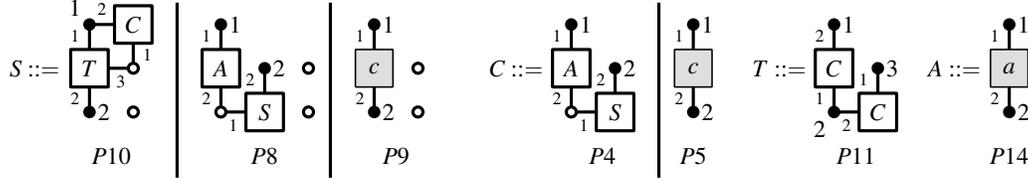


Figure 9: Removal of productions $P8$ and $P4$

Both *rhs* of productions $P4$ and $P8$ are composed by a terminal and a non-terminal hyperedge. We introduce a new label A and a its handle A^\bullet along with the production $P14$ (Fig. 9). We then add the productions $P12$ and $P13$ resulting from the substitution of the terminal hyperedges labelled with a by the non-terminal hyperedges labelled with A . Productions $P4$ and $P8$ are then removed from the grammar. □

From this point on, if not explicitly specified, we always refer to an *HRG* as an *HRG in CNF*. We stress that the input of the method must be already provided in this form, that is, the time required for the transformation is not taken into account during the evaluation of the time complexity.

In order to complete the adaptation of the grammar we propose a more suitable short-hand representation of the productions that only extracts the necessary information, so, for each $p = (A, R) \in P$ and $i \in \mathbb{N}_0$ we use the following notations:

- $A \xrightarrow{p} BC, i$ for a *non-terminal* production where $B, C \in N$ are the labels of the marked hyperedges $e_\alpha, e_\beta \in R$ with $\text{mark}(e_\alpha) = \alpha$, $\text{mark}(e_\beta) = \beta$ and $i = |V_R \setminus \text{ext}_R|$.
- $A \xrightarrow{p} a, i$ for a *terminal* production where $a \in \Sigma$ is the label of the marked hyperedge $e_\alpha \in R$ and $i = |V_R \setminus \text{ext}_R|$.
- $A \xrightarrow{p} \lambda, i$ for a *terminal* production where $E_R = \emptyset$ and $i = |V_R \setminus \text{ext}_R|$.

Matrix M_2 (Tab. 1) shows the short-hand representation of the productions of the grammar in Figure 4. Considering a second input n as the size of the hypergraph to be generated, we are ready to describe a pair of algorithms (**Pre, Gen**) for the random sampling of a hypergraph H from a grammar G . Such a hypergraph is sampled in $L_n^A(G)$, where $A \in N$ is the non-terminal we begin the sampling from. If $A = S$ and G is n -unambiguous, H is sampled uniformly at random among all the hypergraphs in $L_n(G)$.

3.1 Pre-processing phase

The Pre-processing phase is used to construct a pair of matrices M_1, M_2 needed in the generation phase. Let $G = (N, \Sigma, P, S, (\text{mark}_p)_{p \in P})$ be an *HRG*, let $n \in \mathbb{N}$ be the size of the hypergraph $H \in L_n(G)$ we would like to generate, then the algorithm **Pre** (Alg. 1) produces the structures required for the generation.

A as input. For each step we show the probability of the production p to be chosen and the choice of the split and its probability if $p \in P^N$. Since G is non-ambiguous, the first step shows that $|L_{12}(G)| = 3920$, that is, there are 3920 unique size-12-hypergraphs to choose from, each having a different ordered derivation tree. Figure 11 shows the tree t for which $yield(t) = H$, so that $trav(t)$, or equivalently $lmd(H)$, corresponds to the unique sequence of productions applied by the generation algorithm to produce H . In the figure are also indicated the starting symbol A and the replaced hyperedges e_α and e_β , respectively on the edges connecting the left and right child of each node. The proof of termination of the Generation algorithm is based on the assumption that the input grammar is non-contracting:

Algorithm 1: Pre - Pre-processing phase

<p>Input: (G, n), where $G = (N, \Sigma, P, S, (mark_p)_{p \in P})$ and $n \in \mathbb{N}, n \geq 1$ Output: $\langle M_1, M_2 \rangle$</p> <hr style="border: 0.5px solid black;"/> <pre> for 1 ≤ ℓ ≤ n do foreach A ∈ N do A[ℓ] := 0; end foreach p ∈ P do p[ℓ] := 0; end end foreach A \xrightarrow{p} a, i ∈ P_Σ do A \xrightarrow{p} a, i[i+1] := 1; end </pre>	<pre> foreach A \xrightarrow{p} λ, i ∈ P_Σ do A \xrightarrow{p} λ, i[i] := 1; end for 1 ≤ ℓ ≤ n do foreach A ∈ N do foreach p ∈ P^A do A[ℓ] := A[ℓ] + p[ℓ]; end end foreach A \xrightarrow{p} BC, i ∈ P_N do for 1 ≤ k < ℓ do A \xrightarrow{p} BC, i[ℓ+i] := A \xrightarrow{p} BC, i[ℓ+i] + B[k] · C[ℓ-k]; end end end end </pre>
--	--

Algorithm 2: Gen - Generation phase

Input: $(G, \langle M_1, M_2 \rangle, \bar{A}, \bar{n})$, where $G = (N, \Sigma, P, S, (mark_p)_{p \in P})$, $\langle M_1, M_2 \rangle := \mathbf{Pre}(G, n)$, $\bar{A} \in N$
and $\bar{n} \in \mathbb{N}, 1 \leq \bar{n} \leq n + type(\bar{A})$
Output: $H \in L_{\bar{n}}^{\bar{A}}(G)$

$\ell = \bar{n} - type(\bar{A})$
if $\bar{A}[\ell] = 0$ **then**
 | **return** \perp ;
end

Recursively generate H using (\bar{A}, ℓ) as first input as follows:
function $derH(A, \ell)$:
 | $p \leftarrow RNG$ with $p \in P^A$ and probability $p[\ell]/A[\ell]$;
 | **if** $p \in P_T$ **then**
 | | **return** $A^\bullet[e/R]$;
 | **else**
 | | $\ell' = \ell - i$;
 | | $k \leftarrow RNG$ with $0 < k < \ell'$ and probability $B[k] \cdot C[\ell' - k]/(A \xrightarrow{p} BC, i)[\ell]$;
 | | **return** $A^\bullet[e/R[e_\alpha/derH(B, k), e_\beta/derH(C, \ell' - k)]]$;
 | **end**
end function

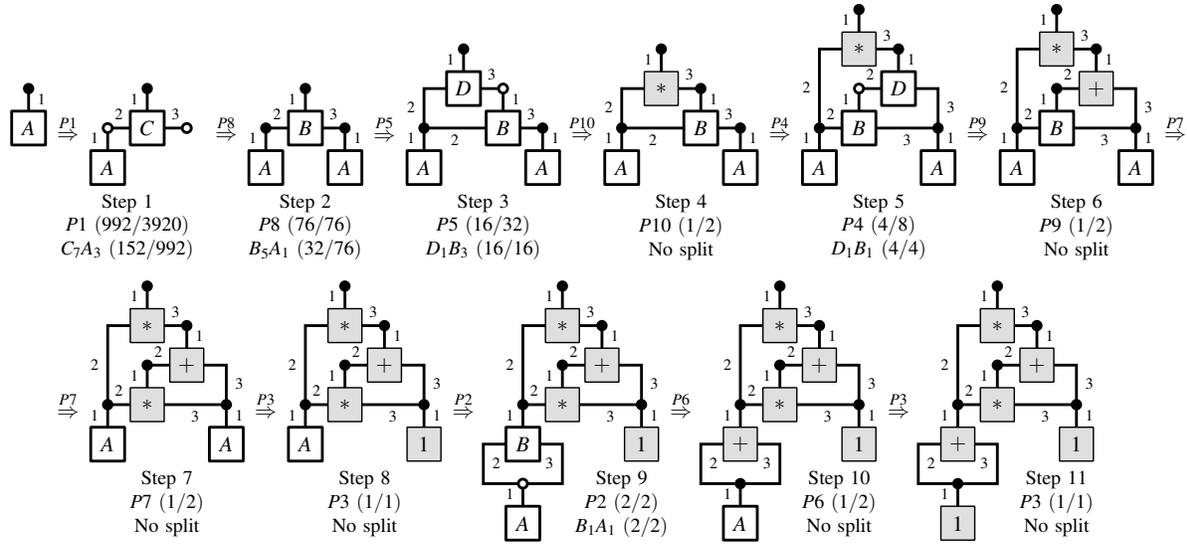


Figure 10: A derivation $d = A^* \Rightarrow_p^* H$ using the grammar of Figure 4

Proof. Let's consider a measure equivalent to the size of a hypergraph $|H|$. To each application of the recursive function **derH** in each step of the algorithm **Gen**, corresponds a direct derivation between two sentential forms $F \Rightarrow F'$ such that $F \leq F'$. Since the grammar is in *CNF*, at each step there are two possible cases:

1. **derH** chooses a non-terminal production. In this case a single hyperedge $e \in F$ is replaced with a hypergraph $R \subseteq F'$ containing 2 hyperedges and 0 or more internal nodes. Clearly $|F| < |F'|$, meaning that the size of the sentential forms gets progressively close to n .
2. **derH** chooses a terminal production. A hyperedge is replaced by a terminal hyperedge or a single node and 0 or more additional internal nodes. In this case $|F| \leq |F'|$. Even if the size is not incremented, being a terminal production, the recursion does not progress any further.

If it is not possible to generate a size- n -hypergraph using the input grammar G the algorithm trivially ends in one step. \square

4 Uniform distribution and time complexity

We now state our first main result, the uniform generation guarantee for Algorithm 2.

Theorem 4.1. *Given a grammar $G = (N, \Sigma, P, S, (mark_p)_{p \in P})$, Algorithm 2 generates from every non-terminal $A \in N$ a size- n -hypergraph $H \in L_n^A(G)$, provided that $L_n^A(G) \neq \emptyset$. If G is n -unambiguous and *RNG* is a uniform random number generator, the hypergraph is chosen uniformly at random.*

Proof. Let G be a n -unambiguous grammar in *CNF*, the recursive function **derH** derives a hypergraph $H \in L_n^A(G)$ simulating $trav(t)$ where $yield(t) = H$ and let $P(c_j)$ denote the probability of the j th choice c made using the *RNG* at each step of the recursion, for a production or a split, according to $lmd(H)$.

Let's remind that for the parallelization, confluence and associativity properties of context-free hyperedge replacement grammars [2], the sequence of replacements associated to a derivation preserves

the result of the derivation, despite of the order in which the replacements are applied. Thus, we are able to discuss each of its steps independently.

By definition, since the grammar is n -unambiguous, for any non-terminal $A \in N$ we know that the set of hypergraphs that can be generated using different productions $p \in P^A$ are pairwise distinct. Otherwise, there would exist $\text{trav}(t') \neq \text{trav}(t'')$ for which $\text{yield}(t') \cong \text{yield}(t'')$.

From algorithm **Pre** (Alg. 1) we know that $\sum_{p \in P^A} p[\ell] = A[\ell]$ and so the probability of the choice c_j of each production in $\text{lmd}(H)$ can be expressed by $P(c_j) = p[\ell]/A[\ell]$. Also, if $p \in P_N$, since the grammar is n -unambiguous the subsets of hypergraphs that can be derived by choosing different splits are also pairwise distinct. For a production $p \in P_N$ then $\sum_{0 < k < \ell} B[k] \cdot C[\ell - k] = A \xrightarrow{p} BC, i[\ell]$, thus a split can be chosen with probability $P(c_j) = B[k] \cdot C[\ell - k]/p[\ell]$.

Knowing that for a lmd , if the grammar is n -unambiguous, both the choices of productions and splits are made from independent sets, considering the corresponding derivation tree t , the probabilities associated to the choice of a node $P(c)$ and the ones associated to its children $P(c')$ and $P(c'')$ are of the form $\frac{m}{q}$, $\frac{m'}{q'}$ and $\frac{m''}{q''}$ with $m, m', m'', q, q', q'' \in \mathbb{N}$ and $q'q'' = m$. Moreover, the probabilities of two consecutive choices $P(c)$ and $P(c')$ are bound to the law of compound probabilities [7], that is, the choice of a node given the choice of its parent is of the form $P(c'|c) = P(c' \cap c)/P(c)$. Then, considering their independence, $P(c'|c) = (P(c')P(c))/P(c) = P(c')$. The same applies for $P(c'')$. The overall probability of the choice of a node and its children is then $P(c)P(c')P(c'') = \frac{m}{q} \frac{m'}{q'} \frac{m''}{q''} = \frac{m'm''}{q}$.

Finally, considering the chain of probabilities described by an lmd , since for \bar{A} $q = |L_{\bar{n}}(G)|$ and for each terminal production $p \in P_{\Sigma}$ $m = 1$, then for each $H \in L_{\bar{n}}(G)$ we can define its probability $P(H)$ to be generated as the productory of independent choices:

$$P(H) = \prod_{j=1}^k P(c_j) = \frac{m_1}{|L_{\bar{n}}(G)|} \cdot \frac{m_2}{q_2} \cdot \frac{m_{k-1}}{q_{k-1}} \cdots \frac{1}{q_k} = \frac{1}{|L_{\bar{n}}(G)|}$$

Each hypergraph $H \in L_{\bar{n}}(G)$ is generated over a uniform distribution given the uniformity of the sampling of the underlying RNG . □

For the complexity analysis we consider the time required by the algorithm **Gen** (Alg. 2) for the generation of the hypergraph and the space required by the algorithm **Pre** (Alg. 1) to store the required data, taking into account that the input grammar is already provided in the correct CNF and the query to the RNG and the replacement operations are performed in unit time. The gaps present in the tables, that are not encountered in string method, are due to the possibility of a production to increase the size of the resulting hypergraph by more than 1 in a single step.

Theorem 4.2. *With the assumptions of Theorem 4.1, the size- n -hypergraph H is generated by **Gen** (Alg. 2) in time $O(n^2)$.*

Proof. The proof of Theorem 4.2 is based on the analysis of the following recurrence relation for the function **derH**: $T(n) \leq cn + \max_{1 \leq k < (n-i)} [T(k) + T(n-k-i)]$, where $T(k)$ and $T(n-k-i)$ are the computational steps required to process the result of the split and i is the number of internal nodes of the current production. In the worst case, we consider that $i = 0$ and that $k = 1$. A simple example is the discrete hypergraph language in which every iteration may generate a terminal hyperedge from e_{α} and the rest of the resulting hypergraph from e_{β} without adding any node. Since the choice of the production is constant, while the choice of a split is linear in n , choosing a split n times leads to a quadratic behavior.

Since $i \ll n$, we may rewrite the recursion as:

$$T(n) \leq cn + \max_{1 \leq k < n} [T(k) + T(n-k)]$$

Then, considering the worst case $k = 1$, for the next step of the recursion we obtain:

$$T(n-1) \leq c(n-1) + \max_{1 \leq k < (n-1)} [T(k) + T(n-k-1)]$$

That is, at each step the choice of a split happens on an input of size $n - 1$. Since this choice requires linear time and it is taken n times, the relation has solution $O(n^2)$. \square

We omit a discussion of the time complexity of the pre-processing phase (Alg. 1) which can be shown to be linear, considering that given a grammar G in *CNF*, being its size $|G|$ constant, for each production p a short form containing the information about the labels and the internal nodes is obtained in constant time.

5 Conclusion

Our main results, presented in Section 4, are that the method generates hypergraphs uniformly at random and in quadratic time. A topic for future work is to modify the generation algorithm so that it runs in linear time and quadratic space, following Mairson's second method in [8].

Another interesting topic is to extend the quasi-polynomial-time approximation algorithm of Gore et al. [5] from strings to hypergraphs. This algorithm guarantees an approximated uniform distribution even for ambiguous grammars.

Our method allows to generate strings uniformly at random in some non-context-free string languages because hyperedge replacement grammars can specify certain *string graph* languages that are not context-free. For example, this applies to the language $\{a^n b^n c^n \mid n \geq 0\}$. Moreover, our method is able to generate strings uniformly at random for a range of inherently ambiguous context-free languages.

The practically most promising application of our generation approach is the testing of programs in arbitrary programming languages that work on graphs. If the inputs of such programs are graphs in a context-free graph language, our method can generate test graphs uniformly at random in the domain of interest. This should allow to refine random testing approaches such as [6, 1].

References

- [1] T. Y. Chen, F.-C. Kuo, R. G. Merkel, and T. Tse. Adaptive random testing: The art of test case diversity. *Journal of Systems and Software*, 83(1):60–66, 2010. SI: Top Scholars.
- [2] B. Courcelle. An axiomatic definition of context-free rewriting and its application to NLC graph grammars. *Theoretical Computer Science*, 55(2):141–181, 1987.
- [3] F. Drewes, A. Habel, and H.-J. Kreowski. Hyperedge replacement graph grammars. In *Handbook of Graph Grammars and Computing by Graph Transformation*, volume 1, pages 95–162. World Scientific, 1997.

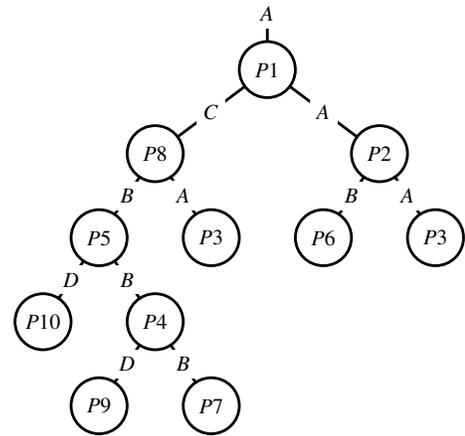


Figure 11: Ordered tree t for the derivation d in Fig. 10

- [4] J. Engelfriet. Context-free graph grammars. In *Handbook of Formal Languages*, volume 3, pages 125–213. Springer, 1997.
- [5] V. Gore, M. Jerrum, S. Kannan, Z. Sweedyk, and S. Mahaney. A quasi-polynomial-time algorithm for sampling words from a context-free language. *Information and Computation*, 134(1):59–74, 1997.
- [6] R. Hamlet. Random testing. In J. J. Marciniak, editor, *Encyclopedia of software Engineering*. Wiley, 1 edition, 1994.
- [7] M. le Marquis de Laplace. Théorie analytique des probabilités. In *Œuvres Complètes de Laplace*, volume 7, pages 181–192. Gauthier-Villars, Imprimeur-Librairie, 3 edition, 1820.
- [8] H. G. Mairson. Generating words in a context-free language uniformly at random. *Information Processing Letters*, 49(2):95–99, 1994.
- [9] D. Plump. Term graph rewriting. In *Handbook of Graph Grammars and Computing by Graph Transformation*, volume 2, pages 3–61. World Scientific, 1999.