



Boltzmann samplers for random generation of lambda terms

Pierre Lescanne

► To cite this version:

| Pierre Lescanne. Boltzmann samplers for random generation of lambda terms. 2014. ensl-00979074v2

HAL Id: ensl-00979074

<https://hal-ens-lyon.archives-ouvertes.fr/ensl-00979074v2>

Submitted on 28 Apr 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Boltzmann samplers

for random generation of lambda terms

Pierre Lescanne
University of Lyon,
École normale supérieure de Lyon,
LIP (UMR 5668 CNRS ENS Lyon UCBL INRIA)
46 allée d'Italie, 69364 Lyon, France
pierre.lescanne@ens-lyon.fr

April 28, 2014

Abstract

Randomly generating structured objects is important in testing and optimizing functional programs, whereas generating random λ -terms is more specifically needed for testing and optimizing compilers. For that a tool called `QuickCheck` has been proposed, but in this tool the control of the random generation is left to the programmer. Ten years ago, a method called Boltzmann samplers has been proposed to generate combinatorial structures. In this paper, we show how Boltzmann samplers can be developed to generate random λ -terms, but also other random data structures like trees. These samplers rely on a critical value which parameters the main random selector and which is exhibited here with explanations on how it is computed. Haskell programs are proposed to show how samplers are actually implemented.

Keywords: lambda calculus, combinatorics, functional programming, test, random generator, Boltzmann sampler

1 Introduction

Claessen and Hughes [5] ask a fundamental question:

[H]ow would one choose a random closed λ -term with a uniform distribution?

Actually generating random lambda terms and more specifically generating random typable λ -terms are fundamental for debugging functional programming compilers [22], but also for answering many questions concerning λ -calculus and functional programming. We will see that this is possible using an apparatus called *Boltzmann samplers* (Duchon et al. [9]) and based on the ideas of

statistical physics initiated by the Austrian physicist Ludwig Boltzmann (1844-1906). The key issue is the ability to generate random combinatorial structures with some flexibility on the size. In this paper the combinatorial structures are λ -terms, but also since we are interested in functional programming they are other data structures, like *binary trees* (the data type `Tree` in [5]) or *1-2-trees*. These latter structures will allow us to introduce the concept of Boltzmann samplers. Assume a program generating random λ -terms is built. Somewhere the program chooses whether it generates a variable, an abstraction or an application. This is done according to a probability distribution. Why the chosen probability distribution is respectively

<i>variable</i>	<i>abstraction</i>	<i>application</i>
0.3703026	0.25939476	0.3703026

will be explained further. For now let us say that there is a *critical value* $\rho \approx 0.509308127$, which plays a main role in counting λ -terms and that the above values are respectively $(1 - \rho^2)/2$ and ρ^2 . Once we know how to generate large plain λ -terms, we are able to address the random generation of large typable λ -terms.

This paper is based on four milestone papers.

- i [8] Nicolaas Govert de Bruijn (1972). *Lambda calculus notation with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser theorem*. *Indagationes mathematicae*, **34**(5), 381–392.
- ii [5] Koen Claessen and John Hughes. *QuickCheck: a lightweight tool for random testing of Haskell programs*. In Martin Odersky and Philip Wadler, editors, *ICFP*, pages 268–279. ACM, 2000. This paper received the *Most Influential ICFP Paper Award* for the year 2000.
- iii [9] Philippe Duchon, Philippe Flajolet, Guy Louchard, and Gilles Schaeffer. *Boltzmann samplers for the random generation of combinatorial structures*. *Combinatorics, Probability & Computing*, 13(4-5):577–625, 2004.
- iv [25] John Tromp. *Binary lambda calculus and combinatory logic*. In Marcus Hutter, Wolfgang Merkle, and Paul M. B. Vitányi, editors, *Kolmogorov Complexity and Applications*, volume 06051 of *Dagstuhl Seminar Proceedings*. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, 2006.

The first paper established what is now called de Bruijn indices which play an important role here in counting λ -terms up to α -conversion. But Nicolaas de Bruijn is better known outside the Functional Programming community as one of the pioneer of modern combinatorics [7] (see [14] fourth cover pages). We will use indirectly his combinatorics results here.

The above quotation of Claessen and Hughes is extracted from the introduction of the second paper [5] and is completed by¹

We have chosen to put distribution under the human tester's control.

However unlike Claessen and Hughes, the philosophy of our paper is that we have chosen to put the distribution under a strict mathematical control, based on Boltzmann samplers and probability theory as introduced in the third paper [9]. For that we use a counting method introduced by John Tromp in [25]. This amazing paper deals with algorithmic complexity theory [19] and proposes to replace Turing machines by λ -calculus in this theory. Therefore it provides a very simple and elegant concrete definition of descriptonal complexity (Kolmogorov complexity). For this Tromp proposes to represent λ -terms by binary chains and in this framework he defines a very small self interpreter of size 210. His coding of λ -terms in $\{0, 1\}^*$ is as follows²:

$$\begin{aligned}\widehat{n} &= 1^n 0 \\ \widehat{\lambda M} &= 00\widehat{M} \\ \widehat{M N} &= 01\widehat{M}\widehat{N}\end{aligned}$$

Representing λ -terms by binary chains confers to them a natural notion of size. This counting is realistic³ as it gives variables a weight depending on the distance from their binder. In other words a variable which is deep in the stack is heavier (larger) than a shallow one .

Moreover our paper relies on the book of Flajolet and Sedgewick [11] which is the reference on generating functions. Our paper is also the application of a recent article containing results on counting binary λ -terms [13].

In this paper we speak about random generation of λ -terms. This applies in testing and optimizing functional programming compilers [22, 21], but it is clear that this applies to other combinatorial objects having bound variables, like imperative programs with a block structure [26].

A survey of Boltzmann approach

Recall that a Boltzmann sampler generates uniformly random objects with a tolerance in the size of the generated objects. In other words, the sampler generates the objects in a cloud around a given size.

¹See appendix for the whole quotation.

²Not exactly since Tromp starts indices at 0 like [16] and we starts indices at 1 like de Bruijn [8].

³In a previous submission [12] where we chose to give de Bruijn indices size 0, a referee said: *"If the authors want to use the de Bruijn representation, another interesting experiment could be done: rather than to count variables as size 0, they should be counted using their unary representation. This would penalize deep lexical scoping, which is not a bad idea since 'local' terms are much easier to understand and analyze than deep terms"*.



An orographic cloud⁴

Before a Boltzmann sampler for combinatorial structures is defined, the generating function associated with these structures is first expressed. Its coefficients count the structures by size. In our case, i.e., trees and λ -terms, the expressions of the generating functions contain a square root of a polynomial, whose smallest root in module is a real which we call the critical value. This is the radius of convergence of the associated series. A Boltzmann sampler is derived from this series and based on a value x which allows generating random structures with a mean and a standard deviation which depend on x . Actually if the user has a specific size in mind, she provides the desired mean and a calculation returns the x that yields this mean. But in our case we assume that we attempt to generate large structures of size, say at least 100, perhaps 1000 or 1000000 or more. In our domain, trees and λ -terms, the choice of x is not any choice, but x shall be set to the critical value, for mathematical reasons that are explained in [9]. This way we get a uniform random sampler for all big trees or λ -terms.

Now let us focus on trees and λ -terms and assume that the structures are distributed in kinds. For instance, for λ -terms, these kinds are variables, abstractions and applications. A random selector among these kinds is built using probabilities computed using the critical value. Then the selector is embedded in a sampler which calls itself recursively. But if any freedom is left to the sampler, it may generate huge structures as well as small ones. To avoid the pitfall of generating huge objects, the sampler has an upper limit of structures it can generate. Furthermore after the generation a sieve retains only the structures that are large enough, yielding the tolerance threshold (the cloud) we mentioned previously.

Structure of the paper The paper is structured as follows. Since the problem of counting λ -terms and presenting their Boltzmann samplers is not straightforward, we present first in Section 2 the generating functions for trees: binary trees and 1-2 trees. Then we present the generating function for counting λ -terms in Section 3. In Section 4 we present a less efficient method for generating λ -terms not based on Boltzmann samplers. In the following sections we address Boltzmann samplers. First in Section 5 we present the general perspective of Boltzmann samplers. Then in Section 6 we specialize on trees: binary trees and 1-2 trees. Section 7 is devoted to Boltzmann samplers for λ -terms. Until that point the paper addresses only the problem of counting and generating plain

⁴Public Domain Picture, origin: Wikimedia, source: *U.S. National Oceanic and Atmospheric Administration*

λ -terms. In Section 8 we survey the harder problem of counting and generating closed terms for which we cannot propose a Boltzmann sampler, since the generating function is not defined by a unique equation.

2 Counting trees

Before presenting how to count λ -terms, we will show in this section how to count trees. We will first show how to count binary trees, then how to count 1-2 trees aka Motzkin trees.

2.1 Binary trees

A binary tree (the data type `Tree` in [5]) is either a leaf or a compound tree made of a node and two subtrees. It can be described as follows:

$$\mathcal{B} = \square + \mathcal{B} \times \bullet \times \mathcal{B}.$$

We assume that leaves and nodes have size 1. General techniques of generating functions ([11] page 738) leads to the following equation for the generating function

$$B(z) = \sum_{n \geq 0} B_n z^n$$

whose coefficients count binary trees of size n :

$$B(z) = z + zB(z)^2$$

or

$$zB(z)^2 - B(z) + z = 0$$

which yields:

$$B(z) = \frac{1 - \sqrt{1 - 4z^2}}{2z}.$$

$B(z)$ is not defined for $z \in \mathbb{R}$ larger than $\frac{1}{2}$. We say that $B(z)$ has a singularity $\rho_{\mathcal{B}}$ which is equal to $\frac{1}{2}$. We are going to use this fact in our sampler of random binary trees.

2.2 Motzkin Trees

We are going to consider another family of trees which have analogies with lambda terms, namely 1-2 trees ([11] pages 81). They are also called *Motzkin trees*. They have two kinds of internal nodes, unary nodes and binary nodes. The class of Motzkin trees is described as

$$\mathcal{M} = \square + \bullet \times \mathcal{M} + \mathcal{M} \times \bullet \times \mathcal{M}.$$

```

unrankToo :: Int → Integer → Term
unrankToo n k
  | k == (trompoo n) = Index $ fromIntegral (n - 1)
  | k <= (trompoo (n - 2)) = Abs (unrankToo (n - 2) k)
  | otherwise = unrankApp (n - 2) 0 (k - trompoo (n - 2))
  where unrankApp n j r | r <= tjnj = let (dv, rm) = (r - 1) 'divMod' tnj
                                     in App (unrankToo j (dv + 1))
                                             (unrankToo (n - j) (rm + 1))
        | otherwise = unrankApp n (j + 1) (r - tjnj)
        where tnj = trompoo (n - j)
              tjnj = (trompoo j) * tnj

```

Figure 1: The unranking function of λ -terms in Haskell

which means that a Motzkin tree is either a leaf of size 1 or a tree of size $n + 1$ rooted on another tree of size n or a tree of size $n_1 + n_2 + 1$ rooted on two trees of size n_1 and n_2 . The generating function

$$M(z) = \sum_{n=0}^{\infty} M_n z^n$$

where M_n counts the number of Motzkin trees of size n , is a solution of the equation:

$$M(z) = z + zM(z) + zM(z)^2$$

or

$$zM(z)^2 - (1 - z)M(z) + z = 0.$$

Hence

$$M(z) = \frac{1 - z - \sqrt{1 - 2z - 3z^2}}{2z}.$$

It is given by the smallest root in module of the polynomial $1 - 2z - 3z^2$, namely $\rho_{\mathcal{M}} = 1/3$ (the other root -1 is larger in module).

3 Counting lambda terms

We now consider λ -terms up to α -conversion. To count equivalence classes modulo α -conversion it is convenient to count canonical representatives of classes modulo α . For that the best way is to count terms with de Bruijn indices, since those terms are unique representatives of classes modulo α .

We work with de Bruijn indices starting at 1. Sometimes we call those indices improperly “variables”. For instance $\lambda 1$ is the representative of the α -conversion equivalence class of the term $\lambda x.x$ and $\lambda \lambda 1 2$ is the representative of the α -conversion equivalence class of the term $\lambda x.\lambda y.y x$. We assume that the

size of the index i is $i + 1$ and the size of an abstraction is 2 and the size of an application is 2 as well. Therefore we get the equation for the numbers⁵ of terms of size n :

$$\begin{aligned} S_{\infty,0} &= S_{\infty,1} = 0, \\ S_{\infty,n+2} &= 1 + S_{\infty,n} + \sum_{k=0}^n S_{\infty,k} S_{\infty,n-k}. \end{aligned}$$

This means that there is no term of size 0 or 1 and for terms of size $n+2$, there is one term which is an index namely $n+1$, plus $S_{\infty,n}$ terms that are abstractions plus $\sum_{k=0}^n S_{\infty,k} S_{\infty,n-k}$ terms that are applications of terms of size k on terms of size $n-k$, for k from 0 to n . Sequence $(S_{\infty,n})_{n \in \mathbb{N}}$ can be found in the *On-line Encyclopedia of Integer Sequences* with the entry number **A114851**. Its first 20 values are:

$$0, 0, 1, 1, 2, 2, 4, 5, 10, 14, 27, 41, \\ 78, 126, 237, 399, 745, 1292, 2404, 4259.$$

Recall some results of Grygiel-Lescanne [13]. Let $S_{\infty}(z)$ denote the generating function for the sequence $(S_{\infty,n})_{n \in \mathbb{N}}$, that is

$$\begin{aligned} S_{\infty}(z) &= \sum_{n=0}^{\infty} S_{\infty,n} z^n \\ &= \sum_{n=0}^{\infty} z^{n+2} + \sum_{n=0}^{\infty} S_{\infty,n} z^{n+2} \\ &\quad + \sum_{n=0}^{\infty} \left(\sum_{k=0}^n S_{\infty,k} S_{\infty,n-k} \right) z^{n+2} \end{aligned}$$

$S_{\infty}(z)$ fulfills the equation

$$S_{\infty}(z) = \frac{z^2}{1-z} + z^2 S_{\infty}(z) + z^2 S_{\infty}(z)^2.$$

which yields

$$S_{\infty}(z) = \frac{z^3 - z^2 - z + 1 - \sqrt{z^6 + 2z^5 - 5z^4 + 4z^3 - z^2 - 2z + 1}}{2z^2(1-z)}.$$

The radius of convergence of this series i.e., the smallest singularity of the associated analytic function, given as the smallest real root of the polynomial

$$\begin{aligned} z^6 + 2z^5 - 5z^4 + 4z^3 - z^2 - 2z + 1 &= \\ (z-1)(z^5 + 3z^4 - 2z^3 + 2z^2 + z - 1). \end{aligned}$$

⁵The index ∞ is justified by the fact that $S_{\infty,n}$ is the limit of $S_{m,n}$ when m goes to ∞ . See Section 8.

A computation using the mathematical software Sage [24] tells us that this smallest root is approximately

$$\rho_{S_\infty} = 0.509308127.$$

Knowledge of the singularity ρ_{S_∞} (which we will write just ρ in what follows) is crucial for computing quantities related to the sequence $(S_{\infty,n})_{n \in \mathbb{N}}$. Especially it plays a key role in computing parameters of the Boltzmann sampler. The following theorem is in [13].

Theorem 1 (Asymptotic evaluation of $S_{\infty,n}$)
The number of all binary λ -terms of size n satisfies

$$S_{\infty,n} \sim (1/\rho)^n \cdot \frac{C}{n^{3/2}},$$

where $1/\rho \doteq 1.963447954$ and $C \doteq 1.021874073$.

It says that whereas binary words are of size 2^n , binary λ -terms are approximately of size 1.963^n for large values of n . For the skeptic reader, we have checked numerically this result with good accuracy [17].

4 Unranking lambda terms

We know how to count the λ -terms, which means basically that we know how to order the terms, assigning a number to each term (ranking). We may use this fact the other way around and build a λ term from its rank. This function is called *unranking* and is implemented in Haskell as *unrankToo* (see Figure 1). The function *trompoo* n^6 implements $S_{\infty,n}$. It works as follows. Given a number k ($1 \leq k \leq S_{\infty,n}$) we consider which interval it belongs to. If $k = S_{\infty,n}$, then k corresponds to a de Bruijn index, namely the de Bruijn index $n - 1$. If $k \leq S_{\infty,n-2}$, then k corresponds to an abstraction, namely the abstraction of the term obtained by unranking k in the interval $[1..S_{\infty,n-2}]$. If $S_{\infty,n-2} < k < S_{\infty,n}$, then k corresponds to an application of two terms which are then computed.

Unranking has at least three applications:

- counting typable λ -terms,
- generating random λ -terms and
- generating random typable λ -terms.

4.1 Counting typable λ -terms

We know very little about the combinatorial properties of simply typable λ -terms. For instance, we do not know any formula for counting them, we have no idea of their distribution among plain λ -terms and we have no information

⁶The notation *oo* mimics ∞ .

on their ratio among plain λ -terms. We propose to count how many terms of size n are typable. We proceed as follows, we generate all the plain terms of size n and we filter those that are typable and we count them. This process is relatively tedious and is limited by the number of plain terms. Here are the numbers $T_{\infty,n}$ of typable terms of size n we computed:

n	$T_{\infty,n}$
0	0
1	0
2	1
3	1
4	2
5	2
6	3
7	5
8	8
9	13
10	22
11	36
12	58
13	103
14	177
15	307
16	535
17	949
18	1645
19	2936
20	5207

n	$T_{\infty,n}$
21	9330
22	16613
23	29921
24	53588
25	96808
26	174443
27	316267
28	572092
29	1040596
30	1888505
31	3441755
32	6268500
33	11449522
34	20902152
35	38256759
36	70004696
37	128336318
38	235302612
39	432050796
40	793513690
41	1459062947
42	2683714350

We did not go further than 42. Notice that

$$S_{\infty,42} = 7\,395\,529\,009 \quad \text{and} \quad S_{\infty,43} = 14\,023\,075\,765.$$

In words this means that there are more than seven billions of terms of size 42 and about 14 billions of terms of size 43 which we have to traverse to know how many are simply typable. We conjecture [17] that like the $S_{\infty,n}$'s, the $T_{\infty,n}$'s increase like 1.963447954^n and the ratio $T_{\infty,n}/S_{\infty,n}$ is polynomial.

4.2 Generating random λ -terms (plain and typable)

Thanks to unranking we have a very natural method for generating random λ -terms [20] (called by Duchon et al. [9] the *recursive method*). It works as follows: to generate a random λ -terms of size n one generates first a random number say k in the interval $[1..S_{\infty,n}]$ and then by unranking k one creates the associated λ -term. The size of the term we can generate this way is limited by the size of the number $S_{\infty,n}$ we can manipulate.

To generate a random typable λ -term of size n , one repeats the process of generating plain terms until one gets a typable λ -term. The process is inefficient for two reasons: one has to repeat the generation of plain terms with a lot of waste and one has to handle huge numbers. But unlike Boltzmann samplers, unranking generates random λ -terms of a precise size. However in applications like compiler testing, we seldom need to generate terms of a very precise size. An approximation on the size is enough.

5 Boltzmann samplers

In this section we present the notion of Boltzmann samplers. For more detail the reader is invited to look at [9].

A *Boltzmann sampler* returns a random object in a given class \mathcal{C} according to a given index x . By convention a Boltzmann sampler for \mathcal{C} depending on the index x is written $\Gamma C(x)$. In **Haskell** we will write it *gammaC gen x*, because we need a generator of random numbers to be passed as parameter.

Before specializing the notion of Boltzmann sampler to count specifically lambda terms of size n assume we consider a generic class \mathcal{C} of combinatorial structures, where objects of size n are counted by numbers C_n . Assume in addition that the class \mathcal{C} admits the generating function C , which means that

$$C(x) = \sum_{n=0}^{\infty} C_n x^n.$$

Uniform probability distribution assigns to each $\gamma \in \mathcal{C}_n$ the probability:

$$\mathbb{P}_{\mathcal{C}_n}(\gamma) = \frac{1}{C_n}$$

We want to generate random objects with some flexibility on the size. In other words we want the objects to be generated in some cloud around a given size n , that is so that the size N of the objects lies in some interval $(1-\varepsilon)n \leq N \leq (1+\varepsilon)n$ for some factor $\varepsilon > 0$ which we call “tolerance”. Such a method is called the *approximate-size* uniform random generation.

The *Boltzmann models* assign to any object $\gamma \in \mathcal{C}$ the following probability:

$$\mathbb{P}_{\mathcal{C},x}(\gamma) = \frac{1}{C(x)} \cdot x^{|\gamma|}.$$

The size of an object in a Boltzmann model is a random variable N . The probability of drawing an object of size n under the model of index x is

$$\mathbb{P}_{\mathcal{C}}(N = n) = \frac{C_n x^n}{C(x)}.$$

This is a probability since

$$\sum_{n \geq 0} \mathbb{P}_{\mathcal{C}}(N = n) = \frac{1}{C(x)} \sum_{n \geq 0} C_n x^n = 1$$

The random variable N has a *first moment* and a *second moment* [9]:

$$\mathbb{E}_x(N) = x \frac{C'(x)}{C(x)} \quad \mathbb{E}_x(N^2) = \frac{x^2 C''(x) + x C'(x)}{C(x)}.$$

and a *standard deviation*:

$$\begin{aligned} \sigma_{\mathcal{C}_n}(x) &= \sqrt{\mathbb{E}_x(N^2) - \mathbb{E}_x(N)^2} \\ &= \sqrt{\frac{x^2 C''(x) + x C'(x)}{C(x)} - x^2 \frac{C'(x)^2}{C(x)^2}} \end{aligned}$$

One builds a Boltzmann generator for a class \mathcal{C} according to the recursive construction of the class \mathcal{C} .

Disjoint union

If the class is built as $\mathcal{C} = \mathcal{A} + \mathcal{B}$, then $C_n = A_n + B_n$ and $C(z) = A(z) + B(z)$. The probability of the occurrence of some object is

$$\mathbb{P}_{\mathcal{C},x}(\gamma \in \mathcal{A}) = \frac{A(x)}{C(x)}, \quad \mathbb{P}_{\mathcal{C},x}(\gamma \in \mathcal{B}) = \frac{B(x)}{C(x)}$$

Assume we have defined a Haskell data type

```
data KindC = IsA | IsB
```

A generator for a Bernoulli variable is given by the Haskell function:

```
bern :: Double -> StdGen -> (KindC, StdGen)
bern x gen = let (p, g) = randomR (0, 1) gen
             in if p < A(x)/C(x) then (IsA, g)
             else (IsB, g)
```

The class \mathcal{C} is represented by the type *TypeC*:

```
data TypeC = A TypeA | B TypeB
```

where *TypeA* and *TypeB* are two types.⁷ Assume two functions *gammaA* and *gammaB*. Then

```
gammaC :: StdGen -> Double -> (C, StdGen)
gammaC gen x = let (class, g) = bern x gen
               in case class of
                 IsA -> let (a, g1) = gammaA g x
                       in (A a, g1)
                 IsB -> let (b, g1) = gammaB g x
                       in (B b, g1)
```

In what follows we will also consider not only a disjoint union of two classes $\mathcal{C} = \mathcal{A} + \mathcal{B}$, but a disjoint union of three classes like $\mathcal{C} = \mathcal{A} + \mathcal{B} + \mathcal{D}$. We call *randSel* the function which replaces *bern* in this case.

⁷It can be the case that *TypeA* and *TypeB* are *TypeC* itself or are defined recursively using *TypeC*.

Cartesian product

If the class is built as $\mathcal{C} = \mathcal{A} \times \mathcal{B}$, then the generating function satisfies $C(z) = A(z) \cdot B(z)$, since

$$C(z) = \sum_{\langle \alpha, \beta \rangle \in \mathcal{A} \times \mathcal{B}} z^{|\alpha|+|\beta|}.$$

$\gamma = \langle \alpha, \beta \rangle \in \mathcal{C}$ has the probability:

$$\mathbb{P}_{\mathcal{C},x}(\gamma) = \frac{x^{|\gamma|}}{C(x)} = \frac{x^{|\alpha|}}{A(x)} \cdot \frac{x^{|\beta|}}{B(x)}.$$

Assume we have defined a Haskell data type

```
data TypeC = Times TypeA TypeB
```

In this case the Boltzmann sampler is

```
gammaC :: Double -> StdGen -> (TypeC, StdGen)
gammaC x gen = let (a, g1) = gammaA gen x
                 (b, g2) = gammaB g1 x
                 in (g2, Times a b)
```

Sequences

\mathcal{C} can be the class of all finite sequences of elements of a class \mathcal{A} . Its associated type is:

```
data TypeC = Nil + Times TypeA TypeC
```

This means that \mathcal{C} is the solution of the equation

$$\mathcal{C} = \mathbf{1} + \mathcal{A} \times \mathcal{C}$$

where $\mathbf{1}$ denotes the empty sequence. The associated generating function is solution of $C(z) = 1 + A(z) \cdot C(z)$ which yields:

$$C(z) = \frac{1}{1 - A(z)}.$$

This means that *Nil* is selected with probability $1/C(x) = 1 - A(x)$ and a non empty sequence is selected with probability $A(x)$. A Boltzmann sampler for sequences is

```
gammaC :: Double -> StdGen -> (TypeC', StdGen)
gammaC x gen = let (p, g) = randomR (0, 1) gen
                 in if p < A(x)
                   then let (a, g1) = gammaA x g
                           (c, g2) = gammaC x g1
                           in (Times a c, g2)
                   else (Nil, g)
```

The case $x = \rho_{\mathcal{C}}$

Assume that the generating function we consider is of the form:

$$C(x) = \frac{P_C(x) - \sqrt{Q_C(x)}}{R_C(x)}$$

where $P_C(x)$, $Q_C(x)$ and $R_C(x)$ are three polynomials and where $\rho_{\mathcal{C}}$ is such that $Q_C(\rho_{\mathcal{C}}) = 0$ and where $Q_C(x) \neq 0$ and $R_C(x) \neq 0$ for $0 \leq x \leq \rho_{\mathcal{C}}$. Those properties are fulfilled for the three generating functions $B(x)$, $M(x)$ and $S_{\infty}(x)$ we have seen. Notice that

$$C(\rho_{\mathcal{C}}) = \frac{P_C(\rho_{\mathcal{C}})}{R_C(\rho_{\mathcal{C}})}$$

is finite. On another hand

$$\begin{aligned} C'(x) &= \frac{P'_C(x)}{R_C(x)} - \frac{Q'_C(x)}{2\sqrt{Q_C(x)}R_C(x)} \\ &\quad - \frac{(P_C(x) - \sqrt{Q_C(x)})R'_C(x)}{R_C(x)^2} \end{aligned}$$

shows that

$$\lim_{x \rightarrow \rho_{\mathcal{C}}} C'(x) = \infty.$$

Hence

$$\lim_{x \rightarrow \rho_{\mathcal{C}}} E_x(N) = \lim_{x \rightarrow \rho_{\mathcal{C}}} \frac{x C'(x)}{C(x)} = \infty.$$

Therefore if we choose x to be $\rho_{\mathcal{C}}$, the size of the generated structures will be distributed all over the natural numbers.

6 Boltzmann samplers for trees

6.1 Boltzmann samplers for Motzkin trees

Given a number x , the *mean value* formula (or *first moment*) $\mathbb{E}_{x, \mathcal{M}}(N) = xM'(x)/M(x)$ allows us by solving the equation $xM'(x)/M(x) = n$ to tell which value of x provides a sampler with n as mean value for the size. For instance, $xM'(x)/M(x) = 100$ returns $x_{100} = 0.33330833286456574$ it says that the sampler *gammaM gen x_{100}* (see below) will generate a random term with a mean value 100. Similarly since $x_{600} = 0.3333326388880542$, *gammaM gen x_{600}* will generate a random term with mean value 600. For a small value, say 10 we get $x_1 = 0.3308286281723805$. The standard deviation $\sigma_{\mathcal{M}_n}(x)$ can be instantiated to x_{10} or x_{100} or to x_{600} . Actually

$$\sigma_{\mathcal{M}_n}(x_{10}) \approx 25 \quad \sigma_{\mathcal{M}_n}(x_{100}) \approx 816 \quad \sigma_{\mathcal{M}_n}(x_{600}) \approx 1200.$$

Added to the fact that x_{100} and x_{600} are very close to the $\rho_{\mathcal{M}} = 1/3$ (which is called the *critical value*) it appears that a Boltzmann sampler for large values

of n is obtained by parameterizing $gammaM$ with $\rho_{\mathcal{M}}$, namely $1/3$. See [9] Section 7.2 for a mathematical justification of this choice.

Since there are three components in the addition, we generalize the Bernoulli generator and then we specialize it to the case of Motzkin trees, producing a function we call $randSelM$. Let us call

$$\frac{M_0(z)}{M(z)} = \frac{z}{M(z)}$$

the first component of $(z + zM(z) + zM(z)^2)/M(z)$,

$$\frac{M_1(z)}{M(z)} = z$$

the second component of $(z + zM(z) + zM(z)^2)/M(z)$ and

$$\frac{M_2(z)}{M(z)} = zM(z)$$

the third component of $(z + zM(z) + zM(z)^2)/M(z)$. We define the following Haskell functions:

```

m z = (1 - z - sqrt(1 - 2 * z - 3 * z * z)) / (2 * z)
p1 x = x / m x
p2 x = p1 x + x

```

One may notice that $p2 x + x * m x == 1$ for every value of x , which means that when one adds all the probabilities of the events one gets 1.

```

data KindMotzkin = IsLeaf | IsUnary | IsBinary

```

```

randSelM :: StdGen -> Double -> (KindMotzkin, StdGen)
randSelM gen x = let (p, g) = randomR (0, 1) gen
                 in if p < p1 x then (IsLeaf, g)
                 else if p < p2 x then (IsUnary, g)
                 else (IsBinary, g)

```

Then we define the data type

```

data MTree = MNil | MU MTree | MB MTree MTree

```

for Motzkin trees and we get the Boltzmann sampler for Motzkin trees:

```

gammaM :: StdGen -> Double -> (MTree, StdGen)
gammaM gen x = let (kind, g) = randSelM gen x
                 in case kind of
                 IsLeaf -> (MNil, g)
                 IsUnary -> let (t1, g1) = gammaM g x
                             in (MU t1, g1)
                 IsBinary -> let (t1, g1) = gammaM g x
                             (t2, g2) = gammaM g1 x
                             in (MB t1 t2, g2)

```

```

ceiledGammaM :: StdGen → (EMTree, Int, StdGen)
ceiledGammaM gen =
  let (kind, g0) = selOneThird gen
  in case kind of
    IsLeaf → (OK MNil, 1, g0)
    IsUnary → let (et1, n1, g1) = ceiledGammaM g0
               in case et1 of
                 OK t1 → if n1 <= upLimit
                        then (OK (MU t1), n1 + 1, g1)
                        else (Error, upLimit, g1)
                 Error → (Error, upLimit, g1)
    IsBinary → let (et1, n1, g1) = ceiledGammaM g0
                 (et2, n2, g2) = ceiledGammaM g1
               in case et1 of
                 OK t1 → case et2 of
                         OK t2 → if n1 + n2 + 1 <= upLimit
                                then (OK (MB t1 t2), n1 + n2 + 2, g2)
                                else (Error, upLimit, g2)
                         Error → (Error, upLimit, g2)
                 Error → (Error, upLimit, g2)

```

Figure 2: The function *ceiledGammaM*

An interesting feature should be noticed in this program. Indeed the function *gammaM gen x* calls *gammaM g x* and should not terminate since we can prove in no way that the first arguments namely *gen* versus *g* decrease. Thus the termination is only probabilistic and depends on the choice of *x*. The implementation of *randSelM* makes *gammaM gen x* to terminate with probability 1, thanks to the whole theory of Boltzmann samplers. For instance if the generator is buggy and takes its values only in the interval $(\frac{1}{3}, 1]$ the program will not terminate.

Samplers for large Motzkin trees

Note that $M(\rho_{\mathcal{M}}) = 1$. For large Motzkin trees the right choice is to take $x = \rho_{\mathcal{M}} = 1/3$. This gives the probability

- for drawing a leaf $\frac{\rho_{\mathcal{M}}}{M(\rho_{\mathcal{M}})} = \frac{1}{3}$,
- for drawing a tree rooted by a unary node $\rho_{\mathcal{M}} = \frac{1}{3}$,
- for drawing a tree rooted by a binary node $\rho_{\mathcal{M}} M(\rho_{\mathcal{M}}) = \frac{1}{3}$.

Actually since choosing $\rho_{\mathcal{M}}$ generates large trees, it is advisable to limit the size of the tree to be generated during the process. Therefore, we consider what Duchon et al. call *ceiled random generation*. We define a constant *upLimit* which

sets a size upper limit that the sampler should respect during the generation. We then define a data type:

```
data EMTree = OK MTree | Error
```

and a selection function *selOneThird* which selects a kind among *IsLeaf*, *IsUnary* and *IsBinary* with probability 1/3, 1/3, 1/3. The function *ceiledGammaM* is defined in Figure 2.

If we want to generate a term between say *a* and *b* we generate terms with an *upLimit* set to *b* until we get a term larger than *a*. Duchon et al. have shown that the complexity of this method is linear. In other words, the generation of a tree with a size between *a* and *b* takes a time $O(b)$.

6.2 Boltzmann samplers for binary trees

The binary trees are implemented in Haskell by

```
data BTree = BNil | BNode BTree BTree
```

If we compute for $B(z)$ the means associated to some values of n , as we did above for $M(z)$, we get:

$$\begin{aligned} x_{10} &= 0.4960783708241233 \\ x_{100} &= 0.4999739685017913 \\ x_{600} &= 0.49999930090217926 \end{aligned}$$

for which we obtain the standard deviations:

$$\sigma_{\mathcal{B}_n}(x_{10}) \approx 22 \quad \sigma_{\mathcal{B}_n}(x_{100}) \approx 970 \quad \sigma_{\mathcal{B}_n}(x_{600}) \approx 14623.$$

There are two kinds of binary trees:

- *leafs* (represented by *BNil*) which correspond to

$$\frac{B_0(z)}{B(z)} = \frac{z}{B(z)}$$

which is the first component of $(z + zB(z)^2)/B(z)$, and

- *nodes* (represented by *BNode t1 t2*) which correspond to

$$\frac{B_1(z)}{B(z)} = zB(z)$$

which is the second component of $(z + zB(z)^2)/B(z)$.

Let us define

$$bz = (1 - \text{sqrt}(1 - 4 * z * z)) / (2 * z)$$

$$q x = x / b x$$

We introduce

$$\text{data } \text{KindBTree} = \text{IsBNil} \mid \text{IsBNode}$$

and

$$\begin{aligned} \text{randSelB} &:: \text{StdGen} \rightarrow \text{Double} \rightarrow (\text{KindBTree}, \text{StdGen}) \\ \text{randSelB gen } x &= \mathbf{let} (p, g) = \text{randomR} (0, 1) \text{ gen} \\ &\quad \mathbf{in} \mathbf{if} p < q x \\ &\quad \quad \mathbf{then} (\text{IsBNil}, g) \\ &\quad \quad \mathbf{else} (\text{IsBNode}, g) \end{aligned}$$

We can then introduce the sampler:

$$\begin{aligned} \text{gammaB} &:: \text{StdGen} \rightarrow \text{Double} \rightarrow (\text{BTree}, \text{StdGen}) \\ \text{gammaB gen } x &= \mathbf{let} (kind, g) = \text{randSelB gen } x \\ &\quad \mathbf{in} \mathbf{case} \text{ kind of} \\ &\quad \quad \text{IsBNil} \rightarrow (\text{BNil}, g) \\ &\quad \quad \text{IsBNode} \rightarrow \mathbf{let} (t1, g1) = \text{gammaB } g x \\ &\quad \quad \quad (t2, g2) = \text{gammaB } g1 x \\ &\quad \quad \mathbf{in} (\text{BNode } t1 \ t2, g2) \end{aligned}$$

Once again we see that the best value for x to generate random large binary trees is $\rho_{\mathcal{B}} = \frac{1}{2}$. Thus $\text{gammaB gen } 0.5$ generates large binary trees. For a similar approach of the same problem with in particular a discussion of *ceiled random generation* see [27]. For other methods for generating random binary trees, see [15], pages 18-19 and [23].

7 Boltzmann samplers for lambda terms

Like for Motzkin trees we consider the equation of the generating function:

$$S_{\infty}(z) = \frac{z^2}{1-z} + z^2 S_{\infty}(z) + z^2 S_{\infty}(z)^2.$$

with three components: the first corresponds to de Bruijn indices, the second to abstractions, the third to applications. Like for Motzkin trees we have to build a random selector among three probabilities. First we describe a data type:

$$\mathbf{data} \text{ KindTerm} = \text{IsVariable} \mid \text{IsAbstraction} \mid \text{IsApplication}$$

a function *soo*

$$\begin{aligned} \text{soo } z &= (z^3 - z^2 - z + 1 - \text{sq}) / (2 * z * z * (1 - z)) \\ &\quad \mathbf{where} \ \text{sq} = \text{sqrt}(z^6 + 2 * (z^5) - 5 * (z^4) + 4 * (z^3) - z^2 - 2 * z + 1) \end{aligned}$$

and two functions:

$$\begin{aligned} p1\ x &= x * x / (1 - x) / so\ x \\ p2\ x &= p1\ x + x^2 \end{aligned}$$

Using Sage we computed the values:

$$\begin{aligned} x_{100} &= 0.5092252666102192 \\ x_{500} &= 0.5093048407797965 \\ x_{600} &= 0.5093058457062517 \\ x_{1000} &= 0.5093073063214039 \end{aligned}$$

which are the choices of x yielding mean values 100, 500, 600 and 1000 respectively.

General samplers of λ -terms

The values of the probabilities for a given x are

- $p_v(x) = \frac{x^2}{(1-x)S_\infty(x)}$ for variables,
- $p_{abs}(x) = x^2$ for abstractions,
- $p_{app}(x) = x^2 S_\infty(x)$ for applications.

For selecting among *IsVariable*, *IsAbstraction* and *IsApplication* we use the following selector written in Haskell

```
randSel :: StdGen -> Double -> (KindTerm, StdGen)
randSel gen x = let (p, g) = randomR (0, 1) gen
                  in if p < pv x then (IsVariable, g)
                     else if p < pv x + pabs x
                          then (IsAbstraction, g)
                          else (IsApplication, g)
```

Samplers for large λ -terms

Like for trees, the best choice of x for generating large λ -terms is ρ which we call *rho* in Haskell:

```
rho :: Double
rho = 0.5093081270242373.
```

whose square is

```
rhosquare = rho * rho
```

which yields $\rho^2 = 0.25939476825293667$. Notice that since ρ is a root of the polynomial below the square root, $S_\infty(\rho) = \frac{1-\rho^2}{2\rho^2}$. The values of the probabilities for selecting among variables, abstractions and applications are:

- $p_v(\rho) = \frac{2\rho^4}{(1-\rho)(1-\rho^2)}$ for variables,
- $p_{abs}(\rho) = \rho^2$ for abstractions,
- $p_{app}(\rho) = \frac{1-\rho^2}{2}$ for applications.

Let us simplify $\frac{2\rho^4}{(1-\rho)(1-\rho^2)}$ into $\frac{1-\rho^2}{2}$ by computing the difference:

$$\begin{aligned} \frac{2\rho^4}{(1-\rho)(1-\rho^2)} - \frac{1-\rho^2}{2} &= \frac{4\rho^4 - (1-\rho^2)^2(1-\rho)}{2(1-\rho)(1-\rho^2)} \\ &= \frac{\rho^5 + 3\rho^4 - 2\rho^3 + 2\rho^2 + \rho - 1}{2(1-\rho)(1-\rho^2)} \\ &= 0. \end{aligned}$$

Therefore we get the result announced in the introduction namely that

- $p_v(\rho) = \frac{1-\rho^2}{2} \approx 0.3703026$ for variables,
- $p_{abs}(\rho) = \rho^2 \approx 0.25939476$ for abstractions,
- $p_{app}(\rho) = \frac{1-\rho^2}{2} \approx 0.3703026$ for applications.

We build the random selector:

```
pvrho = (1 - rhosquare)
p2rho = pvrho + rhosquare
```

```
randSelRho :: StdGen -> (KindTerm, StdGen)
randSelRho gen = let (p, g) = randomR (0, 1) gen
                  in if p < pvrho then (IsVariable, g)
                  else if p < p2rho
                     then (IsAbstraction, g)
                     else (IsApplication, g)
```

We generate a Boltzmann sampler for indices (or variables). A de Bruijn index is a natural number i.e., a sequence of 1's. Hence we take the Boltzmann sampler for sequences with $A(z) = z$ and $A(\rho) = \rho$.

```
gammaV :: StdGen -> (Integer, StdGen)
gammaV gen = let (a, g) = randomR (0, 1) gen
              in if a < rho
                 then let (n, g1) = gammaV g
                        in (n + 1, g1)
                 else (1, g)
```

We want to generate terms that are below a certain *upLimit*. Thus when the limit is passed we generate an error. For that we create a type with error.

```
data ETerm = OK Term | Error
```

This method recalls the **sized** generation of [5]. A Boltzmann sampler for large λ -terms which is given:

```

ceiledGammaSoo :: StdGen → (ETerm, Int, StdGen)
ceiledGammaSoo gen =
  let (kind, g0) = randSelRho gen
  in case kind of
    IsVariable → let (n, g1) = gammaV g0
      in if fromIntegral n <= upLimit
        then (OK (Index n), fromIntegral (n + 1), g1)
        else (Error, upLimit, g1)
    IsAbstraction → let (et1, n1, g1) = ceiledGammaSoo g0
      in case et1 of
        OK t1 → if n1 + 2 <= upLimit
          then (OK (Abs t1), n1 + 2, g1)
          else (Error, upLimit, g1)
        Error → (Error, upLimit, g1)
    IsApplication → let (et1, n1, g1) = ceiledGammaSoo g0
      (et2, n2, g2) = ceiledGammaSoo g1
      in case et1 of
        OK t1 → case et2 of
          OK t2 → if n1 + n2 + 2 <= upLimit
            then (OK (App t1 t2), n1 + n2 + 2, g2)
            else (Error, upLimit, g2)
          Error → (Error, upLimit, g2)
        Error → (Error, upLimit, g2)

```

To generate a large plain λ -term, terms are filtered until a term that is large enough is generated. Recall that the method is linear in time complexity. Thus the generation of a term of size 100,000 takes a few seconds, the generation of a term of size one million takes three minutes and the generation of a term of size five millions takes five minutes on a laptop.

To generate large typable λ -terms we generate λ -terms and check their typability. Currently we are able to generate random typable λ -terms of size 500.

8 Ranking and unranking closed lambda terms

Boltzmann samplers work only when one knows the generating function explicitly. Thus if one wants to generate typable closed terms, one has two solutions. First one generates a random typable term using the Boltzmann sampler described above and if the produced term is not closed, one closes it by adding the adequate number of abstractions. Second one uses an unranking method similar to this described in Section 4.2. For that let us recall how one counts terms with at most m free indices [13]. The specific case of the closed terms is the case when the number of free indices is at most 0.

$$S_{m,0} = S_{m,1} = 0,$$

$$S_{m,n+2} = [m \geq n + 1] + S_{m+1,n} + \sum_{k=0}^n S_{m,k} S_{m,n-k}.$$

Given a predicate P , $[P(\vec{x})]$ denotes the Iverson symbol, i.e., $[P(\vec{x})] = 1$ if $P(\vec{x})$ and $[P(\vec{x})] = 0$ if $\neg P(\vec{x})$. Figure 3 gives a program for unranking terms with at most m indices, where $tromp\ m\ n$ is the Haskell function representing $S_{m,n}$.

```

unrankT :: Int -> Int -> Integer -> Term
unrankT m n k
  | m >= n - 1 && k == (tromp m n) = Index $ fromIntegral (n - 1)  -- terms 1n-10
  | k <= (tromp (m + 1) (n - 2)) = Abs (unrankT (m + 1) (n - 2) k)  -- terms 00M
  | otherwise = unrankApp (n - 2) 0 (k - tromp (m + 1) (n - 2))  -- terms 01MN
  where unrankApp n j h
        | h <= tmjtmnj = let (dv, rm) = (h - 1) 'divMod' tmnj
                          in App (unrankT m j (dv + 1)) (unrankT m (n - j) (rm + 1))
        | otherwise = unrankApp n (j + 1) (h - tmjtmnj)
        where tmnj = tromp m (n - j)
              tmjtmnj = (tromp m j) * tmnj

```

Figure 3: *unrankT* function in Haskell.

$S_{0,n}$ counts closed terms and is computed by $tromp\ 0\ n$.

9 Related works

In the introduction we cited papers that are clearly connected to this work. In a recent work, Bacher et al. [1] propose an improved random generation of binary trees and Motzkin trees, based on Rémy algorithm [23] (or algorithm R as Knuth calls it [15]). Instead of growing the trees from the root, they propose like Rémy to grow the trees from inside by an operation called grafting. It is not clear how this can be generalized to λ -terms as this requires “to find a combinatorial interpretation for the holonomic equations [which] is not [...] always possible, and even for simple combinatorial objects this is not elementary” (Conclusion of [1] page 16).

We would like also to mention papers on counting λ -terms [18, 12] and evaluating their combinatorial proprieties namely [2, 6, 3, 4]. Another related paper is [10] which proposes Haskell programs for enumerating structures.

10 Acknowledgments

We would like to thank Katarzyna Grygiel, Bruno Salvy and John Tromp, for fruitful discussions. We are also indebted to Neil Sloane and the developers of

the *On-line Encyclopedia of Integer Sequences* which allowed us to know about sequence **A114851** which was the starting point of this research.

11 Conclusion

Boltzmann samplers are central tools for the uniform generation of random structures, like trees or λ -terms. Two directions are now open for applications: first to integrate the programs proposed here in actual testers and optimizers and second to extend Boltzmann samplers to other kinds of programs, for instance programs with block structures. From the theoretical point of view, more should be known about generating functions for *closed λ -terms* or *λ -terms with fixed bound on the number of variables*. Boltzmann samplers should be designed for such terms which requires to extend the theory. Concerning combinatorial properties of *simply typable λ -terms* many question are left open and seem to be hard. Besides since we are interested in generating typable terms, it could be worth to build random uniform samplers delivering directly typable terms, for instance based on a hashing table built offline.

References

- [1] A. Bacher, O. Bodini, and A. Jacquot. Efficient random sampling of binary and unary-binary trees via holonomic equations. *CoRR*, abs/1401.1140, 2014.
- [2] O. Bodini, D. Gardy, and B. Gittenberger. Lambda-terms of bounded unary height. *2011 Proceedings of the Eighth Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*, 2011.
- [3] O. Bodini, D. Gardy, B. Gittenberger, and A. Jacquot. Enumeration of generalized BCI lambda-terms. *Electr. J. Comb.*, 20(4):P30, 2013.
- [4] O. Bodini, D. Gardy, and A. Jacquot. Asymptotics and random sampling for BCI and BCK lambda terms. *Theor. Comput. Sci.*, 502:227–238, 2013.
- [5] K. Claessen and J. Hughes. QuickCheck: a lightweight tool for random testing of Haskell programs. In M. Odersky and P. Wadler, editors, *ICFP*, pages 268–279. ACM, 2000.
- [6] R. David, K. Grygiel, J. Kozik, C. Raffalli, G. Theyssier, and M. Zaionc. Asymptotically almost all λ -terms are strongly normalizing. *Logical Methods in Computer Science*, 9(1:02):1–30, 2013.
- [7] N. G. de Bruijn. *Asymptotic Methods in Analysis*. North-Holland, Amsterdam, 1958.
- [8] N. G. de Bruijn. Lambda calculus with nameless dummies, a tool for automatic formula manipulation, with application to the Church-Rosser

- theorem. *Proc. Koninkl. Nederl. Akademie van Wetenschappen*, 75(5):381–392, 1972.
- [9] P. Duchon, P. Flajolet, G. Louchard, and G. Schaeffer. Boltzmann samplers for the random generation of combinatorial structures. *Combinatorics, Probability & Computing*, 13(4-5):577–625, 2004.
- [10] J. Duregård, P. Jansson, and M. Wang. Feat: functional enumeration of algebraic types. In J. Voigtländer, editor, *Proceedings of the 5th ACM SIGPLAN Symposium on Haskell, Haskell 2012, Copenhagen, Denmark, 13 September 2012*, pages 61–72. ACM, 2012.
- [11] P. Flajolet and R. Sedgewick. *Analytic Combinatorics*. Cambridge University Press, 2008.
- [12] K. Grygiel and P. Lescanne. Counting and generating lambda terms. *J. Funct. Program.*, 23(5):594–628, 2013.
- [13] K. Grygiel and P. Lescanne. Counting terms in the binary lambda calculus. *CoRR*, abs/1401.0379, 2014.
- [14] D. E. Knuth. *Selected Papers on Analysis of Algorithms*, volume 102 of *CSLI Lecture Notes*. Stanford, California: Center for the Study of Language and Information, 2000.
- [15] D. E. Knuth. *The Art of Computer Programming, Volume 4, Fascicle 4: Generating All Trees, History of Combinatorial Generation (Art of Computer Programming)*. Addison-Wesley Professional, 2006.
- [16] P. Lescanne. From $\lambda\sigma$ to $\lambda\nu$, a journey through calculi of explicit substitutions. In H. Boehm, editor, *Proceedings of the 21st Annual ACM Symposium on Principles Of Programming Languages, Portland (Or., USA)*, pages 60–69. ACM, 1994.
- [17] P. Lescanne. An exercise on streams: convergence acceleration. *CoRR*, abs/1312.4917, 2013.
- [18] P. Lescanne. On counting untyped lambda terms. *Theor. Comput. Sci.*, 474:80–97, 2013.
- [19] M. Li and P. Vitányi. *An introduction to Kolmogorov complexity and its applications (2nd ed.)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1997.
- [20] A. Nijenhuis and H. S. Wilf. *Combinatorial algorithms for computers and calculators*. Computer science and applied mathematics. Academic Press, New York, 1978. First ed. published in 1975 under title: Combinatorial algorithms.

- [21] M. Pałka. Testing an optimising compiler by generating random lambda terms. Licentiatavhandling, Department of Computer Science and Engineering, Chalmers University of Technology and Göteborg University, may 2012.
- [22] M. H. Pałka, K. Claessen, A. Russo, and J. Hughes. Testing an optimising compiler by generating random lambda terms. In *Proceedings of the 6th International Workshop on Automation of Software Test, AST '11*, pages 91–97, New York, NY, USA, 2011. ACM.
- [23] J.-L. Rémy. Un procédé itératif de dénombrement d’arbres binaires et son application à leur génération aléatoire. *ITA*, 19(2):179–195, 1985.
- [24] W. Stein et al. *Sage Mathematics Software (Version 5.8)*. The Sage Development Team, 2012. <http://www.sagemath.org>.
- [25] J. Tromp. Binary lambda calculus and combinatory logic. In M. Hutter, W. Merkle, and P. M. B. Vitányi, editors, *Kolmogorov Complexity and Applications*, volume 06051 of *Dagstuhl Seminar Proceedings*. Internationales Begegnungs- und Forschungszentrum fuer Informatik (IBFI), Schloss Dagstuhl, Germany, 2006.
- [26] X. Yang, Y. Chen, E. Eide, and J. Regehr. Finding and understanding bugs in C compilers. In M. W. Hall and D. A. Padua, editors, *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2011, San Jose, CA, USA, June 4-8, 2011*, pages 283–294. ACM, 2011.
- [27] B. Yorgey. Random binary trees with a size-limited critical Boltzmann sampler. <http://byorgey.wordpress.com/2013/04/25/>, April 2013.

Here is the whole quotation of Claessen and Hughes [5]:

We have chosen to put distribution under the human tester’s control, by defining a test data generation language (also embedded in Haskell), and a way to observe the distribution of test cases. By programming a suitable generator, the tester can not only control the distribution of test cases, but also ensure that they satisfy arbitrarily complex invariants.