


Linear-time exact sampling of sum-constrained random variables



Andrea Sportiello and Frédérique Bassino
CNRS and LIPN, Université Paris Nord, Villetaneuse


GASCom 2018, Athens, Greece
June 19th 2018



The problem

$$\Omega_n = \{\mathbf{x}\} \quad \mathbf{x} = (x_1, \dots, x_n), \quad |\mathbf{x}| := \sum_i x_i, \quad x_i \in \mathbb{N}$$

$$\mu_{n,m}(\mathbf{x}) = \frac{1}{Z} \prod_{i=1}^n f_i(x_i) \times \delta_{|\mathbf{x}|,m}$$




The problem

$$\Omega_n = \{\mathbf{x}\} \quad \mathbf{x} = (x_1, \dots, x_n), \quad |\mathbf{x}| := \sum_i x_i, \quad x_i \in \mathbb{N}$$

$$\mu_{n,m}(\mathbf{x}) = \frac{1}{Z} \prod_{i=1}^n f_i(x_i) \times \delta_{|\mathbf{x}|,m}$$

Problem: Assume that sampling from each distrib. f_i costs $\mathcal{O}(1)$. Find an algorithm that samples from the distribution $\mu_{n,m}$ in average linear time.




The problem

$$\Omega_n = \{\mathbf{x}\} \quad \mathbf{x} = (x_1, \dots, x_n), \quad |\mathbf{x}| := \sum_i x_i, \quad x_i \in \mathbb{N} \quad \leftarrow \text{random vector of integers}$$

$$\mu_{n,m}(\mathbf{x}) = \frac{1}{Z} \prod_{i=1}^n f_i(x_i) \times \delta_{|\mathbf{x}|,m} \quad \leftarrow \text{NOT completely independent (a single linear constraint)}$$

↑
variables are NOT
identically distributed

Problem: Assume that sampling from each distrib. f_i costs $\mathcal{O}(1)$. Find an algorithm that samples from the distribution $\mu_{n,m}$ in average linear time.




The problem

$\Omega_n = \{\mathbf{x}\} \quad \mathbf{x} = (x_1, \dots, x_n), \quad |\mathbf{x}| := \sum_i x_i, \quad x_i \in \mathbb{N} \quad \leftarrow \text{random vector of integers}$

$\mu_{n,m}(\mathbf{x}) = \frac{1}{Z} \prod_{i=1}^n f_i(x_i) \times \delta_{|\mathbf{x}|,m} \quad \leftarrow \text{completely independent: the problem trivialises!}$

\uparrow variables are NOT identically distributed

Problem: Assume that sampling from each distrib. f_i costs $\mathcal{O}(1)$. Find an algorithm that samples from the distribution $\mu_{n,m}$ in average linear time.




The problem

$$\Omega_n = \{\mathbf{x}\} \quad \mathbf{x} = (x_1, \dots, x_n), \quad |\mathbf{x}| := \sum_i x_i, \quad x_i \in \mathbb{N} \quad \leftarrow \text{random vector of integers}$$

$$\mu_{n,m}(\mathbf{x}) = \frac{1}{Z} \prod_{i=1}^n f_{\mathbf{x}}(x_i) \times \delta_{|\mathbf{x}|,m} \quad \leftarrow \text{NOT completely independent (a single linear constraint)}$$

↑
variables are identically distributed:
doable by using permutation symmetry [L. Devroye, 2012]

Problem: Assume that sampling from each distrib. f_i costs $\mathcal{O}(1)$. Find an algorithm that samples from the distribution $\mu_{n,m}$ in average linear time.




The problem

$$\Omega_n = \{\mathbf{x}\} \quad \mathbf{x} = (x_1, \dots, x_n), \quad |\mathbf{x}| := \sum_i x_i, \quad x_i \in \mathbb{N} \quad \leftarrow \text{random vector of integers}$$

$$\mu_{n,m}(\mathbf{x}) = \frac{1}{Z} \prod_{i=1}^n f_i(x_i) \times \delta_{|\mathbf{x}|,m} \quad \leftarrow \text{NOT completely independent (a single linear constraint)}$$

↑
variables are NOT
identically distributed

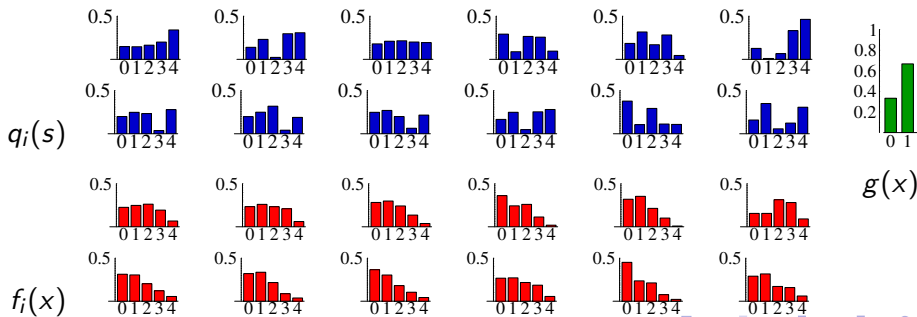
Problem: Assume that sampling from each distrib. f_i costs $\mathcal{O}(1)$. Find an algorithm that samples from the distribution $\mu_{n,m}$ in average linear time.



Our solution

$$\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{N}^n, \quad \mu_{n,m}(\mathbf{x}) = \frac{1}{Z} \prod_{i=1}^n f_i(x_i) \times \delta_{|\mathbf{x}|,m}$$

Our solution: **positive decomposition**. Assume that there exists $g(x) \in \{\text{Bern}_b, \text{Poiss}, \text{Geom}_b\}$, and $\{q_i(s)\}_{1 \leq i \leq n; s \in \mathbb{N}}$ real positive, such that $f_i(x) = \sum_s q_i(s) g^{*s}(x)$. Then our new algorithm does it!



The new algorithm in a nutshell

Our new trick is based on the following ideas:

- ▶ Rejection algorithms have an extra factor in their complexity, on the scale of the inverse of the acceptance rate. In order to have the optimal complexity scaling, **you need the average acceptance rate not to scale with the size n .**
- ▶ **Positive decomposition** gives $f_i(x) = \sum_s q_i(s)g^{*s}(x)$.
As a result the measure $\mu_{n,m}(\mathbf{x}) = \frac{1}{Z} \prod_{i=1}^n f_i(x_i) \times \delta_{|\mathbf{x}|,m}$ is a marginal of a measure **in two sets of variables**:
$$\mu_{n,m}(\mathbf{x}, \mathbf{s}) = \frac{1}{Z} \prod_{i=1}^n (q_i(s_i) g^{*s_i}(x_i)) \times \delta_{|\mathbf{x}|,m}.$$
- ▶ You can **first sample \mathbf{s}** , with measure $\mu_1(\mathbf{s}) = \prod_{i=1}^n q_i(s_i)$, then **accept this vector \mathbf{s}** with rate $a(\mathbf{s}) \propto g^{*|\mathbf{s}|}(m)$, and finally **sample \mathbf{x}** with measure $\mu_2(\mathbf{x} | \mathbf{s}) = \prod_{i=1}^n g^{*s_i}(x_i)$.
- ▶ **The acceptance rate is high because**, although $g^{*|\mathbf{s}|}(m) = \mathcal{O}(n^{-\frac{1}{2}})$, we have $g^{*|\mathbf{s}|}(m) / \max_N(g^{*N}(m)) = \Theta(1)$.

We need something different...

Devroye uses the **permutation symmetry** of the random variables, and samples first the total number of x_i 's equal to any given y (with a rejection scheme), then their reordering.

This is troublesome (use of float approximations for multinomial coefficients, need for extra tricks if the f_i 's do not have finite support, ...), and, most importantly, **cannot be done here**, as the variables are **not** identically distributed...

The obvious rejection scheme

A first algorithm is obtained by neglecting the linear global constraint. Assume (as usual) that $m = \sum_i \mathbb{E}[f_i]$ and $\sigma^2 := \sum_i \text{Var}[f_i]$ are $\Theta(n)$. Up to a Lagrange multiplier, we can assume w.l.o.g. that $\mathbb{E}(|\mathbf{x}|) = m$.

This 'Boltzmann sampling' rejection algorithm would give:

Algorithm : Naïve rejection sampling	<i>complexity</i> $\sim n^{3/2}$
---	----------------------------------

```
begin
|   repeat
|   |   |x| = 0;
|   |   for i ← 1 to n do
|   |   |   x_i ← f_i; |x| += x_i;
|   |   |   ← complexity ~ n;
|   |   until |x| = m
|   |   return (x_1, ..., x_n)
|   |   ← complexity ~ √n;
end
```

Shannon complexity bound

We have seen that the naïve algorithm has complexity $\sim n^{\frac{3}{2}}$.

This seems bad. But how bad exactly? How good can we possibly do?

Let us try to determine the **intrinsic minimal complexity** of this problem.

As we have seen in Olivier's talk, the **time complexity** is defined only up to a multiplicative constant, while for the **random-bit complexity** also the overall constants do matter.

Of course, the second one is a lower bound to the first one.

The intrinsic minimal random-bit complexity of an exact sampling problem is given by the Shannon entropy of the associated measure:

$$S[\mu] = - \sum_{\mathbf{x} \in \Omega_n} \mu(\mathbf{x}) \ln \mu(\mathbf{x})$$

Shannon complexity bound

Simple fact 1: if $\mathbf{x} = (x_1, \dots, x_n)$ and $\mu_{\text{iid}}(\mathbf{x}) = f_1(x_1) \cdots f_n(x_n)$,

$$S[\mu_{\text{iid}}] = \sum_i S[f_i]$$

Simple fact 2: if also $p(s) = \mathbb{P}(|\mathbf{x}| = s)$ and $\mu_s(\mathbf{x}) = \frac{\mu_{\text{iid}}(\mathbf{x})}{p(s)} \cdot \delta_{|\mathbf{x}|,s}$,

$$S[\mu_{\text{iid}}] - S[p] = \sum_s p(s) S[\mu_s] = \mathbb{E}(S[\mu_s])$$

A bit more subtle: Suppose that $\text{Var}[p] = \Theta(n)$, and p has the same value for mode and average, $s^* = \mathbb{E}_p(s) = \text{argmax}(p(s))$.

Then $\mathbb{E}(S[\mu_s]) \leq S[\mu_{s^*}]$, and in turns

$$S[\mu_{\text{iid}}] - S[p] \leq S[\mu_{s^*}] \leq S[\mu_{\text{iid}}]$$

so that

$$S[\mu_{s^*}] = \sum_i S[f_i] + \Theta(\ln n)$$

Shannon complexity bound

Simple fact 1: if $\mathbf{x} = (x_1, \dots, x_n)$ and $\mu_{\text{iid}}(\mathbf{x}) = f_1(x_1) \cdots f_n(x_n)$,

$$S[\mu_{\text{iid}}] = \sum_i S[f_i]$$

Simple fact 2: if also $p(s) = \mathbb{P}(|\mathbf{x}| = s)$ and $\mu_s(\mathbf{x}) = \frac{\mu_{\text{iid}}(\mathbf{x})}{p(s)} \cdot \delta_{|\mathbf{x}|,s}$,

$$S[\mu_{\text{iid}}] - S[p] = \sum_s p(s) S[\mu_s] = \mathbb{E}(S[\mu_s])$$

A bit more subtle: Suppose that $\text{Var}[p] = \Theta(n)$, and p has the same value for mode and average, $s^* = \mathbb{E}_p(s) = \text{argmax}(p(s))$.

Then $\mathbb{E}(S[\mu_s]) \leq S[\mu_{s^*}]$, and in turns

$$S[\mu_{\text{iid}}] - S[p] \leq S[\mu_{s^*}] \leq S[\mu_{\text{iid}}]$$

so that

$$S[\mu_{s^*}] = \sum_i S[f_i] \uparrow \Theta(\ln n)$$

our complexity
goal is linear

Do we know cases in which linearity is achievable
by a conceptually-simple algorithm?

Do we know cases in which linearity is achievable
by a conceptually-simple algorithm?

Yes! You just saw this in Olivier's talk! Call this the **BBHL algorithm**.

The problem of uniformly sampling strings in $\{\bullet, \circ\}^n$
with $\#\{\bullet\} = k$ and $\#\{\circ\} = n - k$ is solved by the BBHL algo,
with **linear** time complexity and **optimal** random-bit complexity.

Do we know cases in which linearity is achievable
by a conceptually-simple algorithm?

Yes! You just saw this in Olivier's talk! Call this the **BBHL algorithm**.

The problem of uniformly sampling strings in $\{\bullet, \circ\}^n$
with $\#\{\bullet\} = k$ and $\#\{\circ\} = n - k$ is solved by the BBHL algo,
with **linear** time complexity and **optimal** random-bit complexity.

- ▶ $\text{BBHL}[n, m]$ solves the problem for

$$\mu_{n,m}(\mathbf{x}) = \frac{1}{Z} \prod_i \text{Bern}_b(x_i) \times \delta_{|\mathbf{x}|,m}, \text{ with } b = \frac{m}{n} \in]0, 1[$$

- ▶ $\text{BBHL}[m + n - 1, m]$ solves the problem for

$$\mu_{n,m}(\mathbf{x}) = \frac{1}{Z} \prod_i \text{Geom}_b(x_i) \times \delta_{|\mathbf{x}|,m}, \text{ with } b = \frac{m}{n} \in]0, +\infty[$$

Very good, but this is only for **i.i.d. cases**...

Algorithm : BBHL shuffling algorithm

begin

$a = k, b = n - k, i = 0;$

repeat

$i++;$

$\nu_i \leftarrow \text{Bern}_\beta;$

if $\nu_i = 1$ **then** $a --$ **else** $b --$

until $a < 0$ **or** $b < 0$

complexity $\sim n;$

if $a < 0$ **then** $\bar{\nu} = 0$ **else** $\bar{\nu} = 1;$

for $j \leftarrow i$ **to** n **do**

$\nu_j = \bar{\nu};$

$h \leftarrow \text{RndInt}_j;$

 swap ν_j and ν_h

complexity $\sim \sqrt{n} \ln n;$

return ν

end

The rejection paradigm

Recall the naïve rejection algorithm:

You want to do exact sampling for $\mu(\mathbf{x})$,
when $\mu(\mathbf{x}) \propto \mu_0(\mathbf{x})a(\mathbf{x})$, with $a(\mathbf{x}) \in [0, 1]$,
supposing that you know how to sample from μ_0

Algorithm : Rejection sampling

$$T[\mu] \sim T[\mu_0] \mathbb{E}(a(\mathbf{x}))^{-1}$$

begin

repeat

$\mathbf{x} \leftarrow \mu_0$;

\leftarrow complexity $T[\mu_0]$;

$\alpha \leftarrow \text{Bern}_{a(\mathbf{x})}$;

until $\alpha = 1$

\leftarrow complexity $\mathbb{E}(a(\mathbf{x}))^{-1}$;

return \mathbf{x}

end

The rejection paradigm for decomposed measures

Now assume $\mu(\mathbf{x}) \propto \sum_{\mathbf{y}} \mu_1(\mathbf{y}) \mu_2(\mathbf{x} | \mathbf{y}) a(\mathbf{y})$, with $a(\mathbf{y}) \in [0, 1]$,
supposing that you know how to sample from μ_1 , and $\mu_2(\cdot | \mathbf{y})$

Algorithm : Rejection sampling for decomposed measures

begin

repeat

$\mathbf{y} \leftarrow \mu_1$;

 ← sample a tentative \mathbf{y} with μ_1 ;

$\alpha \leftarrow \text{Bern}_{a(\mathbf{y})}$;

until $\alpha = 1$

 ← accept \mathbf{y} with rate $a(\mathbf{y})$;

$\mathbf{x} \leftarrow \mu_2(\cdot | \mathbf{y})$;

 ← sample \mathbf{x} with $\mu_2(\cdot | \mathbf{y})$;

return \mathbf{x}

end

$$T = \frac{\sum_{\mathbf{y}} \mu_1(\mathbf{y}) (T_1(\mathbf{y}) + a(\mathbf{y}) T_2(\mathbf{y}))}{\sum_{\mathbf{y}} \mu_1(\mathbf{y}) a(\mathbf{y})} = \frac{\mathbb{E}(T_1 + a T_2)}{\mathbb{E}(a)} \leq \frac{T_1^{\max}}{\mathbb{E}(a)} + T_2^{\max}$$

The rejection paradigm for decomposed measures

Now assume $\mu(\mathbf{x}) \propto \sum_{\mathbf{y}} \mu_1(\mathbf{y}) \mu_2(\mathbf{x} | \mathbf{y}) a(\mathbf{y})$, with $a(\mathbf{y}) \in [0, 1]$,
supposing that you know how to sample from μ_1 , and $\mu_2(\cdot | \mathbf{y})$

Algorithm : Rejection sampling for decomposed measures

begin

repeat

$\mathbf{y} \leftarrow \mu_1$;

\leftarrow complexity $T_1(\mathbf{y})$;

$\alpha \leftarrow \text{Bern}_{a(\mathbf{y})}$;

until $\alpha = 1$

\leftarrow complexity $a(\mathbf{y})^{-1}$;

$\mathbf{x} \leftarrow \mu_2(\cdot | \mathbf{y})$;

\leftarrow complexity $T_2(\mathbf{y})$;

return \mathbf{x}

end

$$T = \frac{\sum_{\mathbf{y}} \mu_1(\mathbf{y}) (T_1(\mathbf{y}) + a(\mathbf{y}) T_2(\mathbf{y}))}{\sum_{\mathbf{y}} \mu_1(\mathbf{y}) a(\mathbf{y})} = \frac{\mathbb{E}(T_1 + a T_2)}{\mathbb{E}(a)} \leq \frac{T_1^{\max}}{\mathbb{E}(a)} + T_2^{\max}$$

Positive decomposition provides a decomposed measure

Positive decomposition tells that, for all i ,
$$f_i(x) = \sum_s q_i(s) g^{*s}(x), \text{ with } q_i(s) \geq 0.$$

From the normalisation of the f_i 's and of g ,
it follows that also the $q_i(s)$ are probability distributions.

As a result the measure $\mu_{n,m}(x) = \frac{1}{Z} \prod_{i=1}^n f_i(x_i) \times \delta_{|x|,m}$
is a marginal of a measure in two sets of variables:

$$\mu_{n,m}(x, s) = \frac{1}{Z} \prod_{i=1}^n (q_i(s_i) g^{*s_i}(x_i)) \times \delta_{|x|,m}.$$

This is exactly as in a decomposed measure, with correspondence

sample \mathbf{s}	with measure $\mu_1(\mathbf{s})$		$\mu_1(\mathbf{s}) = \prod_{i=1}^n q_i(s_i)$
accept \mathbf{s}	with rate $a(\mathbf{s})$		$a(\mathbf{s}) \propto g^{* \mathbf{s} }(m)$
sample \mathbf{x}	with measure $\mu_2(\mathbf{x} \mathbf{s})$		$\mu_2(\mathbf{x} \mathbf{s}) = \prod_{i=1}^n g^{*s_i}(x_i)$

Note: although μ_1 and μ_2 depend on the vector \mathbf{s} ,
the rate a only depends on $|\mathbf{s}| = \sum_i s_i$.

Increasing the acceptance rate

The crucial point is that the decomposition allows to **increase the acceptance rate!**

In the 'ordinary' rejection scheme, you accept \mathbf{x} iff a probabilistic event occurs (in our case, $|\mathbf{x}| = m$). If this probability is **intrinsically small** (in our case, $\Theta(n^{-1/2})$), there is nothing you can do.

In the rejection scheme for decomposed measures, the rate $a(\mathbf{s})$ is defined **up to a multiplicative factor**, as long as $\max_{\mathbf{s}} a(\mathbf{s}) \leq 1$.

Here, the obvious choice for $a(\mathbf{s})$ is $a(\mathbf{s}) = g^{*|\mathbf{s}|}(m)$, which is $\mathcal{O}(n^{-\frac{1}{2}})$.

However, we can push it up to $a(\mathbf{s}) = \frac{g^{*|\mathbf{s}|}(m)}{\max_{\mathbf{N}}(g^{*\mathbf{N}}(m))}$.

As we will see, with this choice $\mathbb{E}(a(\mathbf{s})) = \Theta(1)$.

How to sample from $\text{Bern}_{a(\mathbf{s})}$

This idea is not sufficient by itself. Even if you know in advance that, after maximisation, $\mathbb{E}(a(\mathbf{s})) = \Theta(1)$, you still have a problem:

sampling a Bernoulli rnd var with parameter $a(\mathbf{s})$ is difficult if you do not have an analytic expression for $a(\mathbf{s})$.

It is not compulsory to have an analytic expression for $a(\mathbf{s})$ (just think to how the Monte Carlo algorithm:
 $x \leftarrow \text{Rnd}[0, 1]$; $y \leftarrow \text{Rnd}[0, 1]$; **return** $\text{sign}(1 - x^2 - y^2)$
samples $\text{Bern}_{\pi/4}$ without knowing $\pi \dots$)

however, it makes life easier, and in our case **we have it for free** if we choose the base function $g(x)$ for positive decomposition in the list $g(x) \in \{\text{Bern}_b, \text{Poiss}, \text{Geom}_b\}$

How to sample from $\text{Bern}_{a(s)}$

Example with Bernoulli (the other cases are similar)
(just write $a(s)$ for $a(\mathbf{s})$, with $s = |\mathbf{s}|$)

$$a(s) = \frac{g^{*s}(m)}{\max_N(g^{*N}(m))} = \frac{b^m(1-b)^{s-m} \binom{s}{m}}{\max_N(b^m(1-b)^{N-m} \binom{N}{m})}$$

The max is realised for $N = \bar{N} := \lfloor m/b \rfloor$, thus

$$a(s) = (1-b)^{s-\bar{N}} \frac{s!(\bar{N}-m)!}{\bar{N}!(s-m)!}$$

Good news 1: This is easily evaluated to high precision (i.e., calculating d binary digits has complexity $\ll 2^d$), so that the average cost of $\text{Bern}_{a(s)}$ is $\Theta(1)$.

Good news 2: For large m , and $b = \Theta(1)$, $a(s)$ converges to an un-normalised Gaussian centered around \bar{N} , and of variance $\Theta(m)$.

A rough evaluation of the complexity

Recall the basic steps in the rejection algo for our decomposed measure:

sample \mathbf{s}	with measure $\mu_1(\mathbf{s})$		$\mu_1(\mathbf{s}) = \prod_{i=1}^n q_i(s_i)$
accept \mathbf{s}	with rate $a(\mathbf{s})$		$a(\mathbf{s}) = g^{*\mathbf{s}}(m)/g^{*\bar{N}}(m)$
sample \mathbf{x}	with measure $\mu_2(\mathbf{x} \mathbf{s})$		$\mu_2(\mathbf{x} \mathbf{s}) = \prod_{i=1}^n g^{*s_i}(x_i)$

and that this algorithm has complexity

$$T \leq \frac{T_1^{\max}}{\mathbb{E}_{\mu_1}(a(\mathbf{s}))} + T_2^{\max} \quad \text{where } T_1^{\max}, T_2^{\max} = \Theta(n).$$

Under mild CLT hypotheses, the measure on $s = |\mathbf{s}|$ induced by $\mu_1(\mathbf{s})$ is a (normalised) Gaussian centered in \bar{N} , with variance $\sigma_1^2 n$, while $a(\mathbf{s})$ is an un-normalised Gaussian, centered in \bar{N} , with variance $\sigma_2^2 n$:

$$\mathbb{E}(a) \simeq \int dx \frac{1}{\sqrt{2\pi\sigma_1^2 n}} \exp\left[-\frac{x^2}{2n} \left(\frac{1}{\sigma_1^2} + \frac{1}{\sigma_2^2}\right)\right] = \frac{\sigma_2}{\sqrt{\sigma_1^2 + \sigma_2^2}}$$


$$T \lesssim T_1^{\max} \sqrt{1 + (\sigma_1/\sigma_2)^2} + T_2^{\max} = \Theta(n)$$

The precise result

The three fundamental distributions

$$g_{\beta}^{*s}(r) = \begin{cases} \text{Bern}_{\beta}^{*s}(r) = \beta^r(1-\beta)^{s-r} \binom{s}{r} & \beta \in]0, 1[\\ \text{Pois}_s(r) = e^{-s} \frac{s^r}{r!} & \beta = 0 \\ \text{Geom}_{-\beta}^{*s}(r) = |\beta|^r(1+|\beta|)^{-s-r} \binom{s+r-1}{r} & \beta \in]-\infty, 0[\end{cases}$$

are such that g_{α}^{*s} has a **positive decomposition** in g_{β} iff $\alpha \leq \beta$.




For the list of functions $\mathcal{F} = \{f_1, \dots, f_n\}$ in our measure,
call $\beta_{\min}(\mathcal{F})$ the smallest value of β
such that all the f_i 's have a positive decomposition in g_{β} .

The precise result

Then, the largest value for $\mathbb{E}(a)$ that can be achieved within our framework is

$$a_{\max}(\mathcal{F}) := \sqrt{1 - \beta_{\min}(\mathcal{F})} \cdot \sqrt{\frac{\sum_i \mathbb{E}[f_i]}{\sum_i \text{Var}[f_i]}}$$



Examples of application

So, we have constructed our algorithm for the linear-time exact sampling of sum-constrained random variables, in the case in which they are *not* equally distributed.

However, you could just think:

*«who cares about not-equally-distributed variables?
After all, every time I wanted to generate walks, trees, etc.,
I always wanted equally-distributed variables...»*

The point is: examples of this sort may be **hidden beyond some smart bijection**, starting from more customary (and symmetric) problems.

This is well illustrated by two classical examples:

- Set partitions, and Stirling numbers of the second kind
- Permutations with m cycles, and Stirling numbers of the first kind

Set partitions, and Stirling numbers of the second kind

Call $\mathcal{S}_{n,m}^{\text{set}}$ the ensemble of partitions of a set with n (labelled) elements into m (unlabeled) non-empty subsets.

W.l.o.g. we can assume that the set has a total ordering.

Example, for $(n, m) = (28, 9)$, and the set

$\{a, b, c, d, e, f, g, h, i, j, k, l, m, n, o, p, q, r, s, t, u, v, w, x, y, z, \alpha, \beta\}$

consider the partition

$\{\{a, g, t\}, \{b, d, m, o, \alpha\}, \{c, j, s, y\}, \{e, h, v\}, \{f, k, q\}, \{i, l, z\},$
 $\{n, p, u\}, \{r, \beta\}, \{w, x\}\}$

Set partitions, and Stirling numbers of the second kind

Call $\mathcal{S}_{n,m}^{\text{set}}$ the ensemble of partitions of a set with n (labelled) elements into m (unlabeled) non-empty subsets.

W.l.o.g. we can assume that the set has a total ordering.

Example, for $(n, m) = (28, 9)$, and the set

$\{a, b, c, d, e, f, g, h, i, j, k, l, m, n, o, p, q, r, s, t, u, v, w, x, y, z, \alpha, \beta\}$


consider the partition

$\{\{a, g, t\}, \{b, d, m, o, \alpha\}, \{c, j, s, y\}, \{e, h, v\}, \{f, k, q\}, \{i, l, z\},$
 $\{n, p, u\}, \{r, \beta\}, \{w, x\}\}$

Although the sets are not labeled, they are canonically ordered, e.g. by their smallest element. As a result, we have a canonical incidence matrix T , with $T_{ij} = 1$ if the element j is in subset i .

Set partitions, and Stirling numbers of the second kind


$$\{\{a, g, t\}, \{b, d, m, o, \alpha\}, \{c, j, s, y\}, \{e, h, v\}, \{f, k, q\}, \{i, l, z\}, \\ \{n, p, u\}, \{r, \beta\}, \{w, x\}\}$$



Call **backbone** $B(T)$ the list of smallest elements in the subsets, here $B = \{a, b, c, e, f, i, n, r, w\}$.

Set partitions, and Stirling numbers of the second kind

$$\{\{a, g, t\}, \{b, d, m, o, \alpha\}, \{c, j, s, y\}, \{e, h, v\}, \{f, k, q\}, \{i, l, z\}, \\ \{n, p, u\}, \{r, \beta\}, \{w, x\}\}$$



Call **backbone** $B(T)$ the list of smallest elements in the subsets, here $B = \{a, b, c, e, f, i, n, r, w\}$. The number of partitions T with $B(T) = B$ is the trivial product: $\prod_{y=1}^m y^{c_y}$, but the quantities c_y are **linearly constrained**: $\sum_y c_y = n - m$.

Set partitions, and Stirling numbers of the second kind

As a result, sampling **uniformly** set partitions in $\mathcal{S}_{n,m}$, which bijectively coincides to sampling uniformly the tableaux T , boils down to sampling the backbone B with the **non-uniform** measure

$$\mu_{n,m}(c_1, \dots, c_m) \propto \prod_{y=1}^m y^{c_y} \times \delta_{|c|, n-m}$$

This is exactly our framework! Introduce an appropriate Lagrange multiplier $\omega^{\sum_y c_y}$, in order to have $\mathbb{E}(|c|) = n - m$ (the good choice is the solution to the equation $\frac{n}{m} = -\frac{\ln(1-\omega)}{\omega}$)

The functions $f_y(c_y)$ are $\text{Geom}_{b_y}(c_y)$, with $b_y = \frac{\omega y}{n - \omega y}$


Now, Geom_a has a positive decomposition in terms of Bern_b

$$\text{Geom}_a(x) = \sum_s \text{Geom}_{\frac{a}{a+b}}(s) \text{Bino}_{s,b}(x)$$

Choosing for simplicity $b = \frac{1}{2}$, our algorithm works, with an average acceptance rate $\mathbb{E}(a) = \sqrt{\frac{e^{-\theta} - 1 + \theta}{2(e^\theta - 1 - \theta)}} \quad (\omega = 1 - e^{-\theta})$


m -cycle permutations, and Stirling numbers of the 1st kind

Call $\mathcal{S}_{n,m}^{\text{cyc}}$ the set of permutations $\sigma \in \mathfrak{S}_n$ with m cycles.
Describe σ through the insertion table associated to its growth,
for example, for $\sigma = ((15)(2638)(4)(7))$



m -cycle permutations, and Stirling numbers of the 1st kind


Call $\mathcal{S}_{n,m}^{\text{cyc}}$ the set of permutations $\sigma \in \mathfrak{S}_n$ with m cycles.
Describe σ through the insertion table associated to its growth,
for example, for $\sigma = ((15)(2638)(4)(7))$



m -cycle permutations, and Stirling numbers of the 1st kind

Call $\mathcal{S}_{n,m}^{\text{cyc}}$ the set of permutations $\sigma \in \mathfrak{S}_n$ with m cycles.


Describe σ through the insertion table associated to its growth,
for example, for $\sigma = ((15)(2638)(4)(7))$



$$\sigma = ((1))$$

m -cycle permutations, and Stirling numbers of the 1st kind


Call $\mathcal{S}_{n,m}^{\text{cyc}}$ the set of permutations $\sigma \in \mathfrak{S}_n$ with m cycles.
Describe σ through the insertion table associated to its growth,
for example, for $\sigma = ((15)(2638)(4)(7))$



$$\sigma = ((1)(2))$$

m -cycle permutations, and Stirling numbers of the 1st kind

Call $\mathcal{S}_{n,m}^{\text{cyc}}$ the set of permutations $\sigma \in \mathfrak{S}_n$ with m cycles.
Describe σ through the insertion table associated to its growth,
for example, for $\sigma = ((15)(2638)(4)(7))$




$$\sigma = ((1)(23))$$

m -cycle permutations, and Stirling numbers of the 1st kind

Call $\mathcal{S}_{n,m}^{\text{cyc}}$ the set of permutations $\sigma \in \mathfrak{S}_n$ with m cycles.


Describe σ through the insertion table associated to its growth,
for example, for $\sigma = ((15)(2638)(4)(7))$



$$\sigma = ((1)(23)(4))$$

m -cycle permutations, and Stirling numbers of the 1st kind


Call $\mathcal{S}_{n,m}^{\text{cyc}}$ the set of permutations $\sigma \in \mathfrak{S}_n$ with m cycles.
Describe σ through the insertion table associated to its growth,
for example, for $\sigma = ((15)(2638)(4)(7))$



$$\sigma = ((15)(23)(4))$$

m -cycle permutations, and Stirling numbers of the 1st kind


Call $\mathcal{S}_{n,m}^{\text{cyc}}$ the set of permutations $\sigma \in \mathfrak{S}_n$ with m cycles.
Describe σ through the insertion table associated to its growth,
for example, for $\sigma = ((15)(2638)(4)(7))$



$$\sigma = ((15)(263)(4))$$

m -cycle permutations, and Stirling numbers of the 1st kind


Call $\mathcal{S}_{n,m}^{\text{cyc}}$ the set of permutations $\sigma \in \mathfrak{S}_n$ with m cycles.
Describe σ through the insertion table associated to its growth,
for example, for $\sigma = ((15)(2638)(4)(7))$



$$\sigma = ((15)(263)(4)(7))$$

m -cycle permutations, and Stirling numbers of the 1st kind


Call $\mathcal{S}_{n,m}^{\text{cyc}}$ the set of permutations $\sigma \in \mathfrak{S}_n$ with m cycles.
Describe σ through the insertion table associated to its growth,
for example, for $\sigma = ((15)(2638)(4)(7))$



$$\sigma = ((15)(2638)(4)(7))$$

m -cycle permutations, and Stirling numbers of the 1st kind

Call $\mathcal{S}_{n,m}^{\text{cyc}}$ the set of permutations $\sigma \in \mathfrak{S}_n$ with m cycles.
Describe σ through the insertion table associated to its growth,
for example, for $\sigma = ((15)(2638)(4)(7))$




$$\sigma = ((15)(2638)(4)(7))$$

Call $B(\sigma) = \{0, 0, 1, 0, 1, 1, 0, 1\}$,
the indicator function of “black rows”
of $T(\sigma)$, the **backbone** of σ .

m -cycle permutations, and Stirling numbers of the 1st kind

Call $\mathcal{S}_{n,m}^{\text{cyc}}$ the set of permutations $\sigma \in \mathfrak{S}_n$ with m cycles.
Describe σ through the insertion table associated to its growth,
for example, for $\sigma = ((15)(2638)(4)(7))$



$$\sigma = ((15)(2638)(4)(7))$$

Call $B(\sigma) = \{0, 0, 1, 0, 1, 1, 0, 1\}$,
the indicator function of “black rows”
of $T(\sigma)$, the **backbone** of σ .

The number of σ 's with backbone
 $B = (x_1, \dots, x_n)$ is $\prod_y (y-1)^{x_y}$,
and we must have $|\mathbf{x}| = m$

Again, **this is exactly our framework!**
just with inhomogeneous Bernoulli variables,
instead of inhomogeneous Geometric variables.

Our problem was the **exact sampling** in linear time from the measure $\mu_{n,m}(\mathbf{x}) = \frac{1}{Z} \prod_{i=1}^n f_i(x_i) \times \delta_{|\mathbf{x}|,m}$

We have provided a solution to this problem in the case in which all the f_i 's have a **positive decomposition** in terms of the same function g_β

$$g_\beta(r) = \begin{cases} \text{Bern}_\beta(r) & \beta \in]0, 1[\\ \text{Poiss}_1(r) & \beta = 0 \\ \text{Geom}_{-\beta}(r) & \beta \in]-\infty, 0[\end{cases}$$

In this case, the rejection scheme has a complexity related to the **average acceptance rate**, which is

$$a_{\max} = \sqrt{1 - \beta} \cdot \sqrt{(\mathbb{E}[f_1] + \dots + \mathbb{E}[f_n]) / (\text{Var}[f_1] + \dots + \text{Var}[f_n])}$$

Applications include the uniform exact sampling of **set partitions** with a prescribed number of subsets (and, in turns, minimal automata), and of **permutations** with a prescribed number of cycles.