Lecture Notes
Graduate Course in

STOCHASTIC SIMULATION

Chalmers University of Technology and

Göteborg University

Patrik Albin, March 2000

Contents

1 Generating Univariate Random Variables in a Computer

Definition. A random number is a random variable with a uniform distribution over the interval $(0,1)$.

Note that if ξ_1 and ξ_2 are independent random numbers, then $1-\xi_1$ and $\big(1-|\xi_1-\xi_2|\big)^2$ are also random numbers (see Exercise 44 below for the second claim).

In Section 4 we learn how to generate sequences of numbers in a computer that with good aproximation are independent random numbers. In Sections 5 and 6 we will see how one can make statistical tests to evaluate how good independent random numbers they really are. First however, we learn how independent random numbers can be used to generate more general random variables.

Definition. Given a function $F:\mathbb{R}\to[0,1]$ we write

$$F^{\leftarrow}(y) = \inf\{x : F(x)\geq y\} \quad \text{for} \quad y\in(0,1).$$

Lemma 1. If $F(x) = F_X(x) = \mathbf{P}\{X\leq x\}$ is a distribution function for a random variable X, then we have

$$F^{\leftarrow}(y) \leq x \;\Leftrightarrow\; y \leq F(x).$$

Proof of $\Rightarrow$: For $F(x)<y$ the fact that F is non-decreasing and right-continuous gives

$$F(z)<y \quad \text{for} \quad z\in(-\infty,x+\varepsilon], \quad \text{for some} \quad \varepsilon>0.$$

It follows that

$$F^{\leftarrow}(y) \;=\; \inf\big\{z : F(z)\geq y\big\} \;>\; x.$$

In other words, $F(x)<y \Rightarrow F^{\leftarrow}(y)>x$, which is the same thing as $F^{\leftarrow}(y)\leq x \Rightarrow y\leq F(x)$. $\Box$

Exercise 1. Prove the implication $\Leftarrow$ in Lemma 1.

THEOREM 1. Let $F=F_X$ be the distribution function for a random variable X, and let ξ be a random number. The random variable

$$\eta = F^{\leftarrow}(\xi) \qquad \text{has distribution function} \qquad F_\eta(x) = \mathbf{P}\{\eta\leq x\} = F(x).$$

Proof: By application of Lemma 1 we get

$$F_\eta(x) \;=\; \mathbf{P}\big\{F^{\leftarrow}(\xi)\leq x\big\} \;=\; \mathbf{P}\big\{\xi\leq F(x)\big\} \;=\; F_\xi(F(x)) \;=\; F(x). \quad \Box$$

Of course, a strictly increasing and continuous distribution function F is invertible with invers function $F^{-1} = F^{\leftarrow}$.

Remark. Often explicit expressions for F and or $F^{\leftarrow}$ are not known: This is for example the case with the normal distribution. Therefore Theorem 1 is not a universal solution to the problem of generating random variables (without further ramifications).

Corollary 1. Let F be a continuous distribution function. Suppose that there exist real numbers $-\infty \leq a < b \leq \infty$ such that

- $F(x) \in (0,1)$ for $x \in (a,b)$,
- $F:(a,b) \to (0,1)$ is strictly increasing, and
- $F:(a,b) \to (0,1)$ is onto [that is, F takes all values in the interval $(0,1)$].

The function $F:(a,b) \to (0,1)$ is invertible with invers function $F^{-1}:(0,1) \to (a,b)$, given by $F^{-1}(y) = F^{\leftarrow}(y)$ for $y \in (0,1)$. Hence, if ξ is a random number, the random variable

$$\eta = F^{-1}(\xi) \quad \text{has distribution function} \quad F_\eta(x) = F(x).$$

Corollary 1 might seem complicated, but is in reality not more difficult that this:

Example 1. The $\mathrm{Exp}(\lambda)$-distribution has distribution function $F(x) = 1 - \mathrm{e}^{-\lambda x}$ for $x \in (0,1)$. This distribution function has invers $F^{-1}(y) = -\lambda^{-1}\ln(1-y)$ for $y \in (0,1)$. If ξ is a random number the random variable $\eta = F^{-1}(\xi) = -\lambda^{-1}\ln(1-\xi)$ will therefore be $\exp(\lambda)$-distributed.

Exercise 2. Explain how Theorem 1 works to generate a random variable that is $\mathrm{Bin}(1,p)$-distributed.

Exercise 3. A function $F:\mathbb{R} \to [0,1]$ is a distribution function (for some random variable) if (and only if) F is non-decreasing and right-continuous with "endpoints" $\lim_{x\to\infty} F(x) = 1$ and $\lim_{x\to-\infty} F(x) = 0$: Explain how a possible proof of this statement connects to Theorem 1.

Definition. For a real number x we let $\lfloor x \rfloor$ denote the unique integer k that satisfies $k < x \leq k+1$.

Of course, one can write $\lfloor x \rfloor = \mathrm{int}(x^-)$, where $\mathrm{int}(y)$ is the "usual" integer part of a real number y, that is the unique integer ℓ that satisfies $\ell \leq y < \ell+1$.

THEOREM 2. Let $f(x) = f_X(x) = \mathbf{P}\{X=x\}$ be the density function for a discrete random variable X with (finitely or infinitely many) possible values $\{1, 2, \ldots\}$. Let ξ be a random number. The random variable

$$\eta = \sum_{k\geq 1} \left\lfloor \xi + 1 - \sum_{j=1}^{k-1} f(j) \right\rfloor \qquad \text{is discrete with density function} \qquad f_\eta = f.$$

Proof by direct calculation: Since the terms $\zeta_k = \lfloor \xi + 1 - \sum_{j=1}^{k-1} f(j) \rfloor$ of the sum take values in $\{0,1\}$ and are non-decreasing (when k increases), and since $\zeta_1 = 1$, we have

$$\begin{aligned}
\mathbf{P}\{\eta = n\} &= \mathbf{P}\{\zeta_1 = \ldots = \zeta_n = 1,\ \zeta_{n+1} = \zeta_{n+2} = \ldots = 0\} \\
&= \mathbf{P}\{\zeta_n = 1,\ \zeta_{n+1} = 0\} \\
&= \mathbf{P}\Big\{ \big\lfloor \xi + 1 - \textstyle\sum_{j=1}^{n-1} f(j) \big\rfloor = 1,\ \big\lfloor \xi + 1 - \textstyle\sum_{j=1}^{n} f(j) \big\rfloor = 0 \Big\} \\
&= \mathbf{P}\Big\{ \xi + 1 - \textstyle\sum_{j=1}^{n-1} f(j) > 1,\ \xi + 1 - \textstyle\sum_{j=1}^{n} f(j) \leq 1 \Big\} \\
&= \mathbf{P}\Big\{ \sum_{j=1}^{n-1} f(j) < \xi \leq \sum_{j=1}^{n} f(j) \Big\} \\
&= F_\xi\Big(\sum_{j=1}^{n} f(j) \Big) - F_\xi\Big(\sum_{j=1}^{n-1} f(j) \Big) \\
&= f(n). \quad \square
\end{aligned}$$

Remark: Since $f(0) = 0$ in Theorem 2, one can alternatively put $\eta = \sum_{k\geq 1} \lfloor \xi + 1 - \sum_{j=0}^{k-1} f(j) \rfloor$.

It is easy to modify Theorem 1 so that it can be used to generate a discrete random variable X with possible values $\{x_1, x_2, \ldots\}$ other than $\{1, 2, \ldots\}$.

Exercise 4. Give an alternative proof of Theorem 2 by means of explaining how the theorem follows from an application of Theorem 1.

Exercise 5. Often the formula $\eta = \sum_{k\geq 1} \mathrm{int}(\xi + 1 - \sum_{j=0}^{k-1} f(j))$ is used instead of that in the theorem: Explain why also this formula gives a random variable with the desired density function $f_\eta = f$.

Corollary 2. (THE TABLE METHOD) Let f be the density function of a discrete random variable with (finitely or infinitely many) possible values $\{x_1, x_2, \ldots\}$. If ξ is a random number, then the random variable

$$\eta = \begin{cases} x_1 & \text{if} \quad 0 < \xi \leq f(x_1) \\ x_2 & \text{if} \quad f(x_1) < \xi \leq f(x_1) + f(x_2) \\ x_3 & \text{if} \quad f(x_1) + f(x_2) < \xi \leq f(x_1) + f(x_2) + f(x_3) \\ & \ \vdots \end{cases}$$

will have density function $f_\eta(x_i) = f(x_i)$.

Proof: It is obvious that the statement of the corollary holds. Alternatively, the corollary follows from a comparison with the proof of Theorem 2. □

Exercise 6. Let (X,Y) be an $\mathbb{R}^2$-valued random variable such that there exists a 'nice version' of the conditional probability

$$F^{(y)}(x) \;=\; \mathbf{P}\{X\le x|Y=y\} \quad \text{for} \;\; (x,y)\in\mathbb{R}^2.$$

Shew that if ξ is a random number independent of Y, then we have

$$\Big((F^{(Y)})^{\leftarrow}(\xi), Y\Big) =_{\text{distribution}} (X,Y).$$

How can one generate a random variable which is distributed as (X,Y) using two independent random numbers ξ_1 and ξ_2?

2 More on Generating Random Variables in a Computer

2.1 Rejection Sampling

Definition. An $\mathbb{R}^n$-valued random variable X has uniform distribution over a (non-null measurable) set $O\subseteq\mathbb{R}^n$ if

$$\mathbf{P}\{X\in o\} = \frac{\text{Volume}(o)}{\text{Volume}(O)} \qquad \text{for each (measurable) set} \quad o\subseteq O.$$

THEOREM 3. Let $O\subseteq\mathbb{R}^n$ be a bounded (non-null measurable). Choose a rectangle $R\subseteq\mathbb{R}^n$ such that $O\subseteq R$. Let $\xi_1,\xi_2,\ldots$ be independent identically distributed random variables with a uniform distribution over R, and put $J=\min\{j: \xi_j\in O\}$. The random variable $\eta=\xi_J$ possesses a uniform distribution over O.

Proof: For a subset $o\subseteq O$ we have

$$\begin{aligned}
\mathbf{P}\{\eta\in o\} = \sum_{j=1}^{\infty}\mathbf{P}\{\eta\in o \mid J=j\}\,\mathbf{P}\{J=j\} &= \sum_{j=1}^{\infty}\mathbf{P}\Big\{\xi_j\in o \;\Big|\; \xi_1,\ldots,\xi_{j-1}\notin O,\, \xi_j\in O\Big\}\,\mathbf{P}\{J=j\}\\
&= \sum_{j=1}^{\infty}\tfrac{\mathbf{P}\{\xi_1,\ldots,\xi_{j-1}\notin O,\, \xi_j\in o\}}{\mathbf{P}\{\xi_1,\ldots,\xi_{j-1}\notin O,\, \xi_j\in O\}}\,\mathbf{P}\{J=j\}\\
&= \sum_{j=1}^{\infty}\tfrac{\mathbf{P}\{\xi_j\in o\}}{\mathbf{P}\{\xi_j\in O\}}\,\mathbf{P}\{J=j\}\\
&= \mathbf{P}\{\xi_1\in o\}\,/\,\mathbf{P}\{\xi_1\in O\}\\
&= \big(\text{Vol}(o)/\text{Vol}(R)\big)\;/\;\big(\text{Vol}(O)/\text{Vol}(R)\big)\\
&= \text{Vol}(o)\;/\;\text{Vol}(O). \quad \square
\end{aligned}$$

Of course, it is an easy matter to generate a random variable with a uniform distribution over a rectangle $R = [a_1, b_1] \times \ldots \times [a_n, b_n] \subseteq \mathbb{R}^n$.

Exercise 7. Let X be a random variable with distribution function F, and let $[a, b]$ be an interval such that $0 < \mathbf{P}\{a \leq X \leq b\} < 1$. Explain how independent observations $X_1, X_2, \ldots$ of X can be used to generate a random variable Y with distribution function

$$F_Y(y) = \mathbf{P}\{Y \leq y\} = \mathbf{P}\{X \leq y \mid a \leq X \leq b\} \quad \text{for} \quad y \in [a, b].$$

2.2 Importance Sampling

Consider the problem in Exercise 7 again. Of course, if we can compute the distribution function F_Y of the random variable Y, then there is no need for rejection sampling, and we can instead use the method in Theorem 1. This is importance sampling.

(These methods have fancy names, but amount to little more than a bit of common sense.)

2.3 The Method of Box-Müller

The following theorem is the standard way to generate normal distributed random variables:

THEOREM 4. (Box-Müller) Let ξ_1 and ξ_2 be independent random numbers. The random variable

$$Z = \mu + \sigma\sqrt{-2\ln(\xi_1)}\,\cos(2\pi\xi_2) \qquad \text{is} \quad \mathrm{N}(\mu, \sigma^2)\text{-distributed.}$$

Proof when $\mu = 0$ and $\sigma = 1$: When N_1 and N_2 are independent $\mathrm{N}(0,1)$-distributed random variables we have

$$\text{radius}(N_1, N_2) = \sqrt{N_1^2 + N_2^2} \sim \sqrt{\exp(1/2)}$$

[since $\chi^2(2)$-distribution is the same as $\exp(1/2)$-distribution]. By symmetry, the argument (angle) of (N_1, N_2) is uniformly distributed over $[0, 2\pi]$. Using Example 1, it follows that

$$(N_1, N_2) =_{\text{distribution}} \sqrt{-2\ln(\xi_1)}\,\big(\cos(2\pi\xi_2), \sin(2\pi\xi_2)\big). \quad \square$$

The proof of Theorem 4 indicates that the random numbers ξ_1 and ξ_2 can be used in a slightly more efficient way than suggested in the theorem: It follows from the proof that

$$Z_1 = \mu + \sigma\sqrt{-2\ln(\xi_1)}\,\cos(2\pi\xi_2) \qquad \text{and} \qquad Z_2 = \mu + \sigma\sqrt{-2\ln(\xi_1)}\,\sin(2\pi\xi_2)$$

are independent $\mathrm{N}(\mu, \sigma^2)$-distributed random variables. Hence one can get two normal distributed random variables out of ξ_1 and ξ_2 rather than only one

2.4 Multi-variate Random Variables

We have seen a method to generate random variables in two dimensions in Exercise 6. This method can be extended in an obvious way to higher dimensions. The use of the method requires that conditional distributions can be calculated (analytically or numerically) for one of the vector components of the random variable given the values of the others components.

The following two exercises descrice versions of another general method to generate multi-variate random variables. This method requires that the random variable is absolutely continuous with a density function that is known explicitely:

Exercise 8. Let X be a random variable with values in a bounded (non-null measurable) set $R \subseteq \mathbb{R}^n$. Suppose that X has a bounded density function $f : R \to [0,\infty)$. Let η be a random variable that is uniformly distributed over the set $\{(x,y) \in \mathbb{R}^n \times \mathbb{R} : x \in R,\ y \in (0, f(x))\}$. Prove that the random variable $(\eta_1, \ldots, \eta_n)$ (the vector made up of the first n vector components of η) has density function f.

Exercise 9. Let X be a random variable with values in a (non-null measurable) set $R \subseteq \mathbb{R}^n$. Suppose that X has a density function $f : R \to [0,\infty)$. Let $\{R_y\}_{y\in[0,1)}$ denote a non-decreasing family of bounded (non-null measurable) subsets of R such that $f : R_y \to [0,\infty)$ is bounded for each $y \in [0,1)$, and such that $\bigcup_{y\in[0,1)} R_y = R$. Let ζ be a discrete random variable with values in $\{1,2,\ldots\}$, infinitely many of which have non-zero probability. Let $\{\eta_j\}_{j=1}^{\infty}$ be random variables that are independent of ζ with η_j uniformly distributed over

$$\{(x,y) \in \mathbb{R}^n \times \mathbb{R} : x \in R_{F_\zeta(j)} \backslash R_{F_\zeta(j-1)},\ 0 < y < f(x)/\mathbf{P}\{\zeta = j\}\} \qquad \text{for} \quad j \in \{1,2,\ldots\}.$$

Prove that the random variable $((\eta_\zeta)_1, \ldots, (\eta_\zeta)_n)$ has density function f.

3 Estimation of Expected Values

3.1 Elementary Concepts

In Sections 3.1-3.3, $X_1, \ldots, X_n$ denote independent identically distributed random variables with an unknown expected value $\mathbf{E}\{X_i\} = \mu$ and a finite variance $\mathbf{Var}\{X_i\} = \sigma^2$. We are going to manufacture a confidence interval for μ by means of stochastic simulation.

By the Central Limit Theorem, the sample mean $\overline{X} \equiv \frac{1}{n}(X_1 + \ldots + X_n) \approx_{\text{distribution}} \mathrm{N}(\mu, \sigma^2/n)$. Writing λ_α for the solution to the equation $\Phi(\lambda_\alpha) = \mathbf{P}\{\mathrm{N}(0,1) \le \lambda_\alpha\} = 1-\alpha$, we thus get

$$\mathbf{P}\Big\{\overline{X} - \lambda_{\alpha/2}\,\sigma/\sqrt{n} \le \mu \le \overline{X} + \lambda_{\alpha/2}\,\sigma/\sqrt{n}\Big\} = \mathbf{P}\Big\{\mu - \lambda_{\alpha/2}\,\sigma/\sqrt{n} \le \overline{X} \le \mu + \lambda_{\alpha/2}\,\sigma/\sqrt{n}\Big\} \approx 1-\alpha.$$

If the standard-deviation σ is unknown it is estimated with for example $s = \left[\frac{1}{n-1}\sum_{i=1}^{n}(X_i - \overline{X})^2\right]^{1/2}$. This gives the confidence interval

$$\mathbf{P}\Big\{\overline{X} - \lambda_{\alpha/2}\, s/\sqrt{n} \le \mu \le \overline{X} + \lambda_{\alpha/2}\, s/\sqrt{n}\Big\} \approx 1-\alpha.$$

The particular case when $X_1,\ldots,X_n \sim \mathrm{Bin}(1,p)$, so that $\mu = p$, is very important. Here $\sqrt{\overline{X}(1-\overline{X})}$ is often used to estimate σ instead of s [since $\mathbf{Var}\{X_i\} = p(1-p) = \mu(1-\mu)$].

A quantity that is of apparent interest is the width of the confidence interval

$$b = 2\lambda_{\alpha/2}\,\sigma/\sqrt{n} \qquad \text{or} \qquad b = 2\lambda_{\alpha/2}\, s/\sqrt{n}.$$

Given a certain width b, the number of observations n required will have a quadratic growth as functions of the standard-deviation σ or estimated standard-deviation s;

$$n = \big(2\lambda_{\alpha/2}\,\sigma/b\big)^2 \qquad \text{or} \qquad n = \big(2\lambda_{\alpha/2}\, s/b\big)^2.$$

The idea behind <u>methods for variance reduction</u> is to use the original observations $X_1,\ldots,X_n$ to generate new observations $Z_1,\ldots,Z_n$ with unchanged expected value $\mathbf{E}\{Z_i\} = \mathbf{E}\{X_i\} = \mu$, but with variance $\mathbf{Var}\{Z_i\}$ that is smaller than $\mathbf{Var}\{X_i\} = \sigma^2$. We can use the estimate $\overline{Z}$ to estimate μ, and provided that we really have $\mathbf{Var}\{Z_i\} < \mathbf{Var}\{X_i\} = \sigma^2$, the number of observations n required for a certain interval width b will be smaller than when $\overline{X}$ is used.

3.2 Variance Reduction with Antithetic Variables

Suppose that we can use each observed random variable X_i to create another random variable Y_i in such a way that

$$\mathbf{E}\{Y_i\} = \mathbf{E}\{X_i\} = \mu \qquad \text{and} \qquad \mathbf{Var}\{Y_i\} = \mathbf{Var}\{X_i\} = \sigma^2.$$

Writing $Z_i \equiv \frac{1}{2}(X_i + Y_i)$ we then obtain

$$\mathbf{E}\{Z_i\} = \mu \qquad \text{and} \qquad \mathbf{Var}\{Z_i\} = \frac{1}{2}\big(\sigma^2 + \mathbf{Cov}\{X_i, Y_i\}\big) \le \sigma^2.$$

Particularly useful is the situation in which $\mathbf{Cov}\{X_i, Y_i\} < 0$ (see below), and in that case we call Y_i <u>antithetic variable</u> for X_i.

If the pairs $(X_1, Y_1),\ldots,(X_n, Y_n)$ are independent,

$$\overline{Z} \text{ is an unbiased estimate of } \mu \text{ with variance } \frac{1}{2n}\sum_{i=1}^{n}\big(\sigma^2 + \mathbf{Cov}\{X_i, Y_i\}\big).$$

The estimate $\overline{Z}$ of μ is unbiased and can be used whenever we have $\mathbf{E}\{Y_i\} = \mathbf{E}\{X_i\} = \mu$: One does not have to really prove that $\mathbf{Var}\{Y_i\} = \mathbf{Var}\{X_i\}$ or that Y_i is antithetic for X_i, but it is enough to "suspect" that this is the case. (The worst thing that can happen is that $\mathbf{Var}\{\overline{Z}\} > \mathbf{Var}\{\overline{X}\}$ instead of the desired inequality in the other direction!)

In the context of generation of random variables in a computer, the following trick is often useful for finding antithetic variables: If we generate a random variable X with the help of a random number ξ [using for example the formula $X = F_X^{-1}(\xi)$], then we generate the random variable Y by means of the same scheme, but using $1-\xi$ instead of ξ as our random number [so that $Y = F_X^{-1}(1-\xi)$ in the example above]. The random variables X and Y have the same distribution, and Y will for obvious reasons usually be antithetic for X.

Exercise 10. Consider the so called double exponential distribution function $F(x) = \exp\{-\mathrm{e}^{-x}\}$ for $x \in \mathbb{R}$. The expected value μ of this distribution cannot be calculated analytically. (The expected value happens to be equal to Eulers constant, but the exact value of that constant is not known.) One way to obtain an approximate numerical value for μ is to use stochastic simulation and generate independent random variables $X_1, X_2, \ldots, X_n$ with common distribution function F. The expected value μ is estimated with the sample mean $\overline{X}$. Of course, this procedure becomes more efficient if we can find antithetic variables $Y_1, Y_2, \ldots, Y_n$ to $X_1, X_2, \ldots, X_n$, and instead employ $\overline{Z} = \frac{1}{2}(\overline{X}+\overline{Y})$ to estimate μ. Use this procedure to obtain a confidence intervals for the mean μ, both with and without the use of an antithetic variable. (Recall that one do not have to really prove that a variable is antithetic. Instead one can try a canidate-antithetic variable, and decide to use it if it gives a reduced variance.)

3.3 Variance Reduction with Control Variables

Suppose that we can use each observed random variable X_i to create another random variable Y_i with a known expected value $\mathbf{E}\{Y_i\} = \mu_0$, and such that $\mathbf{Cov}\{X_i, Y_i\} > \frac{1}{2}\mathbf{Var}\{Y_i\}$. For the random variable $Z_i = X_i - (Y_i - \mu_0)$ we then have

$$\mathbf{E}\{Z_i\} = \mu \quad \text{and} \quad \mathbf{Var}\{Z_i\} = \sigma^2 + \mathbf{Var}\{Y_i\} - 2\,\mathbf{Cov}\{X_i, Y_i\} < \sigma^2.$$

In this case we call Y_i a <u>control variable</u> for X_i.

The estimate $\overline{Z}$ of μ is unbiased and can be used for estimation of μ whenever the expected value $\mathbf{E}\{Y_i\} = \mu_0$ is known: It is not necessary to really prove that Y_i is a control variable for X_i, but is enough to "suspect" that this is the case. (The worst thing that can happen is that $\mathbf{Var}\{\overline{Z}\} > \mathbf{Var}\{\overline{X}\}$ instead of the other way around.)

Exercise 11. Again consider the double exponential distribution function $F(x) = \exp\{-\mathrm{e}^{-x}\}$. Employ stochastic simulation to calculate a confidence interval for the expected value μ of F. This time, use an exponentially distributed random variable

Y with parameter 1 (or some other convenient random variable) as control variable for X. Of course, the same random number ξ_i must be used to generate both an observed double exponential random variable X_i and the corresponding control variable Y_i. (Recall that one do not have to prove that a variable is a control variable. Instead one can try it out, and decide to use it if it gives a reduced variance.)

Exercise 12. Let Y_i be a control variable for X_i with expected value $\mathbf{E}\{Y_i\}=\mu_0$. Find the value for of constant c that minimizes the variance $\mathbf{Var}\{X_i-c(Y_i-\mu_0)\}$. How can Y_i be used to construct an improved control variable $\hat{Y}_i$ for X_i [provided that certain covariances (variances) are known, or can be estimated]?

4 Generating Random Numbers in a Computer

Sections 4 through 6 build closely on Sections 3.2 and 3.3, respectively, of "D. Knuth (1998): The Art of Computer Programming, 3rd Ed.". (Note the timely appearance of the third edition.) Donald Knuth is world famous not only within several subdisciplines of computer science, but also in other disiplines, including mathematics and filosofy. Knuth (alone) is the inventor of the TEX-typesetting programme, which argubly is one of the more important achievements of mankind. In addition, Knuth has a lively pedagocial interest, which becomes apparent when reading his book.

The lecture notes are not intended to replace Knuths book, but rather to summarize some important definitions and results therein. The reader is thus recommended to enjoy and take in Knuth's full text. Note especially the following passages on pages 4 and 6, respectively, in the book:

"Many random number generators in use when this chapter was first written were not very good. People have traditionally tended to avoid learning anything about such subroutines; old methods that were comparatively unsatisfactory have been passed down blindly from one programmer to another, until the users have no understanding of the original limitations. We shall see in this chapter that the most important facts about random number generators are not difficult to learn ......" and "...... random numbers should not be generated with a method chosen at random. Some theory should be used."

4.1 Linear Congruential Generators

It is not possible to manufacture "real" independent random numbers in a computer. Instead one has to be content with so called pseudorandom numbers, which are "good approximations" of real random numbers. The way to differ between good and less good generators of pseudorandom numbers (from now on simply called random numbers), is to check to what extent certain "obviously desirable properties" hold (Section 4), and to perform certain statistical tests (Sections 5 and 6).

The most important random number generators are the so called (linear) congruential generators, which give random numbers $U_1, U_2, \ldots$ according to the scheme

$$U_n = X_n/m \qquad \text{where} \qquad X_n = (a\,X_{n-1} + c) \bmod m \qquad \text{for} \quad n \geq 1.$$

Here the module $m \in \mathbb{N}\setminus\{0\}$, the multiplier $a \in \{0, \ldots, m-1\}$, the increment $c \in \{0, \ldots, m-1\}$ and the seed $X_0 \in \{0, \ldots, m-1\}$ are parameters that determine the generator. [Recall that $\ell \bmod m$ denotes the remainder when the non-negative integer ℓ is divided by the positive integer m.]

Although this is not common, one could alternatively take $U_n = X_n/(m-1)$, or possibly $U_n = (X_n + 1/2)/m$. The latter choice have the advantage that the values 0 and 1 are never attained.

Exercise 13. Prove the following two formulas:

$$(a\,b) \bmod m \;=\; ((a \bmod m)\,(b \bmod m)) \bmod m$$

$$(a + b) \bmod m \;=\; ((a \bmod m) + (b \bmod m)) \bmod m.$$

Exercise 14. Prove that a congruential generator with $a \neq 1$ satisfies

$$X_n \;=\; (a^n X_0 + c\,(a^n - 1)/(a-1)) \bmod m \quad \text{for} \quad n \geq 0.$$

(There is not much randomness here, is there?)

The sequence $\{X_n\}_{n\geq 1}$ will ultimately, after an initial transient time-period of length at most $m-2$, be periodic with period-length $\mathcal{P} = \mathcal{P}(m, a, c, X_0)$ at most m: Clearly, it is desirable to have a long period. Hence it is necessary to have a large value of m.

The fact that $\{X_n\}_{n\geq 1}$ is periodic follows from noting that there exist indices $i, j \in \{0, \ldots, m\}$ with $i < j$ such that $X_i = X_j$, because otherwise there would be no "ties" among the numbers $X_0, \ldots, X_m$. (This gives a contradiction, because there are only m different possible values for the X-sequence.) Since every X-value is determined by its predecessor, the fact that $X_i = X_j$ gives $X_{i+k} = X_{j+k}$ for all $k \geq 0$ by iteration. This means that $\{X_n\}_{n\geq 1}$ is periodic with period-length at most $j-i$, which in turn is at most m. Also, the "periodic behavior" starts at (latest at) $n=i$, so that the transient time-period is of length at most $i-1 \leq m-2$.

Exercise 15. Explain why congruential generators with $a=0$ and $a=1$ are poor. What is the period-length of the generator with $a=1$ when $c=1$?

A linear congruential generator with $c = 0$ is for obvious reasons refered to as a multiplicative congruential generator. Such a generator gives random numbers

$$U_n = X_n/m \qquad \text{where} \qquad X_n = (a\,X_{n-1}) \bmod m \qquad \text{for} \quad n \geq 1.$$

4.2 The Module m

We have seen that at large vaule of m is desirable. In addition, m should be chosen in such a way that the generator is efficient from a computational point of view. An especially appropriate choice of the module therfore appears to be $m=w$ where w is the word-length of the computer. Here we have $w=2^e$ in an e-bit computer with binary arithmetic and $m=10^e$ in an e-bit computer with decimal arithmetic. These choices of m are certainly the largest possible, and do have the advantage that the division $X_n = \big(a\,X_{n-1}+c\big) \bmod m$ comes "for free" via automatic truncation of "overflow" (since only the numbers $\{0,\ldots,w-1\}$ can be represented in the computer).

Remark. Of course, typical software are set up to handle integer arithmetic for integers in the range $\{-2^{e-1},\ldots,2^{e-1}-1\}$ or $\{-2^{e-1}+1,\ldots,2^{e-1}\}$, rather than $\{0,\ldots,2^e-1\}$. However, this do not cause much of a problem, and "adaptions" are straightforward. It is merely a matter of "shifting" and "unshifting" numbers a distance 2^{e-1} or $2^{e-1}-1$ to the right and the left, respectively, in order to convert between the two different ways to make use of the word-length 2^e.

The choice $m=w$ has one important drawback: Random numbers from the resulting generator will suffer from "poor randomness" in the sense of short period in the digits that are least significant. In the case when $m = w = 2^e$, for example, the f least significant digits $Y_n = X_n \bmod 2^f$ in the sequence $\{X_n\}_{n\geq 1}$ will have period at most 2^f: This follows from the fact that $Y_n = \big(a\,Y_{n-1}+c\big) \bmod 2^f$.

Exercise 16. Shew that $Y_n = (a\,Y_{n-1}+c) \bmod 2^f$ when $m = w = 2^e$.

Because of the above indicated problems with the choice $m=w$, the alternatives $m=w\pm 1$ are also used: It is possible to take m longer than the actual word-length, because words are equipped with an "overflow-flag", that can be used as an "extra-bit". The division $X_n = \big(a\,X_{n-1}+c\big) \bmod m$ can be carried out comparatively easily also with this m, and the problem with short periods for the least significant digits is no longer present.

Knuth suggests that one makes the choice $m=w$, despite the problem discussed above. However, one should be aware of the existence of this problem, and the limitations it imposes on the use of the random numbers generated. In addition to giving fast generators with simple arithmetic, there is another important advantage with the choice $m=w$: Generators with this module are the ones that are best understood theoretically, and have the most explicitely known features.

In Section 3.6 of his book, Knuth provides historical support for the claim that every random number generator that is so complicated that its performance could not be evaluated (theoretically or by other means) at the time it was introduced, has with use eventually been proven to be not very good. Thus there are historical reasons for "staying simple". Also note Exercise 3.6.6 in Knuth: "*[40]* Look at the subroutine library of each computer installation in your organization, and replace the random number generators by good ones. Try to avoid being too shocked at what you find."

4.3 The Multiplier a

Here we answer the questions what is the maximal length of the period $\mathcal{P} = \mathcal{P}(m, a, c, X_0)$ for a linear congruential random number generator with module m, and how do one achieve that period-length? In other words, given an $m \in \mathbb{N}\backslash\{0\}$, we search for

$$\max\{\mathcal{P}(m, a, c, X_0) : a, c, X_0 \in \{0, \ldots, m-1\}\}$$

$$\text{and} \qquad \operatorname{argmax}\{\mathcal{P}(m, a, c, X_0) : a, c, X_0 \in \{0, \ldots, m-1\}\}.$$

Of course, in the case of a multiplicative generator we instead search for

$$\max\{\mathcal{P}(m, a, 0, X_0) : a, X_0 \in \{0, \ldots, m-1\}\}$$

$$\text{and} \qquad \operatorname{argmax}\{\mathcal{P}(m, a, 0, X_0) : a, X_0 \in \{0, \ldots, m-1\}\}.$$

Recall that we have the general upper bound $\mathcal{P}(m, a, c, X_0) \leq m$. Obviously, the value of the maxima is unique, while the argmax possibly could take several different values [that is, the maximal value (period) could be attained for several different choices of a, c and X_0].

Note that a long period is only one of many desirable properties that should be imposed on a sequence of (pseudo-) random numbers! It is certainly not enough to check just this property, as the following exercise illustrates:

Exercise 17. Given an example of a really poor generator that has period $\mathcal{P} = m$.

Two positive integers are called <u>relatively prime</u> if the only prime number (or other strictly positive integer for that matter) p that divides both of them ["evenly" (without a remainder)] is $p = 1$.

THEOREM 5. (HULL-DOBELL, 1962) A linear congruential generator $X_n = (a\, X_{n-1} + c) \bmod m$ has period-length $\mathcal{P}(m, a, c, X_0) = m$ if and only if the following three conditions hold:

(i) The numbers c and m are relatively prime.

(ii) Each prime number p that divides m will also divide $a-1$.

(iii) If 4 divides m, then 4 also divides $a-1$.

Exercise 18. Shew that $\max\{\mathcal{P}(m,a,c,X_0) : a,c,X_0 \in \{0,\ldots,m-1\}\} = m$ when m is a prime number, and find $\operatorname{argmax}\{\mathcal{P}(m,a,c,X_0) : a,c,X_0 \in \{0,\ldots,m-1\}\}$.

Exercise 19. Use the Hull-Dobell Theorem to shew that a multiplicative generator cannot have period-length m when $m>1$ [so that $\mathcal{P}(m,a,0,X_0) \le \max\{m-1,1\}$]. Alternatively (or in addition), prove this by means of "direct computation".

Definition. (i) Let m and $x \in \{1,\ldots,m-1\}$ be positive and relatively prime integers. The order of x mod m is given by (cf. Exercises 20 and 21 below)

$$\underline{\mathcal{O}(x,m)} \equiv \min\Big\{n \in \{1,\ldots,m-1\} \,:\, x^n \bmod m = 1\Big\}.$$

(ii) Given a positive integer m, the maximal order mod m is defined

$$\underline{\mathcal{O}(m)} \equiv \max\Big\{\mathcal{O}(x,m) \,:\, m \text{ and } x \in \{1,\ldots,m-1\} \text{ are relatively prime}\Big\}.$$

(iii) Given a positive integer m, we say that x is a primitive element mod m when m and $x \in \{1,\ldots,m-1\}$ are relatively prime with $\mathcal{O}(x,m) = \mathcal{O}(m)$, that is, when

$$x = \operatorname{argmax}\Big\{\mathcal{O}(x,m) \,:\, m \text{ and } x \in \{1,\ldots,m-1\} \text{ are relatively prime}\Big\}.$$

Exercise 20. Let m and $x \in \{1,\ldots,m-1\}$ be positive and relatively prime integers. Prove that $x^n \bmod m \neq 0$ for $n \in \{1,\ldots,m-1\}$.

Exercise 21. Let m and $x \in \{1,\ldots,m-1\}$ be positive and relatively prime integers. Prove that if $x^n \bmod m \neq 1$ for $n \in \{1,\ldots,m-1\}$, then we have $x^{n_1} \bmod m \neq x^{n_2} \bmod m$ for distinct $n_1, n_2 \in \{1,\ldots,m-1\}$. Use Exercise 20 to conclude that there exists an $n \in \{1,\ldots,m-1\}$ such that $x^n \bmod m = 1$.

Given positive integers $x_1,\ldots,x_n$, their least common multiple $\operatorname{lcm}\{x_1,\ldots,x_n\}$ is defined as the smallest positive integer that is divided by each of the numbers $x_1,\ldots,x_n$.

THEOREM 6. (Carmichael 1910) Consider a positive integer m with prime number factorization $m = p_1^{e_1} \cdot \ldots \cdot p_t^{e_t}$. The multiplicative congruential generator $X_n = (a\,X_{n-1}) \bmod m$ has maximal period-length $\operatorname{lcm}\{\lambda(p_1^{e_1}),\ldots,\lambda(p_t^{e_t})\}$, where

$\lambda(\cdot)$ is given by

$$\lambda(k) = \begin{cases} 2^{e-1} & \text{for} \quad k = 2^e \;\text{ with }\; e \leq 2 \\ 2^{e-2} & \text{for} \quad k = 2^e \;\text{ with }\; e > 2 \\ (p-1)p^{e-1} & \text{for} \quad k = p^e \;\text{ with }\; p \;\text{ odd} \end{cases} .$$

The maximal period-length is achieved if the following two conditions hold:

(i) The numbers m and X_0 are relatively prime.

(ii) The number a is a primitive element mod m.

Exercise 22. Find the maximal period-length for a multiplicative congruential generator with module m in the cases when m is an odd prime number, and when $m=10^e$ with $e \geq 4$.

On pages 20 and 21 in the text by Knuth, Theorems C and D (together with the text in between), it is described how to find an a that is a primitive element mod m in the important cases when $m=2^e$ and $m=10^e$. These results are quite deep, and are due to Knuth himself.

We conclude this section with a few more citations of Knuth, this time from the summary in Section 3.6 of his book:

"The multiplier a should preferably be chosen between $0.01\,m$ and $0.99\,m$, and its binary or decimal digits should *not* have a simple, regular pattern."

"At most $m/1000$ numbers should be generated; otherwise the future will behave more and more like the past."

"The seed X_0 may be chosen arbitrarily."

The last of these advices of Knuths' is a bit careless: The seed will indeed sometimes affect the perfomance of a linear congruential generator (see Exercises 23 and 24 below). What Knuth really means is that the seed may be chosen arbitrarily for a generator with maximal period-length $\mathcal{P}(m,a,c,X_0) = m$ (see Exercise 25 below).

Exercise 23. Shew with examples that the actual choice of the seed X_0 can affect the performance of a multiplicative generator dramatically.

Exercise 24. Let m be an even positive integer. Will the period-length $\mathcal{P}(m,a, m/2,X_0)$ depend on the choice of the seed X_0?

Exercise 25. Explain with a "direct argument" (which do not rely on the Hull-Dobell Theorem) why it is that the seed may be chosen arbitrarily for a generator with maximal period-length $\mathcal{P}(m,a,c,X_0) = m$. In other words, shew that if $m,a,c \in$

$\{0,\ldots,m-1\}$ satisfy $\mathcal{P}(m,a,c,X_0)=m$ for some $X_0\in\{0,\ldots,m-1\}$, then we have $\mathcal{P}(m,a,c,X_0)=m$ for each $X_0\in\{0,\ldots,m-1\}$.

4.4 Potency

The "output" $X_n=(a\,X_{n-1}+c) \bmod m$ of a linear congruential generator is a polynomial in n when viewed as a function of n (recall Exercise 14). Clearly, it is desirable that the degree of such a polynomial is high. The concept of potency "measures" the degree of polynomial complexity of a linear congruential generator, or in other words, simply the minimal degree of a polynomial that describes X_n as a function of n.

Definition. Consider a linear congruential generator $X_n=(a\,X_{n-1}+c) \bmod m$ that has the maximal period-length $\mathcal{P}(m,a,c,X_0)=m$. The potency of the generator is the number (cf. Exercise 26 below)

$$\underline{\mathcal{S}(m,a)}\equiv\min\Big\{s\in\mathbb{N}\backslash\{0\}:(a-1)^s \bmod m=0\Big\}.$$

Exercise 26. Consider a linear congruential generator $X_n=(a\,X_{n-1}+c) \bmod m$ that has period-length $\mathcal{P}(m,a,c,X_0)=m$. Prove that there exists a positive integer s such that $(a-1)^s \bmod m=0$.

Exercise 27. Consider a linear congruential generator $X_n=(a\,X_{n-1}+c) \bmod m$ that has period-length $\mathcal{P}(m,a,c,X_0)=m$. Let $s=\mathcal{S}(m,a)$ be the potency of the generator. Prove that

$$X_n=c\left(n+\binom{n}{2}(a-1)+\ldots+\binom{n}{\min\{s,n\}}(a-1)^{\min\{s,n\}-1}\right) \bmod m \qquad \text{for}\quad n\geq 0$$

when the seed is $X_0=0$. Further prove that, to each choice of a seed $X_0\in\{0,\ldots,m-1\}$, there exists an integer $n_0=n_0(X_0)\in\{0,\ldots,m-1\}$ such that

$$X_n=c\left((n+n_0)+\binom{n+n_0}{2}(a-1)+\ldots+\binom{n+n_0}{\min\{s,n\}}(a-1)^{\min\{s,n\}-1}\right).$$

5 Empirical Tests of Random Number Generators

An empirical test of a sequnce of (pseudo-) random numbers $\{U_i\}_{i\geq 1}$ is a test procedure that uses established statistical methods to test for deviations from the hypothesis that the dataset $\{U_i\}$ consists of independent observations of a random variable U that is uniformly distributed over $(0,1)$.

Among the empirical tests described below, the collision test and the run test are especially recommended by Knuth. We discuss other weaker tests as well, because they illuminate the problem of testing random number generators, and because one simply should know about them.

5.1 Frequency Methods

Frequency methods test deviations from a uniform distribution using χ^2-tests or Kolmogorov-Smirnov tests of consecutive random numbers $U_1,\ldots,U_n$. Here, contrary to what is usually the case, the number n of random numbers (observations) used in the test should not be too large, because possible "local deficiencies", which are really what one is looking for (since even an only remotely sensible generator will perform well on the "global scale"), tend to cancel out when n is large: As a rule, Knuth suggests that n should not be much larger than 1000.

In a χ^2-test the interval $[0,1)$ is divided into k sub-intervals $[0,\frac{1}{k}),\ldots,[\frac{k-2}{k},\frac{k-1}{k})$, $[\frac{k-1}{k},1)$ in such a way that the following familiar rule of thumb holds

$$n\cdot\mathbf{P}\{\text{a random number takes a value in a certain interval}\} = n/k \geq 5 \ (1)$$

(preferably with a margin). The frequencies $f_1,\ldots,f_k$ for (the total number of) appearances of random numbers in the different intervals are used to compute the

$$\text{test-quantity} \qquad T\equiv\sum_{i=1}^{k}\frac{(f_i-n/k)^2}{n/k}.$$

The test-quantity T is approximately $\chi^2(k-1)$-distributed for n sufficiently large (the rule of thumb ...) when the observations $U_1,\ldots,U_n$ are "real random numbers".

A Kolmogorov-Smirnov test uses the empirical distribution function based on the observations $U_1,\ldots,U_n$ of a random variable U, which is given by

$$\tilde{F}_n(x) = \frac{\#\{i\in\{1,\ldots,n\}: U_i\leq x\}}{n} \qquad \text{(where } \# \text{ stands for cardinality).}$$

Exercise 28. Shew that if η is a random variable that is uniformly distributed over $\{1,\ldots,n\}$, then the random variable U_η has the empirical distribution of U.

Lemma. (The Glivenko-Cantelli Theorem)

$$D_n\equiv\sup_{-\infty<x<\infty}|\tilde{F}_n(x)-F_U(x)|\to_{\text{a.s.}} 0 \qquad \text{as} \quad n\to\infty.$$

Exercise 29. Use the strong law of large numbers to prove the weaker result

$$|\tilde{F}_n(x)-F_U(x)|\to_{\text{a.s.}} 0 \qquad \text{as} \quad n\to\infty \quad \text{for each} \quad x\in\mathbb{R}.$$

Exercise 30. Let $U_1, \ldots, U_n$ be observations of a continuously distributed random variable U. Let $U_{(i)}$ denote the i'th smallest observation among $U_1, \ldots, U_n$. Prove the fact essential for the computation of D_n that

$$D_n = \max\Big\{\max_{1\leq i\leq n}\Big(F_U(U_{(i)}) - \tfrac{i-1}{n}\Big),\ \max_{1\leq i\leq n}\Big(\tfrac{i}{n} - F_U(U_{(i)})\Big)\Big\}.$$

THEOREM 7. (The Kolmogorov-Smirnov Test) Suppose that $U_1, \ldots, U_n$ are independent observations of a random variable U with distribution function F_U. As $n \to \infty$, the test-quantity

$$T \equiv \sqrt{n}\, D_n \to_{\text{weakly}} \tilde{D} \qquad \text{where} \qquad F_{\tilde{D}}(x) = 1 - 2\sum_{k=1}^{\infty}(-1)^{k-1}\mathrm{e}^{-2k^2x^2} \quad \text{for} \quad x > 0.$$

For small values of n, the asymptotic distribution of $\tilde{D}$ cannot be used. However, there are tables providing (exact) percentage points for the distribution of $\sqrt{n}\, D_n$ for n small. The difference between the asympotic percentage points (of $\tilde{D}$) and those of $\sqrt{n}\, D_n$ is bounded by $\frac{1}{3}n^{-1/2}$ plus smaller order terms (as $n \to \infty$).

The test-quantity $\sqrt{n}\, D_n$ does not depend on the actual distribution of U. Hence the Kolmogorov-Smirnov test applies to test whether $U_1, \ldots, U_n$ are independent observations of a random variable U with distribution function F_U for any continuous distribution function F_U.

As is the case with other frequency tests, when using the Kolmogorov-Smirnov test on a random number generator, the sample size n should not be too large (not much more than a 1000 say).

A complete frequency test procedure consists of a number N (10 say, or more) of different frequency tests, each based on a sample size n (not much more than 1000). If the observations are real random numbers (how could they be?), then the corresponding N test-quantities $T_1, \ldots, T_N$ will be independent observations of a χ^2-distribution when it is χ^2-tests that are made, or of the random variable $\sqrt{n}\, D_n$ (that is asymptotically $\tilde{D}$) when Kolmogorov-Smirnov tests are made, respectively.

The final "overall test" is carried out by means of computing the Kolmogorov-Smirnov test-quantity $\sqrt{N}\, D_N$ for the deviation of the observations $T_1, \ldots, T_N$ from the χ^2-distribution and the $\sqrt{n}\, D_n$-distribution in the cases when χ^2-tests and Kolmogorov-Smirnov tests are made, respectively. In Theorem 7, the distribution function F_U should then be replaced with $F_{\chi^2(k-1)}$ and $F_{\sqrt{n}\, D_n}$, respectively.

Note that the scheme suggested above possibly can detect both local and global deviations from the basic assumption about independence and uniform distribution for the random numbers.

5.2 Serial Test

Given random numbers $U_1, \ldots, U_{nN}$, we are going to test deviations from a uniform distribution over the N-dimensional interval $[0,1)^N$ for the random vectors $(U_1, U_2, \ldots, U_N)$, $(U_{N+1}, U_{N+2}, \ldots, U_{2N})$, $\ldots$, $(U_{(n-1)N+1}, U_{(n-1)N+2}, \ldots, U_{nN})$ using χ^2-tests. To that end the interval $[0,1)^N$ is divided into k^N sub-intervals

$$[\tfrac{\ell_1-1}{k}, \tfrac{\ell_1}{k})\times[\tfrac{\ell_2-1}{k}, \tfrac{\ell_2}{k})\times\ldots\times[\tfrac{\ell_N-1}{k}, \tfrac{\ell_N}{k}) \qquad 1 \leq \ell_1, \ldots, \ell_N \leq k,$$

observing the rule of thumb $n/k^N \geq 5$ (1). The frequencies $\{f_{\ell_1,\ldots,\ell_N}\}_{1\leq\ell_1,\ldots,\ell_N\leq k}$ for the appearances of random vectors in the different sub-intervals determine the

$$\text{test-quantity} \qquad T \equiv \sum_{\ell_1,\ldots,\ell_N=1}^{k} \frac{(f_{\ell_1,\ldots,\ell_N} - n/k^N)^2}{n/k^N}.$$

The test-quantity T is approximately $\chi^2(k^N-1)$-distributed for n large when $U_1, \ldots, U_{nN}$ are "real random numbers": This test is called serial test.

A serial test can detect deviations from a uniform distribution, as wells as (very) local deviations from independence. As is the case with fequency tests, the overall test procedure should involve several (at least 10 say) different serial tests. The Kolmogorov-Smirnov test is employed to evaluate the outcome of the test procedure as before.

5.3 Collision Test

The dimension N that can be used in a serial test have to be quite small because of the rule of thumb for the χ^2-test, which means that independence is only tested very locally. This is the reason for the development of the collision test described below. The collision test is an exact test, in the sense that the null distribution (for the test-quantity given the null hypothesis about independence and uniform distributions) can be computed without approximations (asymptotic theory).

Given random numbers $U_1, \ldots, U_{nN}$, we are going to test deviations from a uniform distribution over the N-dimensional interval $[0,1)^N$ for the random vectors $(U_1, U_2, \ldots, U_N)$, $(U_{N+1}, U_{N+2}, \ldots, U_{2N})$, $\ldots$, $(U_{(n-1)N+1}, U_{(n-1)N+2}, \ldots, U_{nN})$. Again we divide the interval $[0,1)^N$ into k^N sub-intervals

$$[\tfrac{\ell_1-1}{k}, \tfrac{\ell_1}{k})\times[\tfrac{\ell_2-1}{k}, \tfrac{\ell_2}{k})\times\ldots\times[\tfrac{\ell_N-1}{k}, \tfrac{\ell_N}{k}) \qquad 1 \leq \ell_1, \ldots, \ell_N \leq k,$$

this time, however, making sure that $n/k^N << 1$ (rather than the opposite inequality required for a serial test). The test-quantity is

$T_n \equiv$ 'the total number of sub-intervals that contains at least one random vector'.

The event $\{T_n = n-\ell\} = \{\omega \in \Omega : T_n(\omega) = n-\ell\}$ occurs precisely when ℓ of the random vectors $\{(U_{(i-1)N+1}, \ldots, U_{iN})\}_{i=1}^n$ that are "trown" into the interval $[0,1)^N$ (consecutively) fall into a sub-interval that has already been hit by an earlier random vector, that is, precisely when ℓ collisions take place. When the random vectors $\{(U_{(i-1)N+1}, \ldots, U_{iN})\}_{i=1}^n$ are uniformly distributed over $[0,1)^N$ (as is the case when $\{U_i\}_{i\geq 1}$ are real random numbers), it is a relatively straightforward combinatorial problem to determine the "null distribution" for T_n:

Exercise 31. The Stirling numbers of the second kind are given by

$$\left\{ \begin{matrix} n \\ \ell \end{matrix} \right\} \equiv \sum_{1 \leq i_0 \leq \ldots \leq i_{n-\ell} \leq \ell} i_0 \cdot \ldots \cdot i_{n-\ell} \, .$$

Prove the recursion formula

$$\left\{ \begin{matrix} n \\ \ell \end{matrix} \right\} = \ell \left\{ \begin{matrix} n-1 \\ \ell \end{matrix} \right\} + \left\{ \begin{matrix} n-1 \\ \ell-1 \end{matrix} \right\}. \qquad \text{(Does this formula ring a bell?)}$$

Exercise 32. Derive a recursive formula where the density function $f_{T_n}(\ell) = \mathbf{P}\{T_n = \ell\}$ of the test-quantity T_n is expressed in terms of the density function $f_{T_{n-1}}$ of T_{n-1} (under the null hypothesis, of course).

Exercise 33. Use Exercises 31 and 32 to give a proof by induction of the fact that

$$f_{T_n}(\ell) = \mathbf{P}\{T_n = \ell\} = \frac{k^N \cdot (k^N-1) \cdot \ldots \cdot (k^N-\ell+1)}{k^{Nn}} \left\{ \begin{matrix} n \\ \ell \end{matrix} \right\} \quad \text{for} \quad \ell \in \{1, \ldots, n\}.$$

The integers k^N, ℓ och n that appear in the expression for the density function of T_n in Exercise 33 can be very large. However, in case this should cause any numerical problems, Knuth provides an algorithm that makes it possible to compute percentage points for T_n avoiding such problems on page 71 in his book.

5.4 Run Test

A run test consists of collecting independent observations $R_1, \ldots, R_n$ of the random variable

$$R^{\uparrow} \equiv \max\{j : U_1 < U_2 < \ldots < U_j\}$$

by means of the recursive procedure

$$R_{i+1} = S_{i+1} - (S_i + 1) \quad \text{where} \quad S_{i+1} = \max\{j \geq S_i + 2 : U_{S_i+2} < U_{S_i+3} < \ldots < U_j\},$$

where $S_0 \equiv -1$. Alternatively, one collects independent observations of the random variable

$$R^{\downarrow} \equiv \max\{j : U_1 > U_2 > \ldots > U_j\},$$

using the scheme

$$R_{i+1} = S_{i+1} - (S_i + 1) \quad \text{where} \quad S_{i+1} = \max\{j \geq S_i + 2 : U_{S_i+2} > U_{S_i+3} > \ldots > U_j\}.$$

A χ^2-test is used to test deviations from (the null hypothesis about) "real random numbers". Because of the rule of thumb, all observations $R_i \geq r_{\max}$ must be collected in a common category, for some choice of $r_{\max}$. Knuth suggests that $r_{\max} = 6$ is used.

Exercise 34. Explain why it is that one do not use the following seemingly "more efficient" procedure (producing more observed runs) to collect observations of runs:

$$R_{i+1} = S_{i+1} - S_i \quad \text{where} \quad S_{i+1} = \max\{j \geq S_i + 1 : U_{S_i+1} < U_{S_i+2} < \ldots < U_j\}.$$

Exercise 35. Determine the common "null distribution" for $R^{\uparrow}$ and $R^{\downarrow}$.

5.5 Further Empirical Tests

Knuth describes several other tests that are more or less equivalent with the series test or the run test. Examples of such tests are gap test, poker test, coupon collectors test, permutation test and maximum-of-t test. There seems to be little to gain from carrying out also these tests.

Knuth discusses a serial correlation test, the idea of which is to test the sample-correlation between two consecutive random numbers for significant deviation from zero. On the personal level, I am a great disbeliver in all statistical methods relying on covariances and correlations, especially when they are employed to make inference about dependence structure, as here. (Obviously, there are a lot of talented people around who belives the opposite.) It is thus with some satisfaction that I note that Knuth states (and even proves) that correlation tests are among the very few weakest tests in the context of testing random number generators (see below). See also Example 2.

Here are a few citations from pages 74-75 in the book by Knuth:

"Of all the test we have discussed, the frequency test and the serial correlation test seem to be the weakest, in the sense that nearly all random number generators pass them. Theoretical grounds for the weakness of these tests are discussed briefly in Section 3.5".

"The run test, on the other hand, is rather strong ...".

"The collision test is also highly recommended, since it has been specially designed to detect the deficiencies of many poor generators ...".

"The need for making several tests has been amply documented."

"As computers become faster, more random numbers are consumed than ever before, and random number generators that once were satisfactory are no longer good enough for sophisticated applications ...".

Note that Knuth gives a referenece for some new developments in empirical testing of random number generators at the very end of page 75.

Example 2. (A VERY POOR TEST OF A VERY POOR GENERATOR) Consider the multiplicative generator $X_i = (65539\, X_{i-1}) \bmod 2^{31}$ with an odd seed, $X_0 = 314159$ say: This is the famous `RANDU`-generator, that was used on 32-bit machines "in most of the world's scientfic computing centras for more than a decade", around the 70's. The "very name `RANDU` is enough to bring dismay into the eyes and stomachs of many computer scientists!" For example, "the generator fails most three-dimensional criteria for randomness, and it should never have been used. Almost any multiplier ... would be better." (Citations are from Section 3.3.4 in Knuth's book.)

If the sequence of random numbers $U_i = X_i/2^{31}$, $i \geq 1$, from this generator consists of independent random variables with a uniform distribution over $(0,1)$, then the sequence has covariance function $r_U(j) = \mathbf{Cov}\{U_i, U_{i+j}\} = \sigma^2\, \delta(j)$ for $j \in \mathbb{Z}$, where $\sigma^2 = \frac{1}{12} = 0.0833...$ is the variance of the uniform distribution. This covariance function has spectral representation

$$r_U(j) = \int_{-1/2}^{1/2} \mathrm{e}^{i2\pi fj}\, \mathcal{R}_U(f)\, df \qquad \text{with spectral density} \qquad \mathcal{R}_U(f) = \sigma^2$$

(the spectral density of discrete-time white noise).

In the framework of methods based on covariances and correlations (second-order methods), it is natural to test the random number generator by means of estimating its spectral density, and check whether the estimate corresponds to $\mathcal{R}_U(f) = \frac{1}{12}$ or not: Given $1\,000\,000$ observations $\{U_i\}_{i=1}^{1000000}$ say, a standard way to estimate $\mathcal{R}_U(f)$ is to use the so called periodogram estimator

$$\mathcal{R}_U^{\star}(f) = \frac{1}{2000} \sum_{k=1}^{2000} \left| \frac{1}{500} \sum_{j=1}^{500} \mathrm{e}^{-i2\pi fj}\, (U_{j+500\,(k-1)} - \overline{U}) \right|^2 \qquad \text{for} \qquad f \in (-1/2, 1/2).$$

(Of course, here the factorization $1\,000\,000 = 2\,000 \times 500$ is quite arbitrary.) In order to detect deficiencies of the generator [that is, deviations from the assumption that $\{U_i\}_{i=1}^{1000000}$ are independent with a uniform distribution over $(0,1)$], one checks for deviations from the "perfect" $\mathcal{R}_U^\star(f)=\frac{1}{12}$ for the estimated spectral density.

However, when all this was implemented in a computer, it turned out that there were no obvious deviations from $\mathcal{R}_U^\star(f) = \frac{1}{12}$ at all (see below). Since we know that the generator is very poor (and "... fails most three-dimensional criteria ..."), a natural conclusion is that the second-order method at hand (which is used very extensively) is a very weak criterion.

Here is the Mathematica-programme that was used to carry out the test:

```
In[1]:= u=Table[0,{1000000}]
In[2]:= b=2^(31); x=314159
In[3]:= For[j=1, j<=1000000, j++, x=Mod[65539*x,b]; u[[j]]=N[x/b]]
In[4]:= m=Sum[u[[j]],{j,1,1000000}]/1000000
In[5]:= Plot[Sum[(Sum[Cos[2*Pi*f*j]*(u[[j+(k-1)*500]]-m),{j,1,500}]
        )^2/500+(Sum[Sin[2*Pi*f*j]*(u[[j+(k-1)*500]]-m),{j,1,500}]
        )^2/500,{k,1,2000}]/2000,{f,-0.5,0.5},AxesLabel->{"f",
        "R(f)"},PlotRange->{-0.01,0.13},Ticks->{{{-0.5,""},{-0.4,
        "-0.4"},{-0.3,""},{-0.2,"-0.2"},{-0.1,""},{0.1,""},{0.2,
        "0.2"},{0.3,""},{0.4,"0.4"},{0.5,""}},{{0.025,""},{0.05,
        "0.05"},{0.075,""},{0.1,"0.10"},{0.125,""}}}]
```

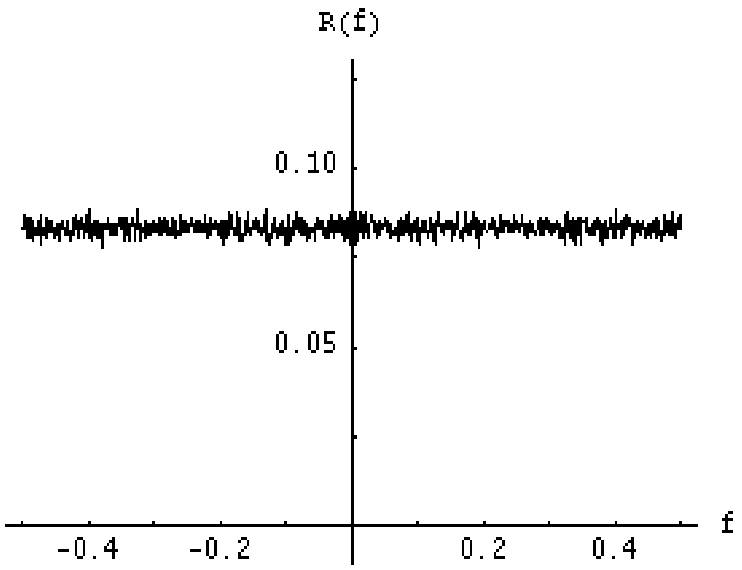

```
Out[5]= - Graphics -
```

6 More on Testing of Random Number Generators

6.1 Theoretical Tests

Theoretical tests use analytical methods (mainly number theory and related fields) to derive upper and lower bounds for the test-quantities that are used in empirical tests, expressed in terms of the parameters m a, c and X_0 of the linear congruential generator. Of course, this means that these empirical tests never have to be carried out in practice (if the bounds are good).

The theory for theoretical tests, that is, the art of deriving sharp theoretical bounds for empirical test-quantities, is difficult, and has not developed sufficiently in order to be really useful. "The results obtained so far are generally for *statistical tests made over the entire period*.": We have already remarked that this tend to make local deficiencies cancel out. In fact, one should never use a generator for more than a small proportion of its entire period (as was pointed out in Section 4).

However, here is one "hybrid-theoretical test" that is very important as well as well-developed:

6.2 Spectral Test

The spectral test is, according to Knuth, "by far the most powerful test known" for detecting poor generators. "Not only do all good generators pass this test, all generators now known to be bad actually *fails* it."

The spectral test is theoretical in the sense that the test-quantity is a well-defined deterministic quantity, which is a (well-defined) function of the parameters m a, c and X_0 of the congruential generator. However, the test is also empirical in the sense that (at present) this function cannot be expressed in closed form, and has to be calculated through extensive computations in a computer. Knuth provides a (flow-plan for a) programme that performs this task in Section 3.3.4.D (pages 102-103) of his book. The reader is likely to find that programme quite complicated!

The spectral test may seem a bit "ad hoc", but there is a soild theoretical base for it, which unfortunately is a bit complicated. See Section 3.3.4.B in Knuths's book for snap-shots of parts of the theory, and the book by Niederreiter mentioned below (if Knuth's text does not hurt enough).

Consider the family $\mathbf{U}$ of consecutive t-tuples $\{(U_i, U_{i+1}, \ldots, U_{i+t-1})\}_{i\geq 1}$ that is generated by the linear congruential generator under consideration. For simplicity, it is assumed that the generator has the maximal period-length $\mathcal{P}(m, a, c, X_0) = m$.

Given a family $\mathbf{H}$ of parallell hyperplanes in t-dimensional hyperspace, the union of which contains $\mathbf{U}$, define

$$d(\mathbf{H}) \equiv \text{'the minimal distance between two hyperplanes in } \mathbf{H}\text{'}.$$

The test-quantity of the spectral test is the precision

$$\nu_t \equiv \Big(\max\{d(\mathbf{H}) : \mathbf{H} \text{ is a family of parallell hyperplanes in } \mathbb{R}^t \text{ that contains } \mathbf{U}\}\Big)^{-1}.$$

Note that a large value of $d(\mathbf{H})$ (for a family $\mathbf{H}$ of parallell planes that contains $\mathbf{U}$) means that $\mathbf{U}$ "lives" on comparatively few and distant planes: This is undesirable, and gives a small precision ν_t [the reciprocal of the largest $d(\mathbf{H})$ possible].

The m different vectors of $\mathbf{U}$ take values in the lattice (≈periodic structure) $\{\frac{0}{m}, \frac{1}{m}, \ldots, \frac{m-1}{m}\}^t$. For a sequence of m vectors with values in such a lattice, we have the following theoretical upper bound for the precision:

$$\nu_t \leq \gamma_t\, m^{1/t}, \qquad \text{where} \quad \gamma_t \quad \text{is a real number of the order of magnitude } 1.$$

See page 109 in Knuth's book for the exact values of γ_t. Because of periodicity, there are only m different vectors in $\mathbf{U}$: This readily gives the rate $m^{1/t}$ of decay with t.

Thus (the upper bound for) the best possible precision will decay with increasing dimensionen t for a sequence of pseudo-random numbers. For "real random numbers", rounded off to take values in the lattice $\{\frac{0}{m}, \frac{1}{m}, \ldots, \frac{m-1}{m}\}^t$, the precision will instead be almost independent of the dimension t.

The lack of "high-dimensionality" for pseudo-random numbers, that is implied by the inequality $\nu_t \leq \gamma_t\, m^{1/t}$, have already been illustrated by means of the concept of potency in Section 4.

In order not to make things worse that they have to be, a generator with high precision should be chosen; preferably the precision should be of the same order of magnitude as is the general (not necessarily attainable) upper bound $\gamma_t\, m^{1/t}$ for the precision: Knuth suggests the following test criterion when using the spectral test:

$$\nu_t \geq 2^{\,30/t} \qquad \text{for} \quad t \in \{2, \ldots, 6\}.$$

Exercise 36. Give a theoretical lower bound for the module m of a linear congruential generator that passes the test criterion suggested by Knuth.

In Section 3.3.4.B of Knuth's book, elementary geometry is used to prove that

$$\nu_t = \min\Big\{|(x_1, \ldots, x_t)| : (x_1, \ldots, x_t) \in \mathbb{N}^t \backslash \{0\},\ (x_1 + a x_2 + \ldots + a^{t-1} x_t) \bmod m = 0\Big\}.$$

This expression can be used to calculate ν_t in a computer. However, first some theory for minimizing discrete expressions in a computer have to be developed, in order to limit the number of different vectors $(x_1, \ldots, x_t)$ that have to be searched. This is done in Section 3.3.4.C of Knuth's book. Finally, as mentioned before, a sketch of a programme for calculating ν_t is given in Section 3.3.4.D.

6.3 Further Reading

The final four chapters in the monograph "H. Niederreiter (1992): Random Number Generation and Quasi-Monte Carlo Methods. SIAM." give a reasonably easy-to-read overview (including an extensive bibliography) of some deeper topics in the art of generating random numbers in a computer. In particular, one chapter deals with non-linear congruential generators, and discuss how they avoid some deficiencies that are associated with linear generators (as for example low precison in the spectral test). The theoretical analysis of these generators is very complicated indeed, and Niederreiter is one authority (of quite few it seems) in the area.

6.4 A Programming Project

Programming Project 1. Construct a linear congruential generator that passes all the empirical tests described in Sections 5.1-5.4, as well as the theoretical spectral test of Section 6.2. At the same time, test the `RANDU`-generator discussed in Example 2 and make comparisons. The test results that are listed by Knuth may not be used, so that the tests really have to be implemented in a computer and carried out in practice. (Of course, the results of Knuth can be used at an "initial stage" to check that programming have been correctly done.)

7 Personal Thoughts on Software for Stochastic Simulation

There are three natural choices of software for completing programming projects in Stochastic Simulation. They are `S`$^+$ (activated with `UNIX`-command `Splus` at the math-computer-system), `Maple V` and `Matlab` (activated by selecting `Maple V` with the left mouse-button), and `Mathematica` (activated by selecting `Mathematica` with the left mouse-button). [A much more advanced interface for `S`$^+$ than you get with this (similar to that of `Mathematica`) do exist.]

`S`$^+$ is a "language" that is very nice to use in simulation contexts. In addition, `S`$^+$ has a very powerful module for "traditional statistics" (not far behind those of for example `SAS` and `SPSS`), which is not the case with `Mathematica`, `Maple V` and `Matlab`.

However, `S`$^+$ is a "university type of language", and does not have the "backing and support" by a multi-national organization that comes with "standard statistical software", like for example `SAS` and SPSS. This means that S^+ cannot be thought of as a replacement for these language (as long as it lacks that backing and support), because it is that very "reliability" a typical ("non-university") user is really looking for. Therefore the "powerful traditional statistics" module of S^+ will not "help the package" as much as perhaps could be expected.

A very serious deficiency of S^+ is that the S^+-programs cannot be compiled, and thus always have to be executed in an interpreting mode. This seriously affects the usefulness of the program in "real" ("non-classroom") simulations: It is simply too slow - by a factor of 1000 or so.

S^+ does not have a serious contender as a programme-package targeted towards statisticians, and covering both traditional statistics as well as stochastic simulation aspects of that subject. Therefore it does not develop in the same way as does `Mathematica`, `Maple V` and `Matlab` (from which you get comercial mail about new features and developments virtually every month in your mail-slot).

As can be heard from their names, `Mathematica`, `Maple V` and `Matlab` are packages intended for mathematicians in general, rather than for statisticians in particular. This means that although they have well-developed modules for traditional statistics, they typically cannot really challenge S^+ in that particular area (or indeed SAS and SPSS).

Looking away from traditional statistics, which is in fact not important at all in the context of stochastic simulation, although very nice to use in programing, S^+ seriously suffers from its inability to render itself to compilation (and the lack of speed following from that). In addition, S^+ does not have the powerful numerical toolboxes that come with the mathematical packages, and which certainly are needed in the context of stochastic simulation. Seemingly, presently, and perhaps unfortunately, these observations make the mathematical packages more natural choices for stochastic simulation than is S^+.

`Maple V` and `Matlab` traditionally have a very strong position among mathematicians, especially when considering numerical mathematics. However, when it comes to formal (analytical) mathematical ability, `Mathematica` appears to be quite superior at present (from what I have made out from discussions with people knowledgeable in this aspect of `Maple V` and `Matlab`). So while `Mathematica` can in fact be used as an aid in some research-type analytical problems, this is not really the case with `Maple V` and `Matlab` (yet).

The surely incomplete as well as personally biased list of evaluation criteria presented above speeks for `Mathematica`. However, that programme does have the disadvantage of being a "one mans invention" (Stephen Wolfram), more than is the case with `Maple V` and `Matlab`. On the other hand, this applies also to TeX (Donald Knuth), and there it has worked well. (I have been shown a photography of Donald Knuth dressed in a `Mathematica` T-shirt.)

Programme-packages are suprisingly easy to learn nowadays. This is a part of the very strong current development aiming at making IT available to everyone. [Obviously, immensely strong economical interests stand behind all this (which does not in anyway imply that it is an undesirable development, rather it is of course very positive).] The effect is that one can learn to use (if not fully master) the programmes discussed above with little effort (compared with the situation 10 to 15 years ago, say). Not only has the programme languages themselves developed strongly towards ease of use, but more importantly, the support for the user via user-friendly interfaces has really dramatically affected the ease of use towards amazing simplicity.

Despite the user-friendliness, of course, there is a great internal momentum against changes of programmes. In particular, at universities a lot of pedagogical material has been developed with a certain programme in mind, and it is very inconvenient to switch to another programme. (Isn't this also how things work on the personal level for most of us as well?)

8 α-stable Random Variables

8.1 Introduction and Brief History of the Subject

Analytical properties of α-stable random variables, and the associated so called "domains of attraction", were a very active area of research during the 40's and 50's. (One important name from this period is the retired professor with our department Harald Bergström.) During the 60's and 70's focus turned to more general infinitely divisible laws, and (more importantly) their associated Lévy processes (that is, processes with stationary and independent increments). The 80's and 90's have seen a rebirth of the interest in α-stability. A few reasons for this are listed below:

(i) It has become increasingly clear that an α-stable stochastic process can be something very much more general than merely an α-stable Lévy process. The study of general α-stable processes offers important and challenging problems. (Holger Rootzén was one of the first names here.)

(ii) The theory for Gaussian processes is a "dead end" in the sense that it is not at all extendable to other kind of processes: The theory is very special and very simple beacuse Gaussian processes are very special and very simple objects. This has to do with the important role of the covariance function [which as a rule is of very little use (if any at all) for more general processes], together with the fact that dependence can be described by means of linear regression (again a unique feature of Gaussians). α-stable processes, on the other hand, are studied with methods that have turned out to be useful also within the much more general class of infinitely divisible processes (which do in fact include virtually every non-pathological process). So, as well as being important on its own, the theory for α-stable processes serve as a bridge to theory for more general processes.

(iii) Modern applications in areas such as weather forecasts, hydrology, telecommunication, finacial markets, and so on, require models that can accomodate phenomena like "heavy tails", "long-range dependence", "self-similarity" and "non-linearity". Traditional models, based on Gaussian processes, either are completely unable to do this (in the cases of "heavy tails" and "non-linearity"), or else are at least comparatively "unwilling" to do so (in the cases of "long-range dependence" and "self-similarity"). α-stable processes, on the other hand, automatically possess three of these properties ("heavy tails", "long-range dependence" and "non-linearity"), and do in addition offer a much more versatile array of self-similar processes that do Gaussians.

The standard reference for α-stable random variables and (especially) α-stable stochastic processes is "G. Samorodnitsky and M. Taqqu (1994): Stable Non-Gaussian Random Processes". This book is quite easy to read, as well as very complete with respect to its content. (I subscribe this rare combination to the "down-to-earth" approach the authors have to research.) A less ambitious reader can see "L. Breiman (1968): Probability", while "M. Ledoux and M. Talagrand (1991): Probability in Banach Spaces" should be a fitting challenge for pure mathematicians.

The standard reference for simulation of α-stable stochastic processes is "A. Janicki and A. Weron (1994): Simulation and Chaotic Behavior of α-stable Stochastic Processes". However, there are a lot of quite serious errors in this book (sometimes on a suprisingly "basic level"), and it must therefore be read with considerable care. (In fact, everything that is used from this book should be double-checked!) Nevertheless, the book is (or was) very important, and has initiated a lot of current research. "A. Janicki, K. Podgórski and A. Weron (1993): Computer Simulation of α-stable Ornstein-Uhlenbeck Processes. In: Stochastic Processes, S. Cambanis, J.K. Ghosh, R.L. Karandikar and P.K. Sen Eds." offers an independent and quite elementary introduction to the book by Janicki and Weron.

8.2 Univariate α-stable Random Variables

Definition. (i) A univariate random variable X is $\underline{\alpha\text{-stable}}$ if given any $m \in \mathbb{N}$ together with independent copies $X_1, \ldots, X_m$ of X, there exists a constant $d_m \in \mathbb{R}$ such that

$$X_1 + \ldots + X_m =_{\text{distribution}} m^{1/\alpha} X + d_m.$$

(ii) A univariate random variable X is $\underline{\text{strictly } \alpha\text{-stable}}$ if given any $m \in \mathbb{N}$ together with independent copies $X_1, \ldots, X_m$ of X, we have

$$X_1 + \ldots + X_m =_{\text{distribution}} m^{1/\alpha} X.$$

The label "stable" alludes to stability under addition (except for a linear scaling).

It turns out that the possible values for α are $\alpha \in (0, 2]$:

Exercise 37. Explain intuitively, using for example the Central Limit Theorem, why it is that $\alpha \in (0, 2]$.

Property 1. (Domains of Attraction) A random variable X is α-stable for some $\alpha \in (0,2]$ if and only if there exist independent identically distributed random variables $Y_1, Y_2, \ldots$ together with real numbers $a_m > 0$ and $d_m \in \mathbb{R}$ such that

$$a_m(Y_1 + \ldots + Y_m) + d_m \rightarrow_{\text{weakly}} X \qquad \text{as} \quad m \to \infty.$$

Property 2. A random variable X is α-stable if and only if

$$\mathbf{E}\{\exp[i\theta X]\} = \begin{cases} \exp\Big\{-\sigma^\alpha|\theta|^\alpha\Big(1 - i\beta\,\mathrm{sign}(\theta)\tan(\tfrac{\pi\alpha}{2})\Big) + i\mu\theta\Big\} & \text{when} \quad \alpha \neq 1 \\ \exp\Big\{-\sigma^\alpha|\theta|^\alpha\Big(1 + i\beta\tfrac{2}{\pi}\,\mathrm{sign}(\theta)\ln|\theta|\Big) + i\mu\theta\Big\} & \text{when} \quad \alpha = 1 \end{cases}$$

for some <u>scale parameter $\sigma > 0$</u>, <u>location parameter $\mu \in \mathbb{R}$</u> and <u>skewness parameter $\beta \in [-1,1]$</u>: This is conveniently denoted <u>$X \sim S_\alpha(\sigma,\beta,\mu)$</u>.

Since the characteristic function of an α-stable random variable apparently is integrable, α-stable random variables possess continuous density functions (see the figures below).

In the case when $\mu = 0$ we write $X \sim S_\alpha(\sigma,\beta)$. Often one writes <u>$S\alpha S$</u> instead of $S_\alpha(\sigma,0)$.

A random variable $X \sim S_\alpha(\sigma,\pm 1)$ is called <u>totally skewed</u>. When $\alpha < 1$ it turns out that $\mathbf{P}\{S_\alpha(\sigma,1) \leq 0\} = 0$ and $\mathbf{P}\{S_\alpha(\sigma,-1) \geq 0\} = 0$. An α-stable random variable that is not totally skewed, or that has $\alpha \geq 1$, possesses a density function that is strictly positive everywhere.

Exercise 38. Shew that $X \sim S_\alpha(\sigma,\beta,\mu) \Rightarrow -X \sim S_\alpha(\sigma,-\beta,-\mu)$. Further, shew that $X \sim S\alpha S$ is symmetric (that is, X and $-X$ are identically distributed).

Exercise 39. Shew that $X \sim S_\alpha(\sigma,\beta)$ is strictly α-stable when $\alpha \neq 1$ for all values of β, and when $\alpha = 1$ for $\beta = 0$.

Many results require that the case when $\alpha = 1$ with $\beta \neq 0$ is treated separately, because of the logarithm in the characteristic function. In the sequel we will therefore always assume that $\beta = 0$ when $\alpha = 1$, so that the upper formula for the characteristic function in Property 2 is valid.

Theory for α-stable random variables is intimately linked with $\mathbb{L}^\alpha$-theory. This means that the cases when $\alpha \geq 1$ and $\alpha < 1$ often have to be treated separately, because the triangle inequality goes in different directions in these two cases. Sometimes the totally skewed case when $\beta = \pm 1$ requires a separate treatment from that when $\beta \neq \pm 1$, because according to Property 4 below, one of the "tails" of the distribution will be "unusually light" when $\beta = \pm 1$.

Exercise 40. Shew that $S_2(\sigma,\beta,\mu) = \mathrm{N}(\mu, 2\sigma^2)$. Further, shew that we have $X \sim S_1(\sigma,0,\mu)$ for a Cauchy distributed random variable X with density function

$$f_X(x) = \frac{\sigma}{\pi\,[(x-\mu)^2 + \sigma^2]} \qquad \text{for} \quad x \in \mathbb{R}.$$

The Gaussian case $\alpha = 2$ is in many aspects quite different from the non-Gaussian case $\alpha < 2$, as has been indicated above. Note the title of the book by Samorodnitsky and Taqqu.

The Gaussian (normal) distribution ($\alpha = 2$) and the Cauchy distribution ($\alpha = 1$ and $\beta = 0$) are, together with the Lévy distribution ($\alpha = 1$ and $\beta = \pm 1$), the only non-degenerate α-stable distributions for which the density functions are known explicitely.

The following graphs of density functions of α-stable random variables have been copied without permission from the book by Janicki and Weron:

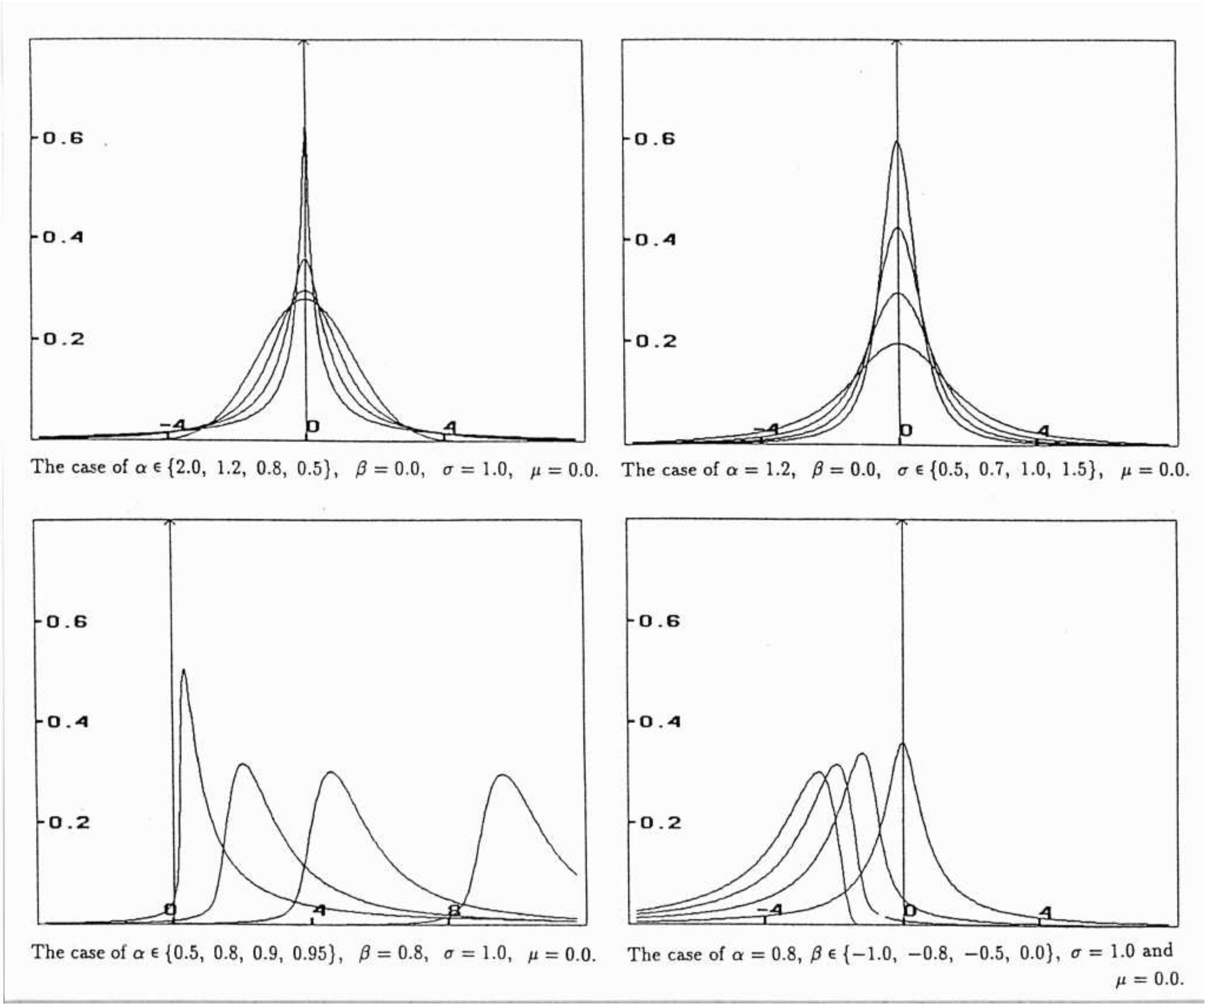

The case of $\alpha \in \{2.0,\ 1.2,\ 0.8,\ 0.5\}$, $\beta = 0.0$, $\sigma = 1.0$, $\mu = 0.0$. The case of $\alpha = 1.2$, $\beta = 0.0$, $\sigma \in \{0.5,\ 0.7,\ 1.0,\ 1.5\}$, $\mu = 0.0$.

The case of $\alpha \in \{0.5,\ 0.8,\ 0.9,\ 0.95\}$, $\beta = 0.8$, $\sigma = 1.0$, $\mu = 0.0$. The case of $\alpha = 0.8$, $\beta \in \{-1.0,\ -0.8,\ -0.5,\ 0.0\}$, $\sigma = 1.0$ and $\mu = 0.0$.

Property 3. **(i)** For independent random variables $X_1 \sim S_\alpha(\sigma_1, \beta_1, \mu_1)$ and $X_2 \sim S_\alpha(\sigma_2, \beta_2, \mu_2)$, we have

$$X_1 + X_2 \sim S_\alpha\Big((\sigma_1^\alpha + \sigma_2^\alpha)^{1/\alpha},\ \frac{\beta_1\sigma_1^\alpha + \beta_2\sigma_2^\alpha}{\sigma_1^\alpha + \sigma_2^\alpha},\ \mu_1 + \mu_2\Big).$$

(ii) For $X \sim S_\alpha(\sigma, \beta, \mu)$ and $a \in \mathbb{R}$, we have $aX \sim S_\alpha(|a|\sigma, \operatorname{sign}(a)\beta, a\mu)$.

Exercise 41. Shew Property 3.

Using so called Tauberian theory (which relates the behavior of the moment generating function at the left endpoint of its region of existence to the behavior of the corresponding distribution function at infinity), one can prove the following important result:

Property 4. For a random variable $X \sim S_\alpha(\sigma,\beta,\mu)$ with $\alpha<2$, we have

$$\lim_{u\to\infty} u^\alpha \mathbf{P}\{X>u\} = C_\alpha(1+\beta)\sigma^\alpha \qquad \text{and} \qquad \lim_{u\to\infty} u^\alpha \mathbf{P}\{X<-u\} = C_\alpha(1-\beta)\sigma^\alpha,$$

where $C_\alpha = [\int_0^\infty x^{-\alpha}\sin(x)\,dx]^{-1}$.

Note that one of the limits is zero when X is totally skewed ($\beta=\pm1$), which means that the "actual tail" decays faster than $u^{-\alpha}$. [In the case when $\alpha<1$ we have already learned that totally skewed distributions are one-sided, so that one of the tails is "empty" already for finite values of u.]

Of course, in the Gaussian case when $\alpha=2$, the tail-behavior is radically different: As $u\to\infty$

$$\mathbf{P}\{S_2(\sigma,\beta,\mu)>u\} = \mathbf{P}\{\mathrm{N}(\mu,2\sigma^2)>u\} = 1-\Phi\left(\frac{u-\mu}{\sqrt{2}\,\sigma}\right) = \frac{\sigma}{\sqrt{\pi}\,(u-\mu)}\exp\left\{-\frac{(u-\mu)^2}{4\,\sigma^2}\right\}[1+o(1)].$$

Exercise 42. Shew that the expected value of an α-stable random variable exists and is finite if and only if $\alpha>1$. Shew that the variance is finite if and only if $\alpha=2$ (that is, in the Gaussian case).

8.3 Generating Univariate α-stable Random Variables in a Computer

THEOREM 8. (Chambers, Mallows and Stuck) Let ξ_1 and ξ_2 be independent random numbers, and put

$$\begin{cases} C \;=\; C(\alpha,\beta) \;=\; \arctan[\beta\tan(\frac{\pi\alpha}{2})] \\ D \;=\; D(\alpha,\beta) \;=\; [\cos(\arctan(\beta\tan(\frac{\pi\alpha}{2})))]^{-1/\alpha} \end{cases}.$$

The following formula gives a random variable $X\sim S_\alpha(\sigma,\beta,\mu)$:

$$X \;=\; \sigma\,\frac{D\,\sin[\pi\alpha(\xi_1-\frac{1}{2})+C]}{[\sin(\pi\xi_1)]^{1/\alpha}}\left(\frac{\cos[(1-\alpha)\pi(\xi_1-\frac{1}{2})-C]}{-\ln(\xi_2)}\right)^{(1-\alpha)/\alpha} + \mu.$$

Exercise 43. Shew that Theorem 8 simplifies to (a version of) the method by Box-Müller (Theorem 4) when $\alpha=2$.

To my experience, the formula in Theorem 8 can cause problems with numerical overflow in very large simulations (involving billions of random numbers). I suspect that what happens is that a linear congruential random number generator with module m and full period-length m, on its way through the possible values $\{\frac{0}{m},\frac{1}{m},\ldots,\frac{m-1}{m}\}$ (or possibly $\{\frac{0}{m-1},\frac{1}{m-1},\ldots,\frac{m-1}{m-1}\}$), takes the value 0 (and possibly 1), once every period. This causes divisions by zero, or that the logarithm of zero is computed. One way to avoid this problem is to let the random numbers generated take the values $\{\frac{1/2}{m}, \frac{3/2}{m}, \ldots, \frac{m-1/2}{m}\}$. The following two exercises give another solution to the problem:

Exercise 44. Let ξ_1 and ξ_2 be independent random numbers. Shew that also $(1-|\xi_1-\xi_2|)^2$ is a random number.

Exercise 45. Explain how Exercise 44 can be used to modify the formula in Theorem 8 so that problems with random numbers that are exactly 0 or 1 do not occur.

The literature provides several representations of α-stable distribution functions and density functions in terms of infinite series. One of these representations could be used together with either Theorem 1 or Exercise 9 as an alternative way to generate α-stable random variables. Such a method would be much slower than that of Theorem 8 (because of the numerical computations of infinite series involved), but could avoid problems with numerical overflow of the type mentioned above.

8.4 Multi-variate α-stable Random Variables

Definition. (i) A random variable X in $\mathbb{R}^n$ is <u>α-stable</u> if given any $m\in\mathbb{N}$ together with independent copies $X_1,\ldots,X_m$ of X, there exists a constant $d_m\in\mathbb{R}^n$ such that

$$X_1+\ldots+X_m =_{\text{distribution}} m^{1/\alpha}X+d_m.$$

(ii) A random variable X in $\mathbb{R}^n$ is <u>strictly α-stable</u> if given any $m\in\mathbb{N}$ together with independent copies $X_1,\ldots,X_m$ of X, we have

$$X_1+\ldots+X_m =_{\text{distribution}} m^{1/\alpha}X.$$

Definition. We say that a family $\{X_i\}_{i\in I}$ of univariate random variables is <u>jointly (strictly) α-stable</u> if the random variable $(X_{i_1},\ldots,X_{i_n})$ is multi-variate (strictly) α-stable for each choice of $n\in\mathbb{R}$ and $i_1,\ldots,i_n\in I$.

In analogy with Property 2, there exists a characterization of the characteristic function of an $\mathbb{R}^n$-valued α-stable random variable. But the multi-variate result is more complicated.

Property 5. (i) A random variable X in $\mathbb{R}^n$ is strictly α-stable if and only if each linear combination of its components $\sum_{i=1}^n a_iX_i$ is univariate strictly α-stable.

(ii) Suppose that $\alpha\geq 1$. A random variable X in $\mathbb{R}^n$ is α-stable if and only if each linear combination of its components $\sum_{i=1}^n a_iX_i$ is univariate α-stable.

In the case when $\alpha<1$, it is <u>not true</u> that a random variable X in $\mathbb{R}^n$ is α-stable if and only if each linear combination of its components $\sum_{i=1}^n a_iX_i$ is univariate α-stable.

Exercise 46. Let $X, X_1, X_2, \ldots$ be jointly α-stable random variables. Shew that $X_n \to_{\mathbf{P}} X$ as $n \to \infty$ if and only if the scale parameter $\sigma_{X_n - X} \to 0$ as $n \to \infty$.

9 α-stable Stochastic Processes

Definition. **(i)** A stochastic process $\{X(t)\}_{t \in T}$ is α-stable if the random variable $(X(t_1), \ldots, X(t_n))$ is α-stable in $\mathbb{R}^n$ for each choice of $n \in \mathbb{N}$ and $t_1, \ldots, t_n \in T$.

(ii) A stochastic process $\{X(t)\}_{t \in T}$ is strictly α-stable if the random variable $(X(t_1), \ldots, X(t_n))$ is strictly α-stable in $\mathbb{R}^n$ for each choice of $n \in \mathbb{N}$ and $t_1, \ldots, t_n \in T$.

In other words, a stochastic process $\{X(t)\}_{t \in T}$ is (strictly) α-stable if the random variables $X(t_1), \ldots, X(t_n)$ are jointly (strictly) α-stable for each choice of $n \in \mathbb{N}$ and $t_1, \ldots, t_n \in T$.

Gaussian intuition is quite useless in the more general α-stable context. Gaussian results often do not have an α-stable analogue, and in the cases they do, proofs are usually quite different.

The most striking and unique feature of Gaussian processes is that they can be specified and thus analyzed by means of their expected values together with the covariance function. For virtually any other stochastic process (than a Gaussian process), and in particular for α-stable processes, it is not possible to specify or analyze the process by such simple means. In fact, the covariance function usually do not even exist for an α-stable processes, and if it does exist it is almost useless.

9.1 α-stable Lévy processes

Definition. A stochastic process $\{L(t)\}_{t \in \mathbb{R}}$ such that $L(0) = 0$ (a.s.) is called a (standard) α-stable Lévy motion with skewness β, if the process has independent and stationary increments that satisfy

$$L(t) - L(s) \sim S_\alpha((t-s)^{1/\alpha}, \beta) \qquad \text{for} \quad -\infty < s < t < \infty.$$

A 2-stable Lévy motion $L(s)$ is a Wiener process with $\mathbf{Var}\{L(1)\} = 2$. The role of the α-stable Lévy motion within the family of α-stable processes is similar to that of the Wiener process within the family of Gaussians.

Unlike the case with the Wiener process ($\alpha = 2$), sample paths of α-stable Lévy motions are not continuous a.s. when $\alpha < 2$, and they have "jump-discontinuities". In fact, they are so called pure-jump processes in the case when $\alpha < 1$.

An α-stable Lévy motion is totally skewed if $\beta = \pm 1$. A totally skewed α-stable Lévy motion with $\alpha < 1$ is non-decreasing when $\beta = 1$, and non-increasing when $\beta = -1$.

Exercise 47. Shew that an α-stable Lévy motion is an α-stable stochastic process.

Exercise 48. Shew that an α-stable Lévy motion is self-similar, that is, that

$$(L(\lambda t_1),\ldots,L(\lambda t_n)) =_{\text{distribution}} (\lambda^{\kappa} L(t_1),\ldots,\lambda^{\kappa} L(t_n))$$

for each choice of $n\in\mathbb{N}$ and $t_1,\ldots,t_n\in\mathbb{R}$, for some (index of self-similarity) $\kappa>0$.

The following graphs of simulated sample paths of α-stable Lévy motions have been copied without permission from the book by Janicki and Weron:

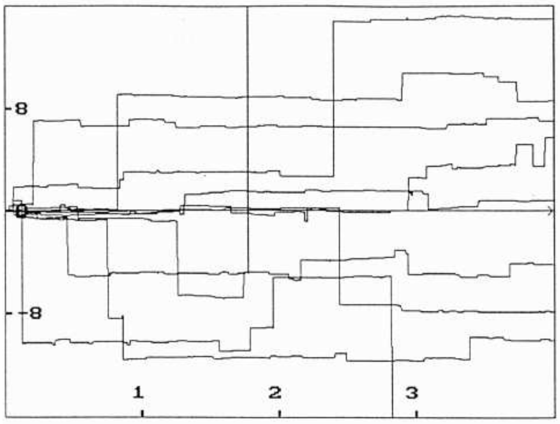

Trajectories of $S\alpha S$ Lévy motion in the case of $\alpha = 0.7$.

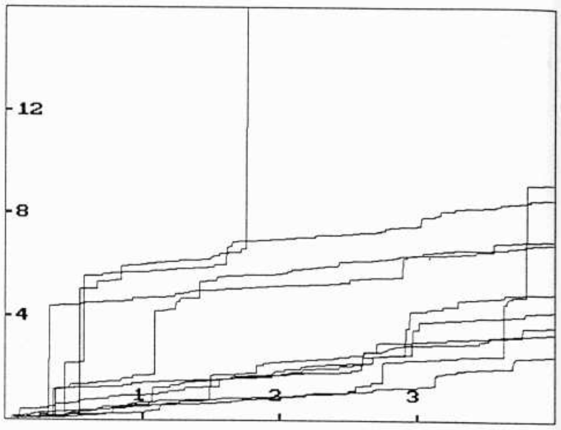

Trajectories of totally skewed Lévy motion in the case of $\alpha = 0.7$.

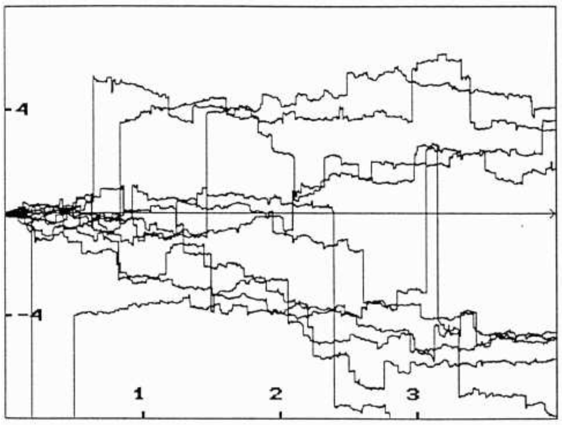

Trajectories of α–stable Lévy motion in the case of $\alpha = 1.2$.

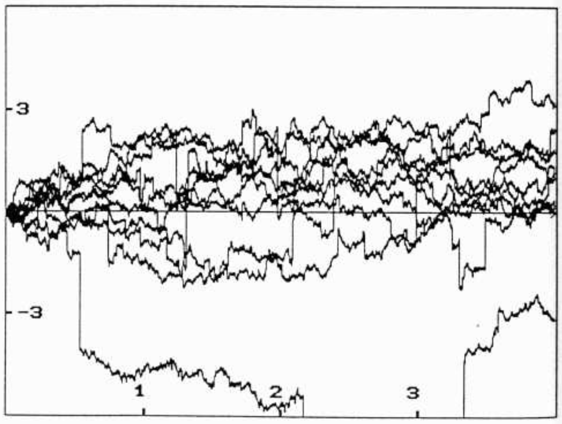

Trajectories of α–stable Lévy motion in the case of $\alpha = 1.7$.

9.2 α-stable Stochastic Integrals

Here we introduce the stochastic integral $\int_{x\in\mathbb{R}} f(x)\,dL(x)$ of a deterministic and measurable function $f:\mathbb{R}\to\mathbb{R}$ with respect to (infinitesimal) increments $dL(\cdot)$ of an α-stable Lévy motion.

Exercise 49. Prove that under suitable technical conditions (that need not be specified), given a constant $\tau>0$, we have

$$\sum_{k=-\infty}^{\infty} f(k\tau)\,\Big(L(k\tau)-L((k-1)\tau)\Big) \sim S_{\alpha}\bigg(\Big(\sum|f(k\tau)|^{\alpha}\tau\Big)^{1/\alpha},\ \beta\,\frac{\sum \mathrm{sign}(f(k\tau))\,|f(k\tau)|^{\alpha}\tau}{\sum|f(k\tau)|^{\alpha}\tau}\bigg).$$

Property 6. The integral $\int f\,dL = \int_{x\in\mathbb{R}} f(x)\,dL(x)$ is well-defined in the sense of convergence in probability for measurable functions $f:\mathbb{R}\to\mathbb{R}$ such that $\int_{x\in\mathbb{R}} |f(x)|^\alpha\, dx < \infty$. In that case we have

$$\int f\,dL \sim S_\alpha\Big(\big[\int_{x\in\mathbb{R}} |f(x)|^\alpha\,dx\big]^{1/\alpha},\ \beta\,\frac{\int_{x\in\mathbb{R}} \operatorname{sign}(f(x))\,|f(x)|^\alpha\,dx}{\int_{x\in\mathbb{R}} |f(x)|^\alpha\,dx}\Big).$$

Under suitable continuity conditions on the function f, Property 6 follows readily from Exercises 46 and 49. The general case with an f that is only measurable with $\int_{x\in\mathbb{R}} |f(x)|^\alpha\,dx < \infty$ is not difficult either, but requires a Lebesgue-type rather than Riemann-type construction of the integral.

Exercise 50. Shew that the random variable $(\int f_1\,dL,\ldots,\int f_n\,dL)$ in $\mathbb{R}^n$ is multivariate α-stable and find its characteristic function. Give a sufficient (non-trivial) condition for the components of the vector to be independent random variables.

9.3 Spectral Representation of α-stable Stochastic Processes

Our next result argubly is the most important result in the theory of α-stable stochastic processes:

Property 7. (Spectral Representation) Let $\{X(t)\}_{t\in T}$ be a strictly α-stable stochastic process. Suppose that either T is countable, or that $T\subseteq\mathbb{R}$ with $X(t)$ continuous in probability. Let $\{L(t)\}_{t\in\mathbb{R}}$ be an α-stable Lévy motion that is totally skewed $(\beta=\pm1)$. There exists a function $f:T\times\mathbb{R}\to\mathbb{R}$ such that

$$(X(t_1),\ldots,X(t_n)) =_{\text{distribution}} \Big(\int f(t_1,\cdot)\,dL,\ldots,\int f(t_n,\cdot)\,dL\Big)$$

for each choice of $n\in\mathbb{N}$ and $t_1,\ldots,t_n\in T$.

Remarks. (i) Of course, (by Property 6) the requirement that the function $f(t,\cdot)$ is measurable with $\int_{x\in\mathbb{R}} |f(t,x)|^\alpha\,dx < \infty$ for each $t\in T$ is implicit in Property 7.

(ii) The requirement that $X(t)$ is continuous in probability means that $X(t)\to_{\mathbf{P}} X(t_0)$ as soon as $t\to t_0$, or equivalently (see Exercise 46), that $\sigma_{X(t)-X(t_0)}\to 0$ as soon as $t\to t_0$: This holds for virtually any non-pathological continuous-time α-stable process that is encountered.

(iii) By means of viewing an α-stable random vector $(X_1,\ldots,X_n)$ as an α-stable stochastic process $\{X_i\}_{i\in\{1,\ldots,n\}}$, we get the representation

$$(X_1,\ldots,X_n) =_{\text{distribution}} \Big(\int f(1,\cdot)\,dL,\ldots,\int f(n,\cdot)\,dL\Big)$$

for a suitable function $f:\{1,\ldots,n\}\times\mathbb{R}\to\mathbb{R}$ (with the integrability property mentioned above).

9.4 Examples of α-stable Stochastic Processes

(i) Let $f:\mathbb{R}\to\mathbb{R}$ be a measurable function with $\int_{x\in\mathbb{R}} f(x)^\alpha\, dx<\infty$. The process

$$X(t) \equiv \int_{s\in\mathbb{R}} f(t-s)\, dL(s) \qquad \text{for} \quad t\in\mathbb{R}$$

is stationary, and is called an α-stable moving average process.

Every non-pathological stationary Gaussian process $(\alpha=2)$ $\{X(t)\}_{t\in\mathbb{R}}$ is a moving average process, because every non-pathological covariance function $r(\cdot)$ has the representation

$$r(t) = \mathbf{Cov}\{X(s), X(s+t)\} = 2 \int_{s\in\mathbb{R}} f(s)f(t+s)\, ds$$

for some choice of a measurable $f(\cdot)$ with $\int_{x\in\mathbb{R}} f(x)^2\, dx<\infty$.

In the case when $\alpha<2$, there are plenty of non-pathological stationary α-stable processes that are not moving average processes.

Exercise 51. Use Exercise 49 together with a limit argument to prove that a 2-stable moving average process $X(t) = \int_{s\in\mathbb{R}} f(t-s)\, dL(s)$ has covariance function $r(t) = 2 \int_{s\in\mathbb{R}} f(s)f(t+s)\, ds$. [Observe that, by Exercise 40, a 2-stable Lévy motion $L(s)$ has incremental variances $\mathbf{Var}\{L(t)-L(s)\} = 2\,|t-s|$.]

Example 3. In the case when $\alpha=2$, the function $f(t) = \sqrt{\lambda}\,\mathrm{e}^{-\lambda t} I_{[0,\infty)}(t)$ makes

$$X(t) = \int f(t-s)\, dL(s) = \int_{-\infty}^{t} \sqrt{\lambda}\,\mathrm{e}^{-\lambda(t-s)}\, dL(s)$$

a stationary Ornstein-Uhlenbeck process with covariance function

$$r(t) = 2 \int_{s\in\mathbb{R}} f(s)f(t+s)\, ds = \ \ldots = \mathrm{e}^{-\lambda|t|}.$$

(ii) As have been mentioned before, a random variable $Y\sim S_\alpha(\sigma,1)$ is non-negative when $\alpha<1$. By Property 2 and some elementary complex analysis, we have

$$\mathbf{E}\{\mathrm{e}^{-\gamma Y}\} = \exp\Big\{-\sigma^\alpha\gamma^\alpha/\cos(\tfrac{\pi\alpha}{2})\Big\} \qquad \text{for} \quad Y\sim S_\alpha(\sigma,1) \quad \text{when} \quad \alpha<1.$$

Let $\{G(t)\}_{t\in T}$ be a zero-mean Gaussian process. Choose an $\alpha\in(0,2)$ and let $Y \in S_{\alpha/2}(\sigma,1)$ be independent of $\{G(t)\}_{t\in T}$. The process

$$X(t) \equiv \sqrt{Y}\, G(t) \qquad \text{for} \quad t\in T$$

is called a sub-Gaussian process.

Exercise 52. Shew that sub-Gaussian processes are α-stable processes.

(iii) A Gaussian process $\{G(t)\}_{t\in\mathbb{R}}$ with mean $\mathbf{E}\{G(t)\}=0$ and covariance function

$$\mathbf{Cov}\{G(s),G(t)\} = \frac{1}{2}\Big(|t|^{2H}+|s|^{2H}-|t-s|^{2H}\Big)\,\mathbf{Var}\{G(1)\} \qquad \text{for some} \quad H\in(0,1)\backslash\{\tfrac{1}{2}\}$$

is called a fractional Brownian motion (fBm.). This process has the representation

$$G(t) = C\int_{-\infty}^{\infty}\Big(((t-s)^+)^{H-1/2}-((-s)^+)^{H-1/2}\Big)\,dL(s) \qquad \text{for} \quad t\in\mathbb{R},$$

where $C=C(H)>0$ is a suitable constant and $\{L(s)\}_{s\in\mathbb{R}}$ is 2-stable Lévy motion (a Wiener-process). Of course, the fact that the process given by the integral really is fBm. is conveniently proven by means of computing its covariance function.

fBm. was introduced in the 60's by Benoit Mandelbrot in his pioneering works about self-similarity, fractals and chaos, and the application of these concepts to climatological data. fBm. is argubly the second most important Gaussian process after the Wiener process.

Exercise 53. Shew that fBm. is a self-similar process with stationary increments.

Linear fractional α-stable motion (lfsm.) is defined by

$$X(t) = C\int_{-\infty}^{\infty}\Big(((t-x)^+)^{H-1/\alpha}-((-x)^+)^{H-1/\alpha}\Big)\,dL(x) \qquad \text{for} \quad t\in\mathbb{R},$$

where $H\in(0,1)\setminus\{1/\alpha\}$ and $C>0$ are constants, and $L(s)$ is an α-stable Lévy motion. lfsm. is the second most important α-stable process after the α-stable Lévy motion.

Exercise 54. Shew that lfsm. is self-similar with stationary increments.

(iv) By a linear α-stable stochastic differential equation (SDE), or a linear α-stable diffusion process, we mean a solution $\{X(t)\}_{t\geq 0}$ to the integral equation

$$X(t) \;=\; X(0)+\int_0^t\Big(a(s)+b(s)X(s)\Big)\,ds+\int_0^t c(s)\,dL(s) \qquad \text{for} \quad t\geq 0.$$

Exercise 55. Shew that the integral equation has the solution

$$X(t) = \Phi_{0,t}\,X(0)+\int_0^t\Phi_{s,t}\,a(s)\,ds+\int_0^t\Phi_{s,t}\,c(s)\,dL(s) \qquad \text{where} \quad \Phi_{s,t}=\exp\Big\{\int_s^t b(r)\,dr\Big\}.$$

A homogeneous Gaussian Ornstein-Uhlenbeck process with drift $\lambda\in\mathbb{R}$ and dispersion $\sigma>0$ can be described as a solution to the equation

$$X(t) \;=\; X(0)-\lambda\int_0^t X(s)\,ds+\sigma\int_0^t dW(s) \qquad \text{for} \quad t\geq 0.$$

where $\{W(s)\}_{t\in\mathbb{R}}$ is a Wiener-processen (2-stable Lévy motion). The process $X(t)$ is Gaussian if the initial value $X(0)$ and $\{W(s)\}_{t\in\mathbb{R}}$ are jointly Gaussian. The process $X(t)$ is stationary if the initial value $X(0)$ is given the stationary distribution and is independent of $\{W(s)\}_{t\in\mathbb{R}}$.

A non-homogeneous Gaussian Ornstein-Uhlenbeck process is a solution to the equation

$$X(t) \;=\; X(0) - \int_0^t \lambda(s)\, X(s)\, ds + \int_0^t \sigma(s)\, dW(s) \qquad \text{for} \quad t \geq 0,$$

where now $\lambda(\cdot)$ and $\sigma(\cdot)$ are functions rather than merely constants. Thus it is natural to say that linear α-stable SDE are non-homogeneous α-stable Ornstein-Uhlenbeck processes.

10 Stochastic Simulation of α-stable SDE

Here we learn how to simulate an α-stable SDE

$$X(t) \;=\; X(0) + \int_0^t \Big(a(s)+b(s)X(s)\Big)\, ds + \int_0^t c(s)\, dL(s) \qquad \text{for} \quad t\in[0,T),$$

where $L(s)$ is an α-stable Lévy-process with skewness β, and $a(\cdot)$, $b(\cdot)$ and $c(\cdot)$ are continuous (at least) functions. Further, $T>0$ is a constant. By Exercise 55, this SDE has a solution that is an α-stable stochastic process [provided that $X(0)$ and $\{L(s)\}_{s\geq 0}$ are jointly α-stable].

Stochastic simulation of a general non-linear α-stable SDE

$$X(t) \;=\; X(0) + \int_0^t d(s, X(s^-))\, ds + \int_0^t e(s, X(s^-))\, dL(s) \qquad \text{for} \quad t\in[0,T),$$

is carried out by means of an obvious modification of the "linear scheme" described below. However, the definition of what the SDE really means is now non-trivial (since it can no longer be solved explicitely as can a linear SDE). Of course, any SDE that is "nice enough" to render itself to simulation must have a unique solution, and the proof of this unicity is in fact pretty much the same thing as an analysis of the precision of the (disrete approximation of the) simulation scheme.

10.1 Discrete Approximation of an α-stable SDE

∘ Divide the interval $[0,T)$ into N sub-intervals $\{[t_{i-1},t_i)\}_{i=1}^N = \{[(i-1)\tau, i\tau)\}_{i=1}^N$, the length τ of which is chosen so that $N\tau = T$.

∘ Choose the initial value $X(0)$. If $X(0)$ is a random variable it could be generated in accordance with the principles of Section 1.

∘ Initialize $\tilde{X}(t_0)=\tilde{X}(0)=X(0)$. Compute the approximations $\tilde{X}(t_1),\ldots,\tilde{X}(t_{N-1})$ of $X(t_1),\ldots,X(t_{N-1})$ by means of the recursive scheme

$$\tilde{X}(t_i) \;=\; \tilde{X}(t_{i-1}) + [a(t_{i-1}) + b(t_{i-1})\tilde{X}(t_{i-1})]\,\tau + c(t_{i-1})\,[L(t_i)-L(t_{i-1})]$$

for $i=1,\ldots,N-1$, where the increments $L(t_i)-L(t_{i-1}) = \int_{t_{i-1}}^{t_i} dL(s)$ are independent $S_\alpha(\tau^{1/\alpha},\beta)$-distributed random variables (generated in accordance with Section 8.3).

∘ The solution $\{X(t)\}_{t\in[0,T)}$ to the SDE is approximated with the process

$$X^\tau(t) \;=\; \textstyle\sum_{j=1}^N I_{[t_{j-1},t_j)}(t)\tilde{X}(t_{j-1}) \;=\; \tilde{X}(t_{i-1}) \quad \text{for} \quad t\in[t_{i-1},t_i) \quad \text{for} \quad i=1,\ldots,N.$$

10.2 A Programming Project

Programming Project 2. Use stochastic simulation to find a confidence interval for the probability

$$p \;=\; \mathbf{P}\{\max_{t\in[0,1)}|X(t)|\ge 1\} \;=\; \mathbf{E}\{I_{\{\max_{t\in[0,1)}|X(t)|\ge 1\}}\}$$

when $X(t)$ is the (solution to) the linear α-stable SDE

$$X(t) \;=\; \int_0^t (-s)\,X(s)\,ds + \int_0^t \mathrm{e}^{-s}\,dL(s) \qquad \text{for} \qquad t\in[0,1),$$

and $L(s)$ is an 1.5-stable Lévy motion with skewness $\beta=0$.

The project should be completed by means of simulating independent observations of the Bin$(1,p)$-distributed random variable $I_{\{\max_{t\in[0,1)}|X^\tau(t)|\ge 1\}}$, and compute a confidence interval for the expected value of that random variable according to the method described in Section 3.

Choose the width of the confidence interval as narrow as is possible with a reasonably time-consuming programme. As always in programming, it is advicable to use quite few "iterations" while programs are developed and tested (making execution times short), and only run "full-length" simulations when programs are ready.

Choose τ small enough to make further decrements of τ have a "neglible effect" on the distribution of $\max_{t\in[0,1)}|X^\tau(t)|$. One way to accomplish this is to check that distances between empirical distribution functions for $\max_{t\in[0,1)}|X^\tau(t)|$, computed for decreasing values of τ, "settle" at a range (of values for the distances) that are small compared with the width of the confidence interval.