Both inequalities can be satisfied simultaneously for all $m \geq 0$. After fixing γ, compute all quantities in the upper bound of Lemma 3.3. Since $C(x)=(C_0+o(1))(|x|+\mu_4^{1/4})^\alpha$ with $C_0=4/(9\pi)$, it is easy to see that

$$\int Cg = (C_0+o(1))E((|X|+\mu_4^{\frac{1}{4}})^\alpha)$$

where X is a random variable with density g, and $\alpha=4(m-1)/(m+4)$. We can choose g such that $E(|X|^\alpha)$ is close to $\mu_4^{\alpha/4}$ (e.g., in Example 3.4, take $r=6$ or larger in the bound for unimodal densities; taking $r=4$ isn't good enough because for $r=4$, $E(|X|^4)=\infty$). Noting next that $\mu_4^{1/4} \sim \sqrt{m}/3^{1/4}$ as $m\to\infty$, we note that $\int Cg$ increases as a constant times $m^{\alpha/2}$. Next, $\int C^{2-\frac{1}{\gamma}}$ increases as a constant times

$$\mu_4^{\frac{1+\alpha(2-\frac{1}{\gamma})}{4}}$$

which in turn increases as $m^{\frac{9}{2}-\frac{2}{\gamma}}$. The upper bound in Lemma 3.3 increases as

$$m^{2-2\gamma+\frac{9\gamma}{2}-2} = m^{\frac{5\gamma}{2}} .$$

The smallest allowable value for γ is $1/\alpha \sim 1/4$. Thus, the upper bound on the expected complexity is of the order of magnitude of $m^{5/8}$. ■

3.4. Exercises.

1. Show that when a characteristic function ϕ is absolutely integrable, then the distribution has a bounded continuous density f. Is the density also uniformly continuous?
2. Construct a symmetric real characteristic function for a distribution with a density, having the property that ϕ takes negative and positive values.
3. Consider symmetric nonnegative characteristic functions ϕ, and define $\nu_{2n}=\int t^{2n}\phi(t)\,dt$.
   - A. Show that $\nu_{2n}{}^{1/(2n)}=o(n)$ implies that $(x^{2n}\nu_{2n})/(2n)!$ is summable for all $x>0$.
   - B. Show that f is unimodal and has a unique mode at 0 (Feller, 1971, p. 528).
   - C. In the alternating series algorithm for this class of densities given in the text, why can we take $b=\mu_1$ or $b=\sigma$ in the formula for the dominating

curve where μ_1 is the first absolute moment for f and σ is the standard deviation for f ?

D. A continuation of part C. If all operations in the algorithm take one unit of time, give a useful sufficient condition on ϕ for the expected time of the algorithm to be finite.

4. The following is an important symmetric nonnegative characteristic function:

$$\phi(t)=\sqrt{\frac{\sqrt{2t}}{\sinh(\sqrt{2t}\,)}}=\frac{1}{\sqrt{1+2\frac{|t|}{3!}+2\frac{|t|^2}{5!}+\cdots}}$$

(see e.g. Anderson and Darling, 1952). Near $t=0$, ϕ varies as $1-|t|/6$. This implies that the first absolute moment is infinite. Find a dominating curve for this particular characteristic function, verify that the density f is determined by its Taylor series about 0, and give all the details of the alternating series method for this distribution.

5. The following characteristic function appears as the limit of a sequence of characteristic functions in mathematical statistics (Anderson and Darling, 1952):

$$\phi(t)=\left(\frac{-2\pi i t}{\cos(\frac{\pi}{2}\sqrt{1+8it}\,)}\right)^{\frac{1}{2}} .$$

Give a finite time random variate generator for this distribution. Ignore efficiency issues (e.g., the expected time is allowed to be infinite).

6. Give the full details of the proof that the expected number of evaluations of ϕ in the series method for generating the sum of m iid uniform $[-1,1]$ random variables (Example 3.6) is $O(m^{(5+\epsilon)/8})$ for all $\epsilon>0$.

7. How can you improve on the expected complexity in Example 3.6?

4. THE SIMULATION OF SUMS.

4.1. Problem statement.

Let X be a random variable with density f on the real line. In this section we consider the problem of the simulation of $S_n=X_1+\cdots+X_n$ where $X_1,\ldots,X_n$ are iid random variables distributed as X. The naive method

Naive method

$S \leftarrow 0$

FOR $i:=1$ TO n DO

  Generate X with density f .

  $S \leftarrow S+X$

RETURN S

takes worst-case or expected time proportional to n depending upon whether X can be generated in constant worst-case or constant expected time. We say that a generator is uniformly fast when the expected time $E(T_n)$ needed to generate S_n satisfies

$$\sup_{n \geq 1} E(T_n) < \infty .$$

This supremum is allowed to depend upon f . Note that the uniformity is with respect to n and not to f . This differs from our standard notion of uniformity over a class of distributions.

In trying to develop uniformly fast generators, we should get a lot of help from the central limit theorem, which states that under some conditions on the distribution of X, the sum S_n, properly normalized, tends in distribution to one of the stable laws. Ideally, a uniformly fast generator should return such a stable random variate most of the time. What complicates matters is that the distribution of S_n is not easy to describe. For example, in a rejection based method, the computation of the value of the density of S_n at one point usually requires time increasing with n. Needless to say, it is this hurdle which makes the problem both challenging and interesting.

In a first approach, we will cheat a bit: recall that if ϕ is the characteristic function of X, then S_n has characteristic function ϕ^n. If we have a uniformly fast generator for the family $\{\phi,\phi^2, \ldots , \phi^n ,\ldots\}$, then we are done. In other words, we reduce the problem to that of the generation of random variates with a given characteristic function, discussed in section 3. The reason why we call this cheating is that ϕ is usually not available, only f .

In the second approach, the problem is tackled head on. We will first derive inequalities which relate the density of S_n to the normal density. In proving the inequalities, we have to rederive a so-called local central limit theorem. The inequalities allow us to design uniformly fast rejection algorithms which return a stable random variate with high probability. The tightness of the bounds allows us to obtain this result despite the fact that the density of S_n can't usually be computed in constant time. When the density can be computed in constant time, the algorithm is extremely efficient. This is the case when the density of S_n has a relatively simple analytic form, as in the case of the exponential density when

S_n is gamma (n).

Other solutions are suggested in the exercises and in later sections, but the most promising generally applicable strategies are definitely the two mentioned above.

4.2. A detour via characteristic functions.

S_n has characteristic function ϕ^n when X has characteristic function ϕ. This fact can be used to generate S_n efficiently provided that all the ϕ_n 's belong to a family of characteristic functions for which a good efficient generator is available.

One such family is the family of Polya characteristic functions dealt with in section IV.6.7. In particular, if ϕ is Polya, so is ϕ^n. Based upon Theorems IV.6.8 and IV.6.9, we can conclude the following:

Theorem 4.1.

If ϕ is a Polya characteristic function, then $X \leftarrow \frac{Y}{Z}$ has characteristic function ϕ^n when Y,Z are independent random variables, Y has the FVP density (defined in Theorem IV.6.9), and Z has distribution function

$$F(s) = 1-\phi^n + sn\,\phi'(s)\phi^{n-1}(s) \quad (s>0)\,.$$

Here ϕ' is the right-hand derivative of ϕ. When F is absolutely continuous, then it has density

$$s^2 n(n-1)\phi'^2(s)\phi^{n-2}(s) + s^2 n\,\phi''(s)\phi^{n-1}(s) \quad (s>0)\,.$$

When ϕ is explicitly given, and it often is, this method should prove to be a formidable competitor. For one thing, we have reduced the problem to one of generating a random variate with an explicitly given distribution function or density, i.e. we have taken the problem out of the domain of characteristic functions.

The principle outlined here can be extended to a few other classes of characteristic functions, but we are still far away from a generally applicable technique, let alone a universal black box method. The approach outlined in the next section is better suited for this purpose.

4.3. Rejection based upon a local central limit theorem.

We assume that f is a zero mean density with finite variance σ^2. Summing n iid random variables with this density is known to give a random variable with approximately normal $(0, n\sigma^2)$ distribution. The study of the closeness of this approximation is the subject of the classical central limit theory. The only things that can be of use to us are precise (i.e., not asymptotic) inequalities which clarify just how close the density of S_n is to the normal $(0, n\sigma^2)$ density. For a smooth treatment, we put two further restrictions on f :

A. The density f has an absolutely integrable characteristic function ϕ. Recall that this implies among other things that f is bounded and continuous.

B. The random variable X has finite third absolute moment not exceeding β: $E(|X|^3) \leq \beta < \infty$.

Condition A allows us to use the simple inversion formula for characteristic functions, while condition B guarantees us that the error term is $O(1/\sqrt{n})$. Densities f satisfying all the conditions outlined above are called **regular**. Clearly, most zero mean densities occurring in practice are regular. There is only one large class of exceptions, the distributions in the domain of attraction of stable laws. By forcing the variance to be finite, we can only have convergence to the normal distribution. In exercise 4.1, which is more a research project than an exercise, the reader is challenged to repeat this section for distributions whose sums converge to symmetric stable laws with parameter $\alpha<2$. For once we will do things backwards, by giving the results and their implications before the proofs, which are deferred to next section.

The fundamental result upon which this entire section rests is the following form of a local central limit theorem:

Theorem 4.2.

Let f be a regular density, and let f_n be the density of $S_n/(\sigma\sqrt{n})$. Let g be the standard normal density. There exist sequences a_n and b_n only depending upon f such that

$$|f_n(x)-g(x)| \leq h_n(x) = \min(a_n, \frac{b_n}{x^2}),$$

and

$$\max(a_n, b_n) = O(\frac{1}{\sqrt{n}}).$$

For a proof and references, see section 4.4. Explicit values for a_n and b_n follow. It is important to note that

$$g-h_n \leq f_n \leq g+h_n ,$$

where $\int h_n = O(1/\sqrt{n})$. In other words, the inequality is eminently suited for use in a rejection algorithm with squeezing. Both g and h_n can be considered as very easy densities from a random variate generation point of view. Furthermore, the obvious rejection algorithm, described in Example II.3.6, has rejection constant $1+\int h_n$ tending to 1 as $n \to \infty$. There is even more good news: if the lower bound is used for squeezing, then the expected number of evaluations of f is at most $2\int h_n = O(1/\sqrt{n}) = o(1)$. The cumbersome part is the evaluation of f_n.

There are essentially two possibilities when it comes to evaluating f_n: first, f_n is explicitly known. This is for example the case when f is an exponential density centered around its mean, and f_n is the density of a linearly transformed gamma (n) density. In the case of the gamma density, we can easily compute the different constants in the bound of Theorem 4.2. as is done in exercise 4.2. Another example for the sums of uniform random variables follows in a separate section.

To compute f_n via convolutions is all but impossible. The only other alternative is to write f_n as a series based upon the inversion formula for ϕ^n, and to apply the series method. Here too the hurdles are formidable.

4.4. A local limit theorem.

It is the purpose of this section to prove Theorem 4.2. The proof is quite long, and is given in full because we require explicit knowledge of the bounding sequence, and a careful derivation of the bounds to keep the constants as small as possible. Local limit theorems of the type needed by us have been derived in a number of papers, see e.g. Inzevitov (1977), Survila (1964) and Maejima (1980). An excellent general reference is Petrov (1975). For example, Survila (1964) has obtained the existence of a constant C depending upon f only such that for regular f,

$$| f_n(x) - g(x) | \leq \frac{C}{1+x^2} .$$

Ibragimov and Linnik (1971) have obtained an upper bound of the type $\frac{C}{\sqrt{n}}$. Note that Survila's bound does not tend to zero with n. The Ibragimov-Linnik upper bound is called a uniform estimate in the local central limit theorem. Such uniform estimates are useless to us because the upper bound when integrated with respect to x is not finite. The bound which we derive here uses well-known tricks of the trade, documented for example in Petrov (1975) and Maejima (1980).

Let us start slowly with a few key lemmas.

Lemma 4.1.

For any real t,

$$| e^{it} - \sum_{j=0}^{n-1} \frac{(it)^j}{j!} | \leq \frac{t^n}{n!} \quad (n \geq 0) .$$

Lemma 4.2.

Let ϕ be the characteristic function for a regular density f. Then the following inequalities are valid:

$$| \phi(t)-1+\frac{\sigma^2 t^2}{2} | \leq \frac{\beta | t |^3}{6} ,$$

$$| \phi'(t)+t\sigma^2 | \leq \frac{\beta}{2} t^2 ,$$

$$| \phi''(t)+\sigma^2 | \leq \beta | t | .$$

Proof of Lemma 4.2.

Since three absolute moments exist, we notice that the first three derivatives of ϕ exist and are continuous functions given by the formulas (Feller, 1971, p. 512)

$$\phi^{(j)}(t) = \int e^{itx} (ix)^j f(x) \, dx \quad (j=0,1,2,3) .$$

Observe that

$$| \phi(t)-1+\frac{\sigma^2 t^2}{2} | \leq \int | e^{itu} -1-itu +\frac{t^2 u^2}{2} | f(u) \, du$$

$$\leq \int | \frac{| t |^3 | u |^3}{6} | f(u) \, du = \frac{\beta}{6} | t |^3 .$$

Next,

$$\phi'(t)+\frac{\sigma^2 t}{2} = \int (e^{itu} -1-itu) iu f(u) \, du .$$

Thus,

$$| \phi'(t)+\frac{\sigma^2 t}{2} | \leq \int | \frac{t^2 u^2}{2} | \, | u | f(u) \, du \leq \frac{\beta}{2} t^2 .$$

Finally,

$$\phi''(t)+\sigma^2 = -\int (e^{itu} -1) u^2 f(u) \, du .$$

Thus,

$$|\phi''(t)+\sigma^2| \leq \int |tu| u^2 f(u)\, du \leq \beta |t| . \blacksquare$$

Lemma 4.3.

Consider the absolute differences

$$A_m(t) = |(1-\frac{t^2}{2n})^m - e^{-\frac{t^2}{2}}| \qquad (m=n-2, n-1, n) .$$

For $t^2 \leq n$, we have

$$A_n(t) \leq \frac{t^4}{4n} e^{-\frac{t^2}{2}} ,$$

$$A_{n-1}(t) \leq \frac{1}{2(n-1)} e^{\frac{1}{2(n-1)}} e^{-\frac{t^2}{2}} ,$$

$$A_{n-2}(t) \leq \frac{2}{n-2} e^{\frac{2}{n-2}} e^{-\frac{t^2}{2}} .$$

If all integrals shown below are over $\{ |t| \leq \sqrt{n} \}$, then we have

$$\int A_n(t)\, dt \leq \frac{3}{4n}\sqrt{2\pi} , \int t^2 A_n(t)\, dt \leq \frac{15}{4n}\sqrt{2\pi} ,$$

$$\int A_{n-1}(t)\, dt \leq \frac{3}{4n}\sqrt{2\pi} + \frac{\sqrt{2\pi}}{2(n-1)} e^{\frac{1}{2(n-1)}} ,$$

$$\int t^2 A_{n-1}(t)\, dt \leq \frac{15}{4n}\sqrt{2\pi} + \frac{\sqrt{2\pi}}{2(n-1)} e^{\frac{1}{2(n-1)}} ,$$

$$\int A_{n-2}(t)\, dt \leq \frac{3}{4n}\sqrt{2\pi} + \frac{2\sqrt{2\pi}}{n-2} e^{\frac{2}{n-2}} ,$$

$$\int t^2 A_{n-2}(t)\, dt \leq \frac{15}{4n}\sqrt{2\pi} + \frac{2\sqrt{2\pi}}{n-2} e^{\frac{2}{n-2}} .$$

Proof of Lemma 4.3.

First,

$$e^{-\frac{t^2}{2}} - (1-\frac{t^2}{2n})^{n-2} \leq e^{-\frac{t^2}{2}} - (1-\frac{t^2}{2n})^{n-1}$$

$$\leq e^{-\frac{t^2}{2}} - (1-\frac{t^2}{2n})^{n}$$

$$\leq e^{-\frac{t^2}{2}}\left(1-e^{-\frac{nt^4}{8n^2(1-\frac{t^2}{2n})}}\right)$$

$$\leq e^{-\frac{t^2}{2}}\left(1-e^{-\frac{nt^4}{4n^2}}\right)$$

$$\leq e^{-\frac{t^2}{2}}\frac{t^4}{4n} .$$

Here we used the inequality $\log(1-u) \geq -u-u^2/(2(1-u)) \geq -u-u^2$ valid for $0 \leq u \leq 1/2$. Since

$$0 \leq e^{-\frac{t^2}{2}}-(1-\frac{t^2}{2n})^n ,$$

the bound for A_n is proved. For the other bounds, consider A_m in general. Clearly,

$$(1-\frac{t^2}{2n})^m - e^{-\frac{t^2}{2}} \leq e^{-\frac{t^2}{2}}\left(e^{t^2(\frac{1}{2}-\frac{m}{2n})-t^4\frac{m}{8n^2}}-1\right) .$$

For $m=n-i$, the exponent is at most $t^2 i/(2n)-t^4(n-i)/(8n^2)$. This function is at most $i^2/(2(n-i))$. By the inequality $e^u-1 \leq ue^u$ valid for $u \geq 0$, we finally conclude that the expression on the right hand side of the last inequality is at most

$$e^{-\frac{t^2}{2}}\frac{i^2}{2(n-i)}e^{\frac{i^2}{2(n-i)}} .$$

This proves all the pointwise inequalities for A_m. The integral inequalities are obtained by integrating the pointwise inequalities over the whole real line (this can only make the upper bounds larger). One needs the facts that for a normal random variable N, $E(N^2)=1$, $E(N^4)=3$, and $E(N^6)=15$. ■

Lemma 4.4.

For regular f, and $|t| \le \frac{3\sigma^3\sqrt{n}}{4\beta}$, we have

$$|\phi^n(\frac{t}{\sigma\sqrt{n}})-e^{-\frac{t^2}{2}}| \le \frac{\beta|t|^3}{3\sigma^3\sqrt{n}} e^{-\frac{t^2}{4}} + |A_n(t)| .$$

Integrated over the given interval for t, we have

$$\int |\phi^n(\frac{t}{\sigma\sqrt{n}})-e^{-\frac{t^2}{2}}| \, dt \le \frac{16\beta}{3\sigma^3\sqrt{n}}+\frac{3}{4n}\sqrt{2\pi} .$$

Proof of Lemma 4.4.

Note that

$$|\phi^n(\frac{t}{\sigma\sqrt{n}})-e^{-\frac{t^2}{2}}| \le |\phi^n(\frac{t}{\sigma\sqrt{n}})-(1-\frac{t^2}{2n})^n| + |A_n(t)| .$$

The last term is taken care of by applying Lemma 4.3. Here we need the fact that the given interval for t is always included in $[-\sqrt{n},\sqrt{n}]$, so that the bounds of Lemma 4.3 are indeed applicable. By Lemma 4.2, the first term can be written as

$$(1-\frac{t^2}{2n})^n \; |(1+\frac{\theta\beta|t|^3}{6\sigma^3 n^{\frac{3}{2}}(1-\frac{t^2}{2n})})^n-1|$$

where $|\theta| \le 1$. Using the fact that $(1+u)^n-1 \le n|u|e^{n|u|}$ for all $n>0$, and all $u \in R$, this can be bounded from above by

$$e^{-\frac{t^2}{2}} \frac{\beta|t|^3}{3\sigma^3\sqrt{n}} e^{\frac{\beta|t|^3}{3\sigma^3\sqrt{n}}} \le e^{-\frac{t^2}{4}} \frac{\beta|t|^3}{3\sigma^3\sqrt{n}} .$$

To obtain the integral inequality, use Lemma 4.3 again, and note that $\int |t|^3 e^{-t^2/4} \, dt = 16$. ■

Lemma 4.5.

For regular f,

$$\sup_x \mid f_n(x)-g(x) \mid \ \leq a_n$$

where

$$a_n = \frac{1}{2\pi}\frac{16\beta}{3\sigma^3\sqrt{n}}(1+\frac{1}{2}e^{-\frac{9\sigma^6 n}{32\beta^2}})$$
$$+\frac{1}{2\pi}\frac{3}{4n}\sqrt{2\pi}+\frac{1}{2\pi}\sup_{\mid t \mid \geq \frac{4\sigma^2}{3\beta}} \mid \phi(t) \mid^{n-1}\sigma\sqrt{n} \int \mid \phi \mid ,$$

f_n is the density of $S_n/(\sigma\sqrt{n})$ and g is the normal density. Also,

$$a_n \sim \frac{8\beta}{3\pi\sigma^3\sqrt{n}}$$

as $n \to \infty$.

Proof of Lemma 4.5.

By the inversion formula for absolutely integrable characteristic functions, we see that

$$2\pi \mid f_n(x)-g(x) \mid \ \leq \int \mid \phi^n(\frac{t}{\sigma\sqrt{n}})-e^{-\frac{t^2}{2}} \mid$$
$$\leq \int_D \mid \phi^n(\frac{t}{\sigma\sqrt{n}})-e^{-\frac{t^2}{2}} \mid \ dt + \int_{D^c} \left(\mid \phi^n(\frac{t}{\sigma\sqrt{n}}) \mid +e^{-\frac{t^2}{2}} \right) dt$$

where D is the interval defined by the condition $\mid t \mid \ \leq \frac{3\sigma^3\sqrt{n}}{4\beta}$, and D^c is the complement of D. The integral over D is bounded in Lemma 4.4 by

$$\frac{16\beta}{3\sigma^3\sqrt{n}}+\frac{3}{4n}\sqrt{2\pi} .$$

The integral over D^c does not exceed

$$\sup_{\mid t \mid \geq \frac{4\sigma^2}{3\beta}} \mid \phi(t) \mid^{n-1}\sigma\sqrt{n} \int \mid \phi \mid \ + \frac{8\beta}{3\sigma^3\sqrt{n}}e^{-\frac{9\sigma^6 n}{32\beta^2}} ,$$

where we used a well-known inequality for the tail of the normal distribution, i.e. $\int_u^\infty g \ \leq g(u)/u$. This concludes the proof of Lemma 4.5. ■

Lemma 4.6.

For regular f, and

$$| t | \leq \frac{3\sigma^3\sqrt{n}}{4\beta} ,$$

we have

$$| \phi^{n-1}(\frac{t}{\sigma\sqrt{n}}) - e^{-\frac{t^2}{2}} | \leq \frac{\beta | t |^3}{3\sigma^3\sqrt{n}} e^{-\frac{t^2}{4}} .$$

Integrated over the given interval for t, we have

$$\int | \phi^{n-1}(\frac{t}{\sigma\sqrt{n}}) - e^{-\frac{t^2}{2}} | \, dt \leq \frac{16\beta}{3\sigma^3\sqrt{n}} + \frac{3}{4n}\sqrt{2\pi} .$$

Proof of Lemma 4.6.

Note that

$$| \phi^{n-1}(\frac{t}{\sigma\sqrt{n}}) - e^{-\frac{t^2}{2}} | \leq | \phi^{n-1}(\frac{t}{\sigma\sqrt{n}}) - (1-\frac{t^2}{2n})^{n-1} | + | A_{n-1}(t) | .$$

The last term is taken care of by applying Lemma 4.3. Here we need the fact that the given interval for t is always included in $[-\sqrt{n}, \sqrt{n}]$, so that the bounds of Lemma 4.3 are indeed applicable. By Lemma 4.2, the first term can be written as

$$(1-\frac{t^2}{2n})^{n-1} \left| \left(1 + \frac{\theta\beta | t |^3}{6\sigma^3 n^{\frac{3}{2}} (1-\frac{t^2}{2n})} \right)^{n-1} - 1 \right|$$

where $| \theta | \leq 1$. Using the fact that $(1+u)^{n-1} - 1 \leq n | u | e^{n | u |}$ for all $n > 0$, and all $u \in R$, this can be bounded from above by

$$\frac{e^{-(n-1)t^2}}{2n} \frac{\beta | t |^3}{3\sigma^3\sqrt{n}} e^{\frac{\beta | t |^3}{3\sigma^3\sqrt{n}}} \leq e^{-(1-\frac{1}{2n})\frac{t^2}{4}} \frac{\beta | t |^3}{3\sigma^3\sqrt{n}} .$$

To obtain the integral inequality, use Lemma 4.3 again, and note that $\int | t |^3 e^{-t^2/4} \, dt = 16$. ■

Lemma 4.7.

Let g be the normal density and let f_n be the density of the normalized sum $S_n/(\sigma\sqrt{n})$ for iid random variables with a regular density f. Let ϕ be the characteristic function for f. Then

$$| f_n(x)-g(x) | \leq \frac{b_n}{x^2} ,$$

where

$$b_n = \frac{4\beta}{3\pi\sigma^3\sqrt{n}} e^{-\frac{9\sigma^6 n}{32\beta^2}} + \frac{\sqrt{8}}{\pi} e^{-\frac{9\sigma^6 n}{64\beta^2}}$$
$$+ \frac{1}{2\pi}\rho^{n-2}\sigma\sqrt{n} \int | \phi | + \frac{1}{2\pi}\rho^{n-3}\sigma^3 n^{\frac{3}{2}} \int t^2 | \phi |$$
$$+ \frac{1}{2\pi}\left(\frac{208\beta}{3\sigma^3\sqrt{n}} + \frac{18\sqrt{2\pi}}{4n} \right)$$
$$+ \frac{1}{\sqrt{4\pi n(n-1)}} + \frac{3}{(n-2)\sqrt{2\pi}}$$
$$+ \frac{1}{n\sigma^2\sqrt{2\pi}} + \frac{\beta}{\sigma^3\sqrt{n}}\left(\frac{1}{\sigma^2\sqrt{8\pi}}+2\right) .$$

Here $\rho = \sup_{|t| \geq \frac{3\sigma^2}{4\beta}} | \phi(t) |$. Note that as $n \to \infty$, $b_n \sim \frac{b}{\sqrt{n}}$ where $b = \frac{\beta}{\sigma^3}\left(\frac{1}{\sigma^2\sqrt{8\pi}}+2+\frac{208}{6\pi}\right)$.

Proof of Lemma 4.7.

As in Lemma 4.5, we define the interval D by the condition $| t | \leq \frac{3\sigma^3\sqrt{n}}{4\beta}$, and let D^c be the complement of D. Let I be the interval defined by $| t | \leq \frac{3\sigma^2}{4\beta}$, and let I^c be the complement of I. By Lemma 4.2, it is easy to see that for $t \in I$, $| \phi(t) | \leq 1-\sigma^2 t^2/4$. Thus,

$$\frac{1}{2\pi}\int_D | \phi^n\left(\frac{t}{\sigma\sqrt{n}}\right)-\phi^{n-1}\left(\frac{t}{\sigma\sqrt{n}}\right) | \ dt$$
$$\leq \frac{1}{2\pi}\int_D | 1-\phi\left(\frac{t}{\sigma\sqrt{n}}\right) | \ | \phi\left(\frac{t}{\sigma\sqrt{n}}\right) |^{n-1} dt$$
$$\leq \frac{1}{2\pi}\int \frac{t^2}{2n} e^{-\frac{(n-1)t^2}{4n}} dt$$
$$= \frac{1}{4\pi n}\sqrt{2\pi}\sqrt{\frac{2n}{n-1}}$$

$$= \frac{1}{\sqrt{4\pi n(n-1)}} .$$

Similarly,

$$\frac{1}{2\pi}\int_D t^2 \mid \phi^n(\frac{t}{\sigma\sqrt{n}})-\phi^{n-2}(\frac{t}{\sigma\sqrt{n}}) \mid \ dt$$

$$\leq \frac{1}{2\pi}\int_D t^2 \mid 1-\phi^2(\frac{t}{\sigma\sqrt{n}}) \mid \ \mid \phi(\frac{t}{\sigma\sqrt{n}}) \mid^{n-2} dt$$

$$\leq \frac{1}{2\pi}\int \frac{t^4}{n} e^{-\frac{(n-2)t^2}{4n}} \ dt$$

$$= \frac{1}{2\pi n}\sqrt{2\pi}\, 3\frac{2n}{n-2}$$

$$= \frac{3}{(n-2)\sqrt{2\pi}} .$$

So far for the preliminary computations. We begin with the observation that

$$x^2(f_n(x)-g(x)) = \frac{1}{2\pi}\int ((t^2-1)e^{-\frac{t^2}{2}}-\phi_n''(t))e^{-itx} \ dt$$

where ϕ_n is the characteristic function corresponding to f_n. Obviously,

$$x^2 \mid f_n(x)-g(x) \mid \ \leq \frac{1}{2\pi}\int \mid (t^2-1)e^{-\frac{t^2}{2}}-\phi_n''(t) \mid \ dt .$$

The second derivative of the n-th power of $\phi(t/(\sigma\sqrt{n}))$ is

$$\frac{n-1}{\sigma^2}\phi'^2\phi^{n-2}+\frac{1}{\sigma^2}\phi''\phi^{n-1} ,$$

where all the omitted arguments are $t/(\sigma\sqrt{n})$. By the triangle inequality, we obtain

$$x^2 \mid f_n(x)-g(x) \mid \ \leq \frac{1}{2\pi}\int \mid (t^2-1)e^{-\frac{t^2}{2}}-\phi_n''(t) \mid \ dt$$

$$\leq \frac{1}{2\pi}\Bigg(\int \mid e^{-\frac{t^2}{2}}-\phi^{n-1}(t/(\sigma\sqrt{n})) \mid \ dt + \int t^2 \mid e^{-\frac{t^2}{2}}-\phi^{n-2}(t/(\sigma\sqrt{n})) \mid \ dt$$

$$+ \int e^{-\frac{t^2}{2}} \mid \frac{n-1}{\sigma^2}\phi'^2(t/(\sigma\sqrt{n}))-t^2 \mid \ dt$$

$$+ \int e^{-\frac{t^2}{2}} \mid \sigma^{-2}\phi''(t/(\sigma\sqrt{n}))+1 \mid \ dt \Bigg)$$

$$= J_1+J_2+J_3+J_4 .$$

From Lemma 4.2, we recall

$$| \frac{\sqrt{n}\,\phi'}{\sigma} + t \,| \le \frac{\beta t^2}{2\sigma^3\sqrt{n}} \,,$$

$$| \frac{\phi''}{\sigma^2} + 1 \,| \le \frac{\beta \,| t |}{\sigma^3\sqrt{n}} \,.$$

Using the fact that $| \phi'(t/(\sigma\sqrt{n}\,)) | \le E(\,|X|\,)/(\sigma\sqrt{n}\,) \le 1/\sqrt{n}$, we have

$$| \frac{n\,\phi'^2}{\sigma^2} - t^2 | \le | \frac{\sqrt{n}\,\phi'}{\sigma} - t \,| \, | \frac{\sqrt{n}\,\phi'}{\sigma} + t \,|$$

$$\le (\frac{1}{\sigma^2} + |\,t\,|\,) \frac{\beta t^2}{2\sigma^3\sqrt{n}} \,.$$

Using the fact that $\int |\,t\,|^i e^{-t^2/2}\,dt$ takes the values $\sqrt{2\pi}$,2,$\sqrt{2\pi}$ and 4 for $i=0,1,2,3$ respectively, we see that

$$J_3 + J_4 \le \int |\,\sigma^{-2}\phi'^2 e^{-\frac{t^2}{2}}\,|\,dt + \frac{1}{\sqrt{2\pi}}\frac{\beta}{2\sigma^5\sqrt{n}} + \frac{2\beta}{\pi\sigma^3\sqrt{n}}$$

$$\le \frac{1}{n\,\sigma^2\sqrt{2\pi}} + \frac{\beta}{\sigma^3\sqrt{n}}(\frac{1}{\sigma^2\sqrt{8\pi}} + 2) \,.$$

This leaves us with J_1 and J_2. Here we will split the integrals over D and D^c. First of all,

$$\frac{1}{2\pi}\left(\int_D |\,e^{-\frac{t^2}{2}} - \phi^{n-1}(t/(\sigma\sqrt{n}\,))\,|\,dt + \int_D t^2\,|\,e^{-\frac{t^2}{2}} - \phi^{n-2}(t/(\sigma\sqrt{n}\,))\,|\,dt \right)$$

$$\le \frac{1}{2\pi}\left(\int_D |\,e^{-\frac{t^2}{2}} - \phi^{n}(t/(\sigma\sqrt{n}\,))\,|\,dt + \int_D t^2\,|\,e^{-\frac{t^2}{2}} - \phi^{n}(t/(\sigma\sqrt{n}\,))\,|\,dt \right)$$

$$+ \frac{1}{2\pi}\left(\int_D |\,\phi^{n-1}(t/(\sigma\sqrt{n}\,)) - \phi^{n}(t/(\sigma\sqrt{n}\,))\,|\,dt \right.$$

$$\left. + \int_D t^2\,|\,\phi^{n-2}(t/(\sigma\sqrt{n}\,)) - \phi^{n}(t/(\sigma\sqrt{n}\,))\,|\,dt \right) .$$

The last two terms were bounded from above earlier on in the proof by

$$\frac{1}{\sqrt{4\pi n\,(n-1)}} + \frac{3}{(n-2)\sqrt{2\pi}} \,.$$

By Lemma 4.4, we have for $t \in D$,

$$|\,\phi^n(\frac{t}{\sigma\sqrt{n}}) - e^{-\frac{t^2}{2}}\,| \le \frac{\beta\,|\,t\,|^3}{3\sigma^3\sqrt{n}} e^{-\frac{t^2}{4}} + |\,A_n(t)\,| \,.$$

Thus, by Lemma 4.3, and the following integrals:

$$\int |\,t\,|^3 e^{-\frac{t^2}{4}}\,dt = 16 \,,$$

$$\int |t|^5 e^{-\frac{t^2}{4}} dt = 192 ,$$

$$\int |t|^4 e^{-\frac{t^2}{2}} dt = 3\sqrt{2\pi} ,$$

$$\int |t|^6 e^{-\frac{t^2}{4}} dt = 15\sqrt{2\pi} ,$$

we have

$$\frac{1}{2\pi}\int_D (1+t^2) \left| \phi^n \left(\frac{t}{\sigma\sqrt{n}}\right) - e^{-\frac{t^2}{2}} \right| dt$$

$$\leq \frac{1}{2\pi}\int (1+t^2)\left(\frac{\beta |t|^3}{3\sigma^3\sqrt{n}} e^{-\frac{t^2}{4}} + \frac{t^4}{4n} e^{-\frac{t^2}{2}} \right) dt$$

$$= \frac{1}{2\pi}\left(\frac{208\beta}{3\sigma^3\sqrt{n}} + \frac{18\sqrt{2\pi}}{4n} \right) .$$

Finally, we have to evaluate the integrals in J_1+J_2 taken over D^c. These are estimated from above by

$$\frac{1}{2\pi}\int_{D^c} (1+t^2) e^{-\frac{t^2}{2}} dt + \frac{1}{2\pi}\rho^{n-2}\sigma\sqrt{n} \int |\phi| + \frac{1}{2\pi}\rho^{n-3}\sigma^3 n^{\frac{3}{2}} \int t^2 |\phi|$$

where $\rho = \sup_{I^c} |\phi|$. The region D^c is defined by the condition $|t| > c$ for some constant c . The first term in the last expression can thus be rewritten as

$$\frac{1}{\pi} \int_{u > c^2/2} ((2u)^{\frac{1}{2}} + \sqrt{2u}) e^{-u} du$$

$$\leq \frac{1}{c\pi} e^{-\frac{c^2}{2}} + \frac{\sqrt{8}}{\pi} e^{-\frac{c^2}{4}}$$

$$= \frac{4\beta}{3\pi\sigma^3\sqrt{n}} e^{-\frac{9\sigma^6 n}{32\beta^2}} + \frac{\sqrt{8}}{\pi} e^{-\frac{9\sigma^6 n}{64\beta^2}} .$$

Collecting bounds gives the desired result. ■

For the bound of Lemma 4.7 to be useful, it is necessary that f not only be regular, but also that its characteristic function satisfy

$$\int t^2 |\phi(t)| \, dt < \infty .$$

This implies that f has two bounded continuous derivatives tending to 0 as $|x| \to \infty$, and in fact

$$f''(x) = -\frac{1}{2\pi}\int e^{-itx}\, t^2\phi(t)\, dt \ .$$

(see e.g. Kawata, 1972, pp. 438-439). This smoothness condition is rather restrictive and can be considerably weakened. The asymptotic bound $b/(x^2\sqrt{n})$ remains valid if $\int t^2 \mid \phi(t) \mid^k < \infty$ for some positive integer k (exercise 4.4). Lemmas 4.5 and 4.7 together are but special cases of more general local limit theorems, such as those found in Maejima (1980) and Inzevitov (1977), except that here we explicitly compute the universal constants in the bounds.

4.5. The mixture method for simulating sums.

When a density f can be written as a mixture

$$f(x) = \sum_{i=1}^{\infty} p_i f_i(x) \ ,$$

where the f_i's are simple densities, then simulation of the sum S_n of n iid random variables with density f can be carried out as follows.

The mixture method for simulating sums

Generate a multinomial $(n, p_1, p_2, \ldots)$ random sequence $N_1, N_2, \ldots$ (note that the N_i's sum to n). Let K be the index of the largest nonzero N_i.

$X \leftarrow 0$

FOR $i := 1$ TO K DO

  Generate S, the sum of N_i iid random variables with common density f_i.

  $X \leftarrow X + S$

RETURN X

The validity of the algorithm is obvious. The algorithm is put in its most general form, allowing for infinite mixtures. A multinomial random sequence is of course defined in the standard way: imagine that we have an infinite number of urns, and that n balls are independently thrown in the urns. Each ball lands with probability p_i in the i-th urn. The sequence of cardinalities of the urns is a multinomial $(n, p_1, p_2, \ldots)$ random sequence. To simulate such a sequence, note that N_1 is binomial (n, p_1), and that given N_1, N_2 is binomial $(n-N_1, p_2/(1-p_1))$, etcetera. If K is the index of the last occupied urn, then it is easy to see that the multinomial sequence can be generated in expected time $O(E(K))$.

The mixture method is efficient if sums of iid random variables with densities f_i are easy to generate. This would for example be the case if f were a finite mixture of stable, gamma, exponential or normal random variables. Perhaps the most intriguing decomposition is that of a unimodal density: every unimodal density can be written as a countable mixture of uniform densities. This statement is intuitively clear, because subtracting a function of the form $cI_{[a,b]}(x)$ from f leaves a unimodal piece on $[a,b]$ and two unimodal tails. This can be repeated for all pieces individually, and at the same time the integral of the left-over function can be made to tend to zero by the judicious choice of rectangular functions (see exercise 4.5). If we can generate sums of iid uniform random variables uniformly fast (with respect to n), then the expected time taken by the mixture method is $O(E(K))$. A few remarks about generating uniform sums are given in the next section.

4.6. Sums of independent uniform random variables.

In this section we consider the distribution of

$$S_n = \sum_{i=1}^{n} U_i ,$$

where $U_1, \ldots, U_n$ are iid uniform [-1,1] random variables. The distribution can be described in a variety of ways:

Theorem 4.3.

The characteristic function of S_n is

$$\left(\frac{\sin(t)}{t}\right)^n .$$

For all $n \geq 2$, the density f_n can be obtained by the inversion formula

$$f_n(x) = \frac{1}{2\pi}\int\left(\frac{\sin(t)}{t}\right)^n \cos(tx)\,dt .$$

This yields

$$f_n(x) = \frac{1}{(i-1)!}\frac{1}{2}\sum_{k=0}^{i-1}(-1)^k \binom{i}{k}(x-(2k-n))^{i-1}$$

where $2i-2-n < x < 2i-n$; $i=1,2, \ldots, n$.

Proof of Theorem 4.3.

The characteristic function is obtained by using the definition. Since the characteristic function of S_n for all $n \geq 2$ is absolutely integrable, f_n can be obtained by the given inversion integral. There is also a direct way of computing the distribution function F_n and density of S_n; its derivation goes back to the nineteenth century (see e.g. Cramer (1951, p. 245)). Different proofs include the geometric approach followed by us in Theorem I.4.4 (see also Hall (1927) and Roach (1963)), an induction argument (Olds, 1952), and an application of the residue theorem (Lusk and Wright, 1982). Taking the derivative of F_n given in Theorem I.4.4 gives the formula

$$\frac{1}{(n-1)!}\left(x_+{}^{n-1} - n\,(x-1)_+{}^{n-1} + \binom{n}{2}(x-2)_+{}^{n-1} - \cdots + (-1)^n \binom{n}{n}(x-n\,)_+{}^{n-1} \right)$$

for the density of the sum of n iid uniform [0,1] random variables. The the density of sums of symmetric uniform random variables is easily obtained by the transformation formula for densities. ■

It is easy to see that the local limit theorems developed in Lemmas 4.5 and 4.7 are applicable to this case. There is one small technical hurdle since the characteristic function of a uniform random variable is not absolutely integrable. This is easily overcome by noting that the square of the characteristic function is absolutely integrable. If we recall the rejection algorithm of section 4.3, we note that the expected number of iterations is $O(1/\sqrt{n}\,)$ and that the expected number of evaluations of f_n is $O(1/\sqrt{n}\,)$. Unfortunately, this is not good enough, since the evaluation of $f_n(x)$ by the last formula of Theorem 4.3 takes time roughly proportional to n for nearly all x of interest. This would yield a global expected time roughly increasing as $\sqrt{n}$. The formula for f_n is thus of limited value. There are two solutions: either one uses the series method based upon a series expansion for f_n which is tailored around the normal density, or one uses a local limit theorem with $O(1/n)$ error by using as main component the normal density plus the first term in the asymptotic expansion which is a normal density multiplied with a Hermite polynomial (see e.g. Petrov, 1975). The latter approach seems the most promising at this point (see exercise 4.6).

4.7. Exercises.

1. Let f be a density, whose normalized sums tend in distribution to the symmetric stable (α) density. Assume that the stable density can be evaluated exactly in one unit of time at every point. Derive first some inequalities for the difference between the density of the normalized sum and the stable density. These non-uniform inequalities should be such that the integral of the error bound with respect to x tends to 0 as $n \to \infty$. Hint: look for error terms of the form $\min(a_n, b_n \mid x \mid^{-c})$ where c is a positive constant, and a_n, b_n are positive number sequences tending to 0 with n. Mimic the derivation of the local limit theorem in the case of attraction to the normal law.

2. **The gamma density.** The zero mean exponential density has characteristic function $\phi = e^{-it}/(1-it)$. In the notation of this chapter, derive for this distribution the following quantities:

   A. $\sigma = 1$, $\beta = \frac{12}{e} - 2$.

   B. $\int \mid \phi \mid = \infty$, $\int \mid \phi \mid^2 = \pi$.

   C. $\sup_{\mid t \mid \geq c} \mid \phi(t) \mid = 1/\sqrt{1+c^2}$ $(c > 0)$.

   Note that the bounds in the local limit theorems are not directly applicable since $\int \mid \phi \mid = \infty$. However, this can be overcome by bounding $\int \mid \phi \mid^n$ by $s \int \mid \phi \mid^2$ where s is the supremum of $\mid \phi \mid$ over the domain of integration, to the power $n-2$. Using this device, derive the rejection constant from the thus modified local limit theorem as a function of n.

3. A continuation of exercise 2. Let f_a be the normalized (zero mean, unit variance) gamma (a) density, and let g be the normal density. By direct means, find sequences a_n, b_n such that for all $a \geq 1$,

$$\mid f_a(x) - g(x) \mid \leq \min(a_n, \frac{b_n}{x^2}),$$

   and compare your constants with those obtained in exercise 2. (They should be dramatically smaller.)

4. Prove the claim that in Lemma 4.7, $b_n \sim b/(x^2\sqrt{n})$ when the condition $\int t^2 \mid \phi(t) \mid \, dt < \infty$ is relaxed to

$$\int t^2 \mid \phi(t) \mid^k \, dt < \infty$$

   where $k > 0$ is a fixed integer.

5. Consider a monotone density f on $[0,\infty)$. Give a constructive completely automatic rule for decomposing this density as a countable mixture of uniform densities, i.e. the decomposition should be obtainable even if f is only given in black box format, and the countable mixture should give us the monotone density again in the sense that the L_1 distance between the two densities is zero (this allows the functions to be different on possibly uncountable sets of zero measure). Can you make a statement about the rate of decrease of p_i for the following subclasses of monotone densities: the log-

concave densities, the concave densities, the convex densities? Prove that when $p_i \leq ce^{-bi}$ for some $b, c > 0$ and all i, then $E(K) = O(\log(n))$, where K is the largest integer in a sample of size n drawn from probability vector $p_1, p_2, \ldots$. Conclude that for important classes of densities, we can generate sums of n iid random variates in expected time $O(\log(n))$.

6. **Gram-Charlier series.** The standard approximation for the density f_n of $S_n/(\sigma\sqrt{n})$ where S_n is the sum of n iid zero mean random variables with second moment σ^2 is g where g is the normal density. The closeness is covered by local central limit theorems, and the errors are of the order of $1/\sqrt{n}$. To obtain errors of the order of $1/n$ it is necessary to user a finer approximation. For example, one could use an extra term in the Gram-Charlier series (see e.g. Ord (1972, p. 26)). This leads to the approximation by

$$\frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} \left(1 + \frac{\mu_3}{6\sigma^3\sqrt{n}} (x^3 - x)\right),$$

where μ_3 is the third moment for f. For symmetric distributions, the extra correction term is zero. This suggests that the local limit theorems of section 4.3 can be improved. For the symmetric uniform density, find constants a, b such that $|f_n - g| \leq \frac{1}{n} \min(a, bx^{-2})$. Use this to design a uniformly fast generator for sums of symmetric uniform random variables.

7. A continuation of the previous exercise. Let $a \in R$ be a constant. Give a random variate generator for the following class of densities related to the Gram-Charlier series approximation of the previous exercise:

$$g(x) = c\left(\frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} (1 + a(x^3 - x))\right)_+ ,$$

where c is a normalization constant.

5. DISCRETE EVENT SIMULATION.

5.1. Future event set algorithms.

Several complex systems evolving in time fall into the following category: they can be characterized by a state, and the state changes only at discrete times. Systems falling into this category include most queueing systems such as those appearing in banks, elevators, computer networks, computer operating systems and telephone networks. Systems not included in this category are those which change state continuously, such as systems driven by differential equations (physical or chemical processes, traffic control systems). In discrete event simulation of such systems, one keeps a subset of all the future events in a future event set, where an event is defined as a change of state, e.g. the arrival or departure of a person in a bank. By taking the next event from the future event set, we can make time advance with big jumps. After having grabbed this event, it is necessary to update the state and if necessary schedule new future events. In other words, the future event set can shrink and grow in its lifetime. What matters is that no event is missed. All future event set algorithms can be summarized as follows:

Future event set algorithm

Time ←0.
Initialize State (the state of the system).
Initialize FES (future event set) by scheduling at least one event.
WHILE NOT EMPTY (FES) DO
    Select the minimal time event in FES, and remove it from FES.
    Time ← time of the selected event, i.e. make time progress.
    Analyze the selected event, and update State and FES accordingly.

For worked out examples, we refer the readers to more specialized texts such as Bratley, Fox and Schrage (1983), Banks and Carson (1984) or Law and Kelton (1982). Our main concern is with the complexity aspect of future event set algorithms. It is difficult to get a good general handle on the time complexity due to the state updates. On the other hand, the contribution to the time complexity of all operations involving FES, the future event set, is amenable to analysis. These operations include

A. INSERT a new event in FES.

B. DELETE the minimal time event from FES.

C. CANCEL a particular event (remove it from FES).

There are two kinds of INSERT: INSERT based upon the time of the event, and INSERT based upon other information related to the event. The latter INSERT is required when a simulation demands information retrieval from the FES other than selection of the minimal time event. This is the case when cancelations can occur, i.e. deletions of events other than the minimal time event. It can always be avoided by leaving the event to be canceled in FES but marking it "canceled", so that when it is selected at some point as the minimal time event, it can immediately be discarded. In most cases we have to use a dual data structure which allows us to implement the operations INSERT, DELETE and either CANCEL or MARK efficiently. Typically, one part of the data structure consists of a dictionary (ordered according to keys used for canceling or marking), and another part is a priority queue (see Aho, Hopcroft and Ullman (1983) for our terminolgy). Since the number of elements in FES grows and shrinks with time, it is difficult to uniformize the analysis. For this reason, sometimes the following assumptions are made:

A. The future event set has n events at all times. This implies that when the minimum time event is deleted, the empty slot is immediately filled by a new event, i.e. the DELETE and INSERT operations always go together.

B. Initially, the future event set has n events, with random times, all iid with common distribution function F on $[0,\infty)$.

C. When an event with event time t is deleted from FES, the new event replacing it in FES has time $t+T$, where T also has distribution function F.

These three assumptions taken together form the basis of the so-called hold model, coined after the SIMULA HOLD operation, which combines our DELETE and INSERT operations. Assumptions B and C are of a stochastic nature to facilitate the expected time analysis. They are motivated by the fact that in homogeneous Poisson processes, the inter-event times are independent exponentially distributed. Therefore, the distribution function F is typically the exponential distribution. The quantity of interest to us is the expected time needed to execute a HOLD operation. Unfortunately, this quantity depends not only upon n, but also on F and the time instant at which the expected time analysis is needed. This is due to the fact that the times of the events in the FES have distributions that vary. It is true that relative to the minimum time in the FES, the distribution of the $n-1$ non-minimal times approaches a limit distribution, which depends upon F and n. Analysis based upon this limit distribution is at times risky because it is difficult to pinpoint in complex systems when the steady state is almost reached. What complicates matters even more is the dependence of the limit distribution upon n. The limit of the limit distribution with respect to n, a double limit of sorts, has density $(1-F(x))/\mu$ $(x>0)$ where μ is the mean for F (Vaucher, 1977). The analyses are greatly facilitated if this limit distribution is used as the distribution of the relative event times in FES. The results of these analyses should be handled with great care. Two extensive reports based upon this model were carried out by Kingston (1985) and McCormack and Sargent (1981). An alternative model was proposed by Reeves (1984). He also works with this limiting distribution, but departs from the HOLD model, in that events are inserted, or scheduled, in the FES according to a homogeneous Poisson process. This implies that the size of the FES is no longer fixed at a given level n, but hovers around a mean value n. It seems thus safer to perform a worst-case time analysis, and to include an expected time analysis only where exact calculations can be carried out. Luckily, for the important exponential distribution, this can be done.

Theorem 5.1.

If assumptions A-C hold, and F is the exponential (λ) distribution, if k HOLD operations have been carried out for any integer k, if $X*$ is the minimal event time in the FES, and $X_1, X_2, \ldots, X_{n-1}$ are the $n-1$ non-minimal event times in the FES (unordered, but in order of their insertion in the FES), then $X_1-X*, \ldots, X_{n-1}-X*$ are iid exponential (λ) random variables.

Proof of Theorem 5.1.

This is best proved inductively. Initially, we have n exponentially distributed times. The assertion is certainly true, by the memoryless property of the exponential distribution. Now, take the minimum time, say M, remove it, and insert the time $M+E$ in the FES, where E is exponential (λ). Clearly, all n times in the FES are now iid with an exponential (λ) distribution on $[M,\infty)$. We are thus back where we started from, and can apply the memoryless property again. ∎

Reeves's model allows for a simple direct analysis for all distribution functions F. Because of its importance, we will briefly study his model in a separate section, before moving on to the description of a few possible data structures for the FES.

5.2. Reeves's model.

In Reeves's model, the FES is initially empty. Insertions occur at random times, which correspond to a homogeneous Poisson process with rate λ. The time of an inserted event is the insertion time plus a delay time which has distribution function F. A few properties will be needed further on, and these are collected in Theorem 5.2:

Theorem 5.2.

Let $0<T_1<T_2<\cdots$ be a homogeneous Poisson process with rate $\lambda>0$ (the T_i's are the insertion times), and let $X_1,X_2,\ldots$ be iid random variables with common distribution function F on $[0,\infty)$. Then

A. The random variables T_i+X_i, $1\leq i$, form a nonhomogeneous Poisson process with rate function $\lambda F(t)$.

B. If N_t is the number of events in FES at time t, then N_t is Poisson $(\lambda\int_0^t(1-F))$. N_t is thus stochastically smaller than a Poisson $(\lambda\mu)$ random variable where $\mu=\int_0^\infty(1-F)$ is the mean for F.

C. Let $V_i, i\leq N_t$, be the event times for the events in FES at time t. Then the random variables V_i-t form a nonhomogeneous Poisson process with rate function $\lambda(F(t+u)-F(u))$, $u\geq 0$.

Proof of Theorem 5.2.

Most of the theorem is left as an exercise on Poisson processes. The main task is to verify the Poisson nature of the defined processes by checking the independence property for nonoverlapping intervals. We will mainly point out how the various rate functions are obtained.

For part A, let L be the number of insertions up to time t, a Poisson (λt) random variable, and let M be the number of T_i+X_i's not exceeding t. Clearly, by the uniform distribution property of homogeneous Poisson processes, M is distributed as

$$\sum_{i=0}^{L} I_{[tU_i+X_i \leq t]} ,$$

where the U_i's are iid uniform [0,1] random variables. Note that this is a Poisson sum of iid Bernoulli random variables. As we have seen elsewhere, such sums are again Poisson distributed. The parameter is λtp where $p=P(tU_1+X_1 \leq t)$. The parameter can be rewritten as

$$\lambda t P(X_1 \leq tU_1) = \lambda t \int_0^1 F(tu)\, du$$

$$= \lambda \int_0^t F(u)\, du .$$

For part B, the rate function can be obtained similarly by writing N_t as a Poisson (λt) sum of iid Bernoulli random variables with success probability $p=P(tU_1+X_1>t)$. This is easily seen to be Poisson $(\lambda \int_0^t (1-F))$. For the second statement of part B, recall that the mean for distribution function F is $\int_0^\infty (1-F)$.

Finally, consider part C. Here again, we argue analogously. Let M be the number of events in FES at time t with event times not exceeding $t+u$. Then M is a Poisson (λt) sum of iid Bernoulli random variables with success parameter p given by

$$P(t \leq tU_1+X_1<t+u) = \int_0^1 (F(tz+u)-F(tz))\, dz$$

$$= \frac{1}{t}\int_0^t (F(z+u)-F(z))\, dz .$$

The statement about the rate function follows directly from this. ∎

The asymptotics in Reeves's model should not be with respect to N_t, the size of the FES, because this oscillates randomly. Rather, it should be with respect to t, the time. The first important observation is that the expected size of the FES at time t is $\lambda \int_0^t (1-F) \to \lambda\mu$ as $t \to \infty$, where μ is the mean for F. If μ is small, the FES is small because events spend only a short time in FES. On the other hand, if $\mu=\infty$, then the expected size of the FES tends to ∞ as $t \to \infty$, i.e. we would need infinite space in order to be able to carry out an unlimited time simulation. The situation is also bad when $\mu<\infty$, although not as bad as in the case $\mu=\infty$: it can be shown (see exercises) that $\limsup_{t \to \infty} N_t = \infty$ almost surely. Thus, in all cases, an unlimited memory would be required. This should be viewed as a serious drawback of Reeves's model. But the insight we gain from his model is invaluable, as we will find out in the next section on linear lists.

5.3. Linear lists.

The oldest and simplest structure for an FES is a linear list in which the elements are kept according to increasing event times. For what follows, it is all but irrelevant whether a linked list implementation or an array implementation is chosen. Deletion is obviously a constant time operation. Insertion of an element in the i-th position takes time proportional to i if we start searching from the front (small event times) of the list, and to $n-i+1$ if we start from the back and n is the cardinality of the FES. We can't say that the time is $\min(i, n-i+1)$ because the value of i is unknown beforehand. Thus, one of the questions to be studied is whether we should start the search from the front or the back.

By Theorem 5.2, part C, we observe that at time t_0, the expected value of the number of events exceeding the currently inserted element (called M_{t_0}) is

$$E(M_{t_0}) = \lambda \int_0^\infty \int_t^\infty (F(t_0+u)-F(u))\, du \; dF(t)$$

$$= \lambda \int_0^\infty (F(t_0+u)-F(u)) \int_0^u dF(t)\, du$$

$$= \lambda \int_0^\infty F(u)(F(t_0+u)-F(u))\, du \; .$$

Here we used a standard interchange of integrals. Since the expected number of elements in the FES is $\lambda \int_0^\infty (F(t_0+u)-F(u))\, du$, the expected value of the number of event times at most equal to the event time of the currently inserted element (called L_{t_0}) is

$$E(L_{t_0}) = \lambda \int_0^\infty (1-F(u))(F(t_0+u)-F(u))\, du \; .$$

We should search from the back when $E(M_{t_0})<E(L_{t_0})$, and from the front otherwise. In an array implementation, the search can always be done by binary search in logarithmic time, but the updating of the array calls for the shift by one position of the entire lower or upper portion of the array. If one imagines a circular array implementation with free wrap-around, of the sort used to implement queues (Standish, 1980), then it is always possible to move only the smaller portion. The same is true for a linked list implementation if we keep pointers to the front, rear and middle elements in the linked list and use double linking to allow for the two types of search. The middle element is first compared with the inserted element. The outcome determines in which half we should insert, where the search should start from, and how the middle element should be updated. The last operation would also require us to keep a count of the number of elements in the linked list. We can thus conclude that for a linear list insertion, we can find an implementation taking time bounded by $\min(M_{t_0},L_{t_0})$. By Jensen's inequality, the expected time for insertion does not exceed

$$\min(E(M_{t_0}),E(L_{t_0})) .$$

The fact that all the formulas for expected values encountered so far depend upon the current time t_0 could deprive us from some badly needed insight. Luckily, as $t_0 \to \infty$, a steady state is reached. In fact, this is the only case studied by Reeves (1984). We summarize:

Theorem 5.3.

In Reeves's model, we have

$$E(M_{t_0}) \uparrow \lambda \int_0^\infty F(1-F) \text{ as } t_0 \to \infty ,$$

$$E(L_{t_0}) \uparrow \lambda \int_0^\infty (1-F)^2 \text{ as } t_0 \to \infty .$$

Proof of Theorem 5.3.

We will only consider the first statement. Note that $E(M_{t_0})$ is monotone $\uparrow$ in t_0, and that for every t_0, the value does not exceed $\lambda \int_0^\infty F(1-F)$. Also, by Fatou's lemma,

$$\liminf_{t_0 \to \infty} E(M_{t_0}) \geq \lambda \int_0^\infty \liminf_{t_0 \to \infty} F(u)(F(t_0+u)-F(u))\, du = \lambda \int_0^\infty F(1-F) . \blacksquare$$

Remark 5.1. Front or back search.

From Theorem 5.3, we deduce that a front search is indicated when $\int(1-F)^2 < \int F(1-F)$. It is perhaps interesting to note that equality is reached for the exponential distribution. Barlow and Proschan (1975) define the NBUE (NWUE) distributions as those distributions for which for all $t > 0$,

$$\int_t^\infty (1-F) \leq (\geq) \; \mu(1-F(t)) ,$$

where μ is the mean for F. Examples of NBUE (new better than used in expectation) distributions include the uniform, normal and gamma distributions for parameter at least one. NWUE distributions include mixtures of exponentials and gamma distributions with parameter at most one. By our original change of integral we note that for NBUE distributions,

$$\lambda \int_0^\infty F(1-F) = \lambda \int_0^\infty \left(\int_t^\infty (1-F) \right) dF(t)$$

$$\leq \lambda\mu \int_0^\infty (1-F(t)) \, dF(t) = \frac{\lambda\mu}{2} .$$

Since the asymptotic expected size of the FES is $\lambda\mu$, we observe that for NBUE distributions, a back search is to be preferred. For NWUE distributions, a front search is better. In all cases, the trick with the median pointer (for linked lists) or the median comparison (for circular arrays) automatically selects the best search mode. ∎

Remark 5.2. The HOLD model.

In the HOLD model, the worst-case insertion time can be as poor as n. For the expected insertion time, the computations are simple for the exponential distribution function. In view of Theorem 5.1, it is easy to see that the expected number of comparisons in a forward scan is $\frac{n+2}{2} - \frac{1}{n+1} = \frac{n}{2} + \frac{n}{n+1}$. The expected number of backward scans is equal to this, by symmetry. For all distributions F having a density, the expected insertion time grows linearly with n (see exercises). ∎

A brief historical remark is in order. Linear lists have been used extensively in the past. They are simple to implement, easy to analyze and use minimal storage. Among the possible physical implementations, the doubly linked list is perhaps the most popular (Knuth, 1969). The asymptotic expected insertion time for front and back search under the HOLD model was obtained by Vaucher (1977) and Englebrecht-Wiggans and Maxwell (1978). Reeves (1984) discusses the same thing for his model. Interestingly, if the size n in the HOLD model is replaced by the asymptotic value of the expected size of the FES, $\lambda\mu$, the two results coincide. In particular, Remark 5.1 applies to both models. The point about NBUE distributions in that remark is due to McCormack and Sargent (1981). The idea of using a median pointer or a median comparison goes back to Pritsker (1976) and Davey and Vaucher (1980). For more analysis involving linear lists, see e.g. Jonassen and Dahl (1975).

The simple linear list has been generalized and improved upon in many ways. For example, a number of algorithms have been proposed which keep an additional set of pointers to selected events in the FES. These are known as multiple pointer methods, and the implementations are sometimes called indexed linear list implementations. The pointers partition the FES into smaller sets containing a few events each. This greatly facilitates insertion. For example, Vaucher and Duval (1975) space pointer events (events pointed to by these pointers) equal amounts of time (Δ) apart. In view of this, we can locate a particular subset of the FES very quickly by making use of the truncation operation. The subset is then searched in the standard sequential manner. Ideally, one would like to have a constant number of events per interval, but this is difficult to enforce. In Reeves's model, the analysis of the Vaucher-Duval bucket structure is easy. We need only concern ourselves with insertions. Furthermore, the time needed to locate the subset (or bucket) in which we should insert is constant. The buckets should be thought of as small linked lists. They actually need not be globally concatenated, but within each list, the events are ordered. The global time interval is divided into intervals $[0,\Delta),[\Delta,2\Delta),\ldots$. Let A_j be the j-th interval, and let $F(A_j)$ denote the probability of the j-th interval. For the sake of simplicity, let us assume that the time spent on an insertion is equal to the number of events already present in the interval into which we need to insert. In any case, ignoring a constant access time, this will be an upper bound on the actual insertion time. The expected number of events in bucket $A_j=[(j-1)\Delta,j\,\Delta)$ under Reeves model at time t is given by

$$\int_{A_j-t} \lambda\,(F(t+u)-F(u))\,du$$

where A_j-t means the obvious thing. Let J be the collection of all indices for which A_j overlaps with $[t,\infty)$, and let B_j be $A_j \cup [t,\infty)$. Then the expected time is

$$\sum_{j\in J} \int_{B_j-t} \lambda\,(F(t+u)-F(u))\,du \;\; F(B_j-t)\,.$$

In Theorem 5.4, we derive useful upper bounds for the expected time.

Theorem 5.4.

Consider the bucket based linear list structure of Vaucher and Duval with bucket width Δ. Then the expected time for inserting (scheduling) an event at time t in the FES under Reeves's model is bounded from above by

A. $\lambda\mu$.

B. $\lambda\Delta$.

C. $\lambda C \mu\Delta$, where C is an upper bound for the density f for F (this point is only applicable when a density exists).

In particular, for any t and F, taking $\Delta \le \frac{c}{\lambda}$ for some constant c guarantees that the expected time spent on insertions is bounded by c.

Proof of Theorem 5.4.

Bound A is obtained by noting that each $F(B_j - t)$ in the sum is at most equal to 1, and that $F(t+u) \le 1$. Bound B is obtained by bounding

$$\int_{B_j - t} \lambda\, (F(t+u) - F(u))\, du$$

by $\lambda\Delta$, and observing that the terms $F(B_j - t)$ summed over $j \in J$ yield the value 1. Finally inequality C uses the fact that $F(B_j - t) \le C\,\Delta$ for all j. ∎

Theorem 5.4 is extremely important. We see that it is possible to have constant expected time deletions and insertions, uniformly over all F, t and λ, provided that Δ is taken small enough. The bound on Δ depends upon λ. If λ is known, there is no problem. Unfortunately, λ has to be estimated most of the time. Recall also that we are in Reeves's idealized model. The present analysis does not extend beyond this model. As a rule of thumb, one can take Δ equal to $1/\lambda$ where λ is the expected number of points inserted per unit of time. This should insure that every bucket has at most one point on the average. Taking Δ too small is harmful from a space point of view because the number of intervals into which the FES is cut up is

$$\left\lceil (\max(Y_i) - t)/\Delta \right\rceil$$

where the Y_i's are the scheduled event times at time t. Taking the expected value, we see that this is bounded from above by

$$1 + \frac{E(\max(Y_1, \ldots, Y_N))}{\Delta} ,$$

where N is Poisson $(\lambda\mu)$. Recall that for an upper bound the Y_i's can be considered as iid random variables with density $(1-F)/\mu$ on $[0,\infty)$. This allows us to get a good idea of the expected number of buckets needed as a function of the expected FES size, or λ. We offer two quantitative results.

Theorem 5.5.

The expected number of buckets needed in Reeves's model does not exceed

$$1+\frac{\sqrt{\frac{\lambda}{3}E(X^3)}}{\Delta},$$

where X has distribution function F. If $\Delta \sim \frac{c}{\lambda}$ as $\lambda \to \infty$ for some constant c, then this upper bound $\sim$

$$\frac{1}{c\sqrt{3}}\sqrt{E((\lambda X)^3)}.$$

Furthermore, if $E(e^{uX})<\infty$ for some $u>0$, and Δ is as shown above, then the expected number of buckets is $O(\lambda \log(\lambda))$.

Proof of Theorem 5.5.

For the first part of the Theorem, we can assume without loss of generality that X has finite third moment. We argue as follows:

$$E(\max(Y_1, \ldots, Y_N)) \le E(\sqrt{\sum_{i \le N} Y_i^2})$$

$$\le \sqrt{E(N)E(Y_1^2)} \quad \text{(Jensen's inequality)}$$

$$= \sqrt{\lambda\mu E(X^3)/(3\mu)} = \sqrt{\lambda E(X^3)/3}.$$

The last step follows from the simple observation that

$$\int_0^\infty x^2 \frac{1-F(x)}{\mu}\, dx = \int_0^\infty x^2 \int_x^\infty \frac{1}{\mu} dF(t)\, dx$$

$$= \int_0^\infty \frac{1}{\mu} \int_0^t x^2 dx \;\; dF(t)$$

$$= \frac{1}{3\mu} E(X^3).$$

The second statement of the Theorem follows in three lines. Let u be a fixed constant for which $E(e^{uX})=a<\infty$. Then, using $X_1, \ldots, X_n$ to denote an iid sample with distribution function F,

$$E(\max(Y_1, \ldots, Y_n)) \le E(\max(X_1, \ldots, X_n))$$

$$\le E(\frac{1}{t}\log(\sum_{i \le N} e^{uX_i}))$$

$$\le \frac{1}{t}\log(E(N)E(e^{uX_i})) = \frac{1}{t}\log(\lambda\mu a).$$

This concludes the proof of Theorem 5.5. ■

Except when F has compact support, the expected number of buckets needed grows superlinearly with λ, when Δ is picked as a constant over λ. The situation is worse when Δ is picked even smaller. This is a good example of the time-space trade-off, because taking Δ larger than $1/\lambda$ effectively decreases the space requirements but slows down the algorithm. However, large Δ's are uninteresting since we will see that there are nonlinear data structures which will run in expected or even worst-case time $O(\log(\lambda))$. Thus, there is no need to study cases in which the Vaucher-Duval structure performs worse than this. Vaucher and Duval (1975) and Davey and Vaucher (1980) circumvent the superlinear (in λ) storage need by collapsing many buckets in one big bucket, called an overflow bucket, or overflow list. Denardo and Fox (1979) consider a hierarchy of bucket structures where bucket width decreases with the level.

Various other multiple pointer structures have been proposed, such as the structures of Franta and Maly (1977, 1978) and Wyman (1976). They are largely similar to the Vaucher-Duval bucket structure. One nice new idea surfacing in these methods is the following. Assume that one wants to keep the cardinality of all sublists about equal and close to a number m, and assume that the FES has about n elements. Therefore, about n/m pointers are needed, which in turn can be kept in a linear list, to be scanned sequentially from left to right or right to left. The time needed for an insertion cannot exceed a constant times $\frac{n}{m}+m$ where the last term accounts for the sequential search into the selected sublist. The optimal choice for m is thus about $\sqrt{n}$, and the resulting complexity of an insertion grows also as $\sqrt{n}$. The difficulty with theses structures is the dynamic balancing of the sublist cardinalities so that all sublists have about m elements. Henriksen (1977) proposes to keep about m events per sublist, but the pointer records are now kept in a balanced binary search tree, which is dynamically adjusted. The complexity of an insertion is not immediately clear since the updating of the pointer tree requires some complicated work. Without the updating, we would need time about equal to $\log(\frac{n}{m})+m$ just to locate the point of insertion of one event. This expression is minimal for constant m ($m=4$ is the usual recommendation for Henriksen's algorithm (Kingston, 1984)). The complexity of insertion without updating is $\Theta(\log(n))$. For a more detailed expected time analysis, see Kingston (1984). In the next section, we discuss $O(\log(n))$ worst-case structures which are much simpler to implement than Henriksen's structure, and perform about equally well in practice.

5.4. Tree structures.

If the event times are kept in a binary search tree, then one would suspect that after a while the tree would be skewed to the right, because elements are deleted from the left and added mostly to the right. Interestingly, this is not always the case, and the explanation parallels that for the forward and backward scanning methods in linear lists. Consider for example an exponential F in the HOLD model. As we have seen in Theorem 5.1, all the relative event times in the FES are iid exponentially distributed. Thus, the binary search tree at every point in time is distributed as for any binary search tree constructed from a random permutation of $1, \ldots, n$. The properties of these trees are well-known. For example, the expected number of comparisons needed for an insertion of a new element, distributed as the n other elements, and independent of it, is $\sim 2\log(n)$ (see e.g. Knuth (1973) or Standish (1980)). The expected time needed to delete the smallest element is $O(\log(n))$. First, we need to locate the element at the bottom left, and then we need to restore the binary tree in case the deleted element had right descendants, by putting the bottom left descendant of these right descendants in its place. Unfortunately, one cannot count on F being exponential, and some distributions could lead to dangerous unbalancing, either to the left or the right. This was for example pointed out by Kingston (1985).

For robust performance, it is necessary to look at worst-case insertion and deletion times. They are $O(\log(n))$ for such structures as the 2-3 tree, the AVL tree and the heap. Of these, the heap is the easiest to implement and understand. The overhead with the other trees is excessive. Suggested for the FES by Floyd in a letter to Fox in the late sixties, and formalized by Gonnet (1976), the heap compares favorably in the extensive timing studies of McCormack and Sargent (1981), Ulrich (1978) and Reeves (1984). However, in isolated applications, it is clearly inferior to the bucket structures (Franta and Maly, 1978). This should come as no surprise since properly designed bucket structures have constant expected time insertions and deletions. If robustness is needed such as in a general purpose software package, the heap structure is warmly recommended (see also Ulrich (1978) and Kingston (1985)).

It is possible to streamline the heap for use in discrete event simulation. The first modification (Franta and Maly, 1978) consists of combining the DELETE and INSERT operations into one operation, the HOLD operation. Since a deletion calls for a replacement of the root of the heap, it would be a waste of effort to replace it by the last element in the heap, fix the heap, then insert a new element in the last position, and finally fix the heap again. In the HOLD operation, the root position can be filled by the new element directly. After this, the heap needs only be fixed once. This improvement is most marked when the number of HOLD operations is relatively large compared to the number of bare DELETE or INSERT operations. A second improvement, suggested by Kingston (1985), consists of using an m-ary heap instead of a binary heap. Good experimental results were obtained by him for the ternary heap. This improvement is based on the fact that insertions are more efficient for large values of m, while deletions become only slightly more time-consuming.

5.5. Exercises.

1. Prove Theorem 5.2.
2. Consider Reeves's model. Show that when $\mu<\infty$, $\limsup_{t\to\infty} N_t = \infty$ almost surely.
3. Show that the gamma (a) ($a\geq 1$) and uniform [0,1] distributions are NBUE. Show that the gamma (a) ($a\leq 1$) distribution is NWUE.
4. Generalize Theorem 5.5 as follows. For $r\geq 1$, the expected number of buckets needed in Reeves's model does not exceed
$$1+\frac{\left(\frac{\lambda}{r+1}E(X^{r+1})\right)^{\frac{1}{r}}}{\Delta},$$
where X has distribution function F. If $\Delta\sim\frac{c}{\lambda}$ as $\lambda\to\infty$ for some constant c, then this upper bound $\sim$
$$\frac{1}{c}\left(\frac{E((\lambda X)^{r+1}}{r+1})\right)^{\frac{1}{r}}.$$
5. Assume that F is the absolute normal distribution function. Prove that if Δ is $1/\lambda$ in the Vaucher-Duval bucket structure, then the expected number of buckets needed is $O(\lambda\sqrt{\log(\lambda)})$ and $\Omega(\lambda\sqrt{\log(\lambda)})$ as $\lambda\to\infty$.
6. In the HOLD model, show that whenever F has a density, the expected time needed for insertion of a new element in an ordered doubly linked list is $\Omega(n)$ and $O(n)$.
7. Consider the binary heap under the HOLD model with an exponential distribution F. Show that the expected time needed for inserting an element at time t in the FES is $O(1)$.
8. Give a heap-based data structure for implementing the operations DELETE, INSERT and CANCEL in $O(\log(n))$ worst-case time.
9. Consider the HOLD model with an ordinary binary search tree implementation. Find a distribution F for which the expected insertion time of a new element at time $t>0$ is $\Omega(\psi(n))$ for some function ψ increasing faster than a logarithm: $\lim_{n\to\infty}\psi(n)/\log(n)=\infty$.

6. REGENERATIVE PHENOMENA.

6.1. The principle.

Many processes in simulation are repetitive, i.e. one can identify a null state, or origin, to which a system evolving in time returns, and given that the system is in the null state at a certain time, the future evolution does not depend upon what has happened up to that point. Consider for example a simple random walk in which at each time unit, one step to the right or left is taken with equal probability 1/2. When the random walk returns to the origin, we start from scratch. The future of the random walk is independent of the history up to the point of return to the origin. In some simulations of such processes, we can efficiently skip ahead in time by generating the waiting time until a return occurs, at least when this waiting time is a proper random variable. Systems in which the probability of a return is less than one should be treated differently.

The gain in efficiency is due to the fact that the waiting time until the first return to the origin is sometimes easy to generate. We will work through the example of the simple random walk in the next section. Regenerative phenomena are ubiquitous: they occur in queueing systems (see section 6.3), in Markov chains, and renewal processes in general. Heyman and Sobel (1982) provide a solid study of many stochastic processes of practical importance and pay particular attention to regenerative phenomena.

6.2. Random walks.

The one-dimensional random walk is defined as follows. Let $U_1,U_2,\ldots$ be iid $\{-1,1\}$-valued random variables where $P(U_1=1)=P(U_1=-1)=\frac{1}{2}$. Form the partial sums

$$S_n = \sum_{i=1}^{n} U_i .$$

Here S_n can be considered as a gambler's gain of coin tossing after n tosses provided that the stake is one dollar; n is the time. Let T be the time until a first return to the origin. If we need to generate S_n, then it is not necessary to generate the individual U_i's. Rather, it suffices to proceed as follows:

$X \leftarrow 0$

WHILE $X \leq n$ DO

  Generate a random variate T (distributed as the waiting time for the first return to the origin).

  $X \leftarrow X+T$

$V \leftarrow X-T$, $Y \leftarrow 0$

WHILE $V < n$ DO

  Generate a random $\{1,-1\}$-valued step U.

  $Y \leftarrow Y+U$, $V \leftarrow V+1$

  IF $Y=0$ THEN $V \leftarrow X-T$ (reset V by rejecting partial random walk)

RETURN Y

The principle is clear: we generate all returns to the origin up to time n, and simulate the random walk explicitly from the last return onwards, keeping in mind that from the last return onwards, the random walk is conditional: no further returns to the origin are allowed. If another return occurs, the partial random walk is rejected. The example of the simple random walk is rather unfortunate in two respects: first, we know that S_n is binomial $(n,\frac{1}{2})$. Thus, there is no need for an algorithm such as the one described above, which cannot possibly run in uniformly bounded time. But more importantly, the method described above is intrinsically inefficient because random walks spend most of their time on one of the two sides of the origin. Thus, the last return to the origin is likely to be $\Omega(n)$ away from n, so that the probability of acceptance of the explicitly generated random walk, which is equal to the probability of not returning to the origin, is $O(\frac{1}{n})$. Even if we could generate T in zero time, we would be looking at an overall expected time complexity of $\theta(n^2)$. Nevertheless, the example has great didactical value.

The distribution of the time of the first return to the origin is given in the following Theorem.

Theorem 6.1.

In a symmetric random walk, the time T of the first return to the origin satisfies

$$P(T=2n) = p_{2n} = \frac{1}{n\, 2^{2n-1}} \binom{2n-2}{n-1} \quad (n \geq 1),$$
$$P(T=2n+1) = 0 \quad (n \geq 0).$$

If q_{2n} is the probability that the random walk returns to the origin at time $2n$, then we have

A. $p_{2n} = q_{2n}/(2n-1)$;

B. $p_{2n} \sim 1/(2\sqrt{\pi} n^{3/2})$;

C. $E(T) = \infty$;

D. $p_{2n} = q_{2n-2} - q_{2n}$;

E. $p_2 = \frac{1}{2}$, $p_{2n+2} = p_{2n}(1-\frac{1}{2n})(1+\frac{1}{n})$.

Proof of Theorem 6.1.

This proof will be given in full, because it is a beautiful illustration of how one can compute certain renewal time distributions via generating functions. We begin with the generating function $G(s)$ for the probabilities $q_{2i} = P(S_{2i}=0)$ where S_{2i} is the value of the random walk at time $2i$. We have

$$G(s) = \sum_i q_{2i} s^i = \sum_{i=1}^{\infty} 2^{-2i} \binom{2i}{i} s^i$$
$$= \sum_{i=1}^{\infty} \binom{-\frac{1}{2}}{i} (-s)^i = \frac{1}{\sqrt{1-s}} - 1 .$$

Let us now relate p_{2n} to q_{2i}. It is clear that

$$q_{2n} = p_{2n} + \sum_{i=1}^{n-1} p_{2n-2i}\, q_{2i} .$$

If $H(s)$ is the generating function for p_{2n}, then we have

$$H(s) = \sum_{n=1}^{\infty} q_{2n} s^n$$
$$= \sum_{n=1}^{\infty} \left(p_{2n} s^n + \sum_{i=1}^{n-1} p_{2n-2i}\, s^{n-i} q_{2i} s^i \right)$$
$$= H(s) + \sum_{i=1}^{\infty} \sum_{n=i+1}^{\infty} p_{2n-2i}\, s^{n-i} q_{2i} s^i$$

$$= H(s)+\sum_{i=1}^{\infty} q_{2i} s^i \sum_{n=1}^{\infty} p_{2n} s^n = H(s)+G(s)H(s) .$$

Therefore,

$$H(s) = \frac{G(s)}{1+G(s)} = 1-\sqrt{1-s} = \sum_{i=1}^{\infty} \binom{\frac{1}{2}}{i} (-1)^{i-1} s^i .$$

Equating the coefficient of s^i with p_{2i} gives the distribution of T. Statement A is easily verified. Statement B follows by using Stirling's formula. Statement C follows directly from B. Finally, D and E are obtained by simple computations. ■

Even though T has a unimodal distribution on the even integers with peak at 2, generation by sequential inversion started at 2 is not recommended because $E(T)=\infty$. We can proceed by rejection based upon the following inequalities:

Lemma 6.1.

The probabilities p_{2n} satisfy for $n \geq 1$,

$$1-\frac{1}{2n} \leq \frac{p_{2n}}{\frac{1}{2\sqrt{\pi}(n-\frac{1}{2})^{\frac{3}{2}}}} \leq e^{\frac{1}{12(2n-1)}} \leq e^{\frac{1}{12}} .$$

Proof of Lemma 6.1.

We rewrite p_{2n} as follows:

$$p_{2n} = \frac{\Gamma(2n-1)}{2n\, 2^{2n-2}\Gamma^2(n)}$$

$$= \frac{e^{-(2n-1)}(2n-1)^{2n-1}\sqrt{2\pi/2n-1}\, e^{\frac{\theta}{12(2n-1)}}}{2n\, 2^{2n-2} e^{-2n}\, n^{2n}\, \frac{2\pi}{n}}$$

$$= \frac{e\,(1-\frac{1}{2n})^{2n-1} e^{\frac{\theta}{12(2n-1)}}}{n\sqrt{2\pi(2n-1)}}$$

for some $0<\theta<1$. An upper bound is provided by

$$= \frac{e^{\frac{1}{12(2n-1)}}}{(n-\frac{1}{2})^{\frac{3}{2}}\sqrt{4\pi}} .$$

A lower bound is provided by

$$= \frac{e\,(1-\frac{1}{2n})^{2n}}{(n-\frac{1}{2})^{\frac{3}{2}}\sqrt{4\pi}}$$

$$\geq \frac{(1+\frac{1}{2n})^{2n}\,(1-\frac{1}{2n})^{2n}}{(n-\frac{1}{2})^{\frac{3}{2}}\sqrt{4\pi}}$$

$$\geq \frac{(1-\frac{1}{4n^2})^{2n}}{(n-\frac{1}{2})^{\frac{3}{2}}\sqrt{4\pi}}$$

$$\geq \frac{(1-\frac{1}{2n})}{(n-\frac{1}{2})^{\frac{3}{2}}\sqrt{4\pi}} . \blacksquare$$

Generation can now be dealt with by truncation of a continuous random variate. Note that $p_{2n} \leq cg(x)$ where

$$cg(x) = \begin{cases} \frac{1}{2} & (n=1\ ,\ n-1<x<n) \\ \frac{e^{\frac{1}{12}}}{\sqrt{4\pi}(x-\frac{1}{2})^{\frac{3}{2}}} & (n>1\ ,\ n-1<x<n) \end{cases}$$

where

$$c = \frac{1}{2}+\frac{2e^{\frac{1}{12}}}{\sqrt{\pi}} .$$

Random variates with density g can quite easily be generated by inversion. The algorithm can be summarized as follows:

Generator for first return to origin in simple random walk

[SET-UP]

$c \leftarrow \frac{1}{2}+\frac{2e^{\frac{1}{12}}}{\sqrt{\pi}}$

[GENERATOR]

REPEAT

    Generate a uniform $[0,c]$ random variate U.

    IF $U \leq \frac{1}{2}$

        THEN RETURN $X \leftarrow 2$

        ELSE

            Generate a uniform $[0,1]$ random variate V.

            $Y \leftarrow \frac{1}{2}+\frac{1}{2-(U-\frac{1}{2})\sqrt{\pi}e^{-\frac{1}{12}}}$ (Y has density g restricted to $[1,\infty)$).

            $T \leftarrow Ve^{\frac{1}{12}}/(\sqrt{\pi}(Y-\frac{1}{2})^{3/2})$ (prepare for rejection)

            $X \leftarrow \lceil Y \rceil$

            $W \leftarrow 1/(\sqrt{\pi}(X-\frac{1}{2})^{3/2})$ (prepare for squeeze steps)

            IF $T/W \leq 1-\frac{1}{2X}$ (quick acceptance)

                THEN RETURN $2X$

                ELSE IF $T/W \leq e^{\frac{1}{12(2X-1)}}$ (quick rejection)

                    THEN IF $T \leq p_{2X}$ THEN RETURN $2X$

UNTIL False

The rejection constant c is a good indicator of the expected time spent before halting provided that p_{2X} can be evaluated in constant time uniformly over all X. However, if p_{2X} is computed directly from its definition, i.e. as a ratio of factorials, then the computation takes time roughly proportional to X. Assume that it is exactly X. Without squeeze steps, the expected time spent computing p_{2X} would be c times $E(X)$ where X has density g. This is ∞ (exercise 6.1). However, with the squeeze steps, the probability of evaluating p_{2X} explicitly for fixed value of X decreases as $\frac{1}{X}$ as $X \rightarrow \infty$. This implies that the overall expected time of the algorithm is finite (exercise 6.2).

6.3. Birth and death processes.

A birth and death process is a process with states 0,1,2,3,..., in which the time spent in state i is distributed as an exponential random variate divided by $\lambda_i+\mu_i$, at which time the system jumps to state $i+1$ (a birth) with probability $\lambda_i/(\lambda_i+\mu_i)$, and to state $i-1$ (a death) otherwise. Simple examples include

A. The Poisson process: $\lambda_i \equiv \lambda > 0$, $\mu_i \equiv 0$. Births correspond essentially to events such as arrivals in a bank.

B. The Yule process: $\lambda_i \equiv \lambda i > 0$, $\mu_i \equiv 0$. Here we also require that at time 0, the state be 1. This is a particular case of a pure birth process. The state can be identified with the size of a given population in which no deaths can occur.

C. The M/M/1 queue: $\lambda_i \equiv \lambda > 0$, $\mu_i \equiv \mu > 0$, $\mu_0 = 0$. Here the state can be identified with the size of a queue, a birth with an arrival, and a death with a departure. The condition $\mu_0=0$ is natural since nobody can leave the queue when the queue is empty.

In all these examples, simulation can often be accelerated by making use of first-passage-time random variables. Formally, we define the first passage time from i to j $(j>i)$, T_{ij}, by

$$T_{ij} = \inf \{ t : X_t = j \mid X_0 = i \} .$$

Here X_t is the state of the system (an integer) at time t, and X_0 is the initial state. Let us consider the M/M/1 queue. The busy period of such a queue is T_{10}. If the system starts in state 0 (empty queue), and we define a system cycle as the minimal time until for the first time another empty queue state is reached after some busy period, i.e. after at least one person has been in the queue, then the system cycle is distributed as $E/\lambda + T_{10}$, where E is an exponential random variate, independent of T_{10}. The only M/M/1 queues of interest to us are those which have with probability one a finite value for T_{10}, i.e. those for which $\mu \geq \lambda$ (Heyman and Sobel, 1982, p. 91). The actual derivation of the distribution of T_{10} would lead us astray. What matters is that we can generate random variates distributed as T_{10} quite easily. This should of course not be done by generating all the arrivals and departures until an empty queue is reached, because the expected time of this method is not uniformly bounded over all values of $\lambda < \mu$. This is best seen by noting that $E(T_{10}) = 1/(\mu-\lambda)$.

The M/M/1 queue provides one of the few instances in which the distribution of the first passage times is analytically manageable. For example, $2\sqrt{\lambda\mu}\,T_{10}$ has density

$$f(x) = e^{-\frac{x}{2}(\xi+\frac{1}{\xi})} I_1(x) \frac{\xi}{x} \quad (x>0) ,$$

where $\xi = \sqrt{\frac{\mu}{\lambda}}$ and I_1 is the Bessel function of the first kind with imaginary argument (see section IX.7.1 for a definition). Direct generation can be carried out based upon the following result.

Theorem 6.2.

When E is exponentially distributed, Y is a random variable with density

$$g(y) = c\frac{\sqrt{y(1-y)}}{\frac{1}{2}(\xi+\frac{1}{\xi})-1+2y} \qquad (0<y<1),$$

where $c=\frac{4\xi}{\pi}$ and $\xi=\sqrt{\frac{\mu}{\lambda}}$, and E,Y are independent, then $E/(\frac{1}{2}(\xi+\frac{1}{\xi})+2Y-1)$ has density f, and $E/(\mu+\lambda+2\sqrt{\mu\lambda}(2Y-1))$ is distributed as T_{10}.

Proof of Theorem 6.2.

This theorem illustrates once again the power of integral representations for densities. By an integral representation of I_1 (Magnus et al, 1966, p. 84),

$$\begin{aligned} f(x) &= e^{-\frac{x}{2}(\xi+\frac{1}{\xi})} I_1(x)\frac{\xi}{x} \\ &= e^{-\frac{x}{2}(\xi+\frac{1}{\xi})} \frac{\xi}{x}\frac{x}{\pi}\int_{-1}^{1} e^{-zx}\sqrt{1-z^2}\,dz \\ &= \int_0^1 (\frac{1}{2}(\xi+\frac{1}{\xi})+2y-1)e^{-x(\frac{1}{2}(\xi+\frac{1}{\xi})+2y-1)}\frac{4\xi}{\pi}\frac{\sqrt{y(1-y)}}{\frac{1}{2}(\xi+\frac{1}{\xi})+2y-1}\,dy \\ &= E((\frac{1}{2}(\xi+\frac{1}{\xi})+2Y-1)e^{-x(\frac{1}{2}(\xi+\frac{1}{\xi})+2Y-1)}) \end{aligned}$$

where Y has density g. Given Y, this is the density of $E/(\frac{1}{2}(\xi+\frac{1}{\xi})+2Y-1)$. ■

Generation of Y can be taken care of very simply by rejection. Note that

$$g(y) \le \begin{cases} c\dfrac{\sqrt{y(1-y)}}{2y} \\ c\dfrac{\sqrt{y(1-y)}}{\frac{1}{2}(\xi+\frac{1}{\xi})-1} \end{cases},$$

where $c=\frac{4\xi}{\pi}$. The top upper bound, proportional to a beta $(\frac{1}{2},\frac{3}{2})$ density integrates to ξ. The bottom upper bound, proportional to a beta $(\frac{3}{2},\frac{3}{2})$ density, integrates to $(\xi/(\xi-1))^2$. One should always pick the bound which has the smallest integral. The cross-over point is at $\xi=\frac{1}{2}(3+\sqrt{5})\approx 2.6$.

Generator for g

CASE

$\xi\leq\frac{3+\sqrt{5}}{2}$:

REPEAT

Generate a uniform [0,1] random variate U.

Generate a beta $(\frac{1}{2},\frac{3}{2})$ random variate Y.

UNTIL $\frac{U}{1-U}\leq\frac{2Y}{\frac{1}{2}(\xi+\frac{1}{\xi})-1}$

$\xi>\frac{3+\sqrt{5}}{2}$:

REPEAT

Generate a uniform [0,1] random variate U.

Generate a beta $(\frac{3}{2},\frac{3}{2})$ random variate Y.

UNTIL $\frac{U}{1-U}\leq\frac{\frac{1}{2}(\xi+\frac{1}{\xi})-1}{2Y}$

RETURN Y

The expected number of iterations is $\min(\xi,(\frac{\xi}{\xi-1})^2)$. This is a unimodal function in ξ, taking the value 1 as $\xi\downarrow 1$ and $\xi\uparrow\infty$. The peak is at $\xi=(3+\sqrt{5})/2$. The algorithm is uniformly fast with respect to $\xi\geq 1$. In the case $\xi=1$ the acceptance condition is automatically satisfied, and the combination of the g generator with the property of Theorem 6.2 is reduced to a generator already dealt with in Theorem IX.7.1.

6.4. Phase type distributions.

Phase type distributions (or simply **PH-distributions**) are the distributions of absorption times in absorbing Markov chains, which are useful in studying queues and reliability problems. We consider only discrete (or: discrete-time) Markov chains with a finite number of states. An absorption state is one which, when reached, does not allow escape. Even if there is at least one absorption state, it is not at all certain that it will ever be reached. Thus, phase type distributions can be degenerate.

Any state can also be "promoted" to absorption state to study the time needed until this state is reached. If we promote the starting state to absorption state immediately after we leave it, then this promotion mechanism can be used to simulate Markov chains by the shortcuts discussed in this section, at least if we can get a good handle on the times until absorption.

Discrete Markov chains can always be simulated by using a simple discrete random variate generator for every state transition (Neuts and Pagano, 1981). This generator is not uniformly fast over all Markov chains with m states and nondegenerate phase type distribution. In the search for uniformly fast generators, simple shortcuts are of little help.

For example, when we are in state i, we could generate the (geometrically distributed) time until we first leave i in constant expected time. The corresponding state can also be generated uniformly fast by a method such as Walker's, because we have a simple conditional discrete distribution with $m-1$ outcomes. This method can be used to eliminate the times spent idling in individual states. It cannot eliminate the times spent in cycles, such as in a Markov chain in which with high probability we stay in a cycle visiting states $i_1, i_2, \ldots, i_k$ in turn. Thus, it cannot possibly be uniformly fast over all Markov chains with m states.

It seems that in this problem, uniform speed does not come cheaply. Some preprocessing involving the transition matrix seems necessary.

6.5. Exercises.

1. Consider the rejection algorithm for the time $2X$ until the first return to the origin in a symmetric random walk given in the text. Show that when the time needed to compute p_{2X} is equal to X, then the expected time taken by the algorithm without squeeze steps is ∞.
2. A continuation of exercise 1. Show that when squeeze steps are added as in the text, then the algorithm halts in finite expected time.
3. **Discrete Markov chains.** Consider a discrete Markov chain with m states and initial state 1. You are allowed to preprocess at any cost, but just once. What sort of preprocessing would you do on the transition matrix so that you can design an algorithm for generating the state S_n at time n in expected time uniformly bounded over n. The expected time is however allowed to increase with m. Hint: can you decompose the transition matrix using a spectral representation so that the n-th power of it can be computed uniformly quickly over all n?
4. **The lost-games distribution.** Let X be the number of games lost before a player is ruined in the classical gambler's ruin problem, i.e. a gambler adds one to his fortune with probability p and loses one unit with probability $1-p$. He starts out with r units (dollars). The purpose of this exercise is to design an algorithm for generating X in expected time uniformly bounded in

r when $p < 1-p$ is fixed. Uniform speed in both r and p would be even better. Notice first that the restriction $p < 1-p$ is needed to insure that X is a proper random variable, i.e. to insure that the player is ruined with probability one.

A. Show that when $p < 1-p$, the player will eventually be ruined with probability one.

B. Show that X has discrete distribution given by

$$P(X=n) = \binom{2n-r}{n} p^{n-r}(1-p)^n \frac{r}{2n-r} \quad (n=r, r+1, \ldots)$$

(Kemp and Kemp, 1968).

C. Suppose that customers arrive at a queue according to a homogeneous Poisson process with parameter λ, that the service time is exponential with parameter $\mu < \lambda$, and that the queue has initially r customers. Show that the number of customers served until the queue first vanishes has the lost-games distribution with parameters r and $p = \lambda/(\lambda+\mu)$.

D. Using Stirling's approximation, determine the general dependence of $P(X=n)$ upon n, and use it to design a uniformly fast rejection algorithm.

For a survey of these and other waiting time mechanisms, see e.g. Patil and Boswell (1975).

7. THE GENERALIZATION OF A SAMPLE.

7.1. Problem statement.

As in section XIV.2, we will discuss an incompletely specified random variate generation problem. Assume that we are given a sample $X_1, \ldots, X_n$ of iid R^d-valued random vectors with common unknown density f, and that we are asked to generate a new independent sample $Y_1, \ldots, Y_m$ of independent random vectors with the same density f. Stated in this manner, the problem is obviously unsolvable, unless we are incredibly lucky.

What one can do is construct a **density estimate** $f_n(x) = f_n(x, X_1, \ldots, X_n)$ of $f(x)$, and then generate a sample of size m from f_n. This procedure has several drawbacks: first of all, f_n is typically not equal to f. Also, the new sample depends upon the original sample. Yet, we have very few options available to us. Ideally, we would like the new sample to appear to be distributed as the original sample. This will be called sample indistinguishability. This and other issues will be discussed in this section. The material appeared originally in Devroye and Gyorfi (1985, chapter 8).

7.2. Sample independence.

There is little that can be done about the dependence between $X_1, \ldots, X_n$ and $Y_1, \ldots, Y_m$ except to hope that for n large enough, some sort of asymptotic independence is obtained. In some applications, sample independence is not an issue at all.

Since the Y_i's are conditionally independent given $X_1, \ldots, X_n$, we need only consider the dependence between Y_1 and $X_1, \ldots, X_n$. A measure of the dependence is

$$D_n = \sup_{A,B} | P(Y \in A, X \in B) - P(Y \in A) P(X \in B) | ,$$

where the supremum is with respect to all Borel sets A of R^d and all Borel sets B of R^{nd}, and where $Y = Y_1$ and X is our shorthand notation for $(X_1, \ldots, X_n)$. We say that the samples are asymptotically independent when

$$\lim_{n \to \infty} D_n = 0 .$$

In situations in which $X_1, \ldots, X_n$ is used to design or build a system, and $Y_1, \ldots, Y_m$ is used to test it, the sample dependence often causes optimistic evaluations. Without the asymptotic independence, we can't even hope to diminish this optimistic bias by increasing n.

The inequality in Theorem 7.1 below provides us with a sufficient condition for asymptotic independence. First, we need the following Lemma.

Lemma 7.1. (Scheffe, 1947).
For all densities f and g on R^d,

$$\int | f - g | = 2 \sup_B | \int_B f - \int_B g | ,$$

where the supremum is with respect to all Borel sets B of R^d.

Proof of Lemma 7.1.

Let us take $B = \{ f > g \}$, and let A be another Borel set of R^d. Because $\int (f - g) = 0$, we see that

$$\int | f - g | = 2 \int_B (f - g) .$$

Thus, we have shown that $\int | f - g |$ is at most twice the supremum over all Borel sets of $| \int_B (f - g) |$. To show the other half of the Lemma, note that if B' denotes the complement of B, then

$$| \int_A (f - g) | = | \int_{A \cap B} (f - g) + \int_{A \cap B'} (f - g) |$$
$$\le \max (\int_{A \cap B} (f - g), \int_{A \cap B'} (f - g))$$

$$\leq \max(\int_B (f-g), \int_{B'} (g-f))$$

$$= \frac{1}{2}\int | f-g | \quad (\text{all } A) . \blacksquare$$

Scheffe's lemma tells us that if we assign probabilities to sets (events) using two different densities, then the maximal difference between the probabilities over all sets is equal to one half of the L_1 distance between the densities. From Lemma 7.1, we obtain

Theorem 7.1.

Let f_n be a density estimate, which itself is density. Then

$$D_n \leq E(\int | f_n - f |) .$$

Proof of Theorem 7.1.

Let $X_1, \ldots, X_{n+1}$ be iid. Then

$$D_n \leq \sup_{A,B} | P(Y\in A, X\in B) - P(X_{n+1}\in A, X\in B) |$$

$$+ \sup_{A,B} | P(X_{n+1}\in A, X\in B) - P(X_{n+1}\in A) P(X\in B) |$$

$$+ \sup_{A,B} | P(X_{n+1}\in A) P(X\in B) - P(Y\in A) P(X\in B) |$$

$$\leq \sup_{A,B} E(I_{X\in B} | \int_A f_n - \int_A f |) + 0 + \sup_A | P(X_{n+1}\in A) - P(Y\in A) |$$

$$\leq \sup_A E(| \int_A f_n - \int_A f |) + \sup_A | \int_A E(f_n) - \int_A f |$$

$$\leq E(\sup_A | \int_A f_n - \int_A f |) + \frac{1}{2}\int | E(f_n) - f |$$

$$= E(\frac{1}{2}\int | f_n - f |) + \frac{1}{2}\int | E(f_n) - f | . \blacksquare$$

We see that for the sake of asymptotic sample independence, it suffices that the expected L_1 distance between f_n and f tends to zero with n. This is also called **consistency**. Consistency does not imply asymptotic independence: just let f_n be the uniform density in all cases, and observe that $D_n \equiv 0$, yet

$\int | f_n - f |$ is a positive constant for all n and all nonuniform f.

7.3. Consistency of density estimates.

A density estimate f_n is **consistent** if for all densities f,

$$\lim_{n \to \infty} E\left(\int | f_n - f |\right) = 0 .$$

Consistency guarantees that the expected value of the maximal error committed by replacing probabilities defined with f with probabilities defined with f_n tends to 0. Many estimates are consistent, see e.g. Devroye and Gyorfi (1985). Parametric estimates, i.e. estimates in which the form of f_n is fixed up to a finite number of parameters, which are estimated from the sample, cannot be consistent because f_n is required to converge to f for all f, not a small subclass. Perhaps the best known and most widely used consistent density estimate is the **kernel estimate**

$$f_n(x) = \frac{1}{nh^d} \sum_{i=1}^{n} K\left(\frac{x - X_i}{h}\right),$$

where K is a given density (or kernel), chosen by the user, and $h > 0$ is a smoothing parameter, which typically depends upon n or the data (Rosenblatt, 1956; Parzen, 1962). For consistency it is necessary and sufficient that $h \to 0$ and $nh^d \to \infty$ in probability as $n \to \infty$ (Devroye and Gyorfi, 1985). How one should choose h as a function of n or the data is the subject of a lot of controversy. Usually, the choice is made based upon the approximate minimization of an error criterion. Sample independence (Theorem 7.1) and sample indistinguishability (next section) suggest that we try to minimize

$$E\left(\int | f_n - f |\right).$$

But even after narrowing down the error criterion, there are several strategies. One could minimize the supremum of the criterion where the supremum is taken over a class of densities. This is called a **minimax strategy**. If f has compact support on the real line and a bounded continuous second derivative, then the best choices for individual f (i.e., not in the minimax sense) are

$$h = Cn^{-\frac{1}{5}},$$

$$K(x) = \frac{3}{4}(1 - x^2) \quad (|x| \le 1),$$

where C is a constant depending upon f only:

$$C = \left(\sqrt{\frac{15}{2\pi}} \frac{\int \sqrt{f}}{\int | f'' |} \right)^{\frac{2}{5}}.$$

The optimal kernel coincides with the optimal kernel for L_2 criteria (Bartlett, 1963). The optimal formula for h, which depends upon the unknown density f, can be estimated from the data. Alternatively, one could compute the formula for a given parametric density, a rough guess of sorts, and then estimate the parameters from the data. For example, if this is done with the normal density as initial guess, we obtain the recommendation to take

$$h = \left(\frac{15e\sqrt{2\pi}}{8n} \right)^{\frac{1}{5}} \hat{\sigma} ,$$

where $\hat{\sigma}$ is a robust estimate of the standard deviation of the normal density (Devroye and Gyorfi, 1985). A typical robust estimate is the so-called quick-and-dirty estimate

$$\hat{\sigma} = \frac{X_{(np)} - X_{(nq)}}{x_p - x_q} ,$$

where x_p, x_q are the p-th and q-th quantiles of the standard normal density, and $X_{(np)}$ and $X_{(nq)}$ are the p-th and q-th quantiles in the data, i.e. the (np)-th and (nq)-th order statistics.

The construction given here with the kernel estimate is simple, and yields fast generators. Other constructions have been suggested in the literature with random variate generation in mind. Often, the explicit form of f_n is not given or needed. Constructions often start from an empirical distribution function based upon $X_1, \ldots, X_n$, and a smooth approximation of this distribution function (obtained by interpolation), which is directly useful in the inversion method. Guerra, Tapia and Thompson (1978) use Akima's (Akima, 1970) quasi-Hermite piecewise cubic interpolation to obtain a smooth monotone function coinciding with the empirical distribution function at the points X_i. Recall that the empirical distribution is the distribution which puts mass $\frac{1}{n}$ at point X_i. Hora (1983) gives another method for the same problem. Butler (1970) on the other hand uses Lagrange's quadratic interpolation on the inverse empirical distribution function to speed random variate generation up even further.

7.4. Sample indistinguishability.

In simulations, one important qualitative measure of the goodness of a method is the indistinguishability of $X_1, \ldots, X_m$ and $Y_1, \ldots, Y_m$ for the given sample size m. Note that we have forced both sample sizes to be the same, although for the construction of f_n we keep on using n points. The indistinguishability could be measured quantitatively by

$$S_{n,m} = \sup_A \left| E(N(A)) - E(M(A) \mid X_1, \ldots, X_n) \right|$$

$$= m \sup_A \left| \int_A f - \int_A f_n \right|$$

$$=\frac{m}{2}\int | f_n - f | .$$

Here, A is a Borel set of R^d, $N(A)$ is the cardinality of A for the original sample (the data, artificially inflated to size m), and $M(A)$ is the cardinality of A for the artificial Y_i sample. By cardinality of a set, we mean the number of data points falling in the set.

When $S_{n,m}$ is smaller than one, then the expected cardinality of a set A with a perfect sample of size m differs by at most one from the conditional expected cardinality of the generated sample of size m. We say that f_n is k-excellent for samples of size m when

$$E(S_{n,m}) \leq k .$$

This is equivalent to asking that the expected L_1 distance between f and f_n is at most $2k/m$. The notion of 1-excellence is very strong. For example, for most nonparametric estimates such as the kernel estimate 1-excellence forces us to use phenomenally large values of n for even moderate values of m. Devroye and Gyorfi (1985) have shown that for all kernel estimates (regardless of choice of K and h), and for all densities f, 1-excellence is not achievable for samples of size $m=1000$ unless $n \geq 4{,}000{,}000$. For $m=10{,}000$, we need $n \geq 1{,}300{,}000{,}000$. For the histogram estimate, the situation is even worse.

But even 1-excellence may not be good enough for one's application. For one thing, no assurances are given as to the discrepancy in moments between the generated sample and the original sample.

7.5. Moment matching.

Some statisticians attach a great deal of importance to the moments of the densities f_n and f. For $d=1$, the i-th **moment mismatch** is defined as

$$M_{n,i} = \int x^i f_n - \int x^i f \qquad (i=1,2,3,\ldots) .$$

Clearly, $M_{n,i}$ is a random variable. Assume that we employ the kernel estimate with a zero mean finite variance (σ^2) kernel K. Then, we have

$$M_{n,1} = \frac{1}{n}\sum_{i=1}^{n}(X_i - E(X_i)) ,$$

$$M_{n,2} = \frac{1}{n}\sum_{i=1}^{n}(X_i^2 - E(X_i^2)) + h^2\sigma^2 .$$

This follows from the fact that f_n is an equiprobable mixture of densities K shifted to the X_i's, each having variance $h^2\sigma^2$ and zero mean. It is interesting to note that the distribution of $M_{n,1}$ is not influenced by h or K. By the weak law of large numbers, $M_{n,1}$ tends to 0 in probability as $n \to \infty$ when f has a finite first moment. The story is different for the second moment mismatch.

Whereas $E(M_{n,1})=0$, we now have $E(M_{n,2})=h^2\sigma^2$, a positive bias. Fortunately, h is usually small enough so that this is not too big a bias. Note further that the variances of $M_{n,1}$, $M_{n,2}$ are equal to

$$\frac{Var(X_1)}{n}, \frac{Var(X_1^2)}{n}$$

respectively. Thus, h and K only affect the bias of the second order mismatch. Making the bias very small is not recommended as it increases the expected L_1 error, and thus the sample dependence and distinguishability.

7.6. Generators for f_n.

For the kernel estimate, generators can be based upon the property that a random variate is distributed as an equiprobable mixture, as is seen from the following trivial algorithm.

Mixture method for kernel estimate

Generate Z, a random integer uniformly distributed on $\{1,2, \ldots, n\}$.
Generate a random variate W with density K.
RETURN X_Z+hW

For Bartlett's kernel $K(x)=\frac{3}{4}(1-x^2)_+$, we suggest either rejection or a method based upon properties of order statistics:

Generator based upon rejection for Bartlett's kernel

REPEAT
    Generate a uniform $[-1,1]$ random variate X and an independent uniform $[0,1]$ random variate U.
UNTIL $U \leq 1-X^2$
RETURN X

The order statistics method for Bartlett's kernel

Generate three iid uniform $[-1,1]$ random variates V_1, V_2, V_3.
IF $|V_3| > \max(|V_1|, |V_2|)$
THEN RETURN $X \leftarrow V_2$
ELSE RETURN $X \leftarrow V_3$

In the rejection method, X is accepted with probability 2/3, so that the algorithm requires on average three independent uniform random variates. However, we also need some multiplications. The order statistics method always uses precisely three independent uniform random variables, but the multiplications are replaced by a few absolute value operations.

7.7. Exercises.

1. **Monte Carlo integration.** To estimate $\int H(x) f(x)\, dx$, where H is a given function, and f is a density, the Monte Carlo method uses a sample of size n drawn from f (say, $X_1, \ldots, X_n$). The naive estimate is

$$\frac{1}{n} \sum_{i=1}^{n} H(X_i) .$$

When n is small, this estimate has a lot of built-in variance. Compute the variance and assume that it is finite. Then construct the **bootstrap estimate**

$$\frac{1}{m} \sum_{i=1}^{m} H(Y_i) ,$$

where the Y_i's are iid random variables with density f_n, the kernel estimate of f based upon $X_1, \ldots, X_n$. The sample size m can be taken as large as the user can afford. Thus, in the limit, one can expect the bootstrap estimate to provide a good estimate of $\int H(x) f_n(x)\, dx$.

A. Show that $|\int Hf - \int Hf_n| \leq 2 (\sup H) \int |f - f_n|$ (Devroye and Gyorfi, 1985).

B. Compare the mean square errors of the naive Monte Carlo estimate and the estimate $\int Hf_n$ (the latter is a limit as $m \to \infty$ of the bootstrap estimate).

C. Compute the mean square error of the bootstrap estimate as a function of n and m, and compare with the naive Monte Carlo estimate. Also consider what happens when you let $m \to \infty$ in the expression for the mean square error.

2. The generators for the kernel estimate based upon Bartlett's kernel in the text use the mixture method. Still for Bartlett's kernel, derive the inversion method with all the details. Hint: note that the distribution function can be written as the sum of polynomials of degree three with compact support, and can therefore be considered as a cubic spline with at most $2n$ breakpoints when there are n data points (Devroye and Gyorfi, 1985).

3. Bratley, Fox and Schrage (1983) consider a density estimate f_n which provides fast generation by inversion. The X_i's are ordered, and f_n is constant on the intervals determined by the order statistics. In addition, in the intervals to the left of the minimum and to the right of the maximum exponential tails are added. The constant pieces and exponentall tails integrate to $1/(n+1)$ over their supports, i.e. all pieces are equally likely to be picked. Rederive their fast inversion algorithm for f_n. Is their estimate asymptotically independent? Show that it is not consistent for any density f. To cure the latter problem, Bratley, Fox and Schrage suggest coalescing breakpoints. Consider coalescing breakpoints by letting f_n be constant on the intervals determined by the k-th, $2k$-th, $3k$-th, $\cdots$ order statistics. How should one define the heights of f_n on these intervals, and how should k vary with n for consistency?

4. For the kernel estimate, show that for any density K, any f, and any sequence of numbers $h > 0$ with $h \to 0$, $nh^d \to \infty$, we have $E(\int |f - f_n|) \to 0$ as $n \to \infty$. Proceed as follows: first prove the statement for continuous f with compact support. Then, using the fact that any measurable function in L_1 can be approximated arbitrarily closely by continuous functions with compact support, wrap up the proof. In the first half of the proof, it is useful to split the integral by considering $|f - E(f_n)|$ separately. In the second half of the proof, you will need an embedding argument, in which you create a sample which with a few deletions can be considered as a sample drawn from f, and with a few different deletions can be considered as a sample drawn from the L_1 approximation of f.