Probabilistic Analysis of Sorting Algorithms
Lecture Notes

Paweł Hitczenko*
Department of Mathematics
Drexel University
Philadelphia, PA 19104, U.S.A.
email: phitczenko@mcs.drexel.edu

*supported in part by NSA grant MSPF-02G-043

1 Introduction

The following are lecture notes to a course "Probabilistic Analysis of Sorting Algorithms" that I gave as a special topics course at Drexel University in the summer of 2002, and at the University of Jyväskylä, Finland, in the fall of the same year. The original intent of the course was to cover large parts of an excellent monograph on the topic, namely, a text by Hosam M. Mahmoud *Sorting: A Distribution Theory*, Wiley 2000. It turned out, however, that his text assumes a solid knowledge of a graduate level probability theory, a rather demanding assumption for most of Drexel's (and, as it turned out, the University of Jyväskylä's) students potentially interested in such a course. As a consequence, almost half of a 10 week long course at Drexel was spent on developing necessary tools from the probability theory (the ultimate goal was the central limit theorem). The notes reflect that: the first part consists of a brief introduction to probability theory. The development was based largely on an excellent monograph by P. Billingsley *Probability and Measure*, Wiley (1995), 3rd Ed. (A.N. Shiryaev's *Probability*, Springer (1995), 2nd Ed. is a good alternative). It should be emphasized, however, that the notes are far from a comprehensive treatment of even basic probability; the goal was to get to the central limit theorem as quickly as possible. Either of the texts just mentioned or, on a non–measure theoretic level, books like S. Ross, *A First Course in Probability*, Macmillan, or M. H. DeGroot, *Probability and Statistics*, Addison–Wesley, should be consulted for a much more complete development. The second part of the notes is closely based on selections from Mahmoud's *Sorting*. Again, we would like to emphasize, that because of a very limited time left for the algorithmic part, the selections are very constrained and even for the topics included here, discussions rather brief. Thus, the original text of Mahmoud's book (or D.E. Knuth's 3rd volume of *The Art of Computer Programming. Sorting and Searching*, Addison–Wesley (1973)) should be consulted for a much more detailed analysis. Since I will refrain from providing historical references and credits in the text, I refer to Mahmoud's texts for references (up to, roughly, 2000). As one exception I would like to mention a paper H.–K. Hwang and R. Neininger "Phase change of limit laws in the quicksort recurrence under varying toll functions" *SIAM J. Computing* (2003), 1475–1501, which contains a quite complete, and perhaps, essentially final discussion of quicksort type of recurrences.

Someone who wishes to learn more about general methods used in discrete mathematics (they are used throughout the notes in a very limited way), will find plenty of methods in *Concrete Mathematics: A Foundation for Computer Science*, Addison–Wesley (1989), by R.L. Graham, D.E. Knuth, and O. Patashnik.

As I mentioned above, one occasion on which this course was taught, was during fall 2002 semester, which I spent as a visiting professor at the University of Jyväskylä, Finland. I would like to thank the Department of the Mathematics and Statistics of the University of Jyväskylä for the invitation and to Christel and Stefan Geiss for their hospitality during that visit.

Finally, I would like to thank Rainer Avikainen for his help with final preparation of the notes in June of 2003, for a careful reading of most of these notes, and for preparing hints to the solutions of the exercises.

2 Complex Numbers

Real numbers, as good as they are, do not suffice for all purposes. For example, some very simple polynomials do not necessarily have real roots. To bypass that problem we need to consider more general system of numbers, namely complex numbers. This is done by introducing the number i, called an *imaginary unit*, by letting $i = \sqrt{-1}$ (or by saying that $i^2 = -1$). We then define a complex number z by

$$z = a + bi,$$

where a and b are real numbers. These are referred to as the *real* and *imaginary* part of z, respectively, and are denoted by $a = \Re(z)$, and $b = \Im(z)$. A complex *conjugate* of z is a complex number $\overline{z}$ defined by $\overline{z} = a - bi$. The *absolute value* of z is a nonnegative number defined by $|z| = \sqrt{a^2 + b^2}$. Arithmetic operations on complex numbers are defined as follows: addition and subtraction are coordinate-wise

$$z_1 \pm z_2 = (a_1 \pm a_2) + (b_1 \pm b_2)i,$$

and multiplication by expanding the product

$$\begin{aligned} z_1 z_2 &= (a_1 + b_1 i)(a_2 + b_2 i) = a_1 a_2 + b_1 b_2 i^2 + (a_1 b_2 + a_2 b_1)i \\ &= (a_1 a_2 - b_1 b_2) + (a_1 b_2 + a_2 b_1)i. \end{aligned}$$

Division is performed by first removing imaginary part from the denominator; for any complex number $z = a + bi$ we have

$$z\overline{z} = (a + bi)(a - bi) = a^2 + b^2 = |z|^2.$$

Hence,

$$\begin{aligned} \frac{a_1 + b_1 i}{a_2 + b_2 i} &= \frac{(a_1 + b_1 i)(a_2 - b_2 i)}{(a_2 + b_2 i)(a_2 - b_2 i)} = \frac{(a_1 a_2 + b_1 b_2) + (b_1 a_2 - a_1 b_2)i}{a_2^2 + b_2^2} \\ &= \frac{a_1 a_2 + b_1 b_2}{a_2^2 + b_2^2} + i\frac{b_1 a_2 - a_1 b_2}{a_2^2 + b_2^2}. \end{aligned}$$

We will need a complex valued function of a real variable, namely, $x \to e^{ix}$. The following provides heuristics behind the definition: consider Taylor series of e^y with $y = ix$. By grouping together even and odd powers and taking into account the defining property of i we get

$$\begin{aligned} e^{ix} &= \sum_{n=0}^{\infty} \frac{(ix)^n}{n!} = \sum_{k=0}^{\infty} \frac{(ix)^{2k}}{(2k)!} + \sum_{k=0}^{\infty} \frac{(ix)^{2k+1}}{(2k+1)!} \\ &= \sum_{k=0}^{\infty} \frac{(-1)^k x^{2k}}{(2k)!} + i\sum_{k=0}^{\infty} \frac{(-1)^k x^{2k+1}}{(2k+1)!}. \end{aligned}$$

We now notice that the first sum is just the Taylor series of $\cos x$ while the second is the that of $\sin x$. In view of this we set: the function $x \to e^{ix}$ is defined by

$$e^{ix} = \cos x + i \sin x.$$

The following lemma lists the basic properties of e^{ix}:

Lemma 2.1 *We have*

(i) $|e^{ix}| = 1$

(ii) $e^{ix} \cdot e^{iy} = e^{i(x+y)}$

(iii) $(e^{ix})' = ie^{ix}$

(iv) $\int e^{ix} dx = -ie^{ix}$

Proof: We begin with (ii). Using the formulas for the cosine and sine of the sum of two angles, we see that:

$$\begin{aligned}
e^{i(x+y)} &= \cos(x+y) + i\sin(x+y) \\
&= \cos x \cos y - \sin x \sin y + i(\sin x \cos y + \sin y \cos x) \\
&= \cos x \cos y + i^2 \sin x \sin y + i(\sin x \cos y + \sin y \cos x) \\
&= (\cos x + i \sin x)(\cos y + i \sin y) \\
&= e^{ix} \cdot e^{iy}
\end{aligned}$$

as required. Part (i) follows from (ii), the fact that $|z| = \sqrt{z\overline{z}}$, and that

$$\overline{e^{ix}} = \cos x - i \sin x = \cos(-x) + i \sin(-x) = e^{-ix}.$$

For part (iii) write

$$\begin{aligned}
(e^{ix})' &= (\cos x + i \sin x)' = -\sin x + i \cos x = i^2 \sin x + i \cos x \\
&= i(\cos x + i \sin x) = ie^{ix}.
\end{aligned}$$

Finally, (iv) follows from (iii) and the fact that that $1/i = -i$. $\square$

3 Probability Space

Let Ω be a set. A family $\mathcal{A}$ of subsets of Ω is called a σ–*algebra* (or a σ–field) of sets if it satisfies the following conditions:

- $\Omega \in \mathcal{A}$,
- for every set A, if $A \in \mathcal{A}$, then its complement A^c is in $\mathcal{A}$,
- if $A_1, A_2, \ldots$ is a sequence such that $\forall i \geq 1$ $A_i \in \mathcal{A}$ then their union $\bigcup_{j=1}^{\infty} A_j \in \mathcal{A}$.

A triple $(\Omega, \mathcal{A}, \mathbf{P})$ is called a probability space if Ω is a set, $\mathcal{A}$ is a σ–field of subsets of Ω, and $\mathbf{P}$ is a function defined on $\mathcal{A}$ satisfying the following conditions (axioms of probability):

- $\forall A \in \mathcal{A}$, $\mathbf{P}(A) \geq 0$,
- $\mathbf{P}(\Omega) = 1$,
- $\mathbf{P}(\bigcup_{j=1}^{\infty} A_j) = \sum_{j=1}^{\infty} \mathbf{P}(A_j)$, whenever $A_1, A_2, \ldots$ is a sequence of *pairwise disjoint* sets in $\mathcal{A}$ (i.e. $A_j \cap A_k = \emptyset$ for all $j \neq k$).

This function $\mathbf{P}$ is called a *probability measure* (or just *probability*) on Ω. Two basic examples are:

1. Discrete uniform space: Let Ω be a finite set and let N be the number of elements in Ω. Let $\mathcal{A}$ be a σ–field of all possible subsets of Ω and let $\mathbf{P}$ be defined by:
$$\forall a \in \Omega, \quad \mathbf{P}(\{a\}) = \frac{1}{N}.$$
Then $(\Omega, \mathcal{A}, \mathbf{P})$ is a probability space and for every $A \subset \Omega$,
$$\mathbf{P}(A) = \frac{\#A}{N},$$
where $\#A$ denotes the number of elements in A. Such probability measure is usually called the discrete uniform measure on Ω.

2. Let Ω be a subset of a real line (typically an interval, or a positive half-line). Let $\mathcal{A}$ be a Borel σ–algebra (i.e. the smallest σ–algebra containing all subintervals of Ω). Let f be a piecewise continuous, nonnegative function on R such that
$$\int_{-\infty}^{\infty} f(x)dx = 1.$$
If $\mathbf{P}$ is defined by
$$\mathbf{P}(A) = \int_A f(x)dx,$$
then $(\Omega, \mathcal{A}, \mathbf{P})$ is a probability space, and the function f is called the probability density function of $\mathbf{P}$.

Note: Because these two examples are by far the most important for us, and in both there is a "canonical" choice of $\mathcal{A}$, from now on the role of $\mathcal{A}$ will be diminished to a point that it will be usually ignored. This will free us of certain technical issues.

Theorem 3.1 *(basic properties of probability)*

(i) $\mathbf{P}(\emptyset) = 0$.

(ii) for any finite sequence of pairwise disjoint sets $A_1, \ldots, A_n$

$$\mathbf{P}(\bigcup_{j=1}^{n} A_j) = \sum_{j=1}^{n} \mathbf{P}(A_j),$$

(iii) for any set $A \in \mathcal{A}$, $\mathbf{P}(A) = 1 - \mathbf{P}(A^c)$.

(iv) for any sets $A, B \in \mathcal{A}$ such that $A \subset B$ we have $\mathbf{P}(A) \leq \mathbf{P}(B)$,

(v) for any sets $A, B \in \mathcal{A}$, $\mathbf{P}(A \cup B) = \mathbf{P}(A) + \mathbf{P}(B) - \mathbf{P}(A \cap B)$.

Proof: To prove (i) set $A_j = \emptyset$, $j = 1, 2, \ldots$. Then A_j's are pairwise disjoint and their union is the empty set. Hence, by the third axiom, we have

$$\mathbf{P}(\emptyset) = \sum_{j=1}^{\infty} \mathbf{P}(\emptyset),$$

which means that we must have $\mathbf{P}(\emptyset) = 0$. Once we know (i), (ii) follows by extending the finite sequence $A_1, \ldots, A_n$ to the infinite one by setting $A_j = \emptyset$, for $j = n+1, n+2, \ldots$. Then $\bigcup_{j=1}^{\infty} A_j = \bigcup_{j=1}^{n} A_j$ and $\sum_{j=1}^{\infty} A_j = \sum_{j=1}^{n} A_j$ and (ii) follows from the third axiom. The remaining properties are direct consequences of axioms and (ii). For example to see (iii) notice that A, A^c are disjoint and their union is Ω. Hence

$$1 = \mathbf{P}(\Omega) = \mathbf{P}(A \cup A^c) = \mathbf{P}(A) + \mathbf{P}(A^c),$$

which gives (iii). Similarly, for (iv) write $B = A \cup (A^c \cap B)$ with A and $A^c \cap B$ disjoint. Thus,

$$\mathbf{P}(B) = \mathbf{P}(A \cup (A^c \cup B)) = \mathbf{P}(A) + \mathbf{P}(A^c \cap B) \geq \mathbf{P}(A),$$

since by the first axiom $\mathbf{P}(A^c \cap B) \geq 0$. A proof of the last part is similar and is omitted. □

The following is a useful observation.

Disjointification: Let $A_1, A_2, \ldots,$ be a sequence of sets. Define their disjointification $B_1, B_2, \ldots$ as follows:

$$\begin{aligned} B_1 &= A_1, \\ B_k &= A_1^c \cap A_2^c \cap \cdots \cap A_{k-1}^c \cap A_k, \qquad \text{for } k \geq 2. \end{aligned}$$

Then, the sets $B_1, B_2, \ldots,$ have the following two properties:

(i) $\mathbf{P}(\bigcup_{j=1}^{\infty} A_j) = \mathbf{P}(\bigcup_{j=1}^{\infty} B_j)$, and

(ii) $B_1, B_2, \ldots,$ are pairwise disjoint.

4 Conditional Probability

Let $(\Omega, \mathcal{A}, \mathbf{P})$ be a probability space and let $A, B \in \mathcal{A}$, $\mathbf{P}(B) > 0$. Then the *conditional probability* of A given B (denoted by $\mathbf{P}(A|B)$) is defined by

$$\mathbf{P}(A|B) = \frac{\mathbf{P}(A \cap B)}{\mathbf{P}(B)}.$$

Note: The above formula is also frequently used in another form, namely, to compute the probability of the intersection of two sets:

$$\mathbf{P}(A \cap B) = \mathbf{P}(A|B)\mathbf{P}(B).$$

We say that a (possibly finite) sequence of pairwise disjoint sets $A_1, A_2, \dots$ is a (measurable) partition of Ω if

(i) $\forall j \geq 1 \; A_j \in \mathcal{A}$,

(ii) $\forall j \geq 1 \; \mathbf{P}(A_j) > 0$, and

(iii) $\bigcup_{j=1}^{\infty} A_j = \Omega$.

For example, if $A \in \mathcal{A}$ is such that $0 < \mathbf{P}(A) < 1$, then A, A^c is a partition of Ω.

Theorem 4.1 *(the law of total probability) Let $A_1, A_2, \dots,$ be a (possibly finite) partition of Ω. Then for every set $A \in \mathcal{A}$ we have*

$$\mathbf{P}(A) = \sum_{j \geq 1} \mathbf{P}(A|A_j)\mathbf{P}(A_j).$$

Proof: Let $B_j = A \cap A_j$. Then B_j's are pairwise disjoint (since A_j's are) and their union is A (because the union of all A_j's is Ω). Therefore,

$$\mathbf{P}(A) = \mathbf{P}(\bigcup_{j \geq 1} B_j) = \sum_{j \geq 1} \mathbf{P}(B_j) = \sum_{j \geq 1} \mathbf{P}(A \cap A_j) = \sum_{j \geq 1} \mathbf{P}(A|A_j)\mathbf{P}(A_j),$$

where the last equality follows from the definition of the conditional probability. $\square$

5 Independence

Two sets, $A, B \in \mathcal{A}$ are *independent* if

$$\mathbf{P}(A \cap B) = \mathbf{P}(A) \cdot \mathbf{P}(B).$$

Two sets that are not independent are called *dependent*.

Note: This definition and the definition of the conditional probability imply that A, B are independent iff $\mathbf{P}(A|B) = \mathbf{P}(A)$ (and thus also $\mathbf{P}(B|A) = \mathbf{P}(B)$). If one thinks of the conditional probability of A given B as the probability of A once we have information contained in B, then the last equation simply says that the information contained in B is totally irrelevant as far as chances of A occurring are concerned.
We also note the following: if A, B are independent then so are A, B^c (and thus also A^c, B^c, and A^c, B). Indeed,

$$\mathbf{P}(A) = \mathbf{P}(A \cap B) + \mathbf{P}(A \cap B^c) = \mathbf{P}(A)\mathbf{P}(B) + \mathbf{P}(A \cap B^c),$$

by independence. Hence,

$$\mathbf{P}(A \cap B^c) = \mathbf{P}(A) - \mathbf{P}(A)\mathbf{P}(B) = \mathbf{P}(A)(1 - \mathbf{P}(B)) = \mathbf{P}(A)\mathbf{P}(B^c),$$

i.e. A and B^c are independent.

For more than two sets the definition becomes as follows: sets $A_1, A_2, \ldots, A_n$ are independent if for any choice of any number of them, the probability of the intersection is equal to the product of probabilities, that is, if

$$\forall\ 2 \le m \le n,\ \forall\ 1 \le j_1 < j_2 < \cdots < j_m \le n,$$
$$\mathbf{P}(A_{j_1} \cap A_{j_2} \cap \cdots \cap A_{j_m}) = \mathbf{P}(A_{j_1}) \cdot \mathbf{P}(A_{j_2}) \cdot \cdots \cdot \mathbf{P}(A_{j_m}).$$

Note: This means, in particular, that three sets A, B, C are independent if *all* of the following hold:

$$\mathbf{P}(A \cap B) = \mathbf{P}(A)\mathbf{P}(B), \quad \mathbf{P}(A \cap C) = \mathbf{P}(A)\mathbf{P}(C),$$
$$\mathbf{P}(B \cap C) = \mathbf{P}(B)\mathbf{P}(C), \quad \mathbf{P}(A \cap B \cap C) = \mathbf{P}(A)\mathbf{P}(B)\mathbf{P}(C).$$

It is not enough that just the last equality holds.

6 Random Variables

A *random variable* X is a real valued function whose domain is a probability space, i.e.

$$X:\ (\Omega, \mathcal{A}, \mathbf{P}) \longrightarrow \mathbf{R}.$$

Note: In general we need more, namely that this function X is *measurable* which means that inverse image of any interval in $\mathbf{R}$ belongs to $\mathcal{A}$, i.e. $X^{-1}(I) \in \mathcal{A}$ for any interval I. Since practically all of the functions we will consider will satisfy that, (for example a random variable defined on a discrete uniform probability space is automatically measurable since in this case $\mathcal{A}$ contains all subsets of Ω) it will simplify our discussion if we just ignore that requirement.

In our study of random variables we are usually interested not in the random variables themselves, but, rather, in their distributions.

Definition: Let X be a random variable on $(\Omega, \mathcal{A}, \mathbf{P})$. Then, its *distribution function* F is a function $F : \mathbf{R} \longrightarrow [0, 1]$ defined by

$$\forall x \in \mathbf{R}, \quad F(x) = \mathbf{P}(X \leq x).$$

Note: Distribution function is frequently called the cumulative distribution function (abbreviated as c.d.f.). Also, sometimes we write F_X to emphasize that F is the c.d.f. of X.

Theorem 6.1 *(basic properties of the c.d.f.)*

(i) *F is increasing in the sense that if $x_1 \leq x_2$ then $F(x_1) \leq F(x_2)$,*

(ii) $\lim_{x\to\infty} F(x) = 1$,

(iii) $\lim_{x\to-\infty} F(x) = 0$,

(iv) *F is right continuous. That is, for every x and for every decreasing sequence (x_n) such that $\lim_{n\to\infty} x_n = x$ we have $\lim_{x\to\infty} F(x_n) = F(x)$.*

Note: *F does not* have to be continuous in general.

Two special classes of random variables are: discrete and continuous. A random variable that takes on countably (or finitely) many values, each with a positive probability, is called a *discrete* random variable. For such a random variable we define its *probability mass function* p as follows:

$$p(x) = \mathbf{P}(X = x).$$

Thus, $p(x)$ is equal to 0, unless x is one of the values taken on by X. Suppose, for simplicity that $x_1 < x_2 < \ldots$ are all the possible values of X and set $p_k = \mathbf{P}(X = x_k)$. Then we have the following relationship:

$$F(x) = \sum_{j:\ x_j \leq x} p_j.$$

Thus the c.d.f. of a discrete random variable is a step function, with steps occurring at the points $x_k,\ k = 1, 2, \ldots$.

We say that a random variable X is continuous if there exists a function $f : \mathbf{R} \longrightarrow \mathbf{R}$ such that

(i) $\forall\ x \in R,\ f(x) \geq 0$.

(ii) $\forall\ x < y$ we have

$$\mathbf{P}(x < X \leq y) = \int_x^y f(t)dt.$$

In that case the function f is called a density of X. Note that the left hand side above can be expressed in terms of the c.d.f. of X as $F(y) - F(x)$. This is because

$$\begin{aligned} \{X \le y\} &= (\{X \le y\} \cap \{X \le x\}) \cup (\{X \le y\} \cap \{X \le x\}^c) \\ &= \{X \le x\} \cup (\{X \le y\} \cap \{X > x\}) \,. \end{aligned}$$

Since the sets $\{X \le x\}$ and $\{X \le y\} \cap \{X > x\}$ are disjoint, for probabilities we get

$$\mathbf{P}(X \le y) = \mathbf{P}(X \le x) + \mathbf{P}(x < X \le y),$$

which means that

$$\mathbf{P}(x < X \le y) = F(y) - F(x).$$

In particular, letting $x \to -\infty$ we obtain

$$\forall y \in \mathbf{R}, \quad F(y) = \int_{-\infty}^{y} f(t)dt.$$

This is similar to a formula expressing the c.d.f. of a discrete random variable in terms of its probability mass function. In any case, the point is that if we know the density (in the continuous case) or the probability mass function (in the discrete case) then we can determine the c.d.f. and vice versa, if we know the c.d.f. of a discrete (or continuous) random variable X, then we can determine the probability mass function (or density) of X.

Examples of the most important random variables (note that all of them are described through their c.d.f's – or densities/probability mass functions):

(I) Discrete:

(i) discrete uniform: A random variable which takes on each of a finitely many values $x_1, x_2, \ldots, x_n$ with equal probability (necessarily $1/n$) is called a *discrete uniform* random variable.

(ii) Bernoulli with parameter p, $0 < p < 1$: A random variable X which is equal to 1 with probability p and is equal to zero with probability $1 - p$, i.e.

$$\mathbf{P}(X = 0) = 1 - p, \quad \mathbf{P}(X = 1) = p,$$

is called *Bernoulli* (with parameter p) random variable.

(iii) Binomial: Let an experiment resulting either with a "success" with probability p or a "failure" (with probability $1 - p$) be repeated independently n times. A random variable X that counts the number of successes in these n trials is called *binomial with parameters* n, p. Thus, for any k, $0 \le k \le n$,

$$\mathbf{P}(X = k) = \binom{n}{k} p^k (1 - p)^{n-k}.$$

(iv) Geometric with parameter p, $0 < p < 1$: Suppose that the experiment in (iii) is repeated independently until a success is obtained for the first time (and then it is stopped). A random variable X that counts how many repetitions are needed until this happens is called a *geometric random variable* (with parameter p). Thus,

$$\mathbf{P}(X = k) = (1-p)^{k-1} \cdot p, \quad \text{for } k = 1, 2, \ldots.$$

Note: sometimes geometric random variable is defined as the number of *failures before the first success*. In that case the formula for the probability mass function becomes:

$$\mathbf{P}(X = k) = (1-p)^{k} \cdot p, \quad \text{for } k = 0, 1, 2, \ldots.$$

(v) Poisson with parameter λ, $\lambda > 0$: Let λ be a positive number. A random variable X whose probability mass function is given by:

$$\mathbf{P}(X = k) = e^{-\lambda} \frac{\lambda^k}{k!}, \quad \text{for } k = 0, 1, 2, \ldots.$$

is called a *Poisson random variable with parameter* λ.

(II) Continuous:

(i) uniform on $[a, b]$: Let $a < b$ be real numbers. A random variable X whose density is given by:

$$f(x) = \begin{cases} \frac{1}{b-a} & \text{if } a \le x \le b \\ 0 & \text{otherwise} \end{cases}$$

is called a *uniform* random variable on the interval $[a, b]$.

(ii) exponential with parameter λ, $\lambda > 0$: let λ be a positive number. A random variable X whose density is given by

$$f(x) = \begin{cases} \lambda e^{-\lambda x} & \text{if } x \ge 0 \\ 0 & \text{otherwise} \end{cases}$$

is called a *exponential* random variable with parameter λ.

(iii) Gaussian (or normal random variable with parameters μ, σ: Let μ be a real number and let σ be a positive number. Then a random variable X whose density is given by

$$f(x) = \frac{1}{\sqrt{2\pi}\sigma} \exp\{-\frac{(x-\mu)^2}{2\sigma^2}\}, \quad \text{for } -\infty < x < \infty,$$

is called a *normal* (or a Gaussian) random variable with parameters μ and σ^2.

7 Expectations of Random Variables

Let X be a random variable on a probability space $(\Omega, \mathcal{A}, \mathbf{P})$. Then the *expected value* $\mathbf{E}X$ is defined by

$$\mathbf{E}X = \begin{cases} \sum_{j\geq 1} x_j \mathbf{P}(X = x_j), & \text{if } X \text{ is discrete with values } x_1, x_2, \ldots \\ \\ \int_{-\infty}^{\infty} x f(x) dx, & \text{if } X \text{ is continuous with density } f(x); \end{cases}$$

provided that the sum

$$\sum_{j\geq 1} |x_j| \mathbf{P}(X = x_j),$$

or, respectively, the integral

$$\int_{-\infty}^{\infty} |x| f(x) dx,$$

is finite.

Example:

(i) if X is a Bernoulli random variable with parameter p, then

$$\mathbf{E}X = 0 \cdot \mathbf{P}(X = 0) + 1 \cdot \mathbf{P}(X = 1) = 0 \cdot (1 - p) + 1 \cdot p = p.$$

(ii) if X is exponential with parameter 1, then its expected value is

$$\int_{-\infty}^{\infty} x f(x) dx = \int_0^{\infty} x e^{-x} dx = \int_0^{\infty} e^{-x} dx = 1,$$

where the next to last equality is justified by integration by parts.

Linearity of expectation: The expected value is linear, that is, if X and Y are two random variables and a, b are two constants then

$$\mathbf{E}(aX + bY) = a\mathbf{E}X + b\mathbf{E}Y.$$

Let X be a random variable on $(\Omega, \mathcal{A}, \mathbf{P})$ and let h be a function (say, piecewise continuous) $h : \mathbf{R} \longrightarrow \mathbf{R}$. Then the composition $Y = h(X) : \Omega \longrightarrow \mathbf{R}$ is again a random variable. Thus, we can talk about its expected value. We have:

$$\mathbf{E}Y = \mathbf{E}h(X) = \begin{cases} \sum_{j\geq 1} h(x_j) \mathbf{P}(X = x_j), & \text{if } X \text{ is discrete} \\ \\ \int_{-\infty}^{\infty} h(x) f(x) dx, & \text{if } X \text{ is continuous with density } f(x); \end{cases}$$

provided that the corresponding sum or integral is finite.

Example:

(i) (linearity) if X is a binomial random variable with parameters n, p, then $X = \sum_{i=1}^{n} X_i$, where

$$X_i = \begin{cases} 1, & \text{if there is a success on the } i\text{th trial}, \\ 0, & \text{otherwise.} \end{cases}$$

We usually write $X_i = I(X_i = 1)$ and call it the indicator function (of the set $\{$a success on the ith trial$\}$). Hence,

$$\mathbf{E}X = \mathbf{E}(\sum_{j=1}^{n} X_j) = \sum_{j=1}^{n} \mathbf{E}X_j = \sum_{j=1}^{n} p = n \cdot p.$$

(ii) if X is uniform on $[0, 1]$ and $h(x) = x^3$, then

$$\mathbf{E}Y = \mathbf{E}X^3 = \int_{-\infty}^{\infty} x^3 f(x) dx = \int_0^1 x^3 dx = \frac{1}{4}.$$

A few functions of particular importance:

(i) $h(x) = x^2$, or more generally $h(x) = x^p$, or $h(x) = |x|^p$, $p > 0$. Then $\mathbf{E}X^p$ and $\mathbf{E}|X|^p$ are called the pth and the absolute pth moment, respectively.

(ii) if for a random variable X, $h(x) = (x - \mathbf{E}X)^2$ then

$$\mathbf{E}h(X) = \mathbf{E}(X - \mathbf{E}X)^2,$$

is called the *variance* of X and is denoted by $\mathrm{var}(X)$, or $D^2(X)$. By squaring out the term under the expected value sign and using linearity we see that $\mathrm{var}(X) = \mathbf{E}X^2 - (\mathbf{E}X)^2$.

(iii) if $h(x) = e^{t \cdot x}$ then the function $t \to M(t) = \mathbf{E}e^{tX}$ is called the *moment generating function* of X. Its domain is the set of all t for which the corresponding expression is finite.

(iv) (outside of real valued functions but barely) for a real t consider $h(x) = e^{itx}$ where i is the imaginary unit. Then the function $\phi(t)$ defined by

$$\phi(t) = \mathbf{E}e^{itX},$$

is the *characteristic function* of X. Note that if $Y = aX + b$ where a, b are constants then

$$\phi_Y(t) = \mathbf{E}e^{it(aX+b)} = e^{itb}\mathbf{E}e^{itaX} = e^{itb}\phi_X(at).$$

Note: The advantage of the characteristic function over the moment generating function is that the former *always* exists and its domain is *always* the whole real line $\mathbf{R}$.

8 Independent Random Variables

Two random variables X and Y are called independent if for all x, y we have

$$\mathbf{P}(\{X \le x\} \cap \{Y \le y\}) = \mathbf{P}(X \le x)\mathbf{P}(Y \le y).$$

The left hand side is usually written as $\mathbf{P}(X \le x, Y \le y)$ and the function $F_{X,Y}(x,y) : \mathbf{R} \times \mathbf{R} \longrightarrow \mathbf{R}$ is called the joint distribution function of X and Y. The equation becomes

$$F_{X,Y}(x,y) = F_X(x) \cdot F_Y(y),$$

i.e. the joint distribution function is the product of individual c.d.f.'s. If we denote the joint probability mass function of two discrete random variables X, Y by

$$p_{X,Y}(x,y) = \mathbf{P}(X = x, Y = y),$$

then if $x_1 < x_2 \ldots$ and $y_1 < y_2 \ldots$ are values taken by X and Y, respectively, then we have

$$p_{X,Y}(x_j, y_k) = \mathbf{P}(X = x_j, Y = y_k).$$

Expressing the latter in terms of the c.d.f.'s by using

$$\{X = x_j, Y = y_k\} = \{X \le x_j\} \cap \{X \le x_{j-1}\}^c \cap \{Y \le y_k\} \cap \{Y \le y_{k-1}\}^c,$$

and using independence we obtain that

$$p_{X,Y}(x,y) = p_X(x)p_Y(y).$$

By a similar argument using the fact that relationship for the joint density is given by

$$f_{X,Y}(x,y) = \frac{\partial^2 F_{X,Y}(x,y)}{\partial x \partial y},$$

at every point x, y at which the joint c.d.f. is differentiable, we see that in the continuous case, if X and Y are independent then

$$f_{X,Y}(x,y) = f_X(x)f_Y(y).$$

As a consequence we obtain that whenever X, Y are independent then

$$\mathbf{E}(X \cdot Y) = \mathbf{E}(X) \cdot \mathbf{E}(Y),$$

whenever $\mathbf{E}|X \cdot Y| < \infty$. What is more, if $h : \mathbf{R} \times \mathbf{R} \longrightarrow$ is a function such that $h(x,y) = h_1(x) \cdot h_2(y)$, and X, Y are independent then

$$\mathbf{E}h(X,Y) = \mathbf{E}h_1(X) \cdot \mathbf{E}h_2(Y),$$

provided $\mathbf{E}|h(X,Y)| < \infty$. For discrete X, Y the argument goes like this:

$$\begin{aligned}
\mathbf{E}h(X,Y) &= \sum_{j\geq 1, k\geq 1} h(x_j, y_k)\mathbf{P}(X = x_j, Y = y_k) \\
&= \sum_{j\geq 1, k\geq 1} h_1(x_j)h_2(y_k)\mathbf{P}(X = x_j)\mathbf{P}(Y = y_k) \\
&= \sum_{j\geq 1}\sum_{k\geq 1} h_1(x_j)\mathbf{P}(X = x_j)h_2(y_k)\mathbf{P}(Y = y_k) \\
&= \sum_{j\geq 1} h_1(x_j)\mathbf{P}(X = x_j)\left(\sum_{k\geq 1} h_2(y_k)\mathbf{P}(Y = y_k)\right) \\
&= \mathbf{E}X\mathbf{E}Y
\end{aligned}$$

where the second equality is by the property of the function h and independence, while the third one is by rearranging the order of summation (that's where the assumption $\mathbf{E}|h(X,Y)| < \infty$ is needed). For continuous variables, the argument is essentially the same - one needs to replace the summation by integration. Two key examples of functions h with that property are functions related to the moment generating function or characteristic function of the sum of independent random variables,

$$h(x,y) = e^{t(x+y)} = e^{tx+ty} = e^{tx}\cdot e^{ty} = h_1(x)\cdot h_2(y),$$

or

$$h(x,y) = e^{it(x+y)} = e^{itx+ity} = e^{itx}\cdot e^{ity} = h_1(x)\cdot h_2(y).$$

Thus, we have the following statement:

Theorem 8.1 *If X, Y are independent random variables then*

(i) the moment generating function of a sum $X + Y$ is the product of the moment generating functions of X and Y (for those t for all three exist)

(ii) the characteristic function $\phi_{X+Y}(t)$ of the sum is the product of the characteristic functions $\phi_X(t)$ of X and $\phi_Y(t)$ of Y, i.e.

$$\phi_{X+Y}(t) = \phi_X(t)\cdot\phi_Y(t), \quad \textit{for all } t \in \mathbf{R}.$$

The concept of independence for more than two random variables (say, $X_1, X_2, \ldots, X_n$) is defined by requiring that for all choices of $x_1, x_2, \ldots, x_n$ the events

$$\{X_1 \leq x_1\},\ \{X_2 \leq x_2\}, \ldots, \{X_n \leq x_n\},$$

are independent. We then have the same property: the joint p.d.f. (or density) is the product of individual p.d.f.'s (or densities) and thus, if

$$h : \mathbf{R}^n \longrightarrow \mathbf{R}$$

has the property that

$$h(x_1, x_2, \ldots, x_n) = h_1(x_1) \cdot h_2(x_2) \cdot \cdots \cdot h_n(x_n),$$

then for independent random variables $X_1, X_2, \ldots, X_n$ we have

$$\mathbf{E}h(X_1, X_2, \ldots, X_n) = \mathbf{E}h_1(X_1) \cdot \mathbf{E}h_2(X_2) \cdot \cdots \cdot \mathbf{E}h_n(X_n).$$

In particular, the characteristic function of the sum of independent random variables is the product of their characteristic functions. That is, if $X_1, X_2, \ldots, X_n$ are independent random variables then

$$\phi_{X_1+X_2+\cdots+X_n}(t) = \phi_{X_1}(t) \cdot \phi_{X_2}(t) \cdots \phi_{X_n}(t).$$

9 Inversion Formula

Theorem 9.1 *(inversion formula) Let X be a random variable with distribution F and characteristic function ϕ. Then for all $a < b$ such that $\mathbf{P}(X = a) = \mathbf{P}(X = b) = 0$, the following holds*

$$F(b) - F(a) = \lim_{T\to\infty} \frac{1}{2\pi} \int_{-T}^{T} \frac{e^{-ita} - e^{-itb}}{it} \phi(t)dt.$$

Proof: First, write Taylor expansion of the exponential function (with the remainder in the integral form):

$$e^{ix} = \sum_{k=0}^{n} \frac{(ix)^k}{k!} + \frac{i^{n+1}}{n!} \int_0^x (x-s)^n e^{is} ds.$$

By integration by parts we have

$$\int_0^x (x-s)^{n-1} e^{is} ds = \frac{x^n}{n} + \frac{i}{n} \int_0^x (x-s)^n e^{is} ds,$$

Hence

$$\int_0^x (x-s)^n e^{is} ds = \frac{n}{i} \left(\int_0^x (x-s)^{n-1} e^{is} ds - \frac{x^n}{n} \right),$$

and since

$$\frac{x^n}{n} = \int_0^x (x-s)^{n-1} ds,$$

combining those two expressions we obtain:

$$\int_0^x (x-s)^n e^{is} ds = \frac{n}{i} \int_0^x (x-s)^{n-1} (e^{is} - 1) ds.$$

Substituting this integral into our Taylor's expansion for e^{ix} we obtain another expansion:

$$e^{ix} = \sum_{k=0}^{n} \frac{(ix)^k}{k!} + \frac{i^n}{(n-1)!} \int_0^x (x-s)^{n-1}(e^{is}-1)ds.$$

If $x > 0$ then the absolute value of the last term on the right is no more than (recall that $|e^{is}| = 1$ for real s)

$$\left|\frac{i^n}{(n-1)!}\right| \cdot \left|\int_0^x (x-s)^{n-1}(e^{is}-1)ds\right| \leq \frac{1}{(n-1)!} \int_0^x (x-s)^{n-1}|e^{is}-1|ds.$$

Since $|e^{is}-1| \leq |e^{is}| + |1| \leq 2$, this is further upper-bounded by

$$\frac{2}{(n-1)!} \int_0^x (x-s)^{n-1} ds = \frac{2x^n}{n!}.$$

Essentially the same computation can be repeated for $x < 0$ and we get

$$\left|\frac{i^n}{(n-1)!} \int_0^x (x-s)^{n-1}(e^{is}-1)ds\right| \leq \frac{2|x|^n}{n!}.$$

Repeating the same argument for the remainder in the first expansion of e^{ix} we get

$$\left|\frac{i^{n+1}}{n!} \int_0^x (x-s)^n e^{is} ds\right| \leq \frac{|x|^{n+1}}{(n+1)!}.$$

Hence, combining those two bounds we derive a useful bound on the expansion of the function e^{ix} (we will need it later !):

$$\left|e^{ix} - \sum_{k=0}^{n} \frac{(ix)^k}{k!}\right| \leq \min\left\{\frac{|x|^{n+1}}{(n+1)!}, \frac{2|x|^n}{n!}\right\},$$

valid for all $n \geq 0$. In particular, letting $n = 0$ we obtain

$$|e^{ix} - 1| \leq \min\{|x|, 2\}.$$

We now return to the proof. For simplicity we will prove it for continuous random variables. So, suppose that X has density f. Thus its characteristic function is $\phi(t) = \int_{-\infty}^{\infty} e^{itx} f(x)dx$, and therefore the expression

$$\frac{1}{2\pi} \int_{-T}^{T} \frac{e^{-ita} - e^{-itb}}{it} \phi(t)dt$$

becomes

$$\frac{1}{2\pi} \int_{-T}^{T} \frac{e^{-ita} - e^{-itb}}{it} \left(\int_{-\infty}^{\infty} e^{itx} f(x)dx\right) dt.$$

Rewriting this as a double integral and multiplying out we get

$$\frac{1}{2\pi}\int_{-T}^{T}\int_{-\infty}^{\infty}\frac{e^{it(x-a)}-e^{it(x-b)}}{it}f(x)dxdt.$$

By our previous estimates

$$\begin{aligned}\left|\frac{e^{it(x-a)}-e^{it(x-b)}}{it}\right| &= \frac{\left|e^{it(x-b)}\left(e^{it(x-a-(x-b))}-1\right)\right|}{|it|}\\ &= \frac{\left|e^{it(b-a)}-1\right|}{|t|}\le\frac{\min\{|t(b-a)|,2\}}{|t|}\le|b-a|,\end{aligned}$$

which is bounded independently of t and x. Therefore, the double integral can be computed as iterated integrals in any order. Exchanging the order of integration we see that it is:

$$\frac{1}{2\pi}\int_{-\infty}^{\infty}\left(\int_{-T}^{T}\frac{e^{it(x-a)}-e^{it(x-b)}}{it}dt\right)f(x)dx.$$

For the inner integral use the fact that $e^{iu}=\cos u+i\sin u$ and linearity of the integral to get

$$\begin{aligned}\frac{1}{i}\int_{-T}^{T}\frac{\cos(t(x-a))}{t}dt &- \frac{1}{i}\int_{-T}^{T}\frac{\cos(t(x-b))}{t}dt\\ &+ \int_{-T}^{T}\frac{\sin(t(x-a))}{t}dt-\int_{-T}^{T}\frac{\sin(t(x-b))}{t}dt.\end{aligned}$$

Since the interval of integration is symmetric about 0 and the first two integrands are odd functions (and thus their integrals are zero) while the last two are even, this expression is equal to

$$2\int_{0}^{T}\frac{\sin(t(x-a))}{t}dt-2\int_{0}^{T}\frac{\sin(t(x-b))}{t}dt.$$

Let us denote for $T\ge 0$

$$S(T)=\int_{0}^{T}\frac{\sin x}{x}dx$$

and let $\text{sign}(u)$ be a function which is $+1$ if $u>0$, is -1 if $u<0$, and is 0 if $u=0$. By change of variables, for any v

$$\int_{0}^{T}\frac{\sin(tv)}{t}dt=\text{sign}(v)\cdot\int_{0}^{T|v|}\frac{\sin y}{y}dy=\text{sign}(v)S(T|v|).$$

Thus, substituting this into inner integral gives the following value for the double integral

$$\int_{-\infty}^{\infty}\left(\frac{\text{sign}(x-a)}{\pi}S(T|x-a|)-\frac{\text{sign}(x-b)}{\pi}S(T|x-b|)\right)f(x)dx.$$

It remains to take the limit as $T \to \infty$. It is known from calculus that

$$\lim_{T\to\infty} \int_0^T \frac{\sin x}{x} dx = \frac{\pi}{2}, \quad \text{that is} \quad \lim_{T\to\infty} S(T) = \frac{\pi}{2}.$$

Therefore the function

$$\frac{\operatorname{sign}(x-a)}{\pi} S(T|x-a|) - \frac{\operatorname{sign}(x-b)}{\pi} S(T|x-b|)$$

converges to 0 if $x < a$ or $x > b$ and to 1 if $a < x < b$. (It also converges to $1/2$ if $x = a$ or $x = b$, but this is irrelevant.) Thus,

$$\lim_{T\to\infty} \int_{-\infty}^{\infty} \left(\frac{\operatorname{sign}(x-a)}{\pi} S(T|x-a|) - \frac{\operatorname{sign}(x-b)}{\pi} S(T|x-b|) \right) f(x)dx$$
$$= \int_a^b f(x)dx = F(b) - F(a).$$

This proves the inversion formula. □

Note: The inversion formula, allows one to recover the distribution function of a random variable from its characteristic function (in principle at least). But, even more importantly, as a consequence, we have the so-called uniqueness theorem, namely if two random variables have the same characteristic function then they have to have the same distribution function as well. Since, obviously, two random variables with the same distribution function have to have the same characteristic function, this is an "if and only if" statement. In other words, characteristic function *uniquely* determines the distribution of a random variable.

10 Convergence in Distribution

Definition. Let X be a random variable with c.d.f. $F(x)$, and let $X_1, X_2 \ldots$ be a sequence of random variables with c.d.f's $F_1(x), F_2(x), \ldots$, respectively. We say that the sequence (X_n), $n \geq 1$ *converges in distribution to* X if

$$F_n(x) \longrightarrow F(x), \quad \text{as} \quad n \to \infty,$$

for all $x \in \mathbf{R}$ at which $F(x)$ is continuous. We write

$$X_n \Longrightarrow X$$

to indicate the convergence in distribution.

Note:

(i) Since F is nondecreasing it can have at most countably many points at which it is not continuous. Furthermore, in the most interesting cases F is, in fact, continuous.

(ii) The definition is really about the convergence of *distribution functions* of random variables, rather than random variables themselves. In particular, the random variables $X_1, X_2, \dots$ need not be defined on the same probability space. For that reason we will often speak of convergence of distributions and write
$$F_n \Longrightarrow F.$$

Example: Let $X_1, X_2, \dots$ be i.i.d. (independent and identically distributed) random variables, each having an exponential distribution with parameter 1 (and thus the common c.d.f. is $F(x) = 1 - e^{-x}$). Let $M_m = \max\{X_1, X_2, \dots, X_n\}$, let F_n be the c.d.f. of M_n and set $b_n = \ln n$. Then we have

$$\mathbf{P}(M_n - b_n \le x) = \mathbf{P}(M_n \le b_n + x) = F_n(b_n + x).$$

Since $\{M_n \le z\} = \{X_j \le z \,\forall 1 \le j \le n\}$ by independence $F_n(z) = F^n(z)$. Thus

$$\begin{aligned} F_n(b_n + x) &= F_n(\ln n + x) = F^n(\ln n + x) = \left(1 - e^{-\ln n - x}\right)^n \\ &= \left(1 - \frac{e^{-x}}{n}\right)^n \longrightarrow e^{-e^{-x}}. \end{aligned}$$

Since
$$\lim_{x \to -\infty} e^{-e^{-x}} = 0, \ \lim_{x \to \infty} e^{-e^{-x}} = 1,$$
and this function is continuous and nondecreasing, it is a c.d.f. (called a maximal value distribution). Thus,
$$M_n - \ln n \Longrightarrow Y,$$
where the c.d.f. of Y is the extreme value distribution.

The following observation is very useful.

Lemma 10.1 *Let $F, F_1, F_2, \dots$ be distribution functions such that $F_n \Longrightarrow F$. Then there exist $Y, Y_1, Y_2 \dots$ defined on the common probability space $(\Omega, \mathcal{A}, \mathbf{P})$ and having distributions $F, F_1, F_2, \dots$ respectively, and such that $Y_n(\omega) \to Y(\omega)$ for all $\omega \in \Omega$.*

Proof: We take as the common probability space the interval $(0, 1)$ with Borel σ-algebra (i.e. σ-algebra generated by all intervals) and the Lebesgue measure (i.e. the unique measure $\mathbf{P}$ satisfying $\mathbf{P}((a, b)) = |b - a|$). Define Y_n and Y by

$$Y_n(\omega) = \inf\{x : \ \omega \le F_n(x)\}, \quad Y(\omega) = \inf\{x : \ \omega \le F(x)\},$$

(it helps to draw a picture; these are essentially inverse functions of F_n's and F, except that F and F_n's need not be strictly increasing and thus may not have the "real" inverses). Then we have $\omega \le F_n(x)$ iff $Y_n(\omega) \le x$ with the similar statement for Y. Hence

$$\mathbf{P}(\{\omega : \ Y_n(\omega) \le x\}) = \mathbf{P}(\{\omega : \ \omega \le F_n(x)\}) = F_n(x),$$

where the last equality holds by the definition of our $(\Omega, \mathcal{A}, \mathbf{P})$. Thus, Y_n has F_n as its c.d.f. and the same argument shows that Y has c.d.f. F. It remains to show that $Y_n(\omega) \to Y(\omega)$ for all $\omega \in (0,1)$ (again, if one thinks of Y_n as inverses of F_n's then the fact that F_n converge implies the the inverses converge, too). For the general case there is a bit more to say: Let $0 < \omega < 1$ and let $\varepsilon > 0$. Pick an x such that

$$Y(\omega) - \varepsilon < x < Y(\omega) \quad \text{and such that } F \text{ is continuous at } x,$$

(since F is nondecreasing it has at most countably many points of discontinuity). Then $F(x) < \omega$ and since $F_n(x)$ converges to $F(x)$ it follows that for sufficiently large n's we will have $F_n(x) < \omega$ as well; hence,

$$Y(\omega) - \varepsilon < x < Y_n(\omega),$$

which implies that $\liminf_n Y_n(\omega) \geq Y(\omega)$. To show that $\limsup_n Y_n(\omega) \leq Y(\omega)$, pick any ω' such that $\omega < \omega'$ and for a given ε choose a y such that

$$Y(\omega') < y < Y(\omega') + \varepsilon \quad \text{and such that } F \text{ is continuous at } y.$$

Then we have

$$\omega < \omega' \leq F(Y(\omega')) \leq F(y),$$

and since $F_n(y)$ converges to $F(y)$ for sufficiently large n's we also have $\omega \leq F_n(y)$ and thus $Y_n(\omega) \leq y < Y(\omega') + \varepsilon$. Since ε is arbitrary, this implies that $\limsup_n Y_n(\omega) \leq Y(\omega')$, whenever $\omega < \omega'$. Letting $\omega' \to \omega$ we conclude that $\limsup_n Y_n(\omega) \leq Y(\omega)$, provided Y is continuous at ω. Since Y is nondecreasing it has at most countably many points of discontinuity. At any such point we can redefine $Y_n(\omega) = Y(\omega) = 0$. This does not change the distributions of Y or Y_n's (since they are changed at at most countably many points) and makes $Y_n(\omega)$ converge to $Y(\omega)$ for all $\omega \in (0,1)$. □

Theorem 10.2 *Let $X, X_1, X_2, \dots$ be random variables with distribution functions $F, F_1, F_2 \dots$, respectively. The following two conditions are equivalent:*

(i) $X_n \Longrightarrow X$.

(ii) $\mathbf{E}f(X_n) \longrightarrow \mathbf{E}f(X)$, *for every bounded, continuous real valued function* f.

Proof of (i)$\Longrightarrow$ (ii). Suppose $F_n \Longrightarrow F$ and consider random variables Y_n and Y given by the above lemma. Let $f : \mathbf{R} \longrightarrow \mathbf{R}$ be a continuous function such that $|f(x)| \leq K$, for $x \in \mathbf{R}$. Since $Y_n(\omega) \to Y(\omega)$ for all $\omega \in (0,1)$ and f is continuous, $f(Y_n(\omega)) \to f(Y(\omega))$ for all $\omega \in (0,1)$ as well. That is,

$$\forall \omega \in (0,1)\ \forall \varepsilon > 0\ \exists k\ \forall n \geq k, \text{ we have } |f(Y_n(\omega)) - f(Y(\omega))| < \varepsilon.$$

Letting $\varepsilon = \frac{1}{m}$, $m = 1, 2, \dots$, and defining the sets

$$A_{n,m} = \{\omega : \ |f(Y_n(\omega)) - f(Y(\omega))| \geq \frac{1}{m}\},$$

the above quantification can be restated as

$$\bigcap_{m\geq 1}\bigcup_{k\geq 1}\bigcap_{n\geq k} A_{n,m}^c = \Omega = (0,1).$$

That means that

$$\forall m \geq 1 \ \bigcup_{k\geq 1}\bigcap_{n\geq k} A_{n,m}^c = \Omega = (0,1)$$

and since the sets

$$\bigcap_{n\geq k} A_{n,m}^c, \quad k = 1, 2, \ldots$$

form an increasing sequence whose union is the whole Ω we must have

$$\lim_{k\to\infty} \mathbf{P}\left(\bigcap_{n\geq k} A_{n,m}^c\right) = 1.$$

This means that for any $\varepsilon > 0$ there exists a k_0 such that for any $k \geq k_0$ we have

$$\mathbf{P}\left(\bigcap_{n\geq k} A_{n,m}^c\right) \geq 1-\varepsilon$$

and thus for the complement

$$\mathbf{P}\left(\bigcup_{n\geq k}\{\omega : \ |f(Y_n(\omega)) - f(Y(\omega))| \geq \frac{1}{m}\}\right) \leq \varepsilon.$$

Therefore,

$$\mathbf{E}f(Y_n) = \mathbf{E}f(Y) + \mathbf{E}\left(f(Y_n) - f(Y)\right)$$

and its enough to show that the second term goes to 0. We have

$$\begin{aligned}
&|\mathbf{E}\left(f(Y_n) - f(Y)\right)| \leq \mathbf{E}\left|f(Y_n) - f(Y)\right| \\
&\quad = \mathbf{E}\left|f(Y_n) - f(Y)\right| I\left(\bigcup_{n\geq k}\{|f(Y_n(\omega)) - f(Y(\omega))| \geq \frac{1}{m}\}\right) \\
&\quad\quad + \mathbf{E}\left|f(Y_n) - f(Y)\right| I\left(\bigcap_{n\geq k}\{|f(Y_n(\omega)) - f(Y(\omega))| < \frac{1}{m}\}\right) \\
&(\text{since } |f(x) - f(y)| \leq |f(x)| + |f(y)| \leq 2K) \\
&\quad \leq 2K\mathbf{P}\left(\bigcup_{n\geq k}\{|f(Y_n(\omega)) - f(Y(\omega))| \geq \frac{1}{m}\}\right) + \frac{1}{m} \\
&(\text{provided } n \geq k_0) \\
&\quad \leq 2K\varepsilon + \frac{1}{m} \\
&(\text{since } m \text{ and } \varepsilon \text{ are arbitrarily large and small, respectively}) \\
&\quad \longrightarrow 0.
\end{aligned}$$

This proves that part. (A moments reflection on the order of choices of various parameters might be helpful: to show that $\mathbf{E}\,|f(Y_n) - f(Y)|$ goes to zero, we have to show that for an arbitrary $\eta > 0$ there exists an n_0 such that this quantity is no more than η for all $n \geq n_0$. For our η pick m such that $1/m < \eta/2$, then pick $\varepsilon > 0$ such that $2K\varepsilon < \eta/2$ (K is given in advance.) For that ε there exists a k_0 with the properties described above. Now we may take $n_0 \geq k_0$.)
Proof of (ii)$\Longrightarrow$(i). Let $x < y$ and define the function $f(t)$ as follows

$$f(t) = \begin{cases} 1 & \text{if } t \leq x \\ \frac{y-t}{y-x} & \text{if } x \leq t \leq y \\ 0 & \text{if } t \geq y \end{cases}$$

Then $F_n(x) \leq \mathbf{E}f(X_n)$ and $\mathbf{E}f(X) \leq F(y)$. Therefore, $\limsup_n F_n(x) \leq F(y)$, and letting $y \to x$ we get $\limsup_n F_n(x) \leq F(x)$, provided F is continuous at x. Similarly, for $u < x$ $F(u) \leq \liminf_n F_n(x)$ and again, letting $u \to x$ we get $F(x) \leq \liminf_n F_n(x)$, provided F is continuous at x. Thus $\lim_n F_n(x) = F(x)$ for any x at which F is continuous. This completes the proof of the theorem. $\square$ We will need a few more results before the main theorem

Theorem 10.3 *(Helly selection theorem) For every sequence of c.d.f's F_n, $n = 1, 2, \ldots$ there exists a subsequence n_k, $k = 1, 2\ldots$ and a nondecreasing, right-continuous function F such that $\lim_{k\to\infty} F_{n_k}(x) = F(x)$ at all points x at which F is continuous.*

Proof: Let $r_1, r_2, \ldots$ be an enumeration of rational numbers. Then $F_n(r_k)$ $n \geq 1$, $k \geq 1$ is a doubly indexed bounded sequence. We will use what is called a diagonal method to find a subsequence n_m such that $F_{n_m}(r_k)$ converges as $m \to \infty$ for all $k \geq 1$. First, since $F_n(r_1)$ is bounded, there exists a subsequence $n_1^1, n_2^1, \ldots$ such that $F_{n_m^1}(r_1)$ converges as $m \to \infty$. Look at the values of F's at r_2 along that subsequence, i.e. at $F_{n_m^1}(r_2)$, $m \geq 1$. It is a bounded sequence and thus there exists a subsequence n_m^2, $m \geq 1$ such that $F_{n_m^2}(r_2)$ converges as $m \to \infty$. Look at the values $F_{n_m^2}(r_3)$ and continue the same argument. This produces an array:

$$\begin{array}{cccc} F_{n_1^1}(r_1) & F_{n_2^1}(r_1) & F_{n_3^1}(r_1) & \ldots \\ F_{n_1^2}(r_2) & F_{n_2^2}(r_2) & F_{n_3^2}(r_2) & \ldots \\ F_{n_1^3}(r_3) & F_{n_2^3}(r_3) & F_{n_3^3}(r_3) & \ldots \\ \cdot & \cdot & \cdot & \\ \cdot & \cdot & \cdot & \\ \cdot & \cdot & \cdot & \end{array}$$

This array has the property that the sequence converges in every row, and that the sequence of indices in every row (except the first one, of course) is a subsequence of those in the previous one. Choose a diagonal (hence the name) $n_m = n_m^m$, $m = 1, 2, \ldots$. Then the sequence $F_{n_m}(r_k)$ converges as $m \to \infty$ for all $k \geq 1$. Indeed, for a given k look at n_m for $m \geq k$. All of them are a subsequence of the indices in the kth row, and since $F_{n_j^k}(r_k)$ converges as $j \to \infty$

it does so along any subsequence as well. Thus, for every rational number r, the limit $\lim_{m\to\infty} F_{n_m}(r)$ exists; let's call it $G(r)$. Now define

$$F(x) = \inf\{G(r) : \ x < r\}.$$

It is obvious from the definition that F is nondecreasing. To show that it is right-continuous, pick any x and for an arbitrary $\varepsilon > 0$ find a rational r,$r > x$ such that $G(r) < F(x) + \varepsilon$. We have to show that $F(y) \to F(x)$ as y tends to x from the right. But, for any y such that $x \le y < r$ we have $F(y) \le G(r) < F(x) + \varepsilon$. Since ε is arbitrary, and $F(x) \le F(y)$ (F is nondecreasing) it follows that $F(y) \to F(x)$ as $y \searrow x$. Finally, to show that $F_{n_m}(x)$ converges to $F(x)$ at any continuity point x of F, for an $\varepsilon > 0$ first pick a y, $y < x$ and such that $F(x) - \varepsilon < F(y)$ and then rational r and s such that $y < r < x < s$ and $G(s) < F(x) + \varepsilon$. We have

$$F(x) - \varepsilon < G(r) \le G(s) < F(x) + \varepsilon, \quad \text{and } F_n(r) \le F_n(x) \le F_n(s).$$

Hence,

$$\limsup_m F_{n_m}(x) \le \limsup_m F_{n_m}(s) = \lim_m F_{n_m}(s) = G(s) < F(x) + \varepsilon,$$

and

$$\liminf_m F_{n_m}(x) \ge \liminf_m F_{n_m}(r) = \lim_m F_{n_m}(r) = G(r) > F(x) - \varepsilon.$$

Since ε is arbitrary, it follows that $\lim_m F_{n_m}(x) = F(x)$. This completes the proof. $\square$

Note: Clearly, we have $0 \le F(x) \le 1$, but F need not be a c.d.f. because we may fail to satisfy $\lim_{x\to\infty} F(x) = 1$, for example. (Consider for instance $F_n(x) = 0$ if $x < n$ and 1 if $x \ge n$; then F given by Helly's theorem is identically 0.) It would, therefore, be nice to have a condition that would guarantee that this F is, indeed a c.d.f. Hence the following concept:

Definition. A sequence of c.d.f's F_n, $n = 1, 2, \dots$ is *tight* if

$$\forall \varepsilon > 0 \ \exists x, y \ \forall n \ge 1 \text{ such that } F_n(x) < \varepsilon, \text{ and } F_n(y) > 1 - \varepsilon.$$

Theorem 10.4 *Let F_n be a sequence of c.d.f's. Then the following conditions are equivalent:*

(i) F_n, $n = 1, 2\dots$ is tight,

(ii) for every subsequence n_k there exist a further subsequence n_{k_m} and a c.d.f. F such that $F_{n_{k_m}} \Longrightarrow F$.

Proof of (i)⟹(ii). (that's all that will be needed). Let F_n, $n \geq 1$, be tight. Pick any subsequence and apply Helly's theorem to that subsequence to get a nondecreasing, right-continuous F such that (a further) subsequence of F_n's converges to F at its continuity points. Call this final subsequence n_j, $j \geq 1$. All we need to do is to show that F is indeed a c.d.f., i.e. that we have

$$\lim_{x \to -\infty} F(x) = 0, \text{ and } \lim_{x \to \infty} F(x) = 1.$$

To this end pick an $\varepsilon > 0$. By tightness, pick x_0 and y_0 such that $F_n(x_0) < \varepsilon$, and $F_n(y_0) > 1 - \varepsilon$ for all n. By monotonicity of F_n's, decreasing x_0 and increasing y_0 if necessary, we can assume that both are continuity points of F. Thus, since $F_{n_j}(x_0) \to F(x_0)$ and $F_{n_j}(y_0) \to F(y_0)$ as $j \to \infty$ we must have $F(x_0) < \varepsilon$ and $F(y_0) > 1 - \varepsilon$, which, since ε was arbitrary, implies that

$$\lim_{x \to \infty} F(x) = 0, \text{ and } \lim_{x \to \infty} F(x) = 1.$$

as required.

$$\lim_{x \to \infty} F(x) = 0, \text{ and } \lim_{x \to \infty} F(x) = 1.$$

Proof of (ii)⟹(i). Suppose F_n, $n \geq 1$, has the property that every subsequence has the further subsequence which converges in distribution to a c.d.f. We will show that F_n, $n \geq 1$, has to be tight. Suppose it is not. Then there exists an $\varepsilon > 0$ such that for all a, b there exists an n such that $F_n(b) - F_n(a) < 1 - \varepsilon$. For that ε choose a subsequence n_k such that $F_{n_k}(k) - F_{n_k}(-k) \leq 1 - \varepsilon$. This subsequence had a further subsequence n_{k_m} such that along that subsequence $F_{n_{k_m}}$ converges to a c.d.f F. Pick a_0, b_0 such that $F(b_0) - F(a_0) > 1 - \varepsilon$, and again, by decreasing a_0 and increasing b_0 if necessary we can assume that both are continuity points of F. Then $F_{n_{k_m}}(b_0) \to F(b_0)$ and $F_{n_{k_m}}(a_0) \to F(a_0)$ as $m \to \infty$, and since for m large enough $-k_m < a_0$ and $k_m > b_0$ we have $F_{n_{k_m}}(-k_m) \leq F_{n_{k_m}}(a_0)$ and $F_{n_{k_m}}(k_m) \geq F_{n_{k_m}}(b_0)$. Therefore,

$$1 - \varepsilon \geq F_{n_{k_m}}(k_m) - F_{n_{k_m}}(-k_m) \geq F_{n_{k_m}}(b_0) - F_{n_{k_m}}(a_0) \longrightarrow F(b_0) - F(a_0).$$

But that would mean that $F(b_0) - F(a_0) \leq 1 - \varepsilon$, which gives a contradiction, and thus proves that F_n, $n \geq 1$, is tight. □ Finally,

Corollary 10.5 *If F_n, $n \geq 1$ is a tight sequence of c.d.f's with the property that each subsequence that converges in distribution converges to the same distribution F, then*

$$F_n \Longrightarrow F.$$

Proof: Suppose it is not true that $F_n \Longrightarrow F$. That means that there is a continuity point of F x_0 such that $F_n(x_0)$ does not converge to $F(x_0)$. But then, there exists an $\varepsilon > 0$ and a subsequence n_k, $k \geq 1$ such that $|F_{n_k}(x_0) - F(x_0)| \geq \varepsilon$. By tightness and previous result, this subsequence has a further subsequence, say, n_{k_m} along which $F_{n_{k_m}}$ converges, and by the assumption it must converge to F. But this is impossible, since $|F_{n_{k_m}}(x_0) - F(x_0)| \geq \varepsilon$. That finishes the proof. □

11 Continuity Theorem

The following result is a major tool in establishing the convergence in distribution

Theorem 11.1 *(continuity theorem) Let $X, X_1, X_2, \ldots$ be random variables with characteristic functions $\phi(t), \phi_1(t), \phi_2(t), \ldots$, respectively. The following two conditions are equivalent*

(i) $X_n \Longrightarrow X$, as $n \to \infty$,

(ii) for all $t \in \mathbf{R}$, $\phi_n(t) \longrightarrow \phi(t)$ as $n \to \infty$.

Proof of (i)$\Longrightarrow$ (ii). This follows from the Theorem 10.2 applied to the functions $x \to \Re(e^{itx})$, and $x \to \Im(e^{itx})$ (both continuous and bounded for every t).
Proof of (i)$\Longrightarrow$ (ii). Let F_n be the c.d.f of X_n. We first show that F_n, $n \geq 1$ is tight. We have

$$\begin{aligned}
\frac{1}{u}\int_{-u}^{u}(1-\phi_n(t))dt &= \frac{1}{u}\int_{-u}^{u}(1-\mathbf{E}e^{itX_n})dt \\
&= \mathbf{E}\left(\frac{1}{u}\int_{-u}^{u}(1-e^{itX_n}\right)dt.
\end{aligned}$$

To compute the inner integral, write $e^{itX_n} = \cos(tX_n) + i\sin(tX_n)$, use the fact that cosine is even and sine is odd, to get

$$\frac{1}{u}\int_{-u}^{u}(1-e^{itX_n})dt = 2 - 2\frac{1}{u}\int_{0}^{u}\cos(tX_n)dt = 2 - 2\frac{\sin(uX_n)}{uX_n}.$$

Therefore,

$$\begin{aligned}
\mathbf{E}\left(\frac{1}{u}\int_{-u}^{u}(1-e^{itX_n})\right)dt &= 2\mathbf{E}(1-\frac{\sin(uX_n)}{uX_n}) \\
&\geq 2\mathbf{E}(1-\frac{\sin(uX_n)}{uX_n})I(|X_n| \geq 2/u) \\
&\geq 2\mathbf{E}(1-\frac{1}{2})I(|X_n| \geq 2/u) = \mathbf{P}(|X_n| \geq \frac{2}{u}) \\
&= \mathbf{P}(X_n \geq \frac{2}{u}) + \mathbf{P}(X_n \leq -\frac{2}{u}) \\
&\geq 1 - F_n(\frac{2}{u}) + F_n(-\frac{2}{u}) \\
&= 1 - \left(F_n(\frac{2}{u}) - F_n(-\frac{2}{u})\right)
\end{aligned}$$

Now $\phi(0) = 1$ and ϕ is continuous at zero, which follows from

$$\begin{aligned}
|\phi(h) - \phi(0)| &= |\mathbf{E}(e^{itX}-1)| \leq \mathbf{E}|e^{itX}-1| \leq \mathbf{E}\min\{|hX|, 2\} \\
&= \mathbf{E}\min\{|hX|, 2\}I(|X| \leq 1/\sqrt{h}) \\
&\quad +\mathbf{E}\min\{|hX|, 2\}I(|X| > 1/\sqrt{h}) \\
&\leq \sqrt{h} + 2\mathbf{P}(|X| > 1/\sqrt{h}) \longrightarrow 0
\end{aligned}$$

as $h \to 0$. Therefore, for any $\varepsilon > 0$ there exists a $u > 0$ such that

$$\frac{1}{u}\int_{-u}^{u}(1-\phi(t))dt < \varepsilon.$$

Since $\phi_n(t) \to \phi(t)$ for all $t \in \mathbf{R}$, and the functions $|1-\phi_n(t)|$ are bounded by 2, it follows that

$$\int_{-u}^{u}(1-\phi_n(t))dt \longrightarrow \int_{-u}^{u}(1-\phi(t))dt < \varepsilon,$$

and hence we will have

$$\frac{1}{u}\int_{-u}^{u}(1-\phi_n(t))dt < \varepsilon,$$

for $n \geq n_0$. Combining this with the earlier computation we see that, for $n \geq n_0$ we have

$$F_n(\frac{2}{u}) - F_n(-\frac{2}{u}) \geq 1-\varepsilon,$$

for $n \geq n_0$. Decreasing u if necessary we may assume that the above inequality holds for n less than n_0 as well so that it holds for all $n \geq 1$. And, since the above inequality and the fact that $0 \leq F_n(x) \leq 1$ imply that

$$F_n(\frac{2}{u}) \geq 1-\varepsilon, \text{ and } F_n(-\frac{2}{u}) \leq \varepsilon,$$

the tightness is proved. Now we can use the first part and the previous corollary to complete the proof: pick any subsequence n_k such that F_{n_k} converges in distribution. Then by the first part $\phi_{n_k}(t)$ converge to $\phi(t)$ and by the uniqueness theorem we must have that $F_{n_k} \Longrightarrow F$. Hence, by the corollary $F_n \Longrightarrow F$. $\square$

12 Lindeberg's Central Limit Theorem

Let for each $n \geq 1$ $X_{n,1}, X_{n,2}, \ldots, X_{n,r_n}$ be independent random variables. Put

$$S_n = X_{n,1} + X_{n,2} + \cdots + X_{n,r_n},$$

and suppose that

$$\mathbf{E}X_{n,k} = 0, \quad \operatorname{var}(X_{n,k}) = \sigma_{n,k}^2, \quad \text{and set } s_n^2 = \sum_{k=1}^{r_n}\sigma_{n,k}^2.$$

Finally suppose that the following condition (usually referred to as *Lindeberg condition*) is satisfied:

$$\forall \varepsilon > 0 \quad \lim_{n\to\infty}\sum_{k=1}^{r_n}\frac{1}{s_n^2}\mathbf{E}X_{n,k}^2 I(|X_{n,k}| \geq \varepsilon s_n) = 0.$$

Then we have:

Theorem 12.1 *Under the above assumptions:*

$$\frac{S_n}{s_n} \Longrightarrow N(0,1),$$

where $N(0,1)$ *denotes the standard normal (i.e. with mean zero and variance 1) random variable.*

Note:

(i) the assumption $\mathbf{E}X_{n,k} = 0$ is inessential since we can always replace $X_{n,k}$ by $X_{n,k} - \mathbf{E}X_{n,k}$.

(ii) the most common application is in the case when $X_1, X_2, \ldots$ is a one sequence of random variables and S_n is a sequence of partial sums, i.e. $X_{nk} = X_k$ and $r_n = n$.

Proof: The plan is to use the continuity theorem. Let's begin with a simple lemma:

Lemma 12.2 *Let* $z_1, \ldots, z_m$ *and* $w_1, \ldots, w_m$ *be complex numbers such that* $|z_j| \le 1$ *and* $|w_j| \le 1$ *for* $j = 1, \ldots, m$. *Then*

$$|z_1 \cdot z_2 \cdots z_m - w_1 \cdot w_2 \ldots w_m| \le \sum_{j=1}^{m} |z_j - w_j|.$$

Proof of Lemma 12.2: The proof is by induction over m. For $m = 1$ the statement is obvious. Assuming it's true for all $1 \le k \le m-1$ write

$$\begin{aligned}
&|z_1 \cdot z_2 \cdots z_m - w_1 \cdot w_2 \ldots w_m| \\
&\quad = |z_1 \cdot z_2 \cdots z_m - z_1 \cdot w_2 \ldots w_m + z_1 \cdot w_2 \ldots w_m - w_1 \cdot w_2 \ldots w_m| \\
&\quad \le |z_1 (z_2 \cdots z_m - w_2 \cdots w_m)| + |(z_1 - w_1) \cdot w_2 \cdots w_m| \\
&\quad \le |z_1| \sum_{j=2}^{m} |z_j - w_j| + |z_1 - w_1| \cdot |w_2| \cdots |w_m| \le \sum_{j=1}^{m} |z_j - w_j|,
\end{aligned}$$

completing the proof of the lemma.

Coming back to the proof of the theorem, let us assume without loss of generality that $s_n = 1$, and let $\phi_{n,k}$ and ϕ be the characteristic functions of $X_{n,k}$, X_n, respectively. By the continuity theorem we need to show that

$$\forall\, t \in \mathbf{R} \quad \phi_n(t) \longrightarrow e^{-t^2/2}.$$

Since we assume that

$$s_n^2 = \sum_{k=1}^{r_n} \sigma_{n,k}^2 = 1$$

it follows that

$$e^{-t^2/2} = e^{-\frac{t^2}{2}\sum_{k=1}^{r_n}\sigma_{n,k}^2} = \prod_{k=1}^{r_n} e^{-\frac{t^2\sigma_{n,k}^2}{2}}.$$

Therefore, by Lemma 12.2

$$\begin{aligned}
|\phi_n(t) - e^{-\frac{t^2}{2}}| &= \left|\prod_{k=1}^{r_n}\phi_{n,k}(t) - \prod_{k=1}^{r_n} e^{-\frac{t^2\sigma_{n,k}^2}{2}}\right| \le \sum_{k=1}^{r_n}\left|\phi_{n,k}(t) - e^{-\frac{t^2\sigma_{n,k}^2}{2}}\right| \\
&= \sum_{k=1}^{r_n}\left|\mathbf{E}e^{itX_{n,k}} - e^{-\frac{t^2\sigma_{n,k}^2}{2}}\right| \\
&= \sum_{k=1}^{r_n}\left|\mathbf{E}\left(1 + itX_{n,k} - \frac{t^2X_{n,k}^2}{2} - e^{-\frac{t^2\sigma_{n,k}^2}{2}} + e^{itX_{n,k}} - 1 - itX_{n,k} + \frac{t^2X_{n,k}^2}{2}\right)\right| \\
&\le \sum_{k=1}^{r_n}\left|\mathbf{E}\left(1 + itX_{n,k} - \frac{t^2X_{n,k}^2}{2}\right) - e^{-\frac{t^2\sigma_{n,k}^2}{2}}\right| \\
&\quad + \sum_{k=1}^{r_n}\left|\mathbf{E}\left(e^{itX_{n,k}} - 1 - itX_{n,k} + \frac{t^2X_{n,k}^2}{2}\right)\right| \\
&= \sum_{k=1}^{r_n}\left|1 - \frac{t^2\sigma_{n,k}^2}{2} - e^{-\frac{t^2\sigma_{n,k}^2}{2}}\right| + \sum_{k=1}^{r_n}\mathbf{E}\left|e^{itX_{n,k}} - 1 - itX_{n,k} + \frac{t^2X_{n,k}^2}{2}\right|
\end{aligned}$$

We will show that each of the two sums goes to zero. For the first sum, we use the fact that for every real (or complex) number z we have

$$|e^z - 1 - z| \le \sum_{j\ge 2}\frac{|z|^j}{j!} = |z|^2\sum_{j=0}^{\infty}\frac{|z|^j}{(j+2)!} \le |z|^2\sum_{j=0}^{\infty}\frac{|z|^j}{j!} = |z|^2 e^{|z|}.$$

Applying the above to each of the terms and using the fact that $\sigma_{n,k}^2 \le s_n^2 = 1$ we get that the sum is bounded by

$$\sum_{k=1}^{r_n}\frac{t^4\sigma_{n,k}^4}{4}e^{\frac{t^2\sigma_{n,k}^2}{2}} \le \frac{t^4}{4}e^{\frac{t^2}{2}}\sum_{k=1}^{r_n}\sigma_{n,k}^4.$$

Lindeberg condition (recall that $s_n^2 = 1$) implies that

$$\max_{1\le k\le r_n}\sigma_{n,k}^2 \longrightarrow 0, \quad \text{as } n\to\infty.$$

Indeed, for every $\varepsilon > 0$ and every $1 \le k \le r_n$

$$\begin{aligned}
\sigma_{n,k}^2 &= \mathbf{E}X_{n,k}^2 = \mathbf{E}X_{n,k}^2 I(|X_{n,k}| < \varepsilon) + \mathbf{E}X_{n,k}^2 I(|X_{n,k}| \ge \varepsilon) \\
&\le \varepsilon + \sum_{k=1}^{r_n}\mathbf{E}X_{n,k}^2 I(|X_{n,k}| \ge \varepsilon).
\end{aligned}$$

By the Lindeberg condition, the second term goes to zero, and since ε is arbitrary, the whole expression vanishes as $n \to \infty$. Hence,

$$\sum_{k=1}^{r_n} \sigma_{n,k}^4 \leq \max_{1 \leq k \leq r_n} \sigma_{n,k}^2 \sum_{k=1}^{r_n} \sigma_{n,k}^2 = \max_{1 \leq k \leq r_n} \sigma_{n,k}^2 \longrightarrow 0, \quad \text{as } n \to \infty,$$

which implies that the first sum goes to zero for every real t. For the second sum, we use the following estimate (see the proof of Theorem 9.1):

$$|e^{ix} - (1 + ix - \frac{1}{2}x^2)| \leq \min\left\{\frac{|x|^3}{3!}, x^2\right\}.$$

Then

$$\begin{aligned}
\mathbf{E}|e^{itX_{n,k}} - 1 - itX_{n,k} + \frac{t^2 X_{n,k}^2}{2}| &\leq \mathbf{E}\min\left\{\frac{|tX_{n,k}|^3}{3!}, t^2 X_{n,k}^2\right\} \\
&\leq \mathbf{E}\min\left\{\frac{|tX_{n,k}|^3}{3!}, t^2 X_{n,k}^2\right\} I(|X_{n,k}| < \varepsilon) \\
&\quad + \mathbf{E}\min\left\{\frac{|tX_{n,k}|^3}{3!}, t^2 X_{n,k}^2\right\} I(|X_{n,k}| \geq \varepsilon) \\
&\leq \mathbf{E}\frac{|t|^3 \varepsilon}{3!} X_{n,k}^2 I(|X_{n,k}| < \varepsilon) + t^2 \mathbf{E} X_{n,k}^2 I(|X_{n,k}| \geq \varepsilon) \\
&\leq \frac{|t|^3 \varepsilon}{3!} \sigma_{n,k}^2 + t^2 \mathbf{E} X_{n,k}^2 I(|X_{n,k}| \geq \varepsilon),
\end{aligned}$$

and by adding the terms we see that the second sum is bounded by

$$\frac{\varepsilon |t|^3}{3} \sum_{k=1}^{r_n} \sigma_{n,k}^2 + t^2 \sum_{k=1}^{r_n} \mathbf{E} X_{n,k}^2 I(|X_{n,k}| \geq \varepsilon),$$

which goes to zero because $\varepsilon > 0$ is arbitrary and because of Lindeberg's condition. The proof is completed. $\square$

The following is a useful corollary to the Lindeberg's CLT. The condition below (referred to as *Lyapunov condition*) assumes that the higher moments than the second exist, but is usually easier to check than Lindeberg condition.

Corollary 12.3 *Let $X_{n,k}$, $n \geq 1$, $1 \leq k \leq r_n$, be random variables such that $\mathbf{E}X_{n,k} = 0$ and $\mathrm{var}(X_{n,k}) = \sigma_{n,k}^2$, and set $s_n^2 = \sum_{k=1}^{r_n} \sigma_{n,k}^2$. Suppose further that the following condition is satisfied for some $p > 2$:*

$$\lim_{n \to \infty} \frac{1}{s_n^p} \sum_{k=1}^{r_n} \mathbf{E}|X_{n,k}|^p = 0.$$

Then

$$\frac{S_n}{s_n} \Longrightarrow N(0, 1).$$

Proof: We have

$$X_{n,k}^2 I(|X_{n,k}| \geq \varepsilon s_n) = 1 \cdot X_{n,k}^2 I\left(\frac{|X_{n,k}|^{p-2}}{\varepsilon^{p-2} s_n^{p-2}} \geq 1\right) \leq \frac{|X_{n,k}|^p}{\varepsilon^{p-2} s_n^{p-2}},$$

so that, summing up over k gives

$$\frac{1}{s_n^2} \sum_{k=1}^{r_n} \mathbf{E} X_{n,k}^2 I(|X_{n,k}| \geq \varepsilon) \leq \frac{1}{\varepsilon^{p-2}} \frac{1}{s_n^p} \sum_{k=1}^{r_n} \mathbf{E}|X_{n,k}|^p,$$

which shows that Lyapunov condition implies Lindeberg's condition. □

13 Random Permutations

Let $\pi = (\pi_1, \pi_2, \ldots, \pi_n)$ be a permutation i.e. an order of $\{1, 2, \ldots, n\}$. There are $n!$ different permutations of $\{1, 2, \ldots, n\}$. By a *random permutation* of $\{1, 2, \ldots, n\}$ we mean a permutation chosen according to the discrete uniform probability measure on the set Π_n of all permutations of $\{1, 2, \ldots, n\}$.

Let $\pi = (\pi_1, \pi_2, \ldots, \pi_n)$ be a permutation of $\{1, 2, \ldots, n\}$. We define a sequence of *sequential ranks*, $sq(\pi_i), i = 1, \ldots, n$ as follows: $sq(\pi_i) = j$ if and only if π_i is the jth smallest element among $\pi_1, \pi_2, \ldots, \pi_i$. For example, if $\pi = (3, 5, 2, 1, 4)$ is a permutation of $\{1, 2, 3, 4, 5\}$, then the sequence of sequential ranks is $sq(\pi) = (1, 2, 1, 1, 4)$.

Lemma 13.1 *There is a one-to-one correspondence between permutations of $\{1, 2, \ldots, n\}$ and n-long feasible sequences of sequential ranks.*

Proof. We first observe that every permutation has a unique sequence of sequential ranks. Thus, it is enough to show that there is at most one permutation corresponding to a sequence of sequential ranks. We will prove it by induction on n. For $n = 1$ the statement is clear since there is only one sequence of sequential ranks and only one permutation. Assuming the statement true for $1 \leq k \leq n-1$ consider a sequence $(s_1, \ldots, s_n)$ of sequential ranks and suppose that $\pi = (\pi_1, \ldots, \pi_n)$ and $\sigma = (\sigma_1, \ldots, \sigma_n)$ are two permutations corresponding to $(s_1, \ldots, s_n)$. Find a position of the letter n in both π and σ; specifically assume $\pi_i = n = \sigma_j$. We claim that necessarily $i = j$. If not, then suppose that $i < j$. But $\pi_i = n$ and $i < j$ means that π_j is not the largest among $\pi_1, \ldots, \pi_j$ and thus $sq(\pi_j) = s_j < j$. On the other hand, since $\sigma_j = n$ we have $sq(\sigma_j) = s_j = j$, a contradiction.

Having proved that the letter n is at the same position in both π and σ we can now remove it. What is left are two permutations $\tilde{\pi}$ and $\tilde{\sigma}$ of $\{1, 2, \ldots, n-1\}$ with sequence of sequential ranks $(s_1, s_2, \ldots, s_{i-1}, s_{i+1}, \ldots, s_n)$. By inductive hypothesis we must have $\tilde{\pi} = \tilde{\sigma}$ and inserting back n on its original place we see that $\pi = \sigma$, thus completing the proof. □

Theorem 13.2 *Let $\pi = (\pi_1, \dots, \pi_n)$ be a random permutation of $\{1, 2, \dots, n\}$. Let $Y_j = sq(\pi_j)$, $j = 1, \dots, n$ be a sequence of sequential ranks. Then $Y_1, \dots, Y_n$ are independent and Y_j is a discrete uniform random variable on $\{1, \dots, j\}$.*

Proof. Clearly, the values of Y_j are in $\{1, \dots, j\}$. To show that Y_j is uniform pick a k such that $1 \le k \le j$. Then, by the law of total probability

$$\begin{aligned}\mathbf{P}(Y_j = k) &= \sum_{m=k}^{n} \mathbf{P}\left(Y_j = k,\ \pi_j = m\right) \\ &= \sum_{m=k}^{n} \mathbf{P}\left(Y_j = k \mid \pi_j = m\right) \cdot \mathbf{P}(\pi_j = m)\end{aligned}$$

To compute the conditional probability notice that, given that $\pi_j = m$, $Y_j = k$ means that exactly $k-1$ numbers less than m are on the first $j-1$ positions. There are $\binom{m-1}{k-1}$ choices of $k-1$ numbers less than m and there are $\binom{j-1}{k-1} \cdot (k-1)!$ ways of arranging those numbers on the first $(j-1)$ positions. Next, we fill the remaining $j-k$ positions among the first $(j-1)$ with numbers larger than m. This can be done in $\binom{n-m}{j-k} \cdot (j-k)!$ ways. Finally, we arrange the remaining $(n-j)$ numbers on the $(n-j)$ unoccupied positions. There are $(n-j)!$ ways to do that. Now using the fact that given $\pi_j = m$, $(\pi_1, \dots, \pi_{j-1}, \pi_{j+1}, \dots, \pi_n)$ is a random permutation of $\{1, \dots, n\} \setminus \{m\}$, and putting the pieces together we see that our conditional probability is

$$\mathbf{P}\left(Y_j = k \mid \pi_j = m\right) = \frac{1}{(n-1)!} \binom{m-1}{k-1} \cdot \binom{j-1}{k-1} \cdot (k-1)! \binom{n-m}{j-k} \cdot (j-k)!(n-j)!$$

which, after using

$$\binom{j-1}{k-1} = \frac{(j-1)!}{(k-1)!(j-k)!}$$

and cancellation, becomes

$$\mathbf{P}\left(Y_j = k \mid \pi_j = m\right) = \frac{(j-1)!(n-j)!}{(n-1)!} \binom{m-1}{k-1} \cdot \binom{n-m}{j-k}.$$

Hence,

$$\begin{aligned}\mathbf{P}(Y_j = k) &= \frac{1}{n} \sum_{m=k}^{n} \frac{(j-1)!(n-j)!}{(n-1)!} \binom{m-1}{k-1} \cdot \binom{n-m}{j-k} \\ &= \frac{(j-1)!(n-j)!}{n!} \sum_{m=k}^{n} \binom{m-1}{k-1} \cdot \binom{n-m}{j-k} \\ &= \frac{1}{j} \cdot \frac{j!(n-j)!}{n!} \sum_{m=k}^{n} \binom{m-1}{k-1} \cdot \binom{n-m}{j-k} \\ &= \frac{1}{j \cdot \binom{n}{j}} \sum_{m=k}^{n} \binom{m-1}{k-1} \cdot \binom{n-m}{j-k}.\end{aligned}$$

and the proof of this part will be complete once we realize that

$$\sum_{m=k}^{n} \binom{m-1}{k-1} \cdot \binom{n-m}{j-k} = \binom{n}{j}.$$

But this is true since the left-hand side counts the number of j-element subsets of $\{1, \ldots, n\}$ by first picking an integer m between k and n and then choosing $k-1$ integers less than m and $j-k$ integers larger than m. Thus, we have proved that Y_j is uniform on $\{1, \ldots, j\}$ for $j = 1, \ldots, n$. Given that, a proof that they are independent is easy. Let $(s_1, \ldots, s_n)$ be a feasible sequence of sequential ranks. Then, since there is only one permutation with that sequence of sequential ranks, we must have

$$\mathbf{P}\left(\bigcap_{j=1}^{n} \{Y_j = s_j\}\right) = \frac{1}{n!}.$$

On the other hand, since Y_j is uniform on $\{1, \ldots, j\}$,

$$\mathbf{P}(Y_j = s_j) = \frac{1}{j}.$$

Hence,

$$\prod_{j=1}^{n} \mathbf{P}(Y_j = s_j) = \prod_{j=1}^{n} \frac{1}{j} = \frac{1}{n!} = \mathbf{P}\left(\bigcap_{j=1}^{n} \{Y_j = s_j\}\right),$$

which shows independence. $\square$

The last theorem has a number of useful consequences two of which are given below. Let $\pi = (\pi_1, \pi_2, \ldots, \pi_n)$ be a permutation. We say that π_k is a *record* if $k = 1$ or $\pi_k > \pi_j$ for all $1 \leq j < k$. A pair (π_i, π_j) is an *inversion* if $i < j$ and $\pi_i > \pi_j$. Clearly, a permutation of n letters can have up to n records and $\binom{n}{2}$ inversions. Let $R_n(\pi)$ be the total number of records in π. For example, if $\pi = (3, 1, 2, 4, 6, 5)$ then $R_6(\pi) = 2$, namely, 4, and 6 are the records. Similarly if $V_n(\pi)$ denotes the total number of inversions then in this example we would have $V_6(\pi) = 3$ since $(3, 1)$, $(3, 2)$, and $(6, 5)$ are the only inversions. Suppose now that π is a random permutation. In that case R_n may be viewed as a random variable on the probability space of all permutations of n letters equipped with the discrete uniform probability measure and the same applies to V_n. What can we say about the distribution of these two random variables?. The following gives a rather complete answer.

Theorem 13.3 *Let π be a random permutation of n distinct letters. Then, as $n \to \infty$, we have*

(i) the total number of records, R_n satisfies

$$\frac{R_n - \ln n}{\sqrt{\ln n}} \Longrightarrow N(0, 1).$$

(ii) the total number of inversions, V_n satisfies

$$\frac{V_n - n^2/4}{n^{3/2}/6} \Longrightarrow N(0,1).$$

Proof of (i): The argument rests on the following simple observation; π_k is a record if and only if $sq(\pi_k) = k$. Hence,

$$R_n = \sum_{k=1}^{n} I(sq(\pi_k) = k) \stackrel{def}{=} \sum_{k=1}^{n} I_k,$$

where, by the previous result, I_k's are independent and $\mathbf{P}(I_k) = 1/k$. It follows that

$$\mathbf{E}R_n = \sum_{k=1}^{n} \frac{1}{k} \sim \ln n, \quad \text{and} \quad \mathrm{var}(R_n) = \sum_{k=1}^{n} \mathrm{var}(I_k) = \sum_{k=1}^{n} \frac{1}{k}\left(1 - \frac{1}{k}\right) \sim \ln n,$$

and the conditions for the CLT hold since the I_k's are bounded.
Proof of (ii): The proof follows the same idea, namely we will link inversions to sequential ranks. Specifically, we have

$$V_n = \sum_{k=1}^{n} W_k,$$

where W_k is the number of $\pi_1, \pi_2, \ldots, \pi_k$ that are larger than π_k. Clearly, $W_k = k - sq(\pi_k)$. Thus, W_k is uniform on $\{1, \ldots, k-1\}$ and they are independent. Therefore,

$$\mathbf{E}V_n = \sum_{k=1}^{n} \mathbf{E}W_k = \sum_{k=1}^{n} \frac{1}{k} \frac{(k-1)k}{2} = \frac{(n-1)n}{4}.$$

Similarly,

$$\mathrm{var}(V_n) = \sum_{k=1}^{n} \mathrm{var}(W_k) = \sum_{k=1}^{n} \frac{(k-1)(k+1)}{12} \sim \frac{n^3}{36}.$$

Finally,

$$\mathbf{E}W_k^3 = \frac{1}{k} \sum_{j=1}^{k-1} j^3 \le \frac{1}{k} \int_0^k x^3 dx = O\left(k^3\right),$$

which implies Lyapunov's sufficient condition for the CLT. □

14 Trees

A *tree* is a hierarchical structure of *node*s arranged in levels. The very top (we tend to picture trees upside down) level has only one node called a *root*. Immediate descendants of any node (linked to it by *edges*) are called the *children* of that node and a node directly above a given node is referred to as a *parent* or a *father*. A node without children is called a *leaf* or an *external* node; any other node is an *internal* node. The *depth*, d_k, of a node k is the distance (counted in terms of levels) between the root and that node. The height h_n of a tree on n nodes is the maximal depth of a node, i.e. $h_n = \max\{d_k : \ k \text{ is a node}\}$ (clearly it is enough to consider leaves only).

We will be exclusively interested in *binary* trees, i.e. trees with the property that every node can have at most two children. In such a tree, there can be at most 2^ℓ nodes on any level ℓ and if this is the case for a particular level we say that that level is *saturated*. If all levels except possibly the last one are saturated, the tree is called *complete*, and if the last level is saturated as well the tree is called *perfect*. It will be convenient to consider a notion of an *extended* binary tree. If T is a binary tree then its extension $\overline{T}$ is obtained by adding nodes so that every original node of T has exactly two children. We will use the symbol $\overline{T}_n$ to indicate that a given extended tree $\overline{T}$ is an extension of a tree on n nodes. Thus, the subscript n refers to the size of the original tree. Also, $\overline{h}_n$ will denote the height of an extension of a tree T on n nodes. Figure 1 depicts an extended, complete binary tree on 6 nodes.

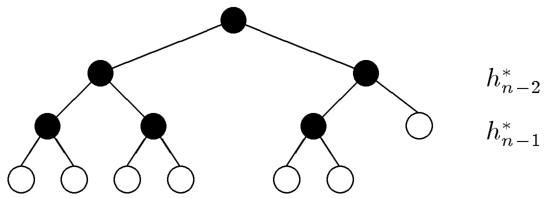

Figure 1: Extended complete binary tree on 6 nodes

We now make a few observations about binary trees that will be useful later:

Lemma 14.1 *Let T_n be a binary tree on n nodes. Then, the height of its extension satisfies*

$$\lceil \lg(n+1) \rceil \le \overline{h}_n \le n,$$

where $\lceil x \rceil$ is the smallest integer no less than x, and lg *denotes the base 2 logarithm.*

Proof: The upper bound is trivial since a height of any tree on n nodes is at most $n-1$ and extending it increases the height by 1. For the lower bound we first note $\overline{h}_n$ is at least as large as h^*_n, the height of an extension of a complete tree on n nodes, so it is enough to show it for such a tree. Since all the levels $0, 1, \ldots, h^*_{n-2}$ are saturated (these are all but the last level of the original, non extended tree), h^*_n is the smallest integer satisfying

$$1 + 2^1 + \cdots + 2^{h^*_n - 1} \ge n,$$

i.e. $2^{h_n^*} - 1 \geq n$, which gives $h_n^* \geq \lg(n+1)$, and since h_n^* is an integer the inequality is also true with the ceiling. $\square$

Lemma 14.2 *For any tree T_n on n nodes $\overline{T}_n$ has exactly $n+1$ nodes leaves.*

Proof: Let ℓ_n be the number of leaves of $\overline{T}_n$. Then $\overline{T}_n$ has exactly $n + \ell_n$ nodes. We will count all the children in $\overline{T}_n$ in two different ways. On one hand, every node except the root is a child; thus there are $n + \ell_n - 1$ children. On the other hand, every internal node of $\overline{T}$ has exactly two children. Since there are exactly n nodes like that, there must be $2n$ children. Thus

$$2n = n + \ell_n - 1,$$

which gives the lemma. $\square$

14.1 Decision Trees

A decision tree is a tree in which internal nodes represent queries and leaves represent outcomes.

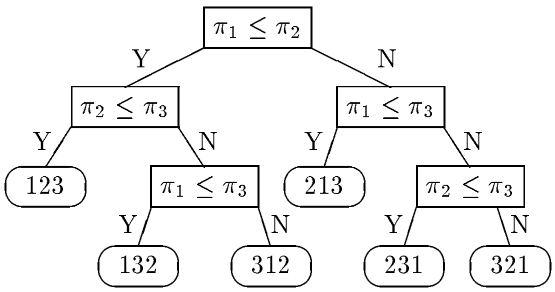

Figure 2: Decision tree for an algorithm sorting three elements

Here is an example of a decision tree which represents an algorithm that sorts a permutation (π_1, π_2, π_3) of $\{1, 2, 3\}$ (see Figure 2). It begins by inquiring if $\pi_1 \leq \pi_2$ and proceeds to compare π_2 with π_3 or π_1 with π_3 according to whether the answer to the first query is "yes" or "no". It then decides what the order is, or continues to the next query. We note that this decision tree is an extended tree (namely, it is an extension of the tree of queries), and that leaves represent all possible permutations of $\{1, 2, 3\}$. Thus there are $6 = 3!$ leaves. The depth of a given leaf is exactly equal to the number of queries needed to find out the permutation represented by that leaf. Thus, the height of this tree is the the largest possible number of queries needed to find any permutation. This is usually called the worst case performance of an algorithm. In this example, 3 is the largest possible number of comparisons needed to find out which of the 6 permutations is π_1, π_2, π_3. Let us now look at a general situation. All of these observations remain true: a decision tree for a particular algorithm would be an extension of a tree of queries, and its height would be the worst case performance of that algorithm. Now comes a key observation: no matter what algorithm is used, the decision tree will have $n!$ leaves (each corresponding to a particular permutation). Thus, by the previous lemma the query tree (of which our decision tree is extension) has $n! - 1$ nodes so its height is at least $\lceil \lg(n! - 1 + 1) \rceil \geq \lg(n!)$. This gives us the following lower bound on any sorting algorithm based on comparisons.

Theorem 14.3 *The worst case performance of any comparison based sorting algorithm is at least $n \lg n - O(n)$ on an input of size n.*

Proof: At this point the proof is no more than a simple calculation. Since the function $\lg x$ is increasing we have

$$\lg(n!) = \sum_{k=1}^{n} \lg n \geq \int_1^{n-1} \lg x dx = (n-1)\lg(n-1) - O(n) = n \lg n - O(n).$$

□

Thus we cannot hope to sort a permutation of $\{1, \ldots, n\}$ any faster than in $n \lg n$ comparisons in the worst case. However, the worst case behavior, although informative does not tell the whole story. We will therefore consider the *average case* behavior. What is meant by that is the following: let a comparison based sorting algorithm be given. Let $C_n(\pi)$ be the number of comparisons required to sort a permutation π. Now, suppose that π is a permutation randomly chosen from the set Π_n of all permutations of $\{1, 2 \ldots, n\}$. Viewed like that, C_n is a random variable defined on the probability space Π_n equipped with the discrete uniform probability measure $\mathbf{P}$, and the average case behavior of our algorithm is simply the expected value $\mathbf{E}C_n$ of C_n. We have

Theorem 14.4 *The average case behavior of any comparison based sorting algorithm is at least $n \lg n - O(n)$.*

Proof: Consider a decision tree of a given sorting algorithm. Label leaves by permutations; then the depth of a given π is the number of comparisons required to sort π. Since the measure $\mathbf{P}$ is uniform on Π_n we get

$$\mathbf{E}C_n = \sum_{\pi \in \Pi_n} d_\pi \mathbf{P}(\pi) = \frac{1}{n!} \sum_{\pi \in \Pi_n} d_\pi,$$

so that the proof will be complete once we show that

$$\sum_{\pi \in \Pi_n} d_\pi \geq n! \lg n!.$$

But this is a general fact, that is true for any extended tree $\overline{T}_k$. As a matter of fact, for a tree T the quantity $\sum d_\ell$, where the sum extends over all leaves of T (that is exactly what we have on the left side) is referred to as the *external path length* and is usually denoted by $X(T)$. We claim

Lemma 14.5 *For any tree T_k on k nodes we have*

$$X(\overline{T}_k) \geq (k+1)\lg(k+1).$$

Proof of Lemma 14.5. We first note that $X(\overline{T}_k) \geq X(T_k^*)$, where T_k^* is a an extension of a complete tree on k nodes. This is because if T_k is not complete then there must be a level with index smaller than h_k that is not saturated. Moving now a node from the bottom level h_k to that level will decrease the external path length in the extended tree. This process can be continued (each time decreasing the external path length) until all the levels, except possibly the last one, are saturated.

It remains to find $X(T_k^*)$. Let h_k^* be the height of T_k^* and L_k be the number of leaves on the h_k^*th level. The leaves are on the last two levels and there are $k+1$ of them altogether. Thus,

$$X(T_k^*) = h_k^* L_k + (h_k^* - 1)(k+1-L_k) = (h_k^* - 1)(k+1) + L_k.$$

But L_k is twice the number of internal nodes on the $(h_k^* - 1)$th level and since all the earlier levels are saturated, it is

$$L_k = 2\left(k - \sum_{j=0}^{h_k^*-2} 2^j\right) = 2\left(k - (2^{h_k^*-1} - 1)\right) = 2(k+1) - 2^{h_k^*}.$$

We now use $h_k^* = \lceil \lg(k+1) \rceil$ to get

$$\begin{aligned} X(T_k^*) &= (\lceil \lg(k+1) \rceil - 1)(k+1) + 2(k+1) - 2^{\lceil \lg(k+1) \rceil} \\ &= (k+1)(\lceil \lg(k+1) \rceil + 1) - 2^{\lceil \lg(k+1) \rceil} \\ &= (k+1)(\lg(k+1) + 1 + x) - 2^{\lg(k+1)+x} \\ &= (k+1)\lg(k+1) + (k+1)(1 + x - 2^x), \end{aligned}$$

where we have set $x = \lceil \lg(k+1) \rceil - \lg(k+1)$. It is now enough to check that $f(x) = 1 + x - 2^x$ is nonnegative on $[0,1]$ which is true since $f(0) = f(1) = 0$ and f is concave on $[0,1]$. $\square$

15 Insertion Sort

Insertion sort algorithm is the simplest and most direct sorting algorithm. The output will be given as a list, $A[1..n]$. We begin with an empty list and at the kth step we will have an already sorted list $A[1..k-1]$ and we will insert the next key,

π_k into its the proper position in that list, creating an ordered list $A[1..k]$. At each stage we will have an *insertion strategy* indicating successive comparisons we are going to make to insert a key. These are conveniently represented by insertion trees; for example a tree in Figure 3 represents a strategy to insert π_7 that consists on first comparing it with $A[4]$ and then, depending on whether π_7 is smaller than $A[4]$ or not, comparing it with $A[3]$ or $A[6]$, and continuing until a proper position is found. In principle, one can follow very different strategies to insert different keys; in practice one does the same type of insertion at every stage. We will insist, however that strategies to insert keys are deterministic. Let T_i be an insertion tree for the ith key and let h_i be the height of T_i. Let X_i be the number of comparisons required at the ith stage. Then, the total number of comparisons, C_n is simply the sum of X_i's,

$$C_n = \sum_{k=1}^{n} X_k,$$

and by the properties of random permutations $X_1, \ldots, X_n$ are independent. The following theorem is a fairly general sufficient condition for the CLT

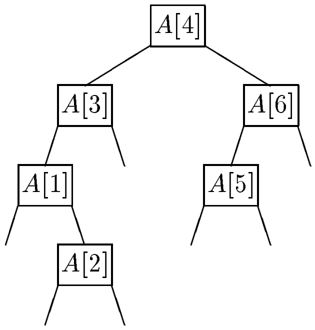

Figure 3: Insertion tree for $A[7]$

Theorem 15.1 *With the above notation, suppose that* $\max\limits_{1 \le k \le n} h_k = o(\sqrt{\mathrm{var}(C_n)})$. *Then, the central limit theorem holds:*

$$\frac{C_n - \mathbf{E}C_n}{\sqrt{\mathrm{var}(C_n)}} \Longrightarrow N(0,1).$$

Note: Our condition $\max\limits_{1 \le k \le n} h_k = o(\sqrt{\mathrm{var}(C_n)})$ is sometimes phrased as $h_n = o(\sqrt{\mathrm{var}(C_n)})$ under the assumption that the heights h_k are nondecreasing (which reasonable strategies would satisfy).

Proof: We will check Lyapunov's condition for the centered random variables $\overline{X}_k = X_k - \mathbf{E}X_k$. Since the number of comparisons to insert a given key is at most the height of a corresponding insertion tree, we have $X_k \leq h_k$; hence $|\overline{X}_k| \leq 2h_k$. Let $p > 2$ be any number. Then,

$$\mathbf{E}|\overline{X}_k|^p = \mathbf{E}|\overline{X}_k|^{p-2}\overline{X}_k^2 \leq (2h_k)^{p-2}\mathbf{E}\overline{X}_k^2.$$

Summing up we obtain

$$\sum_{k=1}^{n} \mathbf{E}|\overline{X}_k|^p \leq 2^{p-2}\left(\max_{1\leq k\leq n} h_k^{p-2}\right)\sum_{k=1}^{n}\mathbf{E}\overline{X}_k^2 = 2^{p-2}\left(\max_{1\leq k\leq n} h_k^{p-2}\right)\mathrm{var}(C_n),$$

in view of our assumption, Lyapunov's condition is evident. □

We will now see how it works on two examples.

15.1 Linear Insertion Sort

Linear sort corresponds to a very naive strategy, namely at the kth stage we start comparing π_k with $A[k-1]$, then (if necessary) with $A[k-2]$ and so on until we encounter an entry smaller than π_k or exhaust the list (perhaps the easiest way to avoid an innesential nuisance caused by exhausting the list is to start not with an empty list but with a list consisting of $A[0] = 0$ before the first insertion). The corresponding insertion tree is depicted in Figure 4.

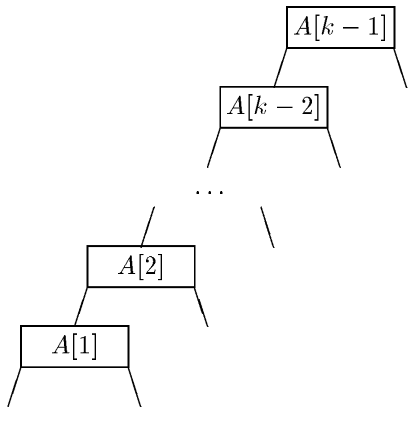

Figure 4: Insertion tree for linear sort

The number of comparisons needed is the number of elements in the list $A[1..k-1]$ that are larger than π_k, plus 1 (to encounter an element smaller than

π_k). In other words,

$$X_k = 1 + (k - sq(\pi_k)) = 1 + W_k,$$

where W_k is uniform on $\{0, \ldots, k-1\}$ as discussed in Theorem 13.3. But this means that X_k is uniform on $\{1, \ldots, k\}$. Hence, by exactly the same computations as in Theorem 13.3 we obtain that $\mathbf{E}C_n \sim n^2/4$ and $\mathrm{var}(C_n) \sim n^3/36$. Since the height of the kth insertion tree satisfies $h_k \leq k \leq n = o(n^{3/2})$ the CLT holds:

Theorem 15.2 *The number of comparisons C_n required to sort a random permutation of $\{1, \ldots, n\}$ satisfies*

$$\frac{C_n - n^2/4}{n^{3/2}/6} \Longrightarrow N(0,1).$$

This is clearly worse performance than the lower bound we know. Quite likely, it is due to a very poor inserting strategy; the fact that at the kth stage $A[1..k-1]$ was sorted was not fully taken advantage of. We will now try to improve on that.

15.2 Binary Insertion Sort

This algorithm as an insertion strategy uses a binary search tree. A binary search on an ordered list is a search that compares an element to be inserted to the "middle" element in a (sorted) list and then, recursively, proceed to the left or right "half" according to whether π_k was smaller or larger than the element it was compared to. Specifically, if a list of length m is to be searched the algorithm compares with the $\lceil m/2 \rceil$th element in that list. It follows that the insertion tree for the kth element, $k \geq 2$, is an extension of a complete binary tree on $k-1$ nodes, and thus has height $\lceil \lg(1 + (k-1)) \rceil$ and it has k leaves. Figure 5 shows an example of a binary insertion tree. Furthermore, the leaves are only on the last two levels $\lceil \lg k \rceil - 1$ and $\lceil \lg k \rceil$. It will be convenient to use the fact $\lceil \lg k \rceil = \lfloor \lg(k-1) \rfloor + 1$. Thus, the number of comparisons to insert π_k is either $\lfloor \lg(k-1) \rfloor$ or one more than that, depending if π_k falls into a leaf on the next to last or the last level. By the properties of random permutations, π_k is equally likely to fall into any of the k leaves. It follows that the number of comparisons X_k satisfies

$$X_k = \lfloor \lg(k-1) \rfloor + B_k,$$

where B_k is a Bernoulli random variable with parameter p_k and p_k is the proportion of leaves on the last level. In particular, $\mathbf{E}B_k = p_k$ and $\mathrm{var}(B_k) = p_k(1-p_k)$.

To compute p_k let L_k denote the number of leaves on the last level, so that $p_k = L_k/k$. This number is twice the number of nodes on the level $\lfloor \lg(k-1) \rfloor$ in the nonextended complete tree on $k-1$ nodes. Since all the previous levels

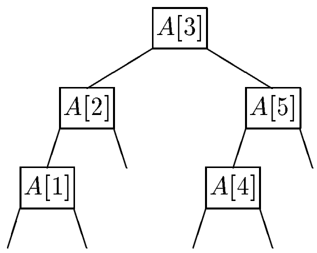

Figure 5: Insertion tree for a binary insertion sort

are saturated we count:

$$\begin{aligned} L_k &= 2\left((k-1) - \left(2^0 + 2^1 + \cdots + 2^{\lfloor \lg(k-1) \rfloor - 1}\right)\right) \\ &= 2\left((k-1) - \left(2^{\lfloor \lg(k-1) \rfloor} - 1\right)\right) \\ &= 2\left(k - 2^{\lfloor \lg(k-1) \rfloor}\right). \end{aligned}$$

Hence the expected value of the number of comparisons is

$$\begin{aligned} \mathbf{E}C_n &= \sum_{k=2}^{n} \mathbf{E}X_k = \sum_{k=1}^{n} \{ \lfloor \lg(k-1) \rfloor + p_k \} \\ &= \sum_{k=2}^{n} \left\{ \lfloor \lg(k-1) \rfloor + \frac{2}{k}(k - 2^{\lfloor \lg(k-1) \rfloor}) \right\} \sim \sum_{k=1}^{n-1} \lg k + O(n) \\ &\sim n \lg n + O(n), \end{aligned}$$

by the same argument as used, for example, in 14.3. Asymptotic evaluation of the variance follows the same pattern. Let $m = \lfloor \lg(n-1) \rfloor$ so that $n = 2^m + r$, where $1 \le r \le 2^m$ Then, by independence of X_k's we have

$$\begin{aligned} \operatorname{var}(C_n) &= \sum_{k=2}^{n} \operatorname{var}(X_k) = \sum_{k=2}^{2^m} \operatorname{var}(X_k) + \sum_{i=1}^{r} \operatorname{var}(X_{2^m+i}) \\ &= \sum_{k=0}^{m-1} \sum_{i=1}^{2^k} \operatorname{var}(X_{2^k+i}) + \sum_{i=1}^{r} \operatorname{var}(X_{2^m+i}) \\ &= \sum_{k=0}^{m-1} \sum_{i=1}^{2^k} p_{2^k+i}(1 - p_{2^k+i}) + \sum_{i=1}^{r} p_{2^m+i}(1 - p_{2^m+i}) \end{aligned}$$

We now recall that

$$p_{2^k+i} = \frac{L_{2^k+i}}{2^k+i},$$

and

$$L_{2^k+i} = 2\left(2^k + i - 2^{\lfloor \lg(2^k+i-1)\rfloor}\right) = 2\left(2^k + i - 2^k\right) = 2i.$$

Substituting yields

$$\begin{aligned}
\mathrm{var}(C_n) &= \sum_{k=0}^{m-1}\sum_{i=1}^{2^k}(p_{2^k+i} - p_{2^k+i}^2) + \sum_{i=1}^{r}(p_{2^m+i} - p_{2^m+i}^2) \\
&= \sum_{k=0}^{m-1}\left\{\sum_{i=1}^{2^k}\frac{2i}{2^k+i} - \sum_{i=1}^{2^k}\left(\frac{2i}{2^k+i}\right)^2\right\} \\
&\quad + \left\{\sum_{i=1}^{r}\frac{2i}{2^m+i} - \sum_{i=1}^{r}\left(\frac{2i}{2^m+i}\right)^2\right\}.
\end{aligned}$$

We now calculate

$$\sum_{i=1}^{2^k}\frac{i}{2^k+i} = \sum_{i=1}^{2^k}\frac{2^k+i}{2^k+i} - 2^k\sum_{i=1}^{2^k}\frac{1}{2^k+i} = 2^k - 2^k(H_{2^{k+1}} - H_{2^k}),$$

where, $H_j = \sum_{i=1}^{j} 1/i$ is the jth harmonic number. Similarly, letting $H_j^{(2)} = \sum_{i=1}^{j} 1/i^2$ we obtain

$$\begin{aligned}
\sum_{i=1}^{2^k}\left(\frac{i}{2^k+i}\right)^2 &= \sum_{i=1}^{2^k}\left(1 - \frac{2^k}{2^k+i}\right)^2 = \sum_{i=1}^{2^k}\left(1 - 2\frac{2^k}{2^k+i} + \frac{2^{2k}}{(2^k+i)^2}\right) \\
&= 2^k - 2^{k+1}(H_{2^{k+1}} - H_{2^k}) + 2^{2k}(H_{2^{k+1}}^{(2)} - H_{2^k}^{(2)}).
\end{aligned}$$

Hence, the term within the first sum for the formula for the variance of C_n is

$$\begin{aligned}
&2(2^k - 2^k(H_{2^{k+1}} - H_{2^k})) - 4(2^k - 2^{k+1}(H_{2^{k+1}} - H_{2^k}) + 2^{2k}(H_{2^{k+1}}^{(2)} - H_{2^k}^{(2)})) \\
&= 3\cdot 2^{k+1}(H_{2^{k+1}} - H_{2^k})) - 2^{k+1} - 2^{2(k+1)}(H_{2^{k+1}}^{(2)} - H_{2^k}^{(2)})).
\end{aligned}$$

The same reasoning applies to the last term and yields

$$3\cdot 2^{m+1}(H_{2^m+r} - H_{2^m})) - 2r - 2^{2(m+1)}(H_{2^m+r}^{(2)} - H_{2^m}^{(2)})).$$

As we will see in a minute, this last term contributes an interesting, albeit not uncommon for binary structures, feature. Using approximations

$$H_j = \ln j + \gamma + O\left(\frac{1}{j}\right) \quad \text{and} \quad H_j^{(2)} = \frac{\pi^2}{6} - \frac{1}{j} + O\left(\frac{1}{j^2}\right),$$

we obtain that

$$\begin{aligned}
\operatorname{var}(C_n) &= \sum_{k=0}^{m-1}\left\{3\cdot 2^{k+1}\ln 2 - 2^{k+1} - 2^{2(k+1)}\left(\frac{2^{k+1}-2^k}{2^{2k+1}}\right) + O(1)\right\} \\
&\quad +3\cdot 2^{m+1}(\ln(2^m+r) - m\ln 2) - 2r \\
&\quad -2^{2(m+1)}\left(\frac{1}{2^m} - \frac{1}{2^m+r}\right) + O(1) \\
&= (6\ln 2 - 4)\sum_{k=0}^{m-1} 2^k + 6\cdot 2^m(\lg(2^m+r) - m)\ln 2 - 2r - 4\cdot 2^m \\
&\quad +4\frac{4^m}{2^m+r} + O(m) \\
&= (6\ln 2 - 6)2^m - 2(r+2^m) + 6\cdot 2^m(\lg(2^m+r) - m)\ln 2 \\
&\quad +4\frac{4^m}{2^m+r} + O(m).
\end{aligned}$$

Now recalling that $n = 2^m + r$ and letting $f_n = \lg(2^m+r) - m$, $0 < f_n \le 1$, so that

$$2^m = 2^{\lg(2^m+r)-f_n} = \frac{n}{2^{f_n}},$$

we see that the variance is

$$\begin{aligned}
\operatorname{var}(C_n) &= (6\ln 2 - 6)2^m + 6\cdot 2^m f_n \ln 2 - 2(r+2^m) \\
&\quad +4\frac{n^2}{(2^m+r)2^{2f_n}} + O(\ln n) \\
&= \frac{6(1+f_n)\ln 2 - 6}{2^{f_n}}\cdot n - 2n + 2^{2(1-f_n)}n + O(\ln n) \\
&= Q(f_n)\cdot n + O(\ln n),
\end{aligned}$$

where

$$Q(x) = \frac{6((1+x)\ln 2 - 1)}{2^x} - 2 + 2^{2(1-x)}.$$

Note that $Q(x)$ is nonconstant and that $Q(0) = Q(1)$, and thus, $\operatorname{var}(C_n)/n$ has an oscillatory behavior, depending on how far n is from a perfect power of 2 (these oscillations are small at the level a bit higher than $0.15\ldots$). This does not change the fact that $\max_{1\le k\le n} h_k = o(\sqrt{\operatorname{var}(C_n)})$, so that we have

Theorem 15.3 *The total number of comparisons C_n required to sort a random permutation of $\{1,\ldots,n\}$ by a binary insertion sort satisfies*

$$\frac{C_n - n\lg n}{\sqrt{Q(f_n)\cdot n}} \Longrightarrow N(0,1).$$

16 Heap Sort

Heap sort is mathematically interesting (and challenging) sorting algorithm. Remarkably, it is designed so that its *worst–case behavior* is of optimal order $O(n \log n)$. As is quite common in such situations, the average case analysis is quite complicated. As a matter of fact, the expected number of comparisons was found, asymptotically, as late as 1993 by Schaeffer and Sedgewick. The asymptotics of the variance is not known at this time, and neither is the limiting distribution function. So, in this section we will focus on the worst–case analysis.

A *heap* (or more precisely a bottom–up heap) is a complete binary tree whose nodes are represented by numbers and are arranged so that the following heap property is satisfied: for any subtree its root is the largest element of that subtree. Since heap is a complete binary tree, all its levels except possibly the last one are saturated. By convention, the last level is filled from left to right. The reason for this is that we often identify a heap on n nodes with a list $H[1..n]$ where $H[1]$ is a root of a heap and for every k the entries $H[2k]$ and $H[2k+1]$ are the children of $H[k]$ (if they are in the list). Our convention about the last level guarantees that there are no "holes" in the list H. A heap corresponding to the list

$$[10, 9, 7, 5, 8, 3, 6, 2, 4, 1]$$

is depicted in figure 6.

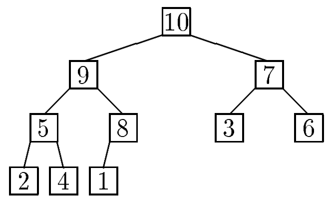

Figure 6: Heap representing $[10, 9, 7, 5, 8, 3, 6, 2, 4, 1]$

Sorting via heap is a two stage process. Let us think of a permutation π as a list $[\pi_1, \ldots, \pi_n]$ and represent it as a tree with π_{2k} and π_{2k+1} as children of π_k. This tree obviously is not a heap in general since it may fail the heap property. The first step is to rearrange the nodes so that the result is a heap (we will call it a heapifying process), represented by a list $H[1..n]$. Once this is done we will take advantage of the heap property: the largest element is at the root. With that given, sorting proceeds according to the following argument: exchange $H[1]$ and $H[n]$ (that's where $H[1]$ belongs) and consider the list $H[1..n-1]$. This list is not a heap in general, since the root may not be the largest element, but once we heapify it, we could repeat the same process over and over again, each time reducing the size of a heap by 1 until we reach size 1, at which point the array will be sorted. It remains to find an efficient way to construct a heap and then to heapify at every stage of the sorting process.

16.1 Construction Stage

Let $\pi = [\pi_1, \ldots, \pi_n]$ be given and let $h_n = \lceil \lg(n+1) \rceil - 1 = \lfloor \lg n \rfloor$ be the height of the corresponding (complete) tree. The construction begins at the very end. Suppose for simplicity that $n = 2k+1$ is odd and look at the subtree of size 3 $[\pi_k, \pi_{2k}, \pi_{2k+1}]$ (if n were even this subtree would have size 2, but as will be seen, this would not affect the process). This subtree does not have to be a heap. In order to heapify it, we need to exchange π_k with the larger of its two children π_{2k} and π_{2k+1}, if necessary. That requires two comparisons (one if n were even) and at most one data move. This results in creating a small heap of height one at the very end of our list. We refer to the process of moving up the larger of the two children as promotion. We now continue in the same fashion until we heapify the last two levels. At this point the nodes on the levels $h_n - 1$ and h_n form heaps of height 1. For example, consider a list $[3, 8, 6, 2, 1, 7, 10, 5, 4, 9]$ which corresponds to a tree in Figure 7. As Figure 8 shows, once the nodes indicated by the arrows are exchanged, the result will be a partially heapified tree with heaps of heights 1 at the bottom. We now move to the level $h_n - 2$. Consider any node at that level and a subtree of height at most 2 rooted at that node. As before, it does not have to be a heap, but we can heapify it by what is usually referred to as a sift down process. First, exchange the root of our subtree with the larger of two children (two comparisons and at most one exchange). The result may not be a heap since a new subtree of height one rooted at the node we just promoted is not guaranteed to be a heap, but if it is not we can complete heapifying process by exchanging with the larger of the two children once again. We then complete heapifying at the $(h_n - 2)$nd level and move to the next one. In our example, this stage is depicted in Figure 9. In general, at every stage we are dealing with a subtree whose all proper subtrees are already heaps, and all we need to do is sift down the root to its proper position.

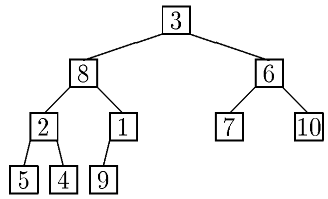

Figure 7: Tree corresponding to $[3, 8, 6, 2, 1, 7, 10, 5, 4, 9]$

Let us see now how many comparisons we need. As a matter of fact it is more convenient to count the number of data moves and remember that there are at most two comparisons for every move. In the worst case scenario at every stage of the sift down process a node may have to be moved all the way down to the bottom of the tree. The height of a tree is $h_n = \lfloor \lg n \rfloor$. Thus, in the first step, a node at the root may have to be moved up to h_n times, each of the two nodes on level one up to $h_n - 1$ times, and so forth. Adding over all nodes

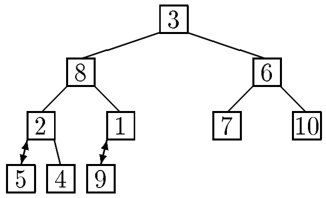

Figure 8: Heaps of heights 1

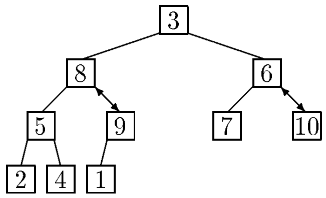

Figure 9: Heaps of heights 2

through the levels $0, 1, \ldots, h_n - 1$ we obtain an upper bound

$$1 \cdot h_n + 2^1 \cdot (h_n - 1) + \cdot + 2^{h_n - 2} \cdot 2 + 2^{h_n - 1} \cdot 1 = \sum_{k=1}^{h_n} k \cdot 2^{h_n - k}$$

$$\leq 2^{h_n} \sum_{k=1}^{\infty} k \cdot 2^{-k} \leq 2^{h_n} \cdot 2 \leq 2^{1 + \lg n} \leq 2n.$$

Remembering that there are at most two comparisons for every data move, yields an upper bound of $4n$ comparisons for the construction stage. This is well below our "benchmark" of $n \log n$ level. It remains to analyze the sorting process.

16.2 Sorting Stage

We consider a heap $H[1..n]$ and will analyze the sorting phase. As we indicated at the beginning we start by exchanging $H[1]$ with $H[n]$ and consider a $n-1$ long list

$$A_1[1..n-1] = [H[n], H[2], \ldots, H[n-1]]\,.$$

We heapify this list by sifting down the element $H[n]$ to its proper position, obtaining a heap $H_1[1..n-1]]$. We now repeat the same procedure; we exchange $H_1[1]$ with $H_1[n-1]$, and then heapify the list

$$A_2[1..n-2] = [H_1[n-1], H_1[2], \ldots, H_1[n-2]]\,,$$

by sifting down $H_1[n-1]$. The process is continued until the length of an unsorted part is reduced to 1. The first iteration of the sift down process for our example is shown on figure 10

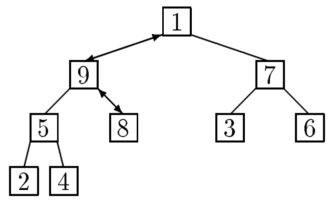

Figure 10: Sift down process

The analysis of the number of moves is very simple. In the worst case scenario at every stage a current root will have to be sifted down all the way to the bottom. When k keys are left, $k = n, n-1, \ldots, 2$, the corresponding tree has height $\lfloor \lg k \rfloor$ so that in the worst case scenario the number of moves is no more than

$$\sum_{k=2}^{n} \lfloor \lg k \rfloor \leq \sum_{k=2}^{n} \lg k \leq \int_{2}^{n+1} \lg x dx = n \lg n + O(1).$$

By a naive argument, the number of comparisons is therefore at most $2n \lg n + O(1)$. Actually one can do better by utilizing the following observation: at every stage all the *proper subtrees* of a tree are heaps, and thus, the nodes of *any* path starting at the level 1 are in a decreasing order. If a specific path was given we could perform a binary insertion in that path (note that that our naive sift down uses effectively a linear insertion, which we know has poor performance). When k keys are left the height of a tree is $\lfloor \lg k \rfloor$ so that the insertion can be done with $\lceil \lg (\lfloor \lg k \rfloor) \rceil$ comparisons (see a discussion in section on a binary insertion sort). It remains to identify a path along which a root will be moved. But that is easy: at every stage a node is moved toward the larger of its two children, and it takes only one comparison (namely, compare the children) to identify which is larger. Thus, for example, in a heap depicted in figure 11 the path of maximal sons is identified by first comparing the two nodes on level 1, then the two children of the larger node, 9, and so on). The path of maximal sons is: 9, 8, 1.

Thus, to sift down the root when k nodes are left requires $h_k = \lfloor \lg k \rfloor$ comparisons to identify the path (called the path of maximal sons) along which the root will be moved, and then $\lceil \lg (\lfloor \lg k \rfloor) \rceil$ comparisons to find its proper position along that path. Summing up over k yields:

$$\sum_{k=2}^{n} \{ \lfloor \lg k \rfloor + \lceil \lg (\lfloor \lg k \rfloor) \rceil \} + O(1) = n \lg n + O(n \lg(\lg n)).$$

Hence we have:

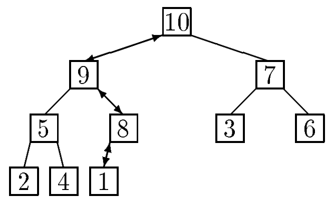

Figure 11: Path of maximal sons

Theorem 16.1 *Heap sort will require no more than $n \lg n + O(n \lg(\lg n))$ comparisons to sort any permutation of n letters.*

Note: One of the reasons for which the average case analysis poses a challenge is that the heapifying process at does not preserve randomness. Specifically, it is known that if π is a random permutation of n distinct keys, then after the construction phase the resulting heap is also random (that is to say, that each of all heaps on n nodes is equally likely to be obtained as a result). However, it is also known, that after sifting down the root not every heap on $n-1$ keys is obtained with the same probability.

17 Merge Sort

Merge sort is an algorithm that tries to utilize what is known as a divide and conquer idea: split the list to be sorted in two sublists, sort each of them and then merge the two (already sorted) lists together. There are two issues to consider here, one is a merging method, the other is the relative lengths of two sublists to be merged. Fortunately, it turns out that if the two lists are of essentially the same length (or more precisely they differ in length by at most one) then the very simplest merging algorithm, called linear merge will produce an optimal result. Let us begin, however, by general considerations. Suppose that $L_1(m)$ and $L_2(k)$ are two sorted lists, of lengths m and k, respectively, that are to be merged. Let $n = m + k$ and we assume without loss of generality, that $m \leq k$. Denote by $S_{m,k}$ the worst–case scenario number of comparisons needed to merge these two lists. Consider a decision tree for a merging algorithm. No matter what merging method is used, the algorithm will have to be able to find each of the possible configurations of the merged lists. There are $\binom{n}{m}$ different ways to do that (since there are that many ways to place the elements of $L_1(m)$ on the n positions). Consequently a decision tree will have that many leaves, and thus the height no less than $\lg \binom{n}{m}$. Suppose now that both lengths are essentially the same, that is $m = \lfloor n/2 \rfloor$ and $k = \lceil n/2 \rceil$. Then, using Stirling's approximation $n! \sim \sqrt{2\pi n}(n/e)^n$ we get

$$\binom{n}{\lfloor n/2 \rfloor} = \frac{n!}{\lfloor n/2 \rfloor ! \lceil n/2 \rceil} \sim \frac{\sqrt{2\pi n}(n/e)^n}{\left(\sqrt{2\pi n/2}(n/(2e))^{n/2}\right)^2} = \frac{\sqrt{2}}{\sqrt{\pi n}} 2^n,$$

and hence

$$S_{\lfloor n/2 \rfloor, \lceil n/2, \rceil} \geq n - \frac{1}{2} \lg n + O(1).$$

It is now time to describe how the linear merge merges two lists. It is convenient to think of removing elements from the lists once their proper position on the merged list is found. With that convention, merging works like this: we just compare the smallest elements in the two lists, the smaller is removed from its list and becomes the first element on the merged list. We then continue in that manner until one of the lists is exhausted, at which point we just transfer the leftover from the other list. For example if L_1 and L_2 are given by

$$L_1 = [3, 6, 8, 9], \quad \text{and} \quad L_2 = [1, 2, 4, 5, 7],$$

linear merge would compare first 3 and 1 (removing 1), then 3 and 2 (removing 2), then 3 and 4 (removing 3 from L_1), then 6 and 4 (removing 4), etc. until, after comparing 8 and 7 (and removing the latter) L_2 became empty. We then transfer 8 and 9 from L_1. How many comparisons does it take? We simply observe that as long as neither of the lists is empty, moving an element to the new list requires one comparison. Thus letting $L_{m,k}$ denote the leftover on the other list at the time that the first becomes empty, we have

$$S_{m,k} = n - L_{m,k} \leq n - 1,$$

since clearly, $1 \leq L_{m,k} \leq k$ (recall that $k \geq m$). This is asymptotically the same as the lower bound for *any* merging if the two lists are balanced. Specifically, we have

$$n - \frac{1}{2} \lg n + O(1) \leq S_{\lfloor n/2 \rfloor, \lceil n/2, \rceil} \leq n - 1.$$

This gives enough information to perform a competent worst case behavior analysis. Letting W_n be the largest number of comparisons required to sort any permutation of $\{1, 2, \ldots, n\}$ by a linear merge sort. We then clearly have the following recurrence:

$$W_n = W_{\lfloor n/2 \rfloor} + W_{\lceil n/2 \rceil} + n - 1.$$

To get a better feel, let us consider first the case when n is a perfect power of 2, say $n = 2^j$. The recurrence becomes

$$W_{2^j} = 2W_{2^{j-1}} + 2^j - 1,$$

and can be easily iterated to yield

$$\begin{aligned} W_{2^j} &= 2W_{2^{j-1}} + 2^j - 1 = 2^2 W_{2^{j-2}} + 2(2^{j-1} - 1) + 2^j - 1 \\ &= 2^3 W_{2^{j-3}} + 2^2(2^{j-2} - 1) + 2(2^{j-1} - 1) + 2^j - 1 \\ &= \cdots = 2^j W_1 + \sum_{i=0}^{j-1} 2^i (2^{j-i} - 1) = j2^j - \sum_{i=0}^{j-1} 2^i = j2^j - (2^j - 1) \\ &= n \lg n - (n-1), \end{aligned}$$

which is of optimal order. What about a general n? The trick is to find a right formulation of the last formula. It turns out that the following is true:

Theorem 17.1 *The worst–case number of comparisons, W_n required to sort any permutation of n distinct keys by a linear merge sort satisfies*

$$W_n = n\lceil \lg n \rceil - (2^{\lceil \lg n \rceil} - 1).$$

Proof: The proof is by an easy induction on n considering separately case of n even or odd to remove the ceilings. For example, if $n = 2\ell$ then

$$\begin{aligned} W_{2\ell} &= 2W_\ell + 2\ell - 1 = 2\left(\ell\lceil \lg \ell \rceil - \left(2^{\lg \ell} - 1\right)\right) + 2\ell - 1 \\ &= 2\ell\lceil \lg 2\ell - 1 \rceil - 2^{1+\lceil \lg \ell \rceil} + 2\ell + 1 = 2\ell\lceil \lg 2\ell \rceil - 2l - 2^{\lceil \lg 2\ell \rceil} + 2\ell + 1 \\ &= 2\ell\lceil \lg 2\ell \rceil - 2^{\lceil \lg 2\ell \rceil} + 1, \end{aligned}$$

and the same argument works for $n = 2\ell + 1$. $\square$

Note: Writing $x_n = \lceil \lg n \rceil - \lg n$ the formula for W_n can be rewritten as

$$W_n = n(x_n + \lg n) - \left(2^{x_n + \lg n} - 1\right) = n \lg n + n f(x_n) + 1,$$

where $f(x) = x - 2^x$, for $0 \le x \le 1$. Thus, the coefficient in front of the linear term has a small periodic oscillation depending how far n is from a perfect power of 2.

18 Quick Sort

Quick sort is an extremely convenient appealing sorting algorithms. It is considered to be one of the algorithms "with greatest influence on the development of science and engineering in the 20th century" by Dongarra and Sullivan in their introduction to *Computing in Science & Engineering* **2** (2000), 22–23. It is based on a very simple (once you know it) but powerful observation. To sort an array, pick an arbitrary element from the list (called pivot) and compare all other elements of the list to the pivot; if a given element is smaller, put it to left of the pivot and if not put it to the right. This process creates two lists (one of them may be empty). While neither of them is sorted, the "left" array consists of elements smaller that the pivot, and the "right" of elements larger than the pivot. Consequently, it suffices to sort each of the two sublists. This is now done recursively. The first question is how to pick a good pivot. If nothing is know a priori about the order, then there is no reason why any particular choice would be better than the other, so choosing pivot uniformly at random from the elements on the array might not be a bad choice. Since we will be sorting a random permutation of $\{1, 2, \ldots, n\}$ the the distribution of the first element is uniform over the range; thus we might as well pivot around the first entry of a given list. Below is a Maple code of a procedure (called `firststep`) that illustrates just that together with an example:

```
firststep:=proc(A)
local L,R,B,k;
 if nops(A)<=1 then RETURN(A)
   else
     L:=[];R:=[]; for k from 2 to nops(A)
    do
     if A[k]<A[1] then L:=[op(L),A[k]]
       else R:=[op(R),A[k]]
     fi;
    od;
        RETURN(L,A[1],R);
 fi;
end;


firststep([6,4,7,9,3,5,2,8,10,1]);

[4, 3, 5, 2, 1], 6, [7, 9, 8, 10];
```

The procedure constructs the two lists L and R and returns them with the pivot $A[1]$ in between. Of course, if we wanted to sort the array A rather than asking for L and R we would have asked for the *sorted* L and R. That is easy enough to do. Here is the code for quick sort

```
qsort := proc (A)
 local L, R, B, k;
  if nops(A) <= 1 then RETURN(A)
     else L := []; R := [];
          for k from 2 to nops(A)
      do
       if A[k] < A[1] then L := [op(L), A[k]]
          else R := [op(R), A[k]]
       end if
      end do;
    B :=[op(qsort(L)), A[1], op(qsort(R))];
    RETURN(B)
  end if
end proc;
```

and here is our example:

```
qsort([6,4,7,9,3,5,2,8,10,1]);

[1, 2, 3, 4, 5, 6, 7, 8, 9, 10];
```

It is perhaps worth reiterating a nice property of random permutations: if A was a random permutation, so was each of the L and R.

The analysis of the number of comparisons is facilitated by the following observation; let C_n be the number of comparisons needed to sort a random permutation of length n. If the value of the pivot were k this number would satisfy the following recurrence:

$$C_n = C_{k-1} + \overline{C}_{n-k} + n - 1,$$

with the initial condition $C(0) = 0$ and with the understanding that the sequences (C_m) and $(\overline{C}_m)$ are independent and they have the same distribution. The reason is that if the pivot is k, then the left and right sublists have, respectively, $k-1$ and $n-k$ elements, and we need $n-1$ comparisons to create these lists. As we said earlier, pivot's value is random, and its distribution is uniform on $\{1, 2, \ldots, n\}$. Thus, letting U_n be such a random variable, and writing $C(k)$ instead of C_k we obtain the following recurrence for the distribution of $C(n)$:

$$C(n) = C(U_n - 1) + \overline{C}(n - U_n) + n - 1, \quad n \geq 1, \quad C(0) = 0.$$

This recurrence is instrumental in the analysis of quick sort. We will first find the expected value and the variance.

18.1 Expected Value and Variance

We have

Theorem 18.1 *Let $C(n)$ be number of comparisons required to sort a random permutation of $\{1, 2, \ldots, n\}$ by quick sort. Then, as $n \to \infty$, we have:*

(i) $\mathbf{E}C(n) = 2n \ln n + O(n)$,

(ii) $\operatorname{var}(C(n)) \sim \left(7 - \frac{2\pi^2}{3}\right) n^2$.

Proof: We begin with part (i). Write

$$\mathbf{E}C(n) = \mathbf{E}\left(C(U_n - 1) + \overline{C}(n - U_n)\right) + n - 1,$$

and then consider the expectation of the first term. We have

$$\begin{aligned}
\mathbf{E}C(U_n-1) &= \sum_{j=1}^{n} \mathbf{E}C(U_n-1)I(U_n=j) \\
&= \sum_{j=1}^{n}\sum_{m} m\mathbf{P}(C(U_n-1)=m, U_n=j) \\
&= \sum_{j=1}^{n}\sum_{m} m\mathbf{P}(C(U_n-1)=m|U_n=j)\mathbf{P}(U_n=j) \\
&= \frac{1}{n}\sum_{j=1}^{n}\sum_{m} m\mathbf{P}(C(U_n-1)=m|U_n=j) \\
&= \frac{1}{n}\sum_{j=1}^{n}\sum_{m} m\mathbf{P}(C(j-1)=m) = \frac{1}{n}\sum_{j=1}^{n}\mathbf{E}C(j-1),
\end{aligned}$$

and, since U_n-1 has the same distribution as $n-U_n$ and C_m the same as $\overline{C}_m$, we obtain exactly the same expression for $\mathbf{E}\overline{C}(n-U_n)$. This gives a recurrence

$$\mathbf{E}C(n) = \frac{2}{n}\sum_{j=1}^{n}\mathbf{E}C(j-1)+n-1 = \frac{2}{n}\sum_{j=0}^{n-1}\mathbf{E}C(j)+n-1.$$

To solve it we multiply both sides by n to get

$$n\mathbf{E}C(n) = 2\sum_{j=0}^{n-1}\mathbf{E}C(j)+n(n-1),$$

write a similar expression for $(n-1)\mathbf{E}C(n-1)$

$$(n-1)\mathbf{E}C(n-1) = 2\sum_{j=0}^{n-2}\mathbf{E}C(j)+(n-1)(n-2),$$

and subtract side by side, to get:

$$n\mathbf{E}C(n)-(n-1)\mathbf{E}C(n-1) = 2\mathbf{E}C(n-1)+n(n-1)-(n-1)(n-2).$$

Rearranging terms and dividing by $n(n+1)$ yields,

$$\begin{aligned}
\frac{\mathbf{E}C(n)}{n+1} &= \frac{\mathbf{E}C(n-1)}{n}+2\frac{n-1}{n(n+1)} \\
&= \frac{\mathbf{E}C(n-2)}{n-1}+2\frac{n-2}{(n-1)n}+2\frac{n-1}{n(n+1)},
\end{aligned}$$

which can be further iterated, and since $\mathbf{E}C(0)=\mathbf{E}C(1)=0$, gives

$$\frac{\mathbf{E}C(n)}{n+1} = 2\sum_{j=2}^{n}\frac{j-1}{j(j+1)} = 2\sum_{j=2}^{n}\left\{\frac{1}{j+1}-\frac{1}{j(j+1)}\right\} = 2\ln n+O(1),$$

which gives

$$\mathbf{E}C(n) = 2n \ln n + O(n),$$

as claimed.

Proof of (ii). The proof of the second part rests on a similar argument. We first introduce the following concepts: for two integer valued random variables, X and Y we let

$$\mathbf{E}(X|Y) = \sum_j I(Y=j)\mathbf{E}(X|Y=j),$$

where

$$\mathbf{E}(X|Y=j) = \sum_m m\mathbf{P}(X=m|Y=j).$$

The second notion is called the *conditional expectation of X given $Y=j$* while the first is called thee *conditional expectation of X given Y*. Both are very common in any undergraduate course in probability. Note that $\mathbf{E}(X|Y)$ is a discrete random variable taking value $\mathbf{E}(X|Y=j)$ on the set $\{Y=j\}$. For that reason we have:

$$\begin{aligned}
\mathbf{E}\left(\mathbf{E}(X|Y)\right) &= \sum_j \mathbf{P}(Y=j)\mathbf{E}(X|Y=j) \\
&= \sum_j \mathbf{P}(Y=j)\sum_m m\mathbf{P}(X=m|Y=j) \\
&= \sum_{j,m} m\mathbf{P}(X=m|Y=j)\mathbf{P}(Y=j) = \sum_{j,m} m\mathbf{P}(X=m, Y=j) \\
&= \mathbf{E}X.
\end{aligned}$$

This is really nothing more than the law of total probability. We now can write:

$$\begin{aligned}
\mathrm{var}(X) &= \mathbf{E}\left(X-\mathbf{E}X\right)^2 = \mathbf{E}\left(X-\mathbf{E}(X|Y)+\mathbf{E}(X|Y)-\mathbf{E}X\right)^2 = \\
&= \mathbf{E}\Big\{(X-\mathbf{E}(X|Y))^2 + (\mathbf{E}(X|Y)-\mathbf{E}X)^2 \\
&\qquad +2(X-\mathbf{E}(X|Y))(\mathbf{E}(X|Y)-\mathbf{E}X)\Big\} \\
&= \mathbf{E}\mathbf{E}(X-\mathbf{E}(X|Y))^2|Y) + \mathbf{E}(\mathbf{E}(X|Y)-\mathbf{E}X)^2 \\
&\qquad +2\mathbf{E}\left((X-\mathbf{E}(X|Y))(\mathbf{E}(X|Y)-\mathbf{E}X)\right).
\end{aligned}$$

Since $\mathbf{E}X = \mathbf{E}\mathbf{E}(X|Y)$ the second term is simply the variance of $\mathbf{E}(X|Y)$. Furthermore, the term inside the first expectation would be the variance of X if the conditional expectations $\mathbf{E}(\,\cdot\,|Y)$ were used rather than "usual" (or unconditional) expectations. This is usually referred to as the *conditional variance given Y* and we will denote it by $\mathrm{var}_Y(\,\cdot\,)$. Thus, the first term can be written as $\mathbf{E}\mathrm{var}_Y(X)$, where $\mathrm{var}_Y(X) = \mathbf{E}((X-\mathbf{E}(X|Y))^2|Y)$. It remains to consider the last term. We claim that it is 0. To see this, insert the conditional expectation given Y to obtain

$$\mathbf{E}\mathbf{E}((X-\mathbf{E}(X|Y))(\mathbf{E}(X|Y)-\mathbf{E}X)|Y),$$

and observe that on each of the sets $\{Y = j\}$ the second factor is constant. Thus, it can be pulled in front of the conditional (but not the unconditional) expectation. Hence the last expression is equal to

$$\mathbf{E}(\mathbf{E}(X|Y) - \mathbf{E}X)\mathbf{E}((X - \mathbf{E}(X|Y))|Y).$$

But now the second factor is 0 since $\mathbf{E}X = \mathbf{E}(\mathbf{E}(X|Y))$. Thus we obtain a formula

$$\mathrm{var}(X) = \mathbf{E}\mathrm{var}_Y(X) + \mathrm{var}(\mathbf{E}(X|Y)),$$

which, as a matter of fact is true for general random variables as well. We will use it with $X = C(U_n - 1) + \overline{C}(n - U_n)$ and $Y = U_n$. One more property of variance that will be used is that it is invariant under shifts, i.e. for any constant a $\mathrm{var}(X + a) = \mathrm{var}(X)$. We now can write:

$$\begin{aligned}
\mathrm{var}(C(n)) &= \mathrm{var}(C(U_n - 1) + \overline{C}(n - U_n) + n - 1) \\
&= \mathrm{var}(C(U_n - 1) + \overline{C}(n - U_n)) \\
&= \mathbf{E}\mathrm{var}_{U_n}(C(U_n - 1) + \overline{C}(n - U_n)) \\
&\qquad + \mathrm{var}(\mathbf{E}(C(U_n - 1) + \overline{C}(n - U_n)|U_n))
\end{aligned}$$

Now, given U_n, $C(U_n - 1)$ and $\overline{C}(n - U_n)$ are independent so that

$$\mathbf{E}\mathrm{var}_{U_n}(C(U_n - 1) + \overline{C}(n - U_n)) = \mathbf{E}\mathrm{var}_{U_n}(C(U_n - 1)) + \mathbf{E}\mathrm{var}_{U_n}(\overline{C}(n - U_n)),$$

and just as before for the conditional expectations, we get

$$\begin{aligned}
\mathbf{E}\mathrm{var}_{U_n}(C(U_n - 1)) &= \sum_{j=1}^{n} \mathbf{E}\mathrm{var}_{U_n}(C(U_n - 1)I(U_n = j)) \\
&= \sum_{j=1}^{n} \mathbf{E}\mathrm{var}_{U_n=j}(C(U_n - 1))\mathbf{P}(U_n = j) \\
&= \frac{1}{n}\sum_{j=1}^{n} \mathbf{E}\mathrm{var}_{U_n=j}(C(U_n - 1)) = \frac{1}{n}\sum_{j=1}^{n} \mathrm{var}(C(j - 1)).
\end{aligned}$$

Once again, the same expression is valid for $\mathbf{E}\mathrm{var}_{U_n}(\overline{C}(n - U_n))$, and letting for simplicity $v_k = \mathrm{var}(C(k))$ we obtain a recurrence for v_n

$$v_n = \frac{2}{n}\sum_{j=0}^{n-1} v_j + \mathrm{var}(\mathbf{E}(C(U_n - 1) + \overline{C}(n - U_n)|U_n)),$$

which is the same form of recurrence that we had for expected value, except that the term $n-1$ there is replaced by a more complicated expression $\mathrm{var}(\mathbf{E}(C(U_n - 1) + \overline{C}(n - U_n)|U_n))$. As we examine how we went about solving that recurrence we realize that most of the argument could be easily carried over to our present situation. In fact, if we consider a recurrence

$$v_n = \frac{2}{n}\sum_{j=0}^{n-1} v_j + t_n$$

for any sequence t_n (usually called the toll function) and repeat exactly the same steps as before, we get that

$$\frac{v_n}{n+1} = \frac{v_{n-1}}{n} + \frac{t_n}{n+1} - \frac{n-1}{n(n+1)} t_{n-1}.$$

Iterating this recurrence, rearranging the terms in the resulting sum, and multiplying by $n+1$ at the end, yields

$$v_n = (n+1)v_0 + t_n + 2(n+1) \sum_{j=2}^{n-1} \frac{t_j}{(j+1)(j+2)}, \tag{1}$$

where, in our case, $v_0 = 0$. In order to proceed we need some information on the values of the tolls t_n. We know from the computations for expectations that $\mathbf{E}(C(U_n - 1)|U_n) \sim 2U_n \ln(U_n)$. Substituting this into our expression for t_n we get (we do have the exact expression for the expected value, so it can be checked by quite lengthy but straightforward calculations, that the error terms are of smaller order; we will omit that in order to keep the analysis as unobstructed as possible)

$$\mathrm{var}(\mathbf{E}(C(U_n - 1) + \overline{C}(n - U_n))|U_n) \sim 4\mathrm{var}(U_n \ln(U_n) + (n - U_n)\ln(n - U_n)).$$

Now we have

$$\begin{aligned}
&U_n \ln(U_n) + (n - U_n)\ln(n - U_n) \\
&\quad = U_n(\ln(U_n) - \ln n) + (n - U_n)(\ln(n - U_n) - \ln n) + U_n \ln n + (n - U_n)\ln n \\
&\quad = U_n \ln\left(\frac{U_n}{n}\right) + (n - U_n)\ln\left(\frac{n - U_n}{n}\right) + n\ln n \\
&\quad = n\left(\frac{U_n}{n}\ln\left(\frac{U_n}{n}\right) + \left(1 - \frac{U_n}{n}\right)\ln\left(1 - \frac{U_n}{n}\right) + \ln n\right).
\end{aligned}$$

As $n \to \infty$ U_n/n becomes effectively U – the uniform random variable on $[0, 1]$ and since $u \ln u$ is bounded on $(0, 1]$, the variance becomes

$$\begin{aligned}
&4n^2\left(\mathrm{var}(U\ln U + (1-U)\ln(1-U)) + O\left(\frac{1}{n}\right)\right) \\
&\quad = 4n^2\Big(\mathbf{E}(U\ln U + (1-U)\ln(1-U))^2 \\
&\qquad -(\mathbf{E}(U\ln U + (1-U)\ln(1-U)))^2\Big) + O(n).
\end{aligned}$$

Now,

$$\mathbf{E}U\ln U = \int_0^1 x\ln x dx = -\frac{1}{4},$$

and

$$\mathbf{E}(U\ln U + (1-U)\ln(1-U))^2 = \frac{5}{6} - \frac{\pi^2}{18}$$

giving

$$t_n = 4n^2\left(\frac{5}{6} - \frac{\pi^2}{18} - \frac{1}{4}\right) + O(n) \sim \left(\frac{7}{3} - \frac{2\pi^2}{9}\right)n^2.$$

This in view of (1) yields

$$v_n \sim \left(\frac{7}{3} - \frac{2\pi^2}{9}\right)n^2 + 2\left(\frac{7}{3} - \frac{2\pi^2}{9}\right)n^2 = \left(7 - \frac{2\pi^2}{3}\right)n^2,$$

and thus completes the proof. $\square$

How about the limiting distribution? This turns out to be more difficult issue than what we have encountered so far and requires new techniques. Since some of these would require a substantial mathematical background we will restricts ourselves to an informal argument (which, however, can be made rigorous).

18.2 Limiting Distribution: Heuristics

As usually, the first step is to normalize; to avoid writing the factor $\sqrt{7 - 2\pi^2/3}$ we divide by n and let

$$C(n)^* = \frac{C(n) - 2n\ln n}{n}.$$

Applying our basic recurrence and performing a sequence of algebraic manipulations, pretty much along the same lines as before, we get

$$\begin{aligned}
C(n)^* &= \frac{C(U_n - 1) + \overline{C}(n - U_n) + n - 1 - 2n\ln n}{n} = \\
&\frac{C(U_n - 1) - 2(U_n - 1)\ln(U_n - 1) + \overline{C}(n - U_n) - 2(n - U_n)\ln(n - U_n)}{n} \\
&\quad + \frac{2(U_n - 1)\ln(U_n - 1) + 2(n - U_n)\ln(n - U_n) + n - 1 - 2n\ln n}{n} \\
&= \frac{C(U_n - 1) - 2(U_n - 1)\ln(U_n - 1)}{U_n - 1}\cdot\frac{U_n - 1}{n} \\
&\quad + \frac{\overline{C}(n - U_n) - 2(n - U_n)\ln(n - U_n)}{n - U_n}\cdot\frac{n - U_n}{n} \\
&\quad + 2\frac{U_n - 1}{n}\ln\left(\frac{U_n - 1}{n}\right) + 2\frac{n - U_n}{n}\ln\left(\frac{n - U_n}{n}\right) + \frac{n - 1 - \ln n}{n} \\
&= C(U_n - 1)^*\cdot\frac{U_n - 1}{n} + \overline{C}(n - U_n)^*\cdot\frac{n - U_n}{n} \\
&\quad + 2\frac{U_n - 1}{n}\ln\left(\frac{U_n - 1}{n}\right) + 2\frac{n - U_n}{n}\ln\left(\frac{n - U_n}{n}\right) + \frac{n - 1 - \ln n}{n}.
\end{aligned}$$

It is now tempting to pass to the limit with n. Since $U_n/n \Longrightarrow U$, assuming that $C(n)$ did converge in distribution to C^*, it would seem that this limit would have to satisfy:

$$C^* \stackrel{d}{=} UC^* + (1 - U)\overline{C}^* + 2U\ln U + 2(1 - U)\ln(1 - U) + 1,$$

where "$\stackrel{d}{=}$" means equality in distribution. This is actually the case, although formal argument is much more involved then what the outlined heuristics suggests. The main reason for being cautious is that, unfortunately, convergence in distribution does not share some of the properties of other types of convergence that we take for granted. For example, the product of two sequences each of which converges in distribution, *does not* have to converge in distribution. Nonetheless, as we mentioned earlier, the result itself is true. Hence,

Theorem 18.2 *The limiting normalized distribution C^* of the number of comparisons that it takes quick sort to sort a random permutation of $\{1, \ldots, n\}$ is the unique distribution with finite second moment which satisfies the following distributional equation:*

$$C^* \stackrel{d}{=} UC^* + (1-U)\overline{C}^* + 2U \ln U + 2(1-U)\ln(1-U) + 1,$$

where $\overline{C}^*$ is an independent copy of C^* and U is a uniform random variable on $[0,1]$, independent of both C^* an $\overline{C}^*$.

Exercises

1. Prove the following properties of probability (those that are not proved in the lecture notes).

(i) $\mathbf{P}(\emptyset) = 0$.

(ii) for any finite sequence of pairwise disjoint sets $A_1, \ldots, A_n$ $\mathbf{P}(\bigcup_{j=1}^{n} A_j) = \sum_{j=1}^{n} \mathbf{P}(A_j)$,

(iii) for any set $A \in \mathcal{A}$, $\mathbf{P}(A) = 1 - \mathbf{P}(A^c)$.

(iv) for any sets $A, B \in \mathcal{A}$ such that $A \subset B$ we have $\mathbf{P}(A) \leq \mathbf{P}(B)$,

(v) for any sets $A, B \in \mathcal{A}$, $\mathbf{P}(A \cup B) = \mathbf{P}(A) + \mathbf{P}(B) - \mathbf{P}(A \cap B)$.

2. Use a concept of disjointification to show the *subadditivity* of the probability: for any (not necessarily pairwise disjoint) sets $A_1, A_2, \ldots$, we have:

$$\mathbf{P}(\bigcup_{j=1}^{\infty} A_j) \leq \sum_{j=1}^{\infty} \mathbf{P}(A_j).$$

3. Show that an exponential random variable with parameter λ has a *memoryless property* i.e. that for all $s, t > 0$ we have:

$$\mathbf{P}(X > s + t | X > t) = \mathbf{P}(X > s).$$

4. Let X be uniform random variable on $[a, b]$ and let c satisfy $a < c < b$. Show that the random variable X conditioned on the event $\{X \leq c\}$ is uniformly distributed on $[a, c]$. (This amounts to finding $\mathbf{P}(X \leq x | X \leq c)$.)

5. Show by appropriate substitution that if X is a normal random variable with parameters μ, σ^2 then the random variable Y defined by

$$Y = \frac{X - \mu}{\sigma}$$

has a normal distribution with parameters $0, 1$.

6. Show that an geometric random variable with parameter p has a *memoryless property* i.e. that for all natural numbers k, m we have:

$$\mathbf{P}(X = k + m | X > m) = \mathbf{P}(X = k).$$

7. Find the characteristic functions of the following random variables:

(i) uniform on the interval $[0, 1]$.

(ii) Poisson with parameter λ.

(iii) binomial with parameters n, p (keep in mind that it is the sum of n independent Bernoulli random variables with parameter p).

(iv) normal with parameters $0, 1$, and then, by change of variables normal with parameters μ, σ^2.

8. Show that

(i) the variance of the sum of two independent random variables is the sum of their variances i.e., that $\text{var}(X+Y) = \text{var}(X) + \text{var}(Y)$ if X, Y are independent (just square out inside the integral sign).

(ii) Extend this to an arbitrary number of independent random variables, i.e. if $X_1, \ldots, X_n$ are independent with finite variances then

$$\text{var}(X_1 + \cdots + X_n) = \text{var}(X_1) + \cdots + \text{var}(X_n).$$

(iii) Find the variance of the binomial distribution with parameters n, p (see a comment for exercise **7**(iii)).

9. Let X and Y be independent Poisson random variables with parameters λ and μ, respectively. Find the characteristic function of the sum $X+Y$ and argue by the uniqueness of the characteristic function that the distribution of $X+Y$ is Poisson with parameter $\lambda + \mu$.

10. Let X_n be a discrete uniform random variable on $[0, 1]$ taking on n values, i.e.

$$\mathbf{P}(X_n = \frac{k}{n}) = \frac{1}{n}, \quad \text{for } k = 0, 1, \ldots, n-1.$$

(i) Find its characteristic function $\phi_n(t)$ (note that the resulting sum is the sum of a geometric progression so it has a closed form).

(ii) Find the limit

$$\phi(t) = \lim_{n \to \infty} \phi_n(t), \quad \text{for all } t \in \mathbf{R},$$

and identify it as a characteristic function of a certain random variable.

11. Let X_n be a binomial random variable with parameters $n, 1/2$.

(i) Find the characteristic function $\phi_n(t)$ of $Y_n = \frac{X_n - n/2}{\sqrt{n}/2}$.

(ii) Find the limit

$$\phi(t) = \lim_{n \to \infty} \phi_n(t), \quad \text{for all } t \in \mathbf{R},$$

and identify it as a characteristic function of a certain random variable.

12. Let X_n be a Bernoulli random variable with parameter p_n, $n \geq 1$ and X a Bernoulli random variable with parameter p. Find a necessary and sufficient condition for $X_n \Longrightarrow X$.

13. Let y_n, $n \geq 1$ be a sequence of points, and let F_n, $n \geq 1$ be a sequence of distributions given by

$$F_n(x) = \begin{cases} 0 & \text{if } x < y_n, \\ 1 & \text{if } x \geq y_n \end{cases}$$

Show that F_n converge in distribution iff $y_n \to y$ and that in this case $F_n \Longrightarrow F$, where

$$F(x) = \begin{cases} 0 & \text{if } x < y, \\ 1 & \text{if } x \geq y \end{cases}$$

14. Let $X, X_1, X_2, \ldots,$ be a sequence of random variables such that $X_n \Longrightarrow X$. Let $c, c_1, c_2, \ldots$ be a sequence of positive numbers such that $\lim_{n\to\infty} c_n = c$ and let $b, b_1, b_2, \ldots$ be a sequence of any numbers such that $\lim_{n\to\infty} b_n = b$. Show that

$$\frac{X_n - b_n}{c_n} \Longrightarrow \frac{X - b}{c}.$$

15. Let X_n, $n \geq 1$ be exponential random variables with parameters λ_n, $n \geq 1$. Show that the sequence of their c.d.f's is tight if and only if there exists a $c > 0$ such that $\lambda_n \geq c$ for all $n \geq 1$.

16. Show that for arbitrary sequence of random variables X_n, $n \geq 1$, there exist constants a_n, $n \geq 1$ such that

$$a_n X_n \Longrightarrow 0.$$

17. (classical central limit theorem). Let X_n, $n \geq 1$ be i.i.d. random variables such that $\mathbf{E}X_n = 0$ and $\mathrm{var}(X_n) = 1$, for all $n \geq 1$. Let $S_n = X_1 + X_2 + \cdots + X_n$, $n \geq 1$. Use continuity theorem (and properties of characteristic functions) to show that

$$\frac{S_n}{\sqrt{n}} \Longrightarrow N(0,1),$$

where $N(0,1)$ denotes the normal random variable with parameters 0 and 1.

18. Let $X_1, X_2, \ldots,$ be a sequence of independent random variables such that $\mathbf{P}(X_n = 1) = 1/n$ and $\mathbf{P}(X_n = 0) = 1 - 1/n$, and let

$$S_n = \sum_{k=1}^{n} X_k.$$

Show that the expected value and the variance of S_n are both asymptotically $\ln n$, and show that

$$\frac{S_n - \ln n}{\sqrt{\ln n}} \Longrightarrow N(0,1).$$

19. Let $X_1, X_2, \ldots,$ be a sequence of independent random variables such that X_k is discrete uniform on $\{0, 1, 2, \ldots, k-1\}$, i.e.

$$\mathbf{P}(X_k = j) = \frac{1}{k}, \quad \text{for } j = 0, 1, \ldots, k-1.$$

Let $S_n = \sum_{k=1}^{n} X_k$. Show that the expected value and the variance of S_n satisfy

$$\mathbf{E}S_n \approx \frac{n^2}{4}, \quad \operatorname{var}(S_n) \approx \frac{n^3}{36},$$

and use Lindeberg's CLT to show that

$$\frac{S_n - n^2/4}{n^{3/2}/6} \Longrightarrow N(0,1).$$

20. Let $(\pi_1, \pi_2, \ldots, \pi_n)$ be a random permutation of n distinct numbers. Let $1 \le i_1 < i_2 < \cdots < i_k \le n$ be fixed. Show that $(\pi_{i_1}, \pi_{i_2}, \ldots, \pi_{i_k})$ is a random permutation of k distinct numbers. (That's extremely useful property of random permutations.)

21. Design a comparison based algorithm that finds a median (the 3rd largest element) in a permutation of 5 distinct numbers (a good one would make no more than 6 comparisons). Assuming that the permutation is random find the expected number of comparisons made by your algorithm.

22. Let T_k be an insertion tree for a kth element in *any* insertion sort algorithm.

(i) Prove that the average number of comparisons needed to insert the jth key is

$$\frac{X(T_j)}{j}.$$

(ii) Prove that for a binary insertion sort each insertion tree is a complete binary tree and thus that the binary insertion sort minimizes the average number of comparisons needed to sort a random permutation by an insertion sort.

23. Find the smallest and the largest number of comparisons needed to sort any permutation of $\{1, \ldots, n\}$ by a linear insertion sort, and find a permutation giving these values. How many comparisons will a binary insertion sort need to sort each of these two permutations?

24. Carry out a step–by–step construction of a heap from the list

$$[1, 12, 8, 4, 7, 10, 9, 3, 13, 2, 6, 11, 5].$$

25. Consider a construction stage of heap sort. Find a permutation of $\{1, \ldots, 10\}$ that gives the worst possible performance during that stage. How many comparisons does it make?

26. Consider the list

$$[x_1, x_2, x_3, 4, 5, 2, 3, 8, 7, 9]$$

which is a permutation of $\{1, \ldots, 10\}$.

(i) Partially out a heap construction for that data (that is you should end up having heaps on all levels except the 0th and the 1st).

(ii) Now, assuming that the above permutation was chosen at random find the expected number of *moves* needed to complete the heap construction.

27. (merge sort) Consider the leftover random variable $L_{m,k}$ discussed in section about merge sort algorithm.

(i) Prove that for a natural number r

$$\mathbf{P}(L_{m,k} \geq r) = \frac{\binom{m+k-r}{k} + \binom{m+k-r}{m}}{\binom{n}{m}},$$

where $m + k = n$.

(ii) Prove that for any random variable X with values in the set of natural numbers, one has

$$\mathbf{E}X = \sum_{r=1}^{\infty} \mathbf{P}(X \geq r),$$

and use it to show that

$$\mathbf{E}L_{m,k} = \frac{k}{m+1} + \frac{m}{k+1}.$$

(iii) Let $n = 2^j$ and let C_n be the number of comparisons needed to sort a random permutation of n distinct numbers by a merge sort (with splits in equal lengths at every stage). Show that

$$\mathbf{E}C_{2^j} = \sum_{i=1}^{j-1} 2^i \left(2^{j-i}\left(1 - \frac{1}{2^{j-i-1}+1}\right)\right),$$

and conclude that in that case $\mathbf{E}C_n \sim n \lg n + O(n)$.

28. The following list is a permutation of the set $\{1, \ldots, 8\}$.

$$[x_1, 2, 6, 8, x_2, 4, 3, 1].$$

Assuming that this permutation was chosen at random find the expected number of comparisons needed to sort this list by quick sort.

29. (FIND) Consider the following FIND algorithm that finds, say, the smallest element in an unsorted list (and which together with quick sort was designed by Hoare): create two lists, L and R just as in quick sort. Since the smallest element has to be in L continue recursively until the list is reduced to a single element or to an empty list. Let C_n be the number of comparisons needed to carry out this search on a list of n distinct elements.

(i) Assuming that the original list is a random permutation, write down a recurrence for C_n.

(ii) Use that recurrence to show that the expected value of C_n satisfies

$$\mathbf{E}C_n = 2n - 2H_n \sim 2n,$$

where $H_n = \sum_{i=1}^{n} 1/i$ is the nth harmonic number.

(iii) By a similar argument, show that

$$\mathrm{var}(C_n) \sim \frac{n^2}{2}.$$

(iv) Now consider the normalized random variable

$$C_n^* = \frac{C_n - 2n}{n}.$$

Use the same heuristics as for quick sort to derive a distributional equation for the limiting distribution C^* of C_n^*, as $n \to \infty$. (It is, in fact, the unique distribution having second moment that satisfies this equation.)

Note: This is NOT the most efficient way to find the smallest (or the largest element on a list. A simple linear scan comparing with a current minimum needs exactly $n-1$ comparisons).

19 Hints to exercises

1. Properties (i)-(iv) are already proved. In (v), notice that for all sets $A_1, A_2 \in \mathcal{A}$, $\mathbf{P}(A_1) = \mathbf{P}(A_1 \cap A_2) + \mathbf{P}(A_1 \cap A_2^c)$.

2. One-line consequence of disjointification and properties of a probability measure.

3. Definitions of conditional probability and exponential random variable.

4. Using only definitions you should get

$$\mathbf{P}(X \le x | X \le c) = \begin{cases} 0, & x < a, \\ \frac{x-a}{c-a}, & a \le x \le c, \\ 1, & x > c. \end{cases}$$

5. Notice that $\mathbf{P}\left(\frac{X-\mu}{\sigma} < x\right) = \mathbf{P}(X < \sigma x + \mu)$, use the definition of normal random variable and make a substitution that gives integral from $-\infty$ to x.

6. Similar to exercise 3. Notice that $\mathbf{P}(X > m)$ is a geometric series and easy to compute in closed form.

7. Parts (i)-(iii) are straightforward computations.

(i)
$$\phi(t) = \frac{e^{it} - 1}{it}$$

(ii)
$$\phi(t) = e^{\lambda(e^{it}-1)}$$

(iii)
$$\phi(t) = (1 + p(e^{it} - 1))^n$$

(iv) Complete square in the exponent, make a change of variable and remember that
$$\int_{-\infty}^{\infty} e^{-x^2/2} dx = \sqrt{2\pi}.$$
You should get
$$\phi(t) = e^{-t^2/2} \text{ with parameters } 0,1,$$
and by change of variables
$$\phi(t) = e^{it\mu-(\sigma t)^2/2} \text{ with parameters } \mu,\ \sigma^2.$$

8.

(i) Remember that $\mathrm{var}(X) = \mathbf{E}(X - \mathbf{E}X)^2 = \mathbf{E}X^2 - (\mathbf{E}X)^2$.

(ii) Result (i) and induction.

(iii) $\mathrm{var}(X) = np(1-p)$.

9. Independence is essential.

10.

(i)
$$\phi_n(t) = \begin{cases} \frac{1-e^{it}}{n(1-e^{it/n})}, & t \neq 0, \\ 1, & t = 0. \end{cases}$$

(ii)

$$\phi(t) = \begin{cases} \frac{e^{it}-1}{it}, & t \neq 0, \\ 1, & t = 0. \end{cases}$$

11.

(i) Using **7**(iii) $\phi_{X_n}(t) = ((1+e^{it})/2)^n$, so that

$$\phi_n(t) = e^{-it\sqrt{n}} \left(\frac{1+e^{i\frac{2t}{\sqrt{n}}}}{2} \right)^n = \left(\frac{e^{-it/\sqrt{n}} + e^{it/\sqrt{n}}}{2} \right)^n .$$

(ii) Expand exponential functions in (i) and use $(1+x/n)^n \to e^x$, for any real x, to get $\phi(t) = e^{-t^2/2}$.

12. The condition is $p_n \to p$.

13. Definition of convergence in distribution.

14. If F_n is the c.d.f. of $\frac{X_n - b_n}{c_n}$, prove that $F_n(x) = F_{X_n}(c_n x + b_n)$, take the limit and show that it's equal to $F(x)$.

15. Use the definition of tight sequence. If the sequence is tight, indirect proof leads you to the contradiction that there always exists n such that $F_n(y) < 1-\varepsilon$. To the opposite direction, requirement $F_N(x) < \varepsilon$ is trivial and $F_n(y) > 1-\varepsilon$ follows easily from the assumption.

16. Since the c.d.f. of $X \equiv 0$ is the unit step function, you must find sequence (a_n) that draws $F_{a_n X_n}$ close to the origin faster than the F_n's get wider (in case they do). For example, pick $a_n = 2^{-n} b_n$, where $b_n > 0$ is a number such that $\max\{F_{X_n}(-b_n), 1 - F_{X_n}(b_n)\} < 1/n$.

17. First, $\phi_{S_n/\sqrt{n}}(t) = (\phi_{X_1}(t/\sqrt{n}))^n = (\mathbf{E} e^{itX_1/\sqrt{n}})^n$. Now use the estimates for the exponential function obtained in the proof of Theorem (9.1) and $\mathbf{E} X_1 = 0$ to get

$$\mathbf{E} \left| e^{itX_1/\sqrt{n}} - \left(1 - \frac{t^2 X_1^2}{2n}\right) \right| \leq \mathbf{E} \min \left\{ \frac{t^2 X_1^2}{n}, \frac{|tX_1|^3}{6n^{3/2}} \right\} = \frac{t^2}{n} \mathbf{E} \min \left\{ X_1^2, \frac{|t||X_1|^3}{6\sqrt{n}} \right\} .$$

The last expectation goes to 0 as $n \to \infty$. Hence

$$\lim_n \phi_{S_n/\sqrt{n}}(t) = \lim_n \left(1 - \frac{t^2}{2n} \right)^n = e^{-t^2/2}.$$

18. It is known that $H_n = \sum_{k=1}^n 1/k$ is asymptotically $\ln n$. Use this and Lindeberg's CLT.

19. Formulas

$$\sum_{j=1}^n j = \frac{n(n+1)}{2}$$

and

$$\sum_{j=1}^{n} j^2 = \frac{n(n+1)(2n+1)}{6}$$

are useful. Lindeberg's sum converges to zero because of the indicator functions.

20. Straightforward consequence of the definition.

21. Take two pairs of numbers, sort both pairs and throw the greatest of the four out. Then take the number that was left out of the original two pairs and put it to the place of the number that was thrown out. Now you need only three more comparisons to find the median.

22.

(i) An insertion tree has already inserted elements as nodes and possible places at which the next element can be inserted as leaves.

(ii) Follows easily from the definition of binary sort.

23. The smallest number of comparisons is $n-1$ and the largest is $\frac{n^2-n}{2}$. The number of comparisons in binary sort does not depend on permutations.

24. Just do it!

25. It takes 15 comparisons to sort the worst case, which is not the obvious one, but quite close to it.

26.

(ii) Probably the easiest way is to go through all possible permutations of the three unknown positions. The expected number of moves is $2\frac{2}{3}$.

27.

(i) Remember that $m+k=n$, and notice that if the r greatest numbers are in one list, then the elements of the other list must be chosen from the $n-r$ smallest numbers.

(ii) The first equation follows from the definition of expected value, once you rearrange the terms in the sum. In the second question you may find the following formula helpful:

$$\sum_{k=0}^{n} \binom{k}{m} = \binom{n+1}{m+1}.$$

(iii) Compute the expected number of comparisons on each merging level, use the previous result and sum up all these terms.

$$\mathbf{E}C_{2^j} = \sum_{i=1}^{j} 2^i \left(2^{j-i}\left(1 - \frac{1}{2^{i-1}+1}\right)\right) = j2^j + O(2^j).$$

28. Choose the pivot to be the first number, and go through the sorting process in the two possible cases. The expected number of comparisons is 16.

29.

(i-iii) These are simplified versions of the corresponding computations in the quick sort section. In (iii), the recurrence relation is

$$v_n = \frac{1}{n}\sum_{j=0}^{n-1} v_j + \underbrace{\mathrm{var}(\mathbf{E}(C(U_n-1)|U_n))}_{=:t_n}.$$

Iteration then gives

$$v_n = t_n + \sum_{j=1}^{n-1} \frac{1}{j+1} t_j.$$

Now use (ii) to conclude that $t_n \sim 4\mathrm{var}(U_n) \sim \frac{n^2}{3}$ and deduce the result from the iterated relation above.

(iv) $C^* \stackrel{d}{=} UC^* + 2U - 1.$