Computational Complexity: A Modern Approach

Draft of a book: Dated January 2007
Comments welcome!

Sanjeev Arora and Boaz Barak
Princeton University
complexitybook@gmail.com

Not to be reproduced or distributed without the authors' permission

This is an Internet draft. Some chapters are more finished than others. References and attributions are very preliminary and we apologize in advance for any omissions (but hope you will nevertheless point them out to us).

Please send us bugs, typos, missing references or general comments to complexitybook@gmail.com — **Thank You!!**

Appendix A

Mathematical Background.

This appendix reviews the mathematical notions used in this book. However, most of these are only used in few places, and so the reader might want to only quickly review Sections A.1, A.2 and A.3, and come back to the other sections as needed. In particular, apart from probability, the first part of the book essentially requires only comfort with mathematical proofs and some very basic notions of discrete math.

The topics described in this appendix are covered in greater depth in many texts and online sources. Almost all of the mathematical background needed is covered in a good undergraduate "discrete math for computer science" course as currently taught at many computer science departments. Some good sources for this material are the lecture notes by Papadimitriou and Vazirani [PV06], Lehman and Leighton [LL06] and the book of Rosen [Ros06].

The mathematical tool we use most often is discrete probability. Alon and Spencer [AS00] is a great resource in this area. Also, the books of Mitzenmacher and Upfal [MU05] and Prabhakar and Raghavan [**prabhakarRa??**] cover probability from a more algorithmic perspective.

Although knowledge of algorithms is not strictly necessary for this book, it would be quite useful. It would be helpful to review either one of the two recent books by Dasgupta et al [DPV06] and Kleinberg and Tardos [KT06] or the earlier text by Cormen et al [CLRS01]. This book does not require prior knowledge of computability and automata theory, but some basic familiarity with that theory could be useful: see Sipser's book [SIP96] for an excellent introduction. See Shoup's book [Sho05] for a computer-science introduction to algebra and number theory.

A.1 Mathematical Proofs

Perhaps *the* mathematical prerequisite needed for this book is a certain level of comfort with mathematical proofs. While in everyday life we might use "proof" to describe a fairly convincing argument, in mathematics a proof is an argument that is convincing *beyond any shadow of a doubt.*[1] For example, consider the following mathematical statement:

> *Every even number greater than* 2 *is equal to the sum of two primes.*

This statement, known as "Goldbach's Conjecture", was conjectured to be true by Christian Goldbach in 1742. In the more than 250 years that have passed since, no one has ever found a counterexample to this statement. In fact, it has been verified to be true for all even numbers from 4 till $100,000,000,000,000,000$. Yet still it is not considered proven, since we have not ruled out the possibility that there is some (very large) even number that cannot be expressed as the sum of two primes.

[1] In a famous joke, as a mathematician and an engineer drive in Scotland they see a white sheep on their left side. The engineer says "you see: all the sheep in Scotland are white". The mathematician replies "All I see is that there exists a sheep in Scotland whose right side is white".

The fact that a mathematical proof has to be absolutely convincing does not mean that it has to be overly formal and tedious. It just has to be clearly written, and contain no logical gaps. When you write proofs try to be clear and concise, rather than using too much formal notation. When you read proofs, try to ask yourself at every statement "am I really convinced that this statement is true?".

Of course, to be absolutely convinced that some statement is true, we need to be certain of what that statement means. This why there is a special emphasis in mathematics on very precise *definitions*. Whenever you read a definition, try to make sure you completely understand it, perhaps by working through some simple examples. Oftentimes, understanding the meaning of a mathematical statement is more than half the work to prove that it is true.

EXAMPLE A.1
Here is an example for a classical mathematical proof, written by Euclid around 300 B.C. Recall that a *prime number* is an integer $p > 1$ whose only divisors are p and 1, and that every number n is a product of prime numbers. Euclid's Theorem is the following:

THEOREM A.2
There exist infinitely many primes.

Before proving it, let's see that we understand what this statement means. It simply means that for every natural number k, there are more than k primes, and hence the number of primes is not finite.

At first, one might think it's obvious that there are infinitely many primes because there are infinitely many natural numbers, and each natural number is a product of primes. However, this is faulty reasoning: for example, the set of numbers of the form 3^n is infinite, even though their only factor is the single prime 3.

To prove Theorem A.2, we use the technique of *proof by contradiction*. That is, we assume it is false and try to derive a contradiction from that assumption. Indeed, assume that all the primes can be enumerated as $p_1, p_2, \ldots, p_k$ for some number k. Define the number $n = p_1 p_2 \cdots p_k + 1$. Since we assume that the numbers $p_1, \ldots, p_k$ are *all* the primes, all of n's prime factors must come from this set, and in particular there is some i between 1 and k such that p_i divides n. That is, $n = p_i m$ for some number m. Thus,

$$p_i m = p_1 p_2 \cdots p_k + 1$$

or equivalently,

$$p_i m - p_1 p_2 \cdots p_k = 1 \,.$$

But dividing both sides of this equation by p_i, we will get a whole number on the left hand side (as p_i is a factor of $p_1 p_2 \cdots p_k$) and the fraction $1/p_i$ on the right hand side, deriving a contradiction. This allows us to rightfully place the QED symbol ■ and consider Theorem A.2 as proven.

A.2 Sets, Functions, Pairs, Strings, Graphs, Logic.

Sets. A *set* contains a finite or infinite number of elements, without repetition or respect to order, for example $\{2, 17, 5\}$, $\mathbb{N} = \{1, 2, 3, \ldots\}$ (the set of natural numbers), $[n] = \{1, 2, \ldots, n\}$ (the set of natural numbers from 1 ro n), $\mathbb{R}$ (the set of real numbers). For a finite set A, we denote by $|A|$ the number of elements in A. Some operations on sets are: **(1)** *union*: $A \cup B = \{x : x \in A \text{ or } x \in B\}$, **(2)** *intersection* : $A \cap B = \{x : x \in A \text{ and } x \in B\}$, and **(3)** *set difference*: $A \setminus B = \{x : x \in A \text{ and } x \notin B\}$.

Functions. We say that f is a *function* from a set A to B, denoted by $f : A \to B$, if it maps any element of A into an element of B. If B and A are finite, then the number of possible functions from A to B is $|B|^{|A|}$. We say that f is *one to one* if for every $x, w \in A$ with $x \neq w$, $f(x) \neq f(w)$. If A,B are finite, the existence of such a function implies that $|A| \leq |B|$. We say that f is *onto* if for every $y \in B$ there exists $x \in A$ such that $f(x) = y$. If A, B are finite, the existence of such a function implies that $|A| \geq |B|$. We say that f is a *permutation* if it is both one-to-one and onto. For finite A, B, the existence of a permutation from A to B implies that $|A| = |B|$.

Pairs and tuples. If A,B are sets, then the $A \times B$ denotes the set of all ordered pairs $\langle a, b \rangle$ with $a \in A, b \in B$. Note that if A, B are finite then $|A \times B| = |A| \cdot |B|$. We can define similarly $A \times B \times C$ to be the set of ordered triples $\langle a, b, c \rangle$ with $a \in A, b \in B, c \in C$. For $n \in \mathbb{N}$, we denote by A^n the set $A \times A \times \cdots \times A$ (n times). We will often use the set $\{0,1\}^n$, consisting of all length-n sequences of bits (i.e., length n strings), and the set $\{0,1\}^* = \cup_{n \geq 0} \{0,1\}^n$ ($\{0,1\}^0$ has a single element: a binary string of length zero, which we call the empty word and denote by ε).

Graphs. A *graph* G consists of a set V of *vertices* (which we often assume is equal to the set $[n] = \{1, \ldots, n\}$ for some $n \in N$) and a set E of *edges*, which consists of unordered pairs (i.e., size two subsets) of elements in V. We denote the edge $\{u, v\}$ of the graph by $\overline{u\,v}$. For $v \in V$, the *neighbors* of v are all the vertices $u \in V$ such that $\overline{u\,v} \in E$. In a *directed graph*, the edges consist of *ordered pairs* of vertices, to stress this we sometimes denote the edge $\langle u, v \rangle$ in a directed graph by $\overrightarrow{u\,v}$. One can represent an n-vertex graph G by its *adjacency matrix* which is an $n \times n$ matrix A such that $A_{i,j}$ is equal to 1 if the edge $\overrightarrow{i\,j}$ is present in G i^{th} and is equal to 0 otherwise. One can think of an undirected graph as a directed graph G that satisfies that for every u, v, G contains the edge $\overrightarrow{u\,v}$ if and only if it contains the edge $\overrightarrow{v\,u}$. Hence, one can represent an undirected graph by an adjecancy matrix that is *symmetric* ($A_{i,j} = A_{j,i}$ for every $i, j \in [n]$).

Boolean operators. A *Boolean variable* is a variable that can be either TRUE or FALSE (we sometimes identify TRUE with 1 and FALSE with 0). We can combine variables via the logical operations AND ($\wedge$), OR ($\vee$) and NOT ($\neg$, sometimes also denoted by an overline), to obtain *Boolean formulae*. For example, the following is a Boolean formulae on the variables u_1, u_2, u_3: $(u_1 \wedge \overline{u}_2) \vee \neg(u_3 \wedge \overline{u}_1)$. The definitions of the operations are the usual: $a \wedge b =$ TRUE if $a =$ TRUE and $b =$ TRUE and is equal to FALSE otherwise; $\overline{a} = \neg a =$ TRUE if $a =$ FALSE and is equal to FALSE otherwise; $a \vee b = \neg(\overline{a} \vee \overline{b})$. If φ is a formulae in n variables $u_1, \ldots, u_n$, then for any *assignment* of values $u \in \{\text{FALSE}, \text{TRUE}\}^n$ (or equivalently, $\{0,1\}^n$), we denote by $\varphi(u)$ the value of φ when its variables are assigned the values in u. We say that φ is *satisfiable* if there exists a u such that $\varphi(u) =$ TRUE.

Quantifiers. We will often use the *quantifiers* $\forall$ (for all) and $\exists$ (exists). That is, if φ is a condition that can be TRUE or FALSE depending on the value of a variable x, then we write $\forall_x \varphi(x)$ to denote the statement that φ is TRUE for *every* possible value that can be assigned to x. If A is a set then we write $\forall_{x \in A} \varphi(x)$ to denote the statement that φ is TRUE for every assignment for x from the set A. The quantifier $\exists$ is defined similarly. Formally, we say that $\exists_x \varphi(x)$ holds if and only if $\neg(\forall_x \neg\varphi(x))$ holds.

Big-Oh Notation. We will often use the big-Oh notation (i.e., $O, \Omega, \Theta, o, \omega$) as defined in Section [**model:sec:oh**].

A.3 Probability theory

A *finite probability space* is a finite set $\Omega = \{\omega_1, \ldots, \omega_N\}$ along with a set of numbers $p_1, \ldots, p_N \in [0,1]$ such that $\sum_{i=1}^{N} p_i = 1$. A random element is selected from this space by choosing ω_i with probability p_i. If x is chosen from the sample space Ω then we denote this by $x \in_R \Omega$. If no distribution is specified then we use the uniform distribution over the elements of Ω (i.e., $p_i = \frac{1}{N}$ for every i).

An *event* over the space Ω is a subset $A \subseteq \Omega$ and the *probability* that A occurs, denoted by $\Pr[A]$, is equal to $\sum_{i:\omega_i \in A} p_i$. To give an example, the probability space could be that of all 2^n possible outcomes of n tosses of a fair coin (i.e., $\Omega = \{0,1\}^n$ and $p_i = 2^{-n}$ for every $i \in [2^n]$) and the event A can be that the number of coins that come up "heads" (or, equivalently, 1) is even. In this case, $\Pr[A] = 1/2$ (exercise). The following simple bound —called the *union bound*—is often used in the book. For every set of events $A_1, A_2, \ldots, A_n$,

$$\Pr[\cup_{i=1}^{n} A_i] \leq \sum_{i=1}^{n} \Pr[A_i]. \tag{1}$$

Inclusion exclusion principle. The union bound is a special case of a more general principle. Indeed, note that if the sets $A_1, \ldots, A_n$ are not *disjoint* then the probability of $\cup_i A_i$ could be smaller than $\sum_i \Pr[A_i]$ since we are overcounting elements that appear in more than one set. We can correct this by substracting $\sum_{i<j} \Pr[A_i \cap A_j]$ but then we might be undercounting, since we subtracted elements that appear in at least 3 sets too many times. Continuing this process we get

Claim A.3 (Inclusion-Exclusion principle)
For every $A_1, \ldots, A_n$,

$$\Pr[\cup_{i=1}^{n} A_i] = \sum_{i=1}^{n} \Pr[A_i] - \sum_{1 \leq i < j \leq n} \Pr[A_i \cap A_j] + \cdots + (-1)^{n-1} \Pr[A_1 \cap \cdots \cap A_n] .$$

Moreover, this is an alternating sum *which means that if we take only the first k summands of the right hand side, then this upper bounds the left-hand side if k is odd, and lower bounds it if k is even.*

We sometimes use the following corollary of this claim: (this is sometimes known as the Bonefforni inequality)

Claim A.4
For every events $A_1, \ldots, A_n$,

$$\Pr[\cup_{i=1}^{n} A_i] \geq \sum_{i=1}^{n} \Pr[A_i] - \sum_{1 \leq i < j \leq n} \Pr[A_i \cap A_j]$$

Random subsum principle. The following fact is used often in the book:

Claim A.5 (The random subsum principle)
For $x, y \in \{0,1\}^n$, denote $x \odot y = \sum_{i=1}^{n} x_i y_i \pmod 2$ (that is, $x \odot y$ is equal to 1 if the number of i's such that $x_i = y_i = 1$ is odd and equal to 0 otherwise). Then for every $y \neq 0^n$,

$$\Pr_{x \in_R \{0,1\}^n}[x \odot y = 1] = \tfrac{1}{2}$$

Proof: Suppose that y_j is nonzero. We can think of choosing x as follows: first choose all the coordinates of x other than the j^{th} and only choose the j^{th} coordinate last. After we choose all the coordinates of x other than the j^{th}, the value $\sum_{i:i \neq j} x_i y_i \pmod 2$ is fixed to be some $c \in \{0,1\}$. Regardless of what c is, with probability $1/2$ we choose $x_j = 0$, in which case $x \odot y = c$ and with probability $1/2$ we choose $x_j = 1$, in which case $x \odot y = 1 - c$. We see that in any case $x \odot y$ will be equal to 1 with probability $1/2$. ■

A.3.1 Random variables and expectations.

A *random variable* is a mapping from a probability space to $\mathbb{R}$. For example, if Ω is as above, the set of all possible outcomes of n tosses of a fair coin, then we can denote by X the number of coins that came up heads.

The *expectation* of a random variable X, denoted by $\mathsf{E}[X]$, is its weighted average. That is, $\mathsf{E}[X] = \sum_{i=1}^{N} p_i X(\omega_i)$. The following simple claim follows from the definition:

Claim A.6 (Linearity of expectation)
For X, Y random variables over a space Ω, denote by $X+Y$ the random variable that maps ω to $X(\omega)+Y(\omega)$. Then,

$$\mathsf{E}[X+Y] = \mathsf{E}[X] + \mathsf{E}[Y]$$

This claims implies that the random variable X from the example above has expectation $n/2$. Indeed $X = \sum_{i=1}^{n} X_i$ where X_i is equal to 1 if the i^{th} coins came up heads and is equal to 0 otherwise. But clearly, $\mathsf{E}[X_i] = 1/2$ for every i.

For a real number α and a random variable X, we define αX to be the random variable mapping ω to $\alpha \cdot X(\omega)$. Note that $\mathsf{E}[\alpha X] = \alpha \mathsf{E}[X]$.

A.3.2 The averaging argument

We list various versions of the "averaging argument." Sometimes we give two versions of the same result, one as a fact about numbers and one as a fact about probability spaces.

Lemma A.7
If $a_1, a_2, \ldots, a_n$ are some numbers whose average is c then some $a_i \geq c$.

Lemma A.8 ("The Probabilistic Method")
If X is a random variable which takes values from a finite set and $\mathsf{E}[X] = \mu$ then the event "$X \geq \mu$" has nonzero probability.

Lemma A.9
If $a_1, a_2, \ldots, a_n \geq 0$ are numbers whose average is c then the fraction of a_i's that are greater than (resp., at least) kc is less than (resp, at most) $1/k$.

Lemma A.10 ("Markov's inequality")
Any non-negative *random variable X satisfies*

$$\mathbf{Pr}\left(X \geq k\mathsf{E}[X]\right) \leq \frac{1}{k}.$$

Corollary A.11
If $a_1, a_2, \ldots, a_n \in [0,1]$ are numbers whose average is $1-\gamma$ then at least $1-\sqrt{\gamma}$ fraction of them are at least $1-\sqrt{\gamma}$.

Can we give any meaningful upper bound on $\mathbf{Pr}[X < c \cdot \mathsf{E}[X]]$ where $c < 1$? Yes, if X is bounded.

Lemma A.12
If $a_1, a_2, \ldots, a_n$ are numbers in the interval $[0,1]$ whose average is ρ then at least $\rho/2$ of the a_i's are at least as large as $\rho/2$.

Proof: Let γ be the fraction of i's such that $a_i \geq \rho/2$. Then $\gamma + (1-\gamma)\rho/2$ must be at least $\rho/2$, so $\gamma \geq \rho/2$. ■ More generally, we have

LEMMA A.13
If $X \in [0,1]$ and $\mathsf{E}[X] = \mu$ then for any $c < 1$ we have

$$\mathbf{Pr}[X \leq c\mu] \leq \frac{1-\mu}{1-c\mu}.$$

EXAMPLE A.14
Suppose you took a lot of exams, each scored from 1 to 100. If your average score was 90 then in at least half the exams you scored at least 80.

A.3.3 Conditional probability and independence

If we already know that an event B happened, this reduces the space from Ω to $\Omega \cap B$, where we need to scale the probabilities by $1/\Pr[B]$ so they will sum up to one. Thus, the probability of an event A *conditioned* on an event B, denoted $\Pr[A|B]$, is equal to $\Pr[A \cap B]/\Pr[B]$ (where we always assume that B has positive probability).

We say that two events A, B are *independent* if $\Pr[A \cap B] = \Pr[A]\Pr[B]$. Note that this implies that $\Pr[A|B] = \Pr[A]$ and $\Pr[B|A] = \Pr[B]$. We say that a set of events $A_1, \ldots, A_n$ are *mutually independent* if for every subset $S \subset [n]$,

$$\Pr[\cap_{i \in S} A_i] = \prod_{i \in S} \Pr[A_i]\,. \tag{2}$$

We say that $A_1, \ldots, A_n$ are *k-wise independent* if (2) holds for every $S \subseteq [n]$ with $|S| \leq k$.

We say that two random variables X, Y are *independent* if for every $x, y \in \mathbb{R}$, the events $\{X = x\}$ and $\{Y = y\}$ are independent. We generalize similarly the definition of mutual independence and k-wise independence to sets of random variables $X_1, \ldots, X_n$. We have the following claim:

CLAIM A.15
If $X_1, \ldots, X_n$ are mutually independent then

$$\mathsf{E}[X_1 \cdots X_n] = \prod_{i=1}^{n} \mathsf{E}[X_i]$$

PROOF:

$$\mathsf{E}[X_1 \cdots X_n] = \sum_x x \Pr[X_1 \cdots X_n = x] =$$

$$\sum_{x_1,\ldots,x_n} x_1 \cdots x_n \Pr[X_1 = x_1 \text{ and } X_2 = x_2 \cdots \text{ and } X_n = x_n] = \text{(by independence)}$$

$$\sum_{x_1,\ldots,x_n} x_1 \cdots x_n \Pr[X_1 = x_1] \cdots \Pr[X_n = x_n] =$$

$$(\sum_{x_1} x_1 \Pr[X_1 = x_1])(\sum_{x_2} x_2 \Pr[X_2 = x_2]) \cdots (\sum_{x_n} x_n \Pr[X_n = x_n]) = \prod_{i=1}^{n} \mathsf{E}[X_i]$$

where the sums above are over all the possible real numbers that can be obtained by applying the random variables or their products to the finite set Ω. ■

A.3.4 Deviation upper bounds

Under various conditions, one can give upper bounds on the probability of a random variable "straying too far" from its expectation. These upper bounds are usually derived by clever use of Markov's inequality.

The *variance* of a random variable X is defined to be $\mathsf{Var}[X] = \mathsf{E}[(X - \mathsf{E}(X))^2]$. Note that since it is the expectation of a non-negative random variable, $\mathsf{Var}[X]$ is always non-negative. Also, using linearity of expectation, we can derive that $\mathsf{Var}[X] = \mathsf{E}[X^2] - (\mathsf{E}[X])^2$. The *standard deviation* of a variable X is defined to be $\sqrt{\mathsf{Var}[X]}$.

The first bound is Chebyshev's inequality, useful when only the variance is known.

LEMMA A.16 (CHEBYSHEV INEQUALITY)
If X is a random variable with standard deviation σ, then for every $k > 0$,

$$\Pr[|X - \mathsf{E}[X]| > k\sigma] \leq 1/k^2$$

PROOF: Apply Markov's inequality to the random variable $(X - \mathsf{E}[X])^2$, noting that by definition of variance, $E[(X - \mathsf{E}[X])^2] = \sigma^2$. ■

Chebyshev's inequality is often useful in the case that X is equal to $\sum_{i=1}^{n} X_i$ for pairwise independent random variables $X_1, \ldots, X_n$. This is because of the following claim, that is left as an exercise:

CLAIM A.17
If $X_1, \ldots, X_n$ are pairwise independent then

$$\mathsf{Var}(\sum_{i=1}^{n} X_i) = \sum_{i=1}^{n} \mathsf{Var}(X_i)$$

The next inequality has many names, and is widely known in theoretical computer science as the *Chernoff bound*. It considers scenarios of the following type. Suppose we toss a fair coin n times. The expected number of heads is $n/2$. How tightly is this number concentrated? Should we be very surprised if after 1000 tosses we have 625 heads? The bound we present is slightly more general, since it concerns n different coin tosses of possibly different expectations (the expectation of a coin is the probability of obtaining "heads"; for a fair coin this is 1/2). These are sometimes known as Poisson trials.

THEOREM A.18 ("CHERNOFF" BOUNDS)
Let $X_1, X_2, \ldots, X_n$ be mutually independent random variables over $\{0, 1\}$ (i.e., X_i can be either 0 or 1) and let $\mu = \sum_{i=1}^{n} \mathsf{E}[X_i]$. Then for every $\delta > 0$,

$$\mathbf{Pr}[\sum_{i=1}^{n} X_i \geq (1+\delta)\mu] \leq \left[\frac{e^{\delta}}{(1+\delta)^{(1+\delta)}}\right]^{\mu}. \tag{3}$$

$$\mathbf{Pr}[\sum_{i=1}^{n} \leq (1-\delta)\mu] \leq \left[\frac{e^{-\delta}}{(1-\delta)^{(1-\delta)}}\right]^{\mu}. \tag{4}$$

Often, we will only use the following corollary:

COROLLARY A.19
Under the above conditions, for every $c > 0$

$$\Pr\left[\left|\sum_{i=1}^{n} X_i - \mu\right| \geq c\mu\right] \leq 2 \cdot e^{-\min\{c^2/4, c/2\}\mu}.$$

In particular this probability is bounded by $2^{-\Omega(\mu)}$ (where the constant in the Ω notation depends on c).

PROOF: Surprisingly, the Chernoff bound is also proved using the Markov inequality. We only prove the first inequality; the second inequality can be proved similarly. We introduce a positive dummy variable t, and observe that

$$\mathsf{E}[\exp(tX)] = \mathsf{E}[\exp(t\sum_i X_i)] = \mathsf{E}[\prod_i \exp(tX_i)] = \prod_i \mathsf{E}[\exp(tX_i)], \tag{5}$$

where $\exp(z)$ denotes e^z and the last equality holds because the X_i r.v.s are independent. Now,

$$\mathsf{E}[\exp(tX_i)] = (1-p_i) + p_i e^t,$$

therefore,

$$\begin{aligned}\prod_i \mathsf{E}[\exp(tX_i)] = \prod_i [1 + p_i(e^t-1)] &\leq \prod_i \exp(p_i(e^t-1)) \\ &= \exp(\sum_i p_i(e^t-1)) = \exp(\mu(e^t-1)),\end{aligned} \tag{6}$$

as $1 + x \leq e^x$. Finally, apply Markov's inequality to the random variable $\exp(tX)$, viz.

$$\mathbf{Pr}[X \geq (1+\delta)\mu] = \mathbf{Pr}[\exp(tX) \geq \exp(t(1+\delta)\mu)] \leq \frac{\mathsf{E}[\exp(tX)]}{\exp(t(1+\delta)\mu)} = \frac{\exp((e^t-1)\mu)}{\exp(t(1+\delta)\mu)},$$

using (5), (6) and the fact that t is positive. Since t is a dummy variable, we can choose any positive value we like for it. Simple calculus shows that the right hand side is minimized for $t = \ln(1+\delta)$ and this leads to the theorem statement. ■

So, if all n coin tosses are fair (Heads has probability $1/2$) then the the probability of seeing N heads where $|N - n/2| > a\sqrt{n}$ is at most $2e^{-a^2/4}$. In particular, the chance of seeing at least 625 heads in 1000 tosses of an unbiased coin is less than 5.3×10^{-7}.

A.3.5 Some other inequalities.

Jensen's inequality.

The following inequality, generalizing the inequality $\mathsf{E}[X^2] \geq \mathsf{E}[X]^2$, is also often useful:

CLAIM A.20
We say that $f : \mathbb{R} \to \mathbb{R}$ is convex *if for every $p \in [0,1]$ and $x, y \in \mathbb{R}$, $f(px + (1-p)y) \leq p \cdot f(x) + (1-p) \cdot f(y)$. Then, for every random variable X and convex function f, $f(\mathsf{E}[X]) \leq \mathsf{E}[f(X)]$.*

Approximating the binomial coefficient

Of special interest is the *Binomial* random variable B_n denoting the number of coins that come up "heads" when tossing n fair coins. For every k, $\Pr[B_n = k] = 2^{-n}\binom{n}{k}$ where $\binom{n}{k} = \frac{n!}{k!(n-k)!}$ denotes the number of size-k subsets of $[n]$. Clearly, $\binom{n}{k} \leq n^k$, but sometimes we will need a better estimate for $\binom{n}{k}$ and use the following approximation:

CLAIM A.21
For every n, $k < n$, $\left(\frac{n}{k}\right)^k \leq \binom{n}{k} \leq \left(\frac{ne}{k}\right)^k$

The best approximation can be obtained via Stirling's formula:

LEMMA A.22 (STIRLING'S FORMULA)
For every n,

$$\sqrt{2\pi n}\left(\frac{n}{e}\right)^n e^{\frac{1}{12n+1}} < n! < \sqrt{2\pi n}\left(\frac{n}{e}\right)^n e^{\frac{1}{12n}}$$

It can be proven by taking natural logarithms and approximating $\ln n! = \ln(1 \cdot 2 \cdots n) = \sum_{i=1}^{n} \ln i$ by the integral $\int_1^n \ln x \, dx = n \ln n - n + 1$. It implies the following corollary:

COROLLARY A.23
For every $n \in \mathbb{N}$ and $\alpha \in [0,1]$,

$$\binom{n}{\alpha n} = (1 \pm O(n^{-1})) \frac{1}{\sqrt{2\pi n \alpha(1-\alpha)}} 2^{H(\alpha)n}$$

where $H(\alpha) = \alpha \log(1/\alpha) + (1-\alpha)\log(1/(1-\alpha))$ and the constants hidden in the O notation are independent of both n and α.

More useful estimates.

The following inequalities can be obtained via elementary calculus:

- For every $x \geq 1$, $\left(1 - \frac{1}{x}\right)^x \leq \frac{1}{e} \leq \left(1 - \frac{1}{x+1}\right)^x$
- For every k, $\sum_{i=1}^{n} i^k = \Theta\left(\frac{n^{k+1}}{k+1}\right)$
- For every $k > 1$, $\sum_{i=1}^{\infty} n^{-k} < O(1)$.
- For every $c, \epsilon > 0$, $\sum_{i=1}^{\infty} \frac{n^c}{(1+\epsilon)^n} < O(1)$.
- For every n, $\sum_{i=1}^{n} \frac{1}{n} = \ln n \pm O(1)$

A.4 Number theory and groups

The *integers* are the set $\mathbb{Z} = \{0, \pm 1, \pm 2, \ldots\}$ while the natural numbers are the subset $\mathbb{N} = \{1, 2, \ldots\}$.[2] A basic fact is that we can *divide* any integer n by an integer k to obtain ℓ, r such that $n = k\ell + r$ and $r \in \{0, \ldots, n-1\}$. If $r = 0$ then we say that k *divides* n and denote this by $k|n$. The *factors* of n are the set of positive integers that divide n.

The *greatest common divisor* of two integers n, m, denoted by $gcd(n,m)$ is the largest integers d such that $d|n$ and $d|m$. We say that n and m are *co-prime* if their greatest common divisor is equal to 1. The following basic facts are not hard to verify:

- If a natural number c divides both n and m then $c|d$.
- The greatest common divisor of n and m is the smallest natural number d such that there exist integers x, y satisfying $nx + my = d$.
- There is a polynomial-time (i.e., polylog(n,m)-time) algorithm that on input n, m outputs the greatest common divisor d of n, m and the integers x, y satisfying $nx + my = d$. (This algorithm is known as *Euclid's Algorithm*.)

A number $p > 1$ is *prime* if its only factors are 1 and p. The following basic facts are known about prime numbers:

- Every natural number n can be written uniquely (up to ordering) as a product of prime numbers. This is called the *prime factorization* of n.
- If $gcd(p,a) = 1$ and $p|ab$ then $p|b$. In particular, if a prime p divides $a \cdot b$ then either $p|a$ or $p|b$.

[2] Some texts include 0 in $\mathbb{N}$; in most cases this does not any difference.

A fundamental question in number theory is how many primes exist. A celebrated result is:

THEOREM A.24 (THE PRIME NUMBER THEOREM (HADAMARD, DE LA VALLÉE POUSSIN 1896)) *Let $\pi(n)$ denote the number of primes between 1 and n then*

$$\pi(n) = \frac{n}{\ln n}(1 \pm o(1))$$

The original proofs of the prime number theorem used rather deep mathematical tools, and in fact people have conjectured that this is *inherently* the case. But in 1949 both Erdös and Selberg (independently) found elementary proofs for this theorem. For most computer science applications, the following weaker statement proven by Chebychev suffices:

THEOREM A.25
$\pi(n) = \Theta(\frac{n}{\log n})$

PROOF: Consider the number $\binom{2n}{n} = \frac{2n!}{n!n!}$. By Stirling's formula we know that $\log \binom{2n}{n} = (1-o(1))2n$ and in particular $n \leq \log \binom{2n}{n} \leq 2n$. Also, all the prime factors of $\binom{2n}{n}$ are between 0 and $2n$, and each factor p cannot appear more than $k = \lfloor \frac{\log 2n}{\log p} \rfloor$ times. Indeed, for every n, the number of times p appears in the factorization of $n!$ is $\sum_i \lfloor \frac{n}{p^i} \rfloor$, since we get $\lfloor \frac{n}{p} \rfloor$ times a factor p in the factorizations of $\{1, \ldots, n\}$, $\lfloor \frac{n}{p^2} \rfloor$ times a factor of the form p^2, etc... Thus the number of times p appears in the factorization of $\binom{2n}{n} = \frac{(2n)!}{n!n!}$ is equal to $\sum_i \lfloor \frac{2n}{p^i} \rfloor - 2\lfloor \frac{n}{p^i} \rfloor$: a sum of at most k elements (since $p^{k+1} > 2n$) each of which is either 0 or 1.

Thus, $\binom{2n}{n} \leq \prod_{\substack{1 \leq p \leq 2n \\ p \text{ prime}}} p^{\lfloor \frac{\log 2n}{\log p} \rfloor}$. Taking logs we get that

$$n \leq \log \binom{2n}{n} \leq \sum_{\substack{1 \leq p \leq 2n \\ p \text{ prime}}} \lfloor \tfrac{\log 2n}{\log p} \rfloor \log p \,,$$

or $\frac{n}{\log n} \leq \pi(2n)$, establishing $\pi(n) = \Omega(\frac{n}{\log n})$.

To prove that $\pi(n) = O(\frac{n}{\log n})$, we define the function $\vartheta(n) = \sum_{\substack{1 \leq p \leq n \\ p \text{ prime}}} \log p$. It suffices to prove that $\vartheta(n) = O(n)$ (exercise!). But since all the primes between $n+1$ and $2n$ divide $\binom{2n}{n}$ at least once, $\binom{2n}{n} \geq \prod_{\substack{n+1 \leq p \leq 2n \\ p \text{ prime}}} p$. Taking logs we get

$$2n \geq \log \binom{2n}{n} \geq \sum_{\substack{n+1 \leq p \leq 2n \\ p \text{ prime}}} \log p = \vartheta(2n) - \vartheta(n) \,,$$

thus getting a recursive equation $\vartheta(2n) \leq \vartheta(n) + 2n$ which solves to $\vartheta(n) = O(n)$. ■

A.4.1 Groups.

A *group* is an abstraction that captures some properties of mathematical objects such as the integers, matrices, functions and more. Formally, a group is a set that has a binary operation, say $\star$, defined on it that is associative and has an inverse. That is, $(G, \star)$ is a group if

1. For every $a, b, c \in G$, $(a \star b) \star c = a \star (b \star c)$
2. There exists a special element $id \in G$ such that $a \star id = a$ for every $a \in G$, and for every $a \in G$ there exists $b \in G$ such that $a \star b = b \star a = id$.

Examples for groups are the integers, with addition being the group operation (and zero the identity element), the non-zero real numbers with multiplication being the group operation(and one the identity element), and the set of functions from a domain A to itself, with function composition being the group operation.

Often, it is natural to use additive ($+$) or multiplicative ($\cdot$) notation to denote the group operation rather than $\star$. In these cases we will use ℓa (or respectively a^ℓ) to denote the result of applying the operation to a ℓ times.

A.4.2 Finite groups

A group is *finite* if it has a finite number of elements. We denote by $|G|$ the number of elements of G. Examples for finite groups are the following:

- The group $\mathbb{Z}_n$ of the integers from 0 to $n-1$ with the operation being addition modulo n. In particular $\mathbb{Z}_2$ is the set $\{0,1\}$ with the XOR operation.
- The group S_n of the permutations on $[n]$, with the operation being function composition.
- The group $(\mathbb{Z}_2)^n$ of n-bit strings with the operation being bitwise XOR. More generally for every two groups G and H, we can define the group $G \times H$ to be a group whose elements are pairs $\langle g, h \rangle$ with $g \in G$ and $h \in H$ and with the group operation corresponding to applying the group operations of G and H componentwise. Similarly, we define G^n to be the group $G \times G \times \cdots \times G$ (n times).
- For every n, we define $\mathbb{Z}_n^*$ to be the set $\{k : 1 \leq k \leq n-1 \ , \ gcd(k,n) = 1\}$ with the operation being multiplication modulo n. Note that if $gcd(k,n) = 1$ then there exist x, y such that $kx + ny = 1$ or in other words $kx = 1 \pmod n$, meaning that x is the inverse of k modulo n. This also means that we can find this inverse in polynomial time using Euclid's algorithm. The size of $\mathbb{Z}_n^*$ is denoted by $\varphi(n)$ and the function φ is known as *Euler's totient function*. Note that if n is prime then $\varphi(n) = n-1$. It is known that for every $n > 6$, $\varphi(n) \geq \sqrt{n}$.

A *subgroup* of G is a subset of G that is itself a group (i.e., closed under the group operation and taking inverses). The following result is often quite useful

THEOREM A.26
If G is a finite group and H is a subgroup of G then $|H|$ divides $|G|$.

PROOF: Consider the family of sets of the form $aH = \{ah : h \in H\}$ for all $a \in G$ (we're using here multiplicative notation for the group). It is easy to see that the map $x \mapsto ax$ is one-to-one and hence $|aH| = |H|$ for every a. Hence it will suffice to show that we can partition G into disjoint sets from this family. Yet this family clearly covers G (as $a \in aH$ for every $a \in G$) and hence it suffices to show that for every a, b either $aH = bH$ or aH and bH are disjoint. Indeed, suppose that there exist $x, y \in H$ such that $ax = by$ then for every element $az \in aH$, we have that $az = (byx^{-1})z$ and since $yx^{-1}z \in H$ we get that $az \in bH$. ■

COROLLARY A.27 (FERMAT'S LITTLE THEOREM)
For every n and $x \in \{1, \ldots, n-1\}$, $x^{\varphi(n)} = 1 \pmod n$. In particular, if n is prime then $x^{n-1} = 1 \pmod n$.

PROOF: Consider the set $H = \left\{x^\ell : \ell \in \mathbb{Z}\right\}$. This is clearly a subgroup of $\mathbb{Z}_n^*$ and hence $|H|$ divides $\varphi(n)$. But the size of H is simply the smallest number k such that $x^k = 1 \pmod n$. Indeed, consider the numbers $1, x, x^2, x^3, \ldots$. Since $\mathbb{Z}_n^*$ is finite, eventually we get i, j such that $x^i = x^j$ for $i < j$, which means that $x^{i-j} = 1 \pmod n$. In other words, this sequence looks like $1, x, x^2, \ldots, x^{k-1}, 1, x, x^2, \ldots$ and so all the elements of H are of the form $x^i \pmod n$ for $i \in [0, k-1]$.

Since $x^{|H|} = 1 \pmod{n}$ obviously taking x to the power $\varphi(n)$ (which is a multiple of $|H|$ yields also 1 modulo n. ■

The *order* of an element x of a group G is the smallest integer k such that x^k is equal to the identity element. The proof above shows that in a finite group G, every element has a finite order and furthermore this order divides the size of G. If an x element of G has order $|G|$ then we say it is a *generator* of G, since in this case the subgroup $\{x, x^1, x^2, \ldots\}$ is all of G.[3] If a group G has a generator then we say that G is *cyclic*. An example for a simple cyclic group is the group $\mathbb{Z}_n$ of the numbers $\{0, \ldots, n-1\}$ with addition modulo n, that is generated by the element 1 (and also by any other element that is co-prime to n— exercise).

A.4.3 The Chinese Remainder Theorem

Let $n = pq$ where p, q are co-prime. The Chinese Remainder Theorem (CRT) says that the group $\mathbb{Z}_n^*$ (multiplicative group modulo n) is isomorphic to the group $Z_p^* \times \mathbb{Z}_q^*$ (pairs of numbers with multiplication done componentwise modulo p and q respectively).

THEOREM A.28
If $n = pq$ where p, q coprime then function f that maps x to $\langle x \pmod{p}, x \pmod{q}\rangle$ is one-to-one on $\mathbb{Z}_n^*$. Furthermore f is an isomorphism in the sense that $f(xy) = f(x)f(y)$ (where multiplication on the left hand side is modulo n and on the right hand side is componentwise modulo p and q respectively).

PROOF: The furthermore part can be easily verified and so we focus on showing that f is one-to-one. We need to show that if $f(x) = f(x')$ then $x = x'$. Since $f(x - x') = f(x) - f(x')$, it suffices to show that if $x = 0 \pmod{p}$ (i.e., $p|x$) and $x = 0 \pmod{q}$ (i.e., $q|x$) then $x = 0 \pmod{n}$ (i.e., $pq|x$). Yet, assume that $p|x$ and write $x = pk$. Then since $gcd(p, q) = 1$ and $q|x$ we know that $q|k$, meaning that $pq|x$. ■

The Chinese Remainder Theorem can be easily generalized to show that for every $n = p_1 p_2 \ldots p_k$, where all the p_i's are co-prime, there is an isomorphism between Z_n^* to $\mathbb{Z}_{p_1}^* \times \cdots \times \mathbb{Z}_{p_k}^*$, meaning that for every n, the group $\mathbb{Z}_n^*$ is isomorphic to a product of groups of the form $\mathbb{Z}_q^*$ for q a prime power (i.e., number of the form p^ℓ for prime p). In fact, it can be generalized even further to show that every Abelian group G is isomorphic to a product $G_1 \times G_2 \cdots G_k$ where all the G_i's are cyclic. (This is a generalization of the CRT because all the groups of the form $\mathbb{Z}_q^*$ for q a power of an odd prime are cyclic, and all groups of the form $\mathbb{Z}_{2^k}^*$ are either cyclic or products of two cyclic groups.)

A.5 Finite fields and groups

A *field* is a set $\mathbb{F}$ that has an addition ($+$) and multiplication ($\cdot$) operations that behave in the expected way: satisfy associative, commutative and distributive laws, have both additive and multiplicative inverses, and neutral elements 0 and 1 for addition and multiplication respectively. In other words, $\mathbb{F}$ is a field if it is an Abelian group with the operation $+$ and an identity element 0, and has an additional operation $\cdot$ such that $\mathbb{F} \setminus \{0\}$ and $\cdot$ forms an Abelian group, and furthermore the two operation satisfy the distributive rule $a(b + c) = ab + ac$.

Familiar fields are the real numbers ($\mathbb{R}$), the rational numbers ($\mathbb{Q}$) and the complex numbers ($\mathbb{C}$), but there are also *finite* fields. Recall that for a prime p, the set $\{0, \ldots, p-1\}$ is an Abelian group with the addition modulo p operation and the set $\{1, \ldots, p-1\}$ is an Abelian group with the multiplication modulo p operation. Hence $\{0, \ldots, p-1\}$ form a field with these two operations, which we denote by GF(p). The simplest example for such a field is the field GF(2) consisting of $\{0, 1\}$ where multiplication is the AND ($\wedge$) operation and addition is the XOR operation.

[3] A more general definition (that works also for infinite groups) is that x is a generator of G if the subgroup $\{x^\ell : \ell \in \mathbb{Z}\}$ is equal to G.

Every finite field $\mathbb{F}$ has a number ℓ such that for every $x \in F$, $x + x + \cdots + x$ (ℓ times) is equal to the zero element of $\mathbb{F}$ (exercise). This number ℓ is called the *characteristic* of $\mathbb{F}$. For every prime q, the characteristic of $\mathrm{GF}(q)$ is equal to q.

A.5.1 Non-prime fields.

One can see that if n is not prime, then the set $\{0, \ldots, n-1\}$ with addition and multiplication modulo n is not a field, as there exist two non-zero elements x, y in this set such that $x \cdot y = n = 0 \pmod{n}$. Nevertheless, there are finite fields of size n for non-prime n. Specifically, for every prime q, and $k \geq 1$, there exists a field of q^k elements, which we denote by $\mathrm{GF}(q^k)$. We will very rarely need to use such fields in this book, but still provide an outline of their construction below.

For every prime q and k there exists an *irreducible* degree k polynomial P over the field $\mathrm{GF}(q)$ (P is irreducible if it cannot be expressed as the product of two polynomials P', P'' of lower degree). We then let $\mathrm{GF}(q^k)$ be the set of all $k-1$-degree polynomials over $\mathrm{GF}(q)$. Each such polynomial can be represented as a vector of its k coefficients. We perform both addition and multiplication modulo the polynomial P. Note that addition corresponds to standard vector addition of k-dimensional vectors over $\mathrm{GF}(q)$, and both addition and multiplication can be easily done in $\mathrm{poly}(n, \log q)$ time (we can reduce a polynomial S modulo a polynomial P using a similar algorithm to long division of numbers). It turns out that no matter how we choose the irreducible polynomial P, we will get the same field, up to renaming of the elements. There is a deterministic $\mathrm{poly}(q, k)$-time algorithm to obtain an irreducible polynomial of degree k over $\mathrm{GF}(q)$. There are also probabilistic algorithms (and deterministic algorithms whose analysis relies on unproven assumptions) that obtain such a polynomial in $\mathrm{poly}(\log q, k)$ time.

For us, the most important example of a finite field is $\mathrm{GF}(2^k)$, which consists of the set $\{0,1\}^k$, with addition being component-wise XOR, and multiplication being polynomial multiplication via some irreducible polynomial which we can find in $\mathrm{poly}(k)$ time. In fact, we will mostly not even be interested in the multiplicative structure of $\mathrm{GF}(2^k)$ and only use the addition operation (i.e., use it as the vector space $\mathrm{GF}(2)^k$, see below).

A.6 Vector spaces and Hilbert spaces

For $\mathbb{F}$ a field and $n \in \mathbb{N}$, we denote by $\mathbb{F}^n$ the set of n-length tuples (or *vectors*) of elements of $\mathbb{F}$. If $\mathbf{u}, \mathbf{v} \in \mathbb{F}^n$ and $x \in \mathbb{F}$ then we denote by $\mathbf{u} + \mathbf{v}$ the vector obtained by componentwise addition of $\mathbf{u}$ and $\mathbf{v}$ and by $x\mathbf{u}$ the vector obtained by multiplying each entry of $\mathbf{u}$ by x.

A set of vectors $\mathbf{u}^1, \ldots, \mathbf{u}^k$ in $\mathbb{F}^n$ is *linearly independent* if the only solution to the equation $x_1\mathbf{u}^1 + \cdots + x_k\mathbf{u}^k = \mathbf{0}$ (where $\mathbf{0}$ denotes the all-zero vector) is $x_1 = x_2 = \cdots = x_k = 0$. It can be shown that if $\mathbf{u}^1, \ldots, \mathbf{u}^k$ are linearly independent then $k \leq n$ (exercise). A set of n linearly independent vectors in $\mathbb{F}^n$ is called a *basis* of $\mathbb{F}^n$. It is not hard to see that if $\mathbf{u}^1, \ldots, \mathbf{u}^n$ is a basis of $\mathbb{F}^n$ then every vector $\mathbf{v} \in \mathbb{F}^n$ can be expressed as a linear combination $\mathbf{v} = \sum_i x_i \mathbf{u}^i$ of the vectors $\mathbf{u}^1, \ldots, \mathbf{u}^n$ and furthermore this expression is unique. The *standard basis* of $\mathbb{F}^n$ is the set $\mathbf{e}^1, \ldots, \mathbf{e}^n$, where $\mathbf{e}^i_j$ is equal to 1 if $j = i$ and to 0 otherwise.

A subset $S \subseteq \mathbb{F}^n$ is called a *subspace* if it is closed under addition and scalar multiplication (i.e., $\mathbf{u}, \mathbf{v} \in S$ and $x, y \in \mathbb{F}$ implies that $x\mathbf{u} + y\mathbf{v} \in S$). The *dimension* of S, denoted by $\dim(S)$ is defined to be the maximum number k such that there are k linearly independent vectors in S. Such a set of $\dim(S)$ linearly independent vectors in S is called a *basis* and one can see that every vector in S can be expressed as a linear combination of the vectors in the basis.

A function $f : \mathbb{F}^n \to \mathbb{F}^m$ is *linear* if $F(\mathbf{u} + \mathbf{v}) = F(\mathbf{u}) + F(\mathbf{v})$. The following facts are not too hard to verify:

- If $\mathbf{u}^1, \ldots, \mathbf{u}^n$ is a basis for $\mathbb{F}^n$ then for every $\mathbf{v} \in \mathbb{F}^n$, $F(\mathbf{v}) = \sum_i x_i F(\mathbf{u}_i)$ where $x_1, \ldots, x_n$ are the elements such that $\mathbf{v} = \sum x_i \mathbf{u}^i$. Thus, to know f's value at every point it suffices to know its value on the basis elements.

- The set $Im(f) = \{f(\mathbf{v}) : \mathbf{v} \in \mathbb{F}^n\}$ is a subspace of $\mathbb{F}^m$.
- The set $Ker(f) = \{\mathbf{v} : f(\mathbf{v}) = 0\}$ is a subspace of $\mathbb{F}^n$.
- $\dim(Im(f)) + \dim(Ker(f)) = n$

A linear function $f : \mathbb{F}^n \to \mathbb{F}^m$ is often described by an $m \times n$ matrix A whose i^{th} column is $f(\mathbf{e}^i)$. The *multiplication* of an $m \times n$ matrix A and an $n \times k$ matrix B is the $n \times k$ matrix $C = AB$ where $C_{i,j} = \sum_{\ell \in [n]} A_{i,\ell} B_{\ell,j}$. One can verify that if A describes a function $f : \mathbb{F}^n \to \mathbb{F}^m$ and B describes a function $g : \mathbb{F}^k \to \mathbb{F}^n$ then C describes the function $h : \mathbb{F}^k \to \mathbb{F}^m$ mapping $\mathbf{v}$ to $f(g(\mathbf{v}))$. It can also be verified that if we identify members of $\mathbb{F}^n$ with $n \times 1$ matrices (i.e., column vectors) then $f(\mathbf{v}) = A\mathbf{v}$.

The *determinant* of an $n \times n$ matrix A, denoted by $\det(A)$ is equal to $\sum_{\sigma \in S_n} (-1)^{sgn(\sigma)} \prod_{i=1}^n A_{i,\sigma(i)}$ where S_n denotes the group of permutations over $[n]$ and $sgn(\sigma)$ is equal to 1 if the number of pairs $\langle i, j \rangle$ such that $i < j$ but $\sigma(i) > \sigma(j)$ is odd, and is equal to 0 otherwise. We have the following two facts:

- $\det(AB) = \det(A)\det(B)$. This can be verified by direct computation.
- If A is an upper diagonal matrix (i.e., $A_{i,j} = 0$ whenever $i > j$) then $\det(A) = A_{1,1} A_{2,2} \cdots A_{n,n}$. Indeed, for a permutation σ to give a non-zero contribution to the determinant in this case it must satisfy $\sigma(i) \geq i$ for every i, which means that it is the identity permutation.

Together these two rules give a polynomial-time algorithm to compute the determinant of a matrix A by following the well known Gaussian elimination algorithm to express A as $E_1 E_2 \cdots E_m D$ where the E_i's are elementary matrices (multiplication by which corresponds to switching two columns, multiplying a column by a field element, or adding one column to another) and the D is upper diagonal. Since the determinant is easy to compute for all these matrices, we can compute the determinant of A as well.

The following lemma relates the determinant of a matrix to the function it represents:

LEMMA A.29
For a function $f : \mathbb{F}^n \to \mathbb{F}^n$ represented by an $n \times n$ matrix A, the following conditions are equivalent:

- *f is one-to-one.*
- $\dim(Im(f)) = n$.
- $\dim(Ker(f)) = 0$.
- $\det(A) \neq 0$.
- *There exists $\mathbf{v} \in \mathbb{F}^n$ such that the equation $A\mathbf{x} = \mathbf{v}$ has exactly one solution.*
- *For every $\mathbf{v} \in \mathbb{F}^n$, the equation $A\mathbf{x} = \mathbf{v}$ has exactly one solution.*

Furthermore, if f is one-to-one then the mapping f^{-1} is linear and is represented by an $n \times n$ matrix A^{-1} whose $(i,j)^{th}$ entry is $\frac{\det(A_{-(i,j)})}{\det(A)}$, where $A_{-(i,j)}$ denotes the $(n-1) \times (n-1)$ matrix obtained by removing the i^{th} row and j^{th} column from A.

A.6.1 Inner product

The vector spaces $\mathbb{R}^n$ and $\mathbb{C}^n$ have an additional structure that is often quite useful.[4] An *inner product* over $\mathbb{C}^n$ to be a function mapping two vectors $\mathbf{u}, \mathbf{v}$ to a complex number $\langle \mathbf{u}, \mathbf{v} \rangle$ satisfying the following conditions:

- $\langle x\mathbf{u} + y\mathbf{w}, \mathbf{v} \rangle = x\langle \mathbf{u}, \mathbf{v} \rangle + y\langle \mathbf{w}, \mathbf{v} \rangle$
- $\langle \mathbf{v}, \mathbf{u} \rangle = \overline{\langle \mathbf{u}, \mathbf{v} \rangle}$ where $\overline{z}$ denotes complex conjugation (i.e., if $z = a + ib$ then $\overline{z} = a - ib$).
- For every $\mathbf{u}$, $\langle \mathbf{u}, \mathbf{u} \rangle$ is a non-negative real number with $\langle \mathbf{u}, \mathbf{u} \rangle = 0$ iff $\mathbf{u} = \mathbf{0}$.

The two examples for inner products we will use are the standard inner product mapping $\mathbf{x}, \mathbf{y} \in \mathbb{C}^n$ to $\sum_{i=1}^n \mathbf{x}_i \overline{\mathbf{y}}_i$ and the expectation or normalized inner product mapping $\mathbf{x}, \mathbf{y} \in \mathbb{C}^n$ to $\frac{1}{n}\sum_{i=1}^n \mathbf{x}_i \overline{\mathbf{y}}_i$. We can also define inner products over the space $\mathbb{R}^n$, in which case we drop the conjugation.

If $\langle \mathbf{u}, \mathbf{v} \rangle = 0$ we say that $\mathbf{u}$ and $\mathbf{v}$ are *orthogonal* and denote this by $\mathbf{u} \perp \mathbf{v}$. We have the following result:

LEMMA A.30
If non-zero vectors $\mathbf{u}^1, \ldots, \mathbf{u}^k$ satisfy $\mathbf{u}^i \perp \mathbf{u}^j$ for all $i \neq j$ then they are linearly independent.

PROOF: Suppose that $\sum_i x_i \mathbf{u}^i = \mathbf{0}$ and consider take an inner product of this vector with itself. We get that

$$0 = \langle \sum_i x_i \mathbf{u}^i, \sum_j x_j \mathbf{u}^j \rangle = \sum_{i,j} x_i \overline{x}_j \langle \mathbf{u}^i, \mathbf{u}^j \rangle = \sum_i |x_i|^2 \langle \mathbf{u}^i, \mathbf{u}^j \rangle \,, \tag{7}$$

where the last equality follows from the fact that $\langle \mathbf{u}^i, \mathbf{u}^j \rangle = 0$ for $i \neq j$. But unless all the x_i's are zero, the righthand side of (7) is strictly positive. (Recall that for a complex number $x = a + ib$, $|x| = \sqrt{a^2 + b^2}$ and $|x|^2 = x\overline{x}$.) ■

A set $\mathbf{u}^1, \ldots, \mathbf{u}^n$ of nonzero vectors in $\mathbb{C}^n$ satisfying $\langle \mathbf{u}^i, \mathbf{u}^j \rangle = 0$ for $i \neq j$ is called an *orthogonal basis* of $\mathbb{C}^n$. If in addition $\langle \mathbf{u}^i, \mathbf{u}^i \rangle = 1$ for all i then we say this is an *orthonormal basis*. An orthonormal basis consists of n linearly independent vectors and hence as its name implies is a basis of $\mathbb{C}^n$, meaning that every vector $\mathbf{v}$ can be expressed as $\mathbf{v} = \sum_i x_i \mathbf{u}^i$. By taking an inner product of this equality with $\mathbf{u}^i$, one can see that $x_i = \langle \mathbf{v}, \mathbf{u}^i \rangle$

The following identity (that can be viewed as a generalization of the Pythagorean theorem) is often useful:

LEMMA A.31 (PARSEVAL'S IDENTITY)
If $\mathbf{u}^1, \ldots, \mathbf{u}^n$ is an orthonormal basis for $\mathbb{C}^n$, then for every $\mathbf{v}$,

$$\langle \mathbf{v}, \mathbf{v} \rangle = \sum_{i=1}^n |x_i|^2 \,,$$

where $x_1, \ldots, x_n$ are the numbers such that $\mathbf{v} = \sum_i x_i \mathbf{u}^i$.

PROOF: As in the proof of Lemma A.30,

$$\langle \mathbf{v}, \mathbf{v} \rangle = \langle \sum_i x_i \mathbf{u}^i, \sum_j x_j \mathbf{u}^j \rangle = \sum_i |x_i|^2 \langle \mathbf{u}^i, \mathbf{u}^i \rangle \,.$$

■

[4] The reason we restrict ourselves to these fields is that they have an infinite characteristic and hence there does not exist a number k such that $\mathbf{v} + \mathbf{v} + \cdots + \mathbf{v} = \mathbf{0}$. You can check that if there is such a number for a field $\mathbb{F}$ then there will not be an inner product over $\mathbb{F}^n$.

A.6.2 Norm

A *norm* of a vector in $\mathbb{C}^n$ is a function mapping a vector $\mathbf{v}$ to a real number $\|\mathbf{v}\|$ satisfying:

- For every $\mathbf{v}$, $\|\mathbf{v}\| \geq 0$ with $\|\mathbf{v}\| = 0$ iff $\mathbf{v} = \mathbf{0}$.
- If $x \in \mathbb{C}$ then $\|x\mathbf{v}\| = |x|\|\mathbf{v}\|$.
- (Triangle inequality) For every $\mathbf{u}, \mathbf{v}$, $\|\mathbf{u} + \mathbf{v}\| \leq \|\mathbf{u}\| + \|\mathbf{v}\|$.

Again, one can make this definition also for real vectors. It's not hard to see that the function $\|\mathbf{v}\| = \sqrt{\langle \mathbf{v}, \mathbf{v} \rangle}$ is a valid norm (called the *Euclidean* or L_2 norm). But there are also other valid norms. In particular, for every $p \geq 1$, the function $\|\mathbf{v}\|_p = (\sum_i |\mathbf{v}_i|^p)^{1/p}$ is also a valid norm. See Note 20.1 for more on these different norms and the relations between them.

A.7 Polynomials

We list some basic facts about univariate polynomials.

THEOREM A.32
A nonzero polynomial of degree d has at most d distinct roots.

PROOF: Suppose $p(x) = \sum_{i=0}^{d} c_i x^i$ has $d+1$ distinct roots $\alpha_1, \ldots, \alpha_{d+1}$ in some field $\mathbb{F}$. Then

$$\sum_{i=0}^{d} \alpha_j^i \cdot c_i = p(\alpha_j) = 0,$$

for $j = 1, \ldots, d+1$. This means that the system $\mathbf{A}\mathbf{y} = \mathbf{0}$ with

$$\mathbf{A} = \begin{pmatrix} 1 & \alpha_1 & \alpha_1^2 & \ldots & \alpha_1^d \\ 1 & \alpha_2 & \alpha_2^2 & \ldots & \alpha_2^d \\ \ldots & \ldots & \ldots & \ldots & \ldots \\ 1 & \alpha_{d+1} & \alpha_{d+1}^2 & \ldots & \alpha_{d+1}^d \end{pmatrix}$$

has a solution $\mathbf{y} = \mathbf{c}$. The matrix $\mathbf{A}$ is a *Vandermonde* matrix, and it can be shown that

$$\det \mathbf{A} = \prod_{i>j} (\alpha_i - \alpha_j),$$

which is nonzero for distinct α_i. Hence rank$\mathbf{A} = d+1$. The system $\mathbf{A}\mathbf{y} = \mathbf{0}$ has therefore only a trivial solution — a contradiction to $\mathbf{c} \neq \mathbf{0}$. ■

This theorem has an interesting corollary:

COROLLARY A.33
For every finite field $\mathbb{F}$, the multiplicative group $\mathbb{F}^*$ is cyclic.

PROOF: The fact that the polynomial $x^k - 1$ has at most k roots implies that the group $\mathbb{F}^*$ has the property **(*)** that for every k the number of elements x satisfying $x^k = 1$ is always at most k. We will prove by induction that every group G satisfying **(*)** is cyclic.

Let $n = |G|$. We consider three cases:

- n is prime. In this case every element of G has either order 1 or order n. Since the only element with order 1 is the identity element, we see that G has an element of order n— G is cyclic.

- $n = p^c$ for some prime p and $c > 1$. In this case if there is no element of order n, then all the orders must divide p^{c-1}. We get $n = p^c$ elements x such that $x^{p^{c-1}} = 1$, violating **(*)**.
- $n = pq$ for co-prime p and q. In this case let H and F be two subgroups of G defined as follows: $H = \{a : a^p = 1\}$ and $F = \{b : b^q = 1\}$. Then $|H| \leq p < n$ and $|F| \leq q < n$ and also as subgroups of G both H and F satisfy **(*)**. Thus by the induction hypothesis both H and F are cyclic and have generators a and b respectively. We claim that ab generates the entire group G. Indeed, let c be any element in G. Since p, q are coprime, there are x, y such that $xq + yp = 1$ and hence $c = c^{xq+yp}$. But $(c^{xq})^p = 1$ and $(c^{yp})^q = 1$ and hence c is a product of an element of H and an element of F, and hence $c = a^i b^j$ for some $i \in \{0, \ldots, p-1\}$ and $j \in \{0, \ldots, q-1\}$. Thus, to show that $c = (ab)^z$ for some z all we need to do is to find z such that $z = i \pmod p$ and $z = j \pmod q$, but this can be done using the Chinese Remainder Theorem.

■

THEOREM A.34
For any set of pairs $(a_1, b_1), \ldots, (a_{d+1}, b_{d+1})$ there exists a unique polynomial $g(x)$ of degree at most d such that $g(a_i) = b_i$ for all $i = 1, 2, \ldots, d+1$.

PROOF: The requirements are satisfied by *Lagrange Interpolating Polynomial*:

$$\sum_{i=1}^{d+1} b_i \cdot \frac{\prod_{j \neq i}(x - a_j)}{\prod_{j \neq i}(a_i - a_j)}.$$

If two polynomials $g_1(x), g_2(x)$ satisfy the requirements then their difference $p(x) = g_1(x) - g_2(x)$ is of degree at most d, and is zero for $x = a_1, \ldots, a_{d+1}$. Thus, from the previous theorem, polynomial $p(x)$ must be zero and polynomials $g_1(x), g_2(x)$ identical. ■

The following elementary result is usually attributed to Schwartz and Zippel in the computer science community, though it was certainly known earlier (see e.g. DeMillo and Lipton [**??**]).

LEMMA A.35
If a polynomial $p(x_1, x_2, \ldots, x_m)$ over $F = GF(q)$ is nonzero and has total degree at most d, then

$$\Pr[p(a_1..a_m) \neq 0] \geq 1 - \frac{d}{q},$$

where the probability is over all choices of $a_1..a_m \in F$.

PROOF: We use induction on m. If $m = 1$ the statement follows from Theorem A.32. Suppose the statement is true when the number of variables is at most $m - 1$. Then p can be written as

$$p(x_1, x_2, \ldots, x_m) = \sum_{i=0}^{d} x_1^i p_i(x_2, \ldots, x_m),$$

where p_i has total degree at most $d - i$. Since p is nonzero, at least one of p_i is nonzero. Let k be the largest i such that p_i is nonzero. Then by the inductive hypothesis,

$$\Pr_{a_2, a_3, \ldots, a_m}[p_i(a_2, a_3, \ldots, a_m) \neq 0] \geq 1 - \frac{d-k}{q}.$$

Whenever $p_i(a_2, a_3, \ldots, a_m) \neq 0$, $p(x_1, a_2, a_3, \ldots, a_m)$ is a nonzero univariate polynomial of degree k, and hence becomes 0 only for at most k values of x_1. Hence

$$\Pr[p(a_1..a_m) \neq 0] \geq (1 - \frac{k}{q})(1 - \frac{d-k}{q}) \geq 1 - \frac{d}{q},$$

and the induction is completed. ■