

AN INTRODUCTION TO MATHEMATICAL CRYPTOGRAPHY
ERRATA FOR THE SECOND EDITION

JEFFREY HOFFSTEIN, JILL PIPHER, JOSEPH H. SILVERMAN

Acknowledgements

We would like to thank the following people who have sent us comments and corrections for the 2nd edition: Jeff Achter, Nicole Andre, Tom Hales, Alex Kontorovich, Benedikt Niemöller, Finn de Ridder, Jeremy Roach, Benjamin Schwendinger,

Page 71, Line 12 (two corrections)

“She can this by first computing c_1^1 ” should be “She can do this by first computing c_1^a ”

Page 89, Theorem 2.31 (Pohlig-Hellman Algorithm)

The rough estimate for the running time leaves out some factors. For each g_i, h_i in Step (1), it seems that one would need to do $N/q_i^{e_i}$ group operations before solving the DLP? And since the baseline algorithm (baby step giant step) counts the number of group operations performed as part of its running time, Pohlig-Hellman should probably do the same as well. This means that we need to add $N/q_i^{e_i}$ to each $S_{q_i^{e_i}}$.

Page 153, Definition

It should possibly be stressed further that in computer science, exponential growth is in terms of the bit-size of the input. Thus in CS, the function N^2 exhibits exponential growth and $(\log N)^2$ exhibits polynomial growth as a function of the bit-size of N .

Page 199, Line 8

“The quantity a ” should be “The quantity A ”. This is a serious typo, since a is Alice’s secret exponent.

Page 167, Line -6

The sentence “We note that $g = 37$ is $\dots p = 18443$ ” is missing a period at the end.

Page 182, Exercise 3.10(b)

The decryption exponent d listed in this exercise is

$$16784693 = e^{-1} \text{ modulo } \frac{(p-1)(q-1)}{2}.$$

This works fine as a decryption exponent. However, it would also be okay to use the alternative decryption exponent

$$36153251 = e^{-1} \text{ modulo } (p-1)(q-1).$$

Page 239, Definition at top of page

The notation $\Pr(X = x)$ is not explicitly defined, although it is reasonably clear from the previous paragraph. But might be worth adding:

$$\Pr(X = x) = \Pr(\{\omega \in \Omega : X(\omega) = x\}),$$

and similarly for $\Pr(X \leq x)$ and $\Pr(X > x)$.

Page 251, Line 6

“Next we choose additional random exponents $z_1, z_2, \dots, z_n$ between 1 and k ” should be “Next we choose additional random exponents $z_1, z_2, \dots, z_n$ between 1 and N ”

Page 264, Definition

Warning: The definition of perfect secrecy is not consistent with some other books. For example, Katz and Lindell’s *Introduction to Modern Cryptography* requires that the given identity hold for all distributions over the message space.

Page 266, Proposition 5.55

It might be better to use $\mathcal{M}^+$, instead of $\mathcal{C}^+$, since $\mathcal{C}$ has been used to denote the space of cipher texts.

Page 270, Property H_3

Property H_3 is incorrect. First, the elements x_{ij} should all be distinct. Secondly, the penultimate displayed formula on page 270 should read

$$\Pr(X = x_{ij}) = \Pr(Y = Z_i) \Pr(Z_i = x_{ij}).$$

(Using the incorrect property as stated in the text, we can construct counterexamples when the random variables Y and Z_i are not independent. For example let Y, Z_0, Z_1 be (dependent) binary choices given by a single toss of an unbiased coin. Then $H(X) = 1$, but the right hand side of the entropy formula would evaluate to 2. This cannot be correct, since we cannot get two bits of entropy out of a single toss.)

Page 296, Exercise 5.45(b)

Possibly add: *Warning*: Perfect secrecy depends on what happens for all c .

Page 321, Remark 6.20

It has been suggested that the criterion for the extra bit should be $\frac{1}{2}p \leq y < p$, instead of $\frac{1}{2}p < y < p$. However, since p is presumably a large (odd) prime, the two formulations are identical.

Page 442, Last displayed equation

b_j^* should be b_j^2 .

Page 454, Exercise 7.3

The encrypted message should be $S = 755$, not $S = 4398$.

SUPPLEMENTARY MATERIAL

Page 24, Section 1.3.2

The history of the square and multiply algorithm is interesting. We would like to thank Carlo Beenakker for providing some information in his answer on MathOverflow:

<http://mathoverflow.net/questions/107708>

The following appears in Donald Knuth, *Seminumerical Algorithms*, volume 2 of *The Art of Computer Programming*, page 441:

“The method is quite ancient; it appeared before 200 B.C. in Pingala’s Hindu classic Chandah-sutra [see B. Datta and A.N. Singh, *History of Hindu Mathematics* 1, 1935]; however, there seem to be no other references to this method outside of India during the next 1000 years. A clear discussion of how to compute 2^n efficiently for arbitrary n was given by al-Uqlidisi of Damascus in 952 A.D.; see *The Arithmetic of al-Uqlidisi* by A.S. Saidan (1975), p. 341-342, where the general ideas are illustrated for $n = 51$. See also al-Biruni’s *Chronology of Ancient Nations* (1879), p. 132-136; this eleventh-century Arabic work had great influence.”

A detailed discussion of the earliest history is in A. Kulkarni, *Recursion and Combinatorial Mathematics in Chandashastra*. [Chandashastra = Chandah-sutra]

<http://arxiv.org/abs/math/0703658>

Page 235, Example 5.26

The Prisoner Paradox can be confusing for students, but since it is not a fundamental part of the course, the authors do not feel it deserves more than half a page. The following material is thus made available as a supplement for instructors and students:

Before Alice gets any information from the jailer, there are three outcomes, each of which has equal probability, so

$$\Pr(\text{Alice released}) = 1/3,$$

$$\Pr(\text{Bob released}) = 1/3,$$

$$\Pr(\text{Carl released}) = 1/3.$$

Next suppose that the jailer tells Alice the name of someone who will stay jailed, but when the jailer has a choice, i.e., when both Bob and Carl will stay jailed, he picks one at random. Then

$$\Pr(\text{Alice released and jailer says “Bob”}) = 1/6,$$

$$\Pr(\text{Alice released and jailer says “Carl”}) = 1/6,$$

$$\Pr(\text{Bob released and jailer says “Carl”}) = 1/3,$$

$$\Pr(\text{Carl released and jailer says “Bob”}) = 1/3.$$

So the fact that the jailer told Alice that Bob will stay jailed means that

$$\Pr(\text{Alice released} \mid \text{jailer says “Bob”}) = \frac{1/6}{1/6 + 1/3} = \frac{1}{3},$$

so Alice’s chances of being released are still $\frac{1}{3}$.

Finally, suppose that the jailer tells Alice the name of someone who will stay jailed, but when the jailer has a choice, he always chooses Bob. Then

$$\Pr(\text{Alice released and jailer says "Bob"}) = 1/3,$$

$$\Pr(\text{Alice released and jailer says "Carl"}) = 0,$$

$$\Pr(\text{Bob released and jailer says "Carl"}) = 1/3,$$

$$\Pr(\text{Carl released and jailer says "Bob"}) = 1/3.$$

So the fact that the jailer told Alice that Bob will stay jailed means that

$$\Pr(\text{Alice released} \mid \text{jailer says "Bob"}) = \frac{1/3}{1/3 + 1/3} = \frac{1}{2},$$

so in this scenario Alice's probability of being released has increased to $\frac{1}{2}$.

Page 238, Section 5.3.4

This section can be difficult for students (and instructors) who have not previously studied probability. It has been suggested that in each of the examples, we explicitly describe the sample space Ω , although once one becomes familiar with the language of probability, this is seldom done. In particular, note that the sample space in Example 4.31 is an infinite set.

Here is expanded text for Examples 4.29 and 4.31.

Example 4.29. The sample space Ω consists of all binary strings $\omega = b_1b_2 \dots b_n$ of length n , where $b_i = 0$ if the i 'th experiment is a failure and $b_i = 1$ if the i 'th experiment is a success. The value of the random variable X at ω is simply $X(\omega) = b_1 + b_2 + \dots + b_n$, which is the number of successes. Using the random variable X , we can express the probability of the event ω as

$$\Pr(\{\omega\}) = p^{X(\omega)}(1-p)^{n-X(\omega)}.$$

(Do you see why this is the correct formula?)

Example 4.31. The sample space Ω consists of all binary strings $\omega = b_1b_2b_3 \dots$, where $b_i = 0$ if the i 'th toss is tails and $b_i = 1$ if the i 'th toss is heads. This is an example of an infinite probability space. The way in which we assign probabilities to events is by specifying a certain number of initial tosses. So for any given finite binary string $\beta_1\beta_2 \dots \beta_n$, we assign a probability

$$\Pr(\{\omega \in \Omega : \omega \text{ starts } \beta_1\beta_2 \dots \beta_n\}) = p^{(\# \text{ of } \beta_i \text{ equal to } 1)}(1-p)^{(\# \text{ of } \beta_i \text{ equal to } 0)}.$$

The random variable X is defined by

$$X(\omega) = X(b_1b_2b_3 \dots) = (\text{smallest } i \text{ such that } b_i = 1).$$

Then

$$\{X = n\} = \{\omega \in \Omega : X(\omega) = n\} = \{\underbrace{000 \dots 00}_{n-1 \text{ zeros}} 1b_{n+1}b_{n+2} \dots\}.$$

Hence

$$f_X(n) = \Pr(X = n) = (1-p)^{n-1}p.$$

Page 340, Example 6.37

We provide a further explanation as to why the divisor of $X - \alpha$ is $2[P] - 2[O]$, and not $[P] - [O]$.

There are various ways to see that $\operatorname{div}(X - \alpha)$ is $2[P] - 2[O]$. In particular, the fact that $(\alpha, 0)$ is a point on E means that α is a root of the cubic used to define E , so the equation of E has the form

$$E : Y^2 = (X - \alpha)(X^2 + aX + b).$$

Further, the polynomials of $X - \alpha$ and $X^2 + aX + b$ have no common roots (this is where we use the nonsingularity of E), so they have no common zeros. And of course, their only pole is the point O . It follows that $\operatorname{div}(X - \alpha)$ and $\operatorname{div}(X^2 + aX + b)$ have no points in common except for O . But

$$2 \operatorname{div}(Y) = \operatorname{div}(X - \alpha) + \operatorname{div}(X^2 + aX + b),$$

which shows that the zeros of $X - \alpha$ appear with even multiplicity. Of course, the only zero of $X - \alpha$ is the point $P = (\alpha, 0)$, which shows that

$$\operatorname{div}(X - \alpha) = 2n[P] - 2n[O]$$

for some integer $n \geq 1$. (Note that the total number of zeros and poles sums to zero.)

There are various ways to prove that n equals 1. For example, we can prove that Y is a local uniformizer at P , i.e., it vanishes to order 1 at P . To do that, we consider the local ring at P , which is the ring

$$R = \left\{ \frac{f(X, Y)}{g(X, Y)} : Y^2 = X^3 + AX + B \text{ and } g(\alpha, 0) \neq 0 \right\}.$$

In other words, we take all rational functions whose denominator does not vanish at P . This is a local ring whose maximal ideal is generated by $X - \alpha$ and Y . But since

$$X - \alpha = \frac{Y^2}{X^2 + aX + b} \quad \text{and} \quad X^2 + aX + b \text{ does not vanish at } P,$$

we see that $X - \alpha$ is in the ideal of R generated by Y , so Y is a uniformizer. Indeed, this shows that $X - \alpha$ generates the same ideal as Y^2 , so

$$\operatorname{ord}_P(X - \alpha) = 2 \operatorname{ord}_P(Y) = 2.$$