

STUDENT MATHEMATICAL LIBRARY
Volume 83

Modern Cryptography and Elliptic Curves

A Beginner's Guide

Thomas R. Shemanske



American Mathematical Society
Providence, Rhode Island

For additional information and updates on this book, visit
www.ams.org/bookpages/stml-83

○ 2017 by the American Mathematical Society

Printed in the United States of America.

Library of Congress Cataloging-in-Publication Data

Names: Shemanske, Thomas R., 1952–

Title: Modern cryptography and elliptic curves : a beginner's guide / Thomas R. Shemanske.

Description: Providence, Rhode Island : American Mathematical Society, [2017] | Series: Student mathematical library ; volume 83 | Includes bibliographical references and index.

Identifiers: LCCN 2017001191 | ISBN 9781470435820 (alk. paper)

Subjects: LCSH: Curves, Elliptic–Textbooks. | Geometry, Algebraic–Textbooks. | Cryptography–Textbooks. | AMS: Number theory – Instructional exposition (textbooks, tutorial papers, etc.). msc | Computer science – Instructional exposition (textbooks, tutorial papers, etc.). msc | Number theory – Elementary number theory – Elementary number theory. msc | Algebraic geometry – Arithmetic problems. Diophantine geometry – Applications to coding theory and cryptography. msc | Number theory – Finite fields and commutative rings (number-theoretic aspects) – Algebraic coding theory; cryptography. msc | Computer science – Theory of data – Data encryption. msc | Number theory – Computational number theory – Factorization. msc | Information and communication, circuits – Communication, information – Cryptography. msc | Number theory – Arithmetic algebraic geometry (Diophantine geometry) – Elliptic curves over global fields. msc | Quantum theory – Axiomatics, foundations, philosophy – Quantum computation. msc

Classification: LCC QA567.2.E44 S534 2017 | DDC 516.3/52–dc23 LC record available at <https://lccn.loc.gov/2017001191>

Contents

Preface	vii
Introduction	ix
Chapter 1. Three Motivating Problems	1
§1.1. Fermat's Last Theorem	3
§1.2. The Congruent Number Problem	5
§1.3. Cryptography	6
Chapter 2. Back to the Beginning	9
§2.1. The Unit Circle: Real vs. Rational Points	10
§2.2. Parametrizing the Rational Points on the Unit Circle	12
§2.3. Finding all Pythagorean Triples	16
§2.4. Looking for Underlying Structure: Geometry vs. Algebra	27
§2.5. More about Points on Curves	34
§2.6. Gathering Some Insight about Plane Curves	38
§2.7. Additional Exercises	43
Chapter 3. Some Elementary Number Theory	45
§3.1. The Integers	46
§3.2. Some Basic Properties of the Integers	47

§3.3.	Euclid's Algorithm	52
§3.4.	A First Pass at Modular Arithmetic	56
§3.5.	Elementary Cryptography: Caesar Cipher	63
§3.6.	Affine Ciphers and Linear Congruences	66
§3.7.	Systems of Congruences	70
Chapter 4.	A Second View of Modular Arithmetic: \mathbb{Z}_n and U_n	73
§4.1.	Groups and Rings	73
§4.2.	Fractions and the Notion of an Equivalence Relation	77
§4.3.	Modular Arithmetic	79
§4.4.	A Few More Comments on the Euler Totient Function	93
§4.5.	An Application to Factoring	95
Chapter 5.	Public-Key Cryptography and RSA	101
§5.1.	A Brief Overview of Cryptographic Systems	102
§5.2.	RSA	107
§5.3.	Hash Functions	114
§5.4.	Breaking Cryptosystems and Practical RSA Security Considerations	123
Chapter 6.	A Little More Algebra	127
§6.1.	Towards a Classification of Groups	128
§6.2.	Cayley Tables	128
§6.3.	A Couple of Non-abelian Groups	131
§6.4.	Cyclic Groups and Direct Products	134
§6.5.	Fundamental Theorem of Finite Abelian Groups	138
§6.6.	Primitive Roots	141
§6.7.	Diffie–Hellman Key Exchange	143
§6.8.	ElGamal Encryption	144
Chapter 7.	Curves in Affine and Projective Space	147
§7.1.	Affine and Projective Space	147
§7.2.	Curves in the Affine and Projective Plane	153

§7.3. Rational Points on Curves	156
§7.4. The Group Law for Points on an Elliptic Curve	159
§7.5. A Formula for the Group Law on an Elliptic Curve	179
§7.6. The Number of Points on an Elliptic Curve	185
Chapter 8. Applications of Elliptic Curves	189
§8.1. Elliptic Curves and Factoring	190
§8.2. Elliptic Curves and Cryptography	196
§8.3. Remarks on a Post-Quantum Cryptographic World	198
Appendix A. Deeper Results and Concluding Thoughts	203
§A.1. The Congruent Number Problem and Tunnell's Solution	203
§A.2. A Digression on Functions of a Complex Variable	209
§A.3. Return to the Birch and Swinnerton-Dyer Conjecture	211
§A.4. Elliptic Curves over \mathbb{C}	212
Appendix B. Answers to Selected Exercises	219
§B.1. Chapter 2	219
§B.2. Chapter 3	231
§B.3. Chapter 4	233
§B.4. Chapter 5	236
§B.5. Chapter 6	238
§B.6. Chapter 7	241
Bibliography	245
Index	249

Preface

A question implicit in many conversations with undergraduates at first-year orientation is, “If I were to take only one mathematics class, what should it be?” Of course, there is no one correct answer. A more interesting question (not that it gets asked) might be, “Is there a class which introduces me to the subject of mathematics?” In answer, some might proffer various discrete/finite math/introduction to proof courses, but I think these are not really an answer to that question. A number of years ago, our department was interested in providing a wide variety of courses which attempted to answer that question, and this book details one such offering.

Therefore, one goal of this book is to present some of the vista of modern mathematics to undergraduates who have just started their mathematical careers, with the intent of enticing (and guiding) them into taking mathematics courses beyond the typical calculus regime. As mathematicians, we have acquired the perspective that mathematics represents not only a broad collection of tools which can be brought to bear to solve a myriad of problems, but more interestingly, that the various subdisciplines in mathematics are interconnected often in surprising ways, adding immensely to the richness and allure of the discipline itself. Undergraduates beginning their careers certainly do not have that perspective, and if they are lucky, they begin to acquire it only near the end of their undergraduate majors. The hope

in writing this book is to provide some of that perspective to students at an early stage in their careers so as both to (re)excite them about mathematical exploration, and to help inform their choices as they go forward in their undergraduate experience.

The focal point for this text is to lead students to understand the arithmetic of elliptic curves over a finite field and some applications of elliptic curves to modern cryptography. Assuming only calculus as prerequisite, there is a great deal of ground to cover, but a wonderful opportunity to demonstrate how many areas of mathematics are intertwined.

That said, this book is not (nor is it intended to be) a typical textbook in many respects. While the topics introduced include material on elementary number theory, abstract algebra, cryptography, affine and projective geometry, the intent is not to present a thorough introduction to any of those subjects; it is meant to generate interest in exploring those subjects in more detail. Excellent books devoted individually to those topics are plentiful, but typically they are aimed at a more mathematically sophisticated audience.

This book aims at a mathematically young audience, one that more likely than not has never seen a substantive mathematical proof. Indeed, the only real prerequisite for this book is some one-variable calculus; the rest of the mathematical topics are introduced on-the-fly. This is also not a standard textbook in another important sense. Instead of presenting succinct proofs of results (as is done in a typical textbook), many proofs are presented more as explorations, including (on occasion) some intentional peeks down blind alleys. That is to say, this book makes a significant effort to teach students about how to produce or discover a proof, by presenting mathematics as an exploration. Indeed, while somewhat of a cliché, the book is a great deal more about the journey than the destination, and it is intended to point to the many branches off the main path to be explored in the future.

In the end, the book seems to serve several purposes. It serves, as initially conceived, as a means of introducing many topics in modern mathematics with interconnections among them to motivate students to take more mathematics. It also seems well suited to serve as an

alternate course introducing proofs and abstract mathematics, which occupies a prominent place in many mathematics programs. And finally, given that one cannot really understand modern cryptography without some of its mathematical underpinnings, this book is well suited to computer science programs which desire to offer a course investigating the practical and implementation sides of cryptography, but which need their students to have some semblance of its necessary mathematical background.

Introduction

This book is written to introduce a student with only single-variable calculus as background to enough mathematics to understand the basics of elliptic curves over finite fields and their applications to modern cryptography. Topics include basic notions in elementary number theory and abstract algebra, aspects of affine and projective geometry, as well cryptography and cryptanalysis. The goal of this book is not so much to provide complete answers to the questions we raise as it is to show the many connections between those questions and areas of mathematics whose further study will provide deeper answers.

In Chapter 1 we give a cursory exposition of three problems in number theory which are connected to elliptic curves: Fermat's Last Theorem, the congruent number problem, and applications of number theory to cryptography.

Chapter 2 is quite broad, recasting problems in number theory as problems amenable to geometric or algebraic interpretation. We look at connections of congruent numbers to Pythagorean triples, and at connections between Pythagorean triples and rational points on the unit circle. We explore in detail how to parametrize the rational points on the unit circle and use the parametrization to produce a simple algorithm to enumerate square-free congruent numbers. Then we begin to look for structure inherent in certain sets. For example, we know the set of points in \mathbb{R}^3 that satisfy $x + y + z = 0$ has the geometric structure of a plane through the origin. The set of rational points that satisfy the same equation does not seem to have geometric structure, but it does still have algebraic structure once we define the

notion of vector space. We give a few examples that characterize the notion of dimension of a space along with the notion of basis, which, while clearly important in their own right, also foreshadow the rank of a finitely generated abelian group, the group of rational points on an elliptic curve.

We talk about rational points on more general curves and give Bachet's duplication formula for the elliptic curves $y^2 = x^3 + k$, $k \neq 0$; this is one of the few places we use some calculus. Beyond that, we work to gain insight into Bézout's theorem concerning the number of points of intersection of two plane curves. We see how the issues of the field of definition and multiplicity affect the answer and hint that this is still not enough to give a complete answer, suggesting a future need to expand our view from affine to projective space.

Chapter 3 is rather traditional, introducing basic concepts in elementary number theory including divisibility, gcd, and division and Euclidean algorithms. We take a first pass at modular arithmetic, noting that congruence is an equivalence relation. We give some simple applications of modular arithmetic, and we use the Caesar cipher both as another application and as a vehicle to introduce some standard terminology in cryptography. We extend Caesar ciphers to affine ones and explore conditions under which affine transformations can function as encryption algorithms. This leads to determining the conditions under which linear congruences can be solved and to determining the number of incongruent solutions. All this is preparatory to the next chapter where we talk about the set of residues modulo n having an algebraic structure.

Chapter 4 is another in which we slowly unravel many important ideas that lead to the characterization of \mathbb{Z}_n (the set of residues modulo n) as a ring and U_n (the set of reduced residues modulo n) as its unit group. We begin by understanding the standard arithmetic operations on the integers as binary operations on the set \mathbb{Z} and how their properties endow \mathbb{Z} with the structure of a commutative ring, passing through the notion of a group on the way. Then we use arithmetic with fractions (\mathbb{Q}) to motivate binary operations on a set of equivalence classes. Armed with that intuition, we define congruence

classes modulo n , and show that there are well-defined binary operations which can be defined on them which make the set of residues \mathbb{Z}_n into a commutative ring with identity. We then show that we can make the set of reduced residues (the units of \mathbb{Z}_n) into an abelian group. We define the Euler totient function ϕ and prove Euler's theorem and Fermat's little theorem, which we will need to justify that RSA (the Rivest–Shamir–Adelman algorithm) functions as intended. We discuss modular exponentiation, and end with an application to factoring, Pollard's $p - 1$ method which serves as the model against which we compare Lenstra's elliptic curve method of factorization.

Chapter 5 begins with a simple description of how public-key cryptography facilitates the creation of a secure connection when making an online purchase. Then we discuss more of the fundamentals of a public-key cryptosystem, followed by a discussion of signatures and authentication. We then begin to make things somewhat more realistic by talking about hash functions and signatures applied to a hash. We discuss the use of hash functions in daily use and specific requirements for hash functions in current use. We discuss preimage resistance problems in relation to the Birthday paradox in probability. The chapter ends with some security considerations for RSA.

Chapter 6 introduces a bit more algebra, including the notion of a cyclic group and the fundamental theorem of finite abelian groups. We use the fundamental theorem to give a proof that for p a prime, the set of reduced residues U_p is cyclic, leading to a discussion of primitive roots. This in turn leads to the notion of discrete logarithms, the Diffie–Hellman key exchange, and ElGamal encryption.

Chapter 7 covers a great deal of ground beginning with a gradual introduction to projective space. We discuss how and why to homogenize a polynomial defining an affine plane curve so as to reveal extra points on the corresponding projective curve. Then we take a significant amount of time to define the group law for the set of points on an elliptic curve, and we abstract from it the algebraic formulas that define the addition law in projective space. We give several examples where we determine the isomorphism class of the abelian group of points of an elliptic curve over a finite field, and we end with Hasse's theorem bounding the number of points on an elliptic curve over the

finite field \mathbb{F}_p . As an application, we show that the probability that a randomly chosen $x \in \mathbb{F}_p$ is the x -coordinate of a point on the elliptic curve is approximately one-half, which we use in the last chapter in discussing an elliptic curve version of Diffie–Hellman and the ElGamal cryptosystem.

In the final chapter we look at applications of all our work thus far. We introduce Lenstra’s elliptic curve method (ECM) of factorization and discuss its analogy with Pollard’s $p - 1$ method. We then talk about how to embed a plaintext message as a point on an elliptic curve and, given that embedding, what would be the appropriate analogs of Diffie–Hellman and ElGamal. We end with some interesting remarks about the NSA’s vision and recommendations regarding cryptography in a post-quantum computer world.

Appendix A completes the discussion of some themes that motivated much of the exposition, but the level of exposition is now far above where it has been in the body of the text. Giving closure to the topic of congruent numbers is Tunnell’s theorem whose solution involves a discussion of Mordell’s theorem on the structure of $E(\mathbb{Q})$, the set of rational points on an elliptic curve, as well as the Birch and Swinnerton-Dyer conjecture. The appendix ends with a brief discussion of elliptic curves over \mathbb{C} , elliptic functions, and the characterization of $E(\mathbb{C})$ as a complex torus.

Appendix B has solutions to the majority of exercises posed in the text.

All code and figures in the text were produced with Sage [S⁺15].

Chapter 1

Three Motivating Problems

The goal of this book is to explain how the set of points on an elliptic curve can be given the structure of an abelian group, and how the arithmetic of elliptic curves over finite fields can be used as a powerful tool in cryptography and cryptanalysis. Perhaps to put it another way, the goal of this book is to help the reader understand the first sentence.

To motivate our study of elliptic curves, we consider three problems in number theory and geometry whose solutions use elliptic curves in an essential, if sometimes subtle, manner. Two of these problems, Fermat's Last Theorem and the congruent number problem, are problems whose statements are completely elementary. They are classical in feel, perhaps almost playful in tone. In contrast the third problem, applications of the theory of elliptic curves to cryptography, is a quite modern subject whose practical importance has grown enormously of late. In 2011 Koblitz et al. [KKM11] wrote that "over a period of sixteen years, elliptic curve cryptography went from being an approach that many people mistrusted or misunderstood to being a public key technology that enjoys almost unquestioned acceptance." A key question for us to address is, What are these elliptic

curves which have had such an impact, and how have they proven to be so important?

At first blush, elliptic curves appear to be nothing special. For us they will just be certain curves given by polynomials of degree 3, but before we even talk about cubics, let's back up a bit. In secondary school it is common to study the properties of lines and conics in the plane as well as the principles of Euclidean geometry. For example, if you were to consider the quadratic equation $y = x^2 + 2$, you might think of the set of all the points (x, y) in the plane which satisfy that equation as a geometric object, the parabola with line of symmetry the y -axis, vertex $(0, 2)$, and opening upward. It might also be interesting to consider what happens when we consider the points (x, y) which satisfy that equation when we restrict or allow the coordinates to come from different domains. For example, we could ask for the *rational points*, that is those points (x, y) in which both $x, y \in \mathbb{Q}$ (are rational numbers) such as $(0, 2)$ or $(1/3, 19/9)$, though you might question why that would be a useful thing to do. We shall answer that question later in the book, but for now let's simply observe that some but not all points on the curve have rational coordinates, e.g., $(\pm\sqrt{2}, 4)$, or $(\sqrt[3]{2}, 2 + \sqrt[3]{4})$ are real, but not rational points. Analogously, we might ask about complex points (points with complex coordinates). Again, why would we do that? Well, at least here, your previous experience in mathematics affords you some insight. If we asked what are the roots of $x^2 + 2$, you would either have said there are none or they are complex (imaginary). Said another way, $(\pm i\sqrt{2}, 0)$ are two complex points on the curve $y = x^2 + 2$ which are not on the real locus, the curve we draw. While we cannot fully appreciate this comment so early in the book, it is perhaps not too surprising that an interest in rational points should somehow be connected to number theory, since rational numbers are the quotients of integers, the domain of number theory. It is also the case that the complex points are often more interesting than the real points since there are more of them and since the set of complex solutions may have an even more interesting structure than we might first think. Indeed, to use the word "structure" in the context of the set of points on a curve is quite intentional, and in some sense it represents the origins of most of the applications we shall discuss.

For us an elliptic curve will be a the set of points (x, y) which satisfy an equation of the form

$$y^2 = x^3 + ax^2 + bx + c,$$

where the cubic, $x^3 + ax^2 + bx + c$, is nonsingular (has distinct roots). Our interests will include a consideration of the set of solutions (x, y) where x, y are restricted to different domains, and indeed we will not only be interested in points whose coordinates are rational, real, or complex numbers, but also those whose coordinates lie in so-called *finite fields*, which we have yet to define. But even we have to admit that despite the build up we have given to elliptic curves so far, the definition seems quite lackluster; in particular, the definition seems to shed no light at all on why elliptic curves should play such a pivotal role in number theory. Yet be assured that they do, and as well provide some of the best schemes for public-key cryptography. Certainly we shall take a closer look at all these things.

Among the three motivating problems, we shall look only briefly at the Fermat problem, slightly more deeply at the congruent number problem, and most deeply at developing an understanding of how elliptic curves are of critical use in cryptography. This emphasis is deliberate, in part because the role elliptic curves play in the solutions of the Fermat and the congruent number problems is more subtle and much more sophisticated, and in part because such an omission provides the opportunity for the reader to do some investigation on her own.

1.1. Fermat's Last Theorem

This theorem, conjectured by Fermat in 1635, states simply that for $n > 2$ the equation $x^n + y^n = z^n$ has no solutions in the integers except when one of the variables is zero. We note that this contrasts sharply with the case of $n = 2$ for which solutions (Pythagorean triples) abound. In fact, Pythagorean triples will play an integral role in the congruent number problem, but before leaping to make that connection, we need to say a few more words about the Fermat theorem.

In 1640, Fermat himself proved that the conjecture was true for $n = 4$ and noted that if $n = km$, the existence of a nontrivial solution to $x^n + y^n = z^n$ implied the existence of nontrivial solutions to $(x^k)^m + (y^k)^m = (z^k)^m$, that is to a Fermat problem whose exponent is a divisor of the original. Fermat's $n = 4$ result and this observation reduces the proof of the Fermat conjecture to showing there are no nontrivial solutions to $x^p + y^p = z^p$, where p is an odd prime.

Until 1839, only the cases $p = 3, 5, 7$ had been resolved; this included Sophie Germain's important work which eventually allowed the conjecture to be proved for all odd primes less than 100. In the 1850s, Ernst Kummer developed techniques to prove the Fermat conjecture for all "regular" primes, which is believed to be an infinite family. Modern computing methods verified the conjecture for primes less than four million.

The first real breakthrough came in 1985 when Gerhard Frey suggested that if there were a counterexample to Fermat's conjecture, it could be used to create an elliptic curve having properties which would provide a counterexample to yet another unproved conjecture due to Taniyama and Shimura. While the Fermat conjecture enjoyed a reputation as a long-standing open problem in mathematics, the Taniyama–Shimura conjecture had deep implications for how the theory of modular forms and elliptic curves fit together. If this later conjecture had been false, it would have been quite disappointing.

In the period 1985–1986, Jean-Pierre Serre showed how the Taniyama–Shimura conjecture together with another smaller conjecture — termed the "epsilon conjecture" — would imply Fermat's theorem. Ken Ribet proved the epsilon conjecture in 1986 reducing the Fermat theorem to a proof of the Taniyama–Shimura conjecture for a special class of elliptic curves. In 1994, Andrew Wiles (after seven or more years of intense work, together with a last minute assist by Richard Taylor) succeeded in proving the required case of the Taniyama–Shimura conjecture, and hence proving Fermat's Last Theorem. And in case you were wondering, the full Taniyama–Shimura conjecture has now been proven as well.

1.2. The Congruent Number Problem

A positive integer is called a *congruent number* if it is the area of a right triangle whose sides all have rational length. For example, 6 is a congruent number since 6 is the area of a 3-4-5 right triangle.

It is also true that 5 is a congruent number, though this is something you might not guess right off. But indeed, 5 is the area of the right triangle with sides: $3/2$, $20/3$, $41/6$. Any reasonable person would agree that one can check the result, but it is certainly mysterious how one would come up with a triangle having those sides. However, with an unexpected solution in hand, it now becomes a much more interesting question to ask which integers are congruent numbers. For example, later in the book we shall see that 157 is a congruent number. Surely it can't be that hard to check that such a small number is or is not a congruent number. But we shall see that the answer to this question is more elusive than it may first appear.

A key observation in characterizing congruent numbers is that if N is a congruent number, then Nt^2 is also for any positive integer t ; indeed if N is the area of a triangle having rational sides a, b, c , then Nt^2 is the area of a triangle with sides at, bt, ct . Let's consider the triangle showing that 5 is a congruent number. It is easily seen that 6 is the common denominator of the rational numbers $3/2$, $20/3$, $41/6$, and from our observation above, since 5 is a congruent number, $5 \cdot 6^2$ is also, being the area of a right triangle with sides $9 = 6 \cdot \frac{3}{2}$, $40 = 6 \cdot \frac{20}{3}$, and $41 = 6 \cdot \frac{41}{6}$. But of course this means that 9, 40, 41 is a Pythagorean triple! Conversely, suppose that A, B, C are a Pythagorean triple, and N is the area of the corresponding right triangle. Write $N = N_0t^2$ where N_0 is square free (1 or the product of distinct primes), $t > 0$. Then N_0 is a congruent number, being the area of a right triangle with rational sides $A/t, B/t, C/t$.

So there is a clear relationship between congruent numbers and Pythagorean triples, which means if we had a way to list all Pythagorean triples, we would know which numbers were congruent numbers. In fact, we will show how to list all the Pythagorean triples! Unfortunately, the congruent numbers that come out of the list do not appear

in any particular order and are often repeated, so this procedure cannot definitely answer whether a given integer is a congruent number. Still it will provide a good deal of insight into the connections between algebra and geometry, so we will spend significant time with it.

As of this writing, the congruent number problem remains open, though many partial results are known. Jerrold Tunnell [**Tun83**] developed a condition based on the arithmetic of elliptic curves (and yet another open conjecture—the Birch and Swinnerton-Dyer conjecture) which provides a beautiful answer to this question. We discuss Tunnell’s approach at the end of the book.

1.3. Cryptography

Cryptography is a subject that has a long and fascinating history and is a matter of critical importance to all of us in an age when so many transactions happen electronically. It is the subject around which essentially all the background material on number theory and algebra that we develop in this book will be focused.

There are many interesting questions of a practical nature which cryptography solves and which we shall examine, but as a teaser, we mention only a few in this introductory chapter. It is not terribly difficult to send private messages to a friend even over an insecure channel. What becomes trickier is when you want to do the same with someone you don’t know. Why would you want to do that? Well, every time you order something online, you want to communicate securely with the vendor so that confidential information (e.g., a credit card number) is not revealed over the insecure web. But how can you (that is, your computer) and your vendor do this? On a different note, how can someone who has received an email from you prove to a third party that the message is indeed from you and not someone forging your address? Or, how can someone be sure a message has not been tampered with (e.g., when a bank receives a message to transfer funds from one account to another)?

All of these are vital questions that modern cryptography answers effectively, and elliptic curves figure prominently in the mix. Of course given any cryptographic system, there are many individuals who do their best to break it, so we will look at standard kinds of

cryptographic attacks on various systems, and how vulnerable each system is to different types of attacks. In the end, elliptic curve cryptography turns out to be among the best public-key cryptosystems currently in use.

Chapter 2

Back to the Beginning

In the first chapter we briefly introduced what will be our main object of study in this book, an elliptic curve. We suggested that in general, curves could be interesting geometrically but they sometimes also contain arithmetic or algebraic information—presuming we could discover how to tease it out of them. One point we made is that it is often revealing to think about those points on a curve each of whose coordinates are rational, or real, or complex. We refer to such points simply as the rational, real, or complex points on a curve. In a rather different vein, we also discussed the connection between congruent numbers and Pythagorean triples. We even suggested that we could list all the Pythagorean triples.

In this chapter, we set as one of our goals to do just that: to determine a way in which to list all the Pythagorean triples—well, at least all the so-called *primitive* Pythagorean triples, but we will talk about that subtlety in a bit. It turns out that the key to listing all the triples is to characterize all the rational points on the unit circle, $x^2 + y^2 = 1$. This represents an important first example of how the algebra and geometry of a curve can combine to answer questions that are seemingly quite unrelated. We shall of course see more examples as the book unfolds.

2.1. The Unit Circle: Real vs. Rational Points

When we look at the equation $x^2 + y^2 = 1$, our minds can be pulled in two different directions. In one direction, we see the geometric object, the unit circle, which displays simultaneously all the (real) solutions (x, y) to the equation. Said another way, the circle is a graphical representation of the real points on the curve $x^2 + y^2 = 1$. In the other direction, we see the underlying algebraic operations: for example, if $x = 1/2$, then $y^2 = 1 - (1/2)^2 = 3/4$, so that $y = \pm\sqrt{3}/2$ which gives us the two real points $(1/2, \pm\sqrt{3}/2)$ on the unit circle.

While the graphical representation of the unit circle allows us to “see” all the points on the unit circle, we need various means of characterizing their coordinates. As a first approach, our experience with calculus provides an outstanding answer to the question of characterizing the real points on the circle: they are parametrized by cosine and sine. That is, each point (x, y) on the unit circle has the form $(\cos \theta, \sin \theta)$ for a uniquely determined $\theta \in [0, 2\pi)$.

This is certainly one reasonable answer if we are interested in the real points. But what about the rational points? Except for the four obvious points, $(0, \pm 1)$ and $(\pm 1, 0)$, we seem to know very little about the values of θ which give rise to rational values of sine and cosine.

So to characterize the rational points, we shall take a different approach. Suppose that $(\frac{a}{b}, \frac{c}{d})$ is a rational point on the unit circle. Then the coordinates (x, y) satisfy the equation $x^2 + y^2 = 1$, so

$$\left(\frac{a}{b}\right)^2 + \left(\frac{c}{d}\right)^2 = 1.$$

Clearing denominators, we obtain

$$(ad)^2 + (bc)^2 = (bd)^2,$$

that is, the rational point $(a/b, c/d)$ on the unit circle corresponds to an integer solution of $x^2 + y^2 = z^2$ with $z \neq 0$, and in fact to a Pythagorean triple (taking absolute values) $|ad|, |bc|, |bd|$, if the rational point $(\frac{a}{b}, \frac{c}{d}) \neq (0, \pm 1), (\pm 1, 0)$. This is certainly an interesting observation. Conversely, any triple, $A, B, C \in \mathbb{Z}$, $C \neq 0$ with $A^2 + B^2 = C^2$ corresponds to the rational point $(A/C, B/C)$ on the unit circle. And note that the constraint $C \neq 0$ is not much of a

constraint, since if $C = 0$, the only possible triple of (real) numbers is $(0, 0, 0)$, which isn't all that arithmetically interesting.

Summary so far.

- Congruent numbers are naturally connected to Pythagorean triples. A Pythagorean triple determines a right triangle whose area is by definition a congruent number, and conversely given a congruent number, we have by definition a right triangle whose sides have rational length. Clearing the denominators for the side lengths produces a Pythagorean triple.
- Pythagorean triples are connected to rational points on the unit circle $x^2 + y^2 = 1$, as we have just seen.
- In the first chapter we said that if we could list all the Pythagorean triples, then all the congruent numbers would eventually be revealed.
- So if we can figure out how to characterize all the rational points on the unit circle, we can get the Pythagorean triples, which will give us the congruent numbers. At least it's a plan. We'll give an algorithm for this later in this chapter.

But before addressing that question, let's digress a bit more. First a general ramble. It may seem a bit circuitous that to list congruent numbers, we have reduced the problem to listing all the Pythagorean triples, and then in turn we have reduced the problem of generating all the Pythagorean triples to characterizing the rational points on the unit circle. This process of reducing the solution of one problem to the solution of another is actually quite common in mathematics. If you want a more dynamic example, consider the progress on the Fermat problem, outlined in the first chapter. But more to the point, it is often this process of transforming one problem into a seemingly quite distinct one that is among the reasons mathematics is perceived as so rich and deep a subject. Occasionally, one might lament the lack of a direct proof as in the Fermat problem, but that is perhaps better left as a question of aesthetics.

A more constrained digression from the task of characterizing the rational points on the unit circle is to consider what might happen if we generalized the problem by asking not simply for the solutions to $x^2 + y^2 = 1$ but for solutions to $x^2 + y^2 = r$, where r is any real number. We might be led to consider solutions to $x^2 + y^2 = 0$ or $x^2 + y^2 = -1$. If we were biased by our perspective above where we sought only solutions in the rational or real numbers, we might be inclined to say the only solution to the first equation is the origin, and that there is no solution for the second. But that is only because our perspective may be a bit narrow. For example if we were looking for solutions in the complex numbers or over a finite field (whatever that is), there would be lots of solutions. Having a flexible perspective will be very useful to us.

2.2. Parametrizing the Rational Points on the Unit Circle

While we have certainly made a good case for why we should be interested in the rational points on the unit circle, we haven't yet come up with a good scheme for enumerating them. We considered using sine and cosine, but not only does characterizing the values of θ for which $\cos \theta$ and $\sin \theta$ are both rational seem difficult, it is not really what we want. We simply want a way to list all the points on the unit circle with rational coordinates.

Perhaps that suggests we go back to the beginning and find a different way to list all the real points on the unit circle. One way that comes to mind is historically quite old and has been used in many contexts—it is the notion of a stereographic projection. Imagine wanting to create a map of the world, and for now let's consider what was classically considered a map, that is, a rendering on a flat surface. The idea was quite simple. Take the globe, and pass a plane through the equator. Draw a line from the north pole to any other point on the globe. That line will pierce the plane in exactly one point. Note that the equator is a circle in our plane, and the northern hemisphere is mapped outside the circle; the southern hemisphere mapped inside the circle.

Today, you might argue that such a representation is rather old fashioned, and indeed, while a stereographic projection has certain interesting properties (it preserves angles), it certainly distorts areas and distances, but that is because we are trying to do too much with one map. Somehow localized projections would suffer less distortion. After all, the Earth is basically a sphere, and every map we look at is a planar rendering of some portion of it. In topology, you make rigorous the notion that the sphere (the Earth) is locally flat, that is if you look at a small enough patch it is essentially the same as a piece of the Euclidean plane. This concept is the germ of the idea behind the notion of a manifold, which begins a foray into differential topology.

But let's return to the issue of characterizing the points on the unit circle in such a way that the rational points are easily recognized. To do so, we use the idea of stereographic projection. Consider Figure 2.1.




Figure 2.1

In analogy with the stereographic projection of the sphere onto a plane through its equator, we consider the projection of the unit circle from the point $(-1, 0)$ onto the line passing through the north and south poles $((0, 1)$ and $(0, -1)$). More specifically, for each point $(a, b) \neq (-1, 0)$ on the unit circle, there is a unique line passing through both points, and that line intersects the y -axis in a unique point we have labeled $(0, t)$. Note that the points of the unit circle in the first and fourth quadrants are mapped to points $(0, t)$ with $t \in [-1, 1]$, while points in the second and third quadrants are mapped to points $(0, t)$ with $|t| \geq 1$.

The analysis we give now is quite standard; we follow the progression of observations in [ST92], though we fill in a bit of background material to make the arguments accessible to the intended audience of this text. We have seen that for each point $(a, b) \neq (-1, 0)$, the line between $(-1, 0)$ and (a, b) crosses the y -axis at a point $(0, t)$. Conversely, by this correspondence, every point $(0, t)$ determines a unique point (a, b) on the unit circle (except for $(-1, 0)$).

So there is a one-to-one correspondence between points $(0, t)$ on the y -axis and points on the unit circle (except for $(-1, 0)$). We will show that the point $(-1, 0)$ corresponds to $t = \pm\infty$.

Let L_t denote the line through $(-1, 0)$ and $(0, t)$; its equation is $y = t(x + 1)$. If (a, b) is the (other) point of intersection of the line and the unit circle, then the coordinates satisfy

$$b = t(a + 1), \text{ and } a^2 + b^2 = 1.$$

Solving simultaneously yields

$$1 - a^2 = b^2 = t^2(a + 1)^2.$$

Now, for a fixed value of t , the equation $1 - a^2 = t^2(a + 1)^2$ can be viewed as a quadratic equation in the variable a whose roots are the x -coordinates of the points of intersection of the line L_t with the circle. Clearly one of them is $a = -1$, so we assume $a \neq -1$; that is, we assume that $1 + a \neq 0$. Now consider the equation

$$1 - a^2 = (1 - a)(1 + a) = t^2(1 + a)^2.$$

Since $1 + a \neq 0$, this implies that

$$1 - a = t^2(1 + a), \text{ or (expanding and regrouping)}$$

$$a(t^2 + 1) = 1 - t^2, \text{ which yields}$$

$$a = \frac{1 - t^2}{1 + t^2}.$$

This is an expression for the x -coordinate of the point of intersection ($x \neq -1$) of the line L_t with the unit circle. Substituting into the equation of the line yields the y coordinate $b = t(1 + a) = \frac{2t}{1 + t^2}$.

So every point on the unit circle (except $(-1, 0)$) has the form

$$(2.1) \quad (a, b) = \left(\frac{1 - t^2}{1 + t^2}, \frac{2t}{1 + t^2} \right)$$

for some value of $t \in \mathbb{R}$. We note that

$$\lim_{t \rightarrow \pm\infty} \left(\frac{1 - t^2}{1 + t^2}, \frac{2t}{1 + t^2} \right) = \lim_{t \rightarrow \pm\infty} \left(\frac{\frac{1}{t^2} - 1}{\frac{1}{t^2} + 1}, \frac{2}{\frac{1}{t} + t} \right) = (-1, 0),$$

suggesting that to characterize all the points on the unit circle, we need all the real numbers ($t \in \mathbb{R}$), plus a little something involving infinity.

Now that we have successfully given a characterization of all the real points on the unit circle, we ask if it is easy to detect the rational points? Let's examine the line L_t a bit more carefully. Now L_t , the line between $(-1, 0)$ and (a, b) , has slope

$$t = \frac{b - 0}{a + 1} = \frac{b}{a + 1},$$

so it is immediate that if (a, b) is a rational point on the unit circle, the slope t is a rational number. What about the converse?

Consider the line L_t given by the equation $y = t(x + 1)$, and assume that t is a rational number. What can be said about the coordinates of the point of intersection (a, b) ? In the work above, we derived an explicit characterization of (a, b) in terms of t , given in equation (2.1), which clearly demonstrates that if t is a rational number, then so are the coordinates a, b . So there is a one-to-one correspondence between the rational points on the circle (except for $(-1, 0)$) and rational values of t .

Exercise. Analogously to what we did above, find a parametrization for the points on the circle $x^2 + y^2 = 2$ and extract a characterization of the rational points.

To start, project from the rational point $(1, 1)$. Note: projecting onto the x or y -axis does not work as expected, as not all lines from $(1, 1)$ to points on the circle intersect those axes. Instead, try to project onto the line $y = -x$.

Exercise. Now consider the issue of rational points on $x^2 + y^2 = 3$. In contrast to the examples above, prove that there are no rational points on this curve, and describe the crucial difference between this example and the one before.

2.3. Finding all Pythagorean Triples

We now use the above formulas to find all Pythagorean triples: positive integers A, B, C with $A^2 + B^2 = C^2$. This is a task that requires a good deal of effort. None of it is terribly deep, but there are many steps, and we work steadily to reduce the problem (that technique again) to a manageable state. We make our first simplification. In enumerating the triples, there is no reason to consider triples with a common divisor. For example, if A, B, C have a common divisor t , then $A = A_0t$, $B = B_0t$, and $C = C_0t$ with $A_0, B_0, C_0 \in \mathbb{Z}$. Moreover, A_0, B_0, C_0 is also a Pythagorean triple since

$$A^2 + B^2 = C^2 \implies A_0^2t^2 + B_0^2t^2 = C_0^2t^2 \implies A_0^2 + B_0^2 = C_0^2,$$

and given the triple A_0, B_0, C_0 , we could recover the original by multiplying all the numbers by t . We call a Pythagorean triple A, B, C *primitive* if there is no common integer divisor of A, B, C other than ± 1 ; we refer to this condition as saying that the integers A, B, C are *relatively prime*. In the next chapter, we shall introduce the notion of a greatest common divisor, so that the triple A, B, C being primitive is equivalent to saying the greatest common divisor of A, B, C is 1, which will be denoted $\gcd(A, B, C) = 1$.

2.3.1. Developing an Algorithm.

Remark 2.1. It is worth noting that if we are given a primitive Pythagorean triple A, B, C , then all pairs of integers $\{A, B\}$, $\{A, C\}$, and $\{B, C\}$ are also relatively prime. Indeed, we show a bit more. We claim that if t is an integer which divides two of the three integers A, B, C forming any Pythagorean triple, then it must divide the third. For example, if $t \mid B$ and $t \mid C$, then $t^2 \mid (C^2 - B^2) = A^2$ which means $t \mid A$. The other two cases are analogous. So if A, B, C is a primitive Pythagorean triple, then the pairs $\{A, B\}$, $\{A, C\}$, and $\{B, C\}$ are all relatively prime. Conversely, if any of the pairs $\{A, B\}$, $\{A, C\}$, or $\{B, C\}$ are relatively prime, then A, B, C is a primitive triple, since a common divisor of all three of A, B, C must obviously divide any two. \square

Thus given a primitive Pythagorean triple A, B, C , we produce the rational point $(A/C, B/C)$ on the unit circle in which the rational numbers $A/C, B/C$ are already in lowest terms. Moreover, the rational point lies in the first quadrant and is not equal to $(0, 1)$ or $(1, 0)$.

In the other direction, using equation (2.1), we can produce the coordinates of a rational point (a, b) on the unit circle (the second point of intersection of the line L_t through $(-1, 0)$ and $(0, t)$) whenever t is a rational number. Let us write $t = m/n$ with m, n relatively prime integers. Since we are interested in only those rational points that will correspond to Pythagorean triples, we can restrict ourselves to rational points in the first quadrant excluding $(0, 1)$ and $(1, 0)$. This means that the corresponding parameter t (slope of the line L_t) satisfies $0 < t < 1$ and, since $t = m/n$, we may also assume that $n > m > 0$; recall we were already assuming that m, n are relatively prime.

Substituting $t = m/n$ into equation (2.1) yields that

$$a = \frac{n^2 - m^2}{n^2 + m^2} \quad \text{and} \quad b = \frac{2mn}{n^2 + m^2}.$$

We claim that there are integers A, B, C so that

$$(2.2) \quad a = \frac{A}{C} = \frac{n^2 - m^2}{n^2 + m^2} \quad \text{and} \quad b = \frac{B}{C} = \frac{2mn}{n^2 + m^2},$$

and that the fractions A/C and B/C are in lowest terms. Well, it is clear that if we put $A_0 = n^2 - m^2$, $B_0 = 2mn$, and $C_0 = n^2 + m^2$, then we can write

$$a = \frac{A_0}{C_0} = \frac{n^2 - m^2}{n^2 + m^2} \quad \text{and} \quad b = \frac{B_0}{C_0} = \frac{2mn}{n^2 + m^2},$$

but generally these fractions are not in lowest terms. For example, if m, n have the same parity (i.e., both are even or odd), then A_0, B_0, C_0 are all even (i.e., have 2 for a common divisor), so the potential snag comes in when we independently reduce the fractions—we need to know that the new denominators are the same.

This is why Remark 2.1 is really important. Since we are assuming that $n > m > 0$, we know that A_0, B_0, C_0 is a Pythagorean triple. Now let t be the greatest common divisor of A_0, C_0 . This means two things: first, if we define $A = A_0/t$ and $C = C_0/t$, then A and C are integers which are relatively prime, so that $A/C = A_0/C_0$ is in lowest terms. Second, by Remark 2.1, since t is a common divisor of A_0 and C_0 , then it is also a divisor of B_0 , so if we put $B = B_0/t$, then we know that B is an integer and $B/C = B_0/C_0$. What's left? We need to know that B/C is in lowest terms. Well, if it were not, there would be a divisor $s > 1$ of both B and C . Recall that since $A_0^2 + B_0^2 = C_0^2$ we have $A^2 + B^2 = C^2$, so the remark would imply that s also was a divisor of A , meaning A/C was not in lowest terms as claimed, which would be a contradiction. So both fractions are reduced and have the same denominator.

Summarizing, we have now written our rational point (a, b) on the unit circle as in equation (2.2):

$$a = \frac{A}{C} = \frac{n^2 - m^2}{n^2 + m^2} \quad \text{and} \quad b = \frac{B}{C} = \frac{2mn}{n^2 + m^2},$$

where the fractions A/C and B/C are in lowest terms, and this in turn means that A, B, C is a primitive Pythagorean triple.

Now we claim a bit more: given a primitive Pythagorean triple A, B, C , one of A or B must be even, and the other odd. Certainly, they cannot both be even, since if they are, then by Remark 2.1, C is also, violating that the triple is primitive. If A and B are both odd, we have to work a bit harder since we have not yet introduced the

notion of modular arithmetic (next chapter), but still the argument is quite elementary.

Integers are even or odd. Those that are even are divisible by 2, but what else is implied? From grade school, we know that if we divide any integer by 2, the only nonnegative remainders are 0 and 1. So the even numbers have remainder 0, and the odd numbers have remainder 1. We prove the desired result that one of A or B must be even, by contradiction; that is we assume that A and B are both odd, and see that we arrive at a conclusion which is impossible. Consequently, our assumption leading to this contradiction is false, and that in turn establishes our original claim.

So we assume that both A and B are odd integers; thus we can write $A = 2r + 1$ and $B = 2s + 1$ for integers r, s . But then

$$\begin{aligned} C^2 &= A^2 + B^2 = (2r + 1)^2 + (2s + 1)^2 = 4r^2 + 4r + 1 + 4s^2 + 4s + 1 \\ &= 2[2r^2 + 2r + 2s^2 + 2s + 1] = 2[2(r^2 + r + s^2 + s) + 1] = 2L, \end{aligned}$$

where L is an odd integer. This says C^2 is an even integer, but the only way for that to happen is for C to be even. But if C is even, then C^2 is a multiple of 4. But our expression above says that $C^2 = 2L$ where L is an odd integer, so this provides a contradiction, and thus we have established that one of A or B is even and one is odd.

Now our goal is to enumerate Pythagorean triples. One obvious duplication which would occur in a list is that if we look at the triple (A, B, C) , we know that (B, A, C) will also be a Pythagorean triple, so again to simplify our efforts, we shall assume that A is odd and B is even.

We are nearing the end of our argument and need to make only a few more observations. Recall that in equation (2.2), we have assumed that m and n are relatively prime integers with $n > m > 0$.

Now since A/C and B/C are already in lowest terms, we must have positive integers λ, μ (here we use that $n > m > 0$) so that

$$\begin{aligned} n^2 - m^2 &= \lambda A, & 2mn &= \mu B, \\ n^2 + m^2 &= \lambda C, & n^2 + m^2 &= \mu C. \end{aligned}$$

Since $C \neq 0$, $n^2 + m^2 = \mu C = \lambda C$ implies $\mu = \lambda$, thus there is a positive integer λ with

$$\lambda A = n^2 - m^2, \quad \lambda B = 2mn, \quad \lambda C = n^2 + m^2.$$

Our final claim is that $\lambda = 1$, which will provide nice formulas for our primitive Pythagorean triples.

For this last claim, we note that $\lambda(C+A) = 2n^2$ while $\lambda(C-A) = 2m^2$. This means that λ divides both $2m^2$ and $2n^2$, but m and n are relatively prime which means λ divides 2, so either $\lambda = 1$ or 2. We need only exclude $\lambda = 2$ as a possibility.

Recall that we are assuming the following: A is odd and B is even, and m and n are relatively prime, so they are not both even. When we established that A and B had different parities, we noted the following about squares of integers. The square of an even integer has the form $(2r)^2 = 4r^2$, a multiple of 4, while the square of an odd integer has the form $(2r+1)^2 = 4(r^2+r)+1$, an integer leaving a remainder of 1 when divided by 4.

So let us assume that $\lambda = 2$ and look for a contradiction which will finish our argument. Since A is an odd integer, we can write it as $A = 2r + 1$, so if $\lambda = 2$, we see that $\lambda A = 4r + 2$.

Now we must have that $\lambda A = n^2 - m^2$ and we know that m and n cannot both be even since they are relatively prime. So $n^2 - m^2$ has one of the following three forms:

$$n^2 - m^2 = \begin{cases} (2k+1)^2 - (2\ell+1)^2 = 4(k^2 + k - \ell^2 - \ell) & \text{if } m, n \text{ are both odd,} \\ (2k+1)^2 - (2\ell)^2 = 4(k^2 + k - \ell^2) + 1 & \text{if } m \text{ is even and } n \text{ is odd,} \\ (2k)^2 - (2\ell+1)^2 = 4(k^2 - \ell^2 - 1) + 3 & \text{if } n \text{ is even and } m \text{ is odd.} \end{cases}$$

But this says that if $\lambda = 2$, λA leaves a remainder of 2 when divided by 4, while $n^2 - m^2 = \lambda A$ can only leave remainders 0, 1, 3, so the case $\lambda = 2$ is precluded, and we have finally that $\lambda = 1$.

Summarizing our persistent efforts, we have the following theorem:

Theorem 2.2. *Every primitive Pythagorean triple A, B, C with B even has the form*

$$A = n^2 - m^2, \quad B = 2mn, \quad C = n^2 + m^2$$

for relatively prime, positive integers $n > m$.

Remark 2.3. There is a small subtlety in the statement of this theorem. It tells us the form of every primitive Pythagorean triple A, B, C with B even. What it does not say (because it is not true) is that every triple given by those formulas is a primitive Pythagorean triple—something we shall see in the examples below. But in terms of listing all the primitive Pythagorean triples, this theorem tells us we will not miss any (with B even). And if we automatically pair the triple (A, B, C) with the one (B, A, C) , we have captured all the primitive triples, as well as a few imprimitive ones. We will then offer a second algorithm which has been tweaked slightly to exclude those triples where A, B, C have a common factor. We will take a few moments to discuss the tweak in the context of all we have developed so far.

2.3.2. Implementing the Algorithm. We provide some code below to generate a list of Pythagorean triples characterized by Theorem 2.2. The theorem determines the triple A, B, C as a function of integers m, n with $1 \leq m < n$ and assuming m, n have no common factors. Then given a triple, the algorithm computes the area ($\text{Area} = \frac{1}{2}AB$) of the corresponding right triangle. Finally, it implicitly writes $\text{Area} = N_0t^2$ where N_0 is square free (that is, $N_0 = 1$ or N_0 is the product of distinct primes) and prints the value of N_0 as the congruent number. As we discussed in Chapter 1, knowing that N_0 is a congruent number tells us that N_0t^2 is also for any positive integer t , so in terms of looking for congruent numbers, it makes the most sense to start with their square-free “core”.

The first bit of code does the computations above that are associated to the values of m, n with $1 \leq m < n < 6$. We shall talk about the function `gcd` in a moment, but for now we assume that it provides the test for whether m and n have no common factors. The code below is written in Sage [S⁺15].

```

for n in range(2,6):
    for m in range(1,n):
        if gcd(m,n) == 1:
            A = n^2 - m^2
            B = 2*m*n
            C = n^2 + m^2
            Area = A*B/2
            Congruent_number = squarefree_part(Area)
            print "(m,n) = (" ,m," ,",n,"),", \
                  "(A,B,C) = (" ,A," ,",B," ,",C,"), "
            print "Congruent number =", \
                  Congruent_number, "\n"

```

The output follows.

```

(m,n) = ( 1 , 2 ), (A,B,C) = ( 3 , 4 , 5 ),
Congruent number = 6

```

```

(m,n) = ( 1 , 3 ), (A,B,C) = ( 8 , 6 , 10 ),
Congruent number = 6

```

```

(m,n) = ( 2 , 3 ), (A,B,C) = ( 5 , 12 , 13 ),
Congruent number = 30

```

```

(m,n) = ( 1 , 4 ), (A,B,C) = ( 15 , 8 , 17 ),
Congruent number = 15

```

```

(m,n) = ( 3 , 4 ), (A,B,C) = ( 7 , 24 , 25 ),
Congruent number = 21

```

```

(m,n) = ( 1 , 5 ), (A,B,C) = ( 24 , 10 , 26 ),
Congruent number = 30

```

```

(m,n) = ( 2 , 5 ), (A,B,C) = ( 21 , 20 , 29 ),
Congruent number = 210

```

```

(m,n) = ( 3 , 5 ), (A,B,C) = ( 16 , 30 , 34 ),
Congruent number = 15

```


$(m,n) = (4, 5)$, $(A,B,C) = (9, 40, 41)$,
Congruent number = 5

Some simple observations include the following.

- (1) The second triple that is produced is certainly not primitive; neither are any of the triples with A even (that is, with m and n both odd).
- (2) We see that the last entry has already given us a proof that 5 is a congruent number.

Now we want to modify the code so that only primitive Pythagorean triples appear; the intent is simply to tweak the code, not to be particularly efficient.

Remark 2.4. The change to the code (that is, where the new code deviates from the statement of the theorem) is to test to see if A and C are relatively prime. By Remark 2.1, if any two of A, B, C are relatively prime, the triple is primitive.

In the code we have jumped ahead of ourselves just a bit by using the function “gcd” which will be introduced formally in the next chapter. But taking a sneak peek now will afford us some perspective and foreshadowing of things to come. The abbreviation “gcd” is shorthand for greatest common divisor, which is something for which we have an intuitive feel: it is the largest common factor, so for example we somehow “know” that the $\text{gcd}(24, 30) = 6$ probably because we thought of $24 = 2^3 \cdot 3$ and $30 = 2 \cdot 3 \cdot 5$, and we see that we can pull exactly one 2 and one 3 from each of the numbers. This intuition comes from the Fundamental Theorem of Arithmetic (Theorem 3.15) that every integer 2 or greater can be factored uniquely into a product of primes. Comparing the factorizations allows us to extract the greatest common divisor, and when two integers a, b have no common divisors other than ± 1 , we say they are relatively prime, and write $\text{gcd}(a, b) = 1$.

What is truly remarkable is that in the next chapter we shall learn how to compute a gcd very quickly without factoring, and this in turn is quite fortuitous because factoring is hard to do, and it is precisely the difficulty of factoring large numbers which makes one of

the most famous public-key encryption schemes (RSA) secure. We shall meet the RSA encryption scheme in Chapter 5.

The new Sage code is as follows.

```
for n in range(2,8):
    for m in range(1,n):
        if gcd(m,n) == 1:
            A = n^2 - m^2
            B = 2*m*n
            C = n^2 + m^2
            if gcd(A,C) == 1:
                Area = A*B/2
                Congruent_number = squarefree_part(Area)
                print "(m,n) = (" ,m, ", " ,n, " ), " , \
                    "(A,B,C) = (" ,A, ", " ,B, ", " ,C, " ), "
                print "Congruent number = " , \
                    Congruent_number, "\n"
```

The new output is shown below.

```
(m,n) = ( 1 , 2 ), (A,B,C) = ( 3 , 4 , 5 ),
Congruent number = 6
```

```
(m,n) = ( 2 , 3 ), (A,B,C) = ( 5 , 12 , 13 ),
Congruent number = 30
```

```
(m,n) = ( 1 , 4 ), (A,B,C) = ( 15 , 8 , 17 ),
Congruent number = 15
```

```
(m,n) = ( 3 , 4 ), (A,B,C) = ( 7 , 24 , 25 ),
Congruent number = 21
```

```
(m,n) = ( 2 , 5 ), (A,B,C) = ( 21 , 20 , 29 ),
Congruent number = 210
```

```
(m,n) = ( 4 , 5 ), (A,B,C) = ( 9 , 40 , 41 ),
Congruent number = 5
```

$(m,n) = (1 , 6)$, $(A,B,C) = (35 , 12 , 37)$,
Congruent number = 210

$(m,n) = (5 , 6)$, $(A,B,C) = (11 , 60 , 61)$,
Congruent number = 330

$(m,n) = (2 , 7)$, $(A,B,C) = (45 , 28 , 53)$,
Congruent number = 70

$(m,n) = (4 , 7)$, $(A,B,C) = (33 , 56 , 65)$,
Congruent number = 231

$(m,n) = (6 , 7)$, $(A,B,C) = (13 , 84 , 85)$,
Congruent number = 546

We make two more observations.

- (1) We see the congruent numbers that are listed do not appear in any particular order. So while even in this short list we see congruent numbers as large as 546, we have no idea whether a number that has not appeared, say 157, is a congruent number. It is true that 157 is a congruent number, but it would take a very long time before it would appear in this list.
- (2) We note that congruent numbers can be repeated in the list.

Here we provide some alternate Sage code (perhaps of more use for larger tables).

```
def CN(nn):
    List=[]
    for n in range(2,nn):
        for m in range(1,n):
            if gcd(m,n) == 1:
                A = n^2 - m^2
                B = 2*m*n
                C = n^2 + m^2
```

```

    if gcd(A,C) == 1:
        Area = A*B/2
        Congruent_number = squarefree_part(Area)
        List.append([m,n,A,B,C,Congruent_number])
return(List)

CNList=CN(8)
table(CNList,align='right', \
      header_row=["$m$", "$n$", "$A$", "$B$", "$C$", "$CN$"])

```

The code above generates the following table.

<i>m</i>	<i>n</i>	<i>A</i>	<i>B</i>	<i>C</i>	<i>CN</i>
1	2	3	4	5	6
2	3	5	12	13	30
1	4	15	8	17	15
3	4	7	24	25	21
2	5	21	20	29	210
4	5	9	40	41	5
1	6	35	12	37	210
5	6	11	60	61	330
2	7	45	28	53	70
4	7	33	56	65	231
6	7	13	84	85	546

Summary. We have given a parametrization of primitive Pythagorean triples which is easily implemented as an algorithm, and which will eventually list all primitive Pythagorean triples, along with every square-free congruent number. The problem is we seem to have to wait an unknown period of time before a given congruent number might appear. Perhaps we need a few more tools that might provide insight into which numbers might or might not be congruent numbers.

Exercise. Find a square-free congruent number not in the list above, and show all the work to obtain it.

2.4. Looking for Underlying Structure: Geometry vs. Algebra

A major theme in this book will be to understand the intrinsic structure of the set of points on an elliptic curve. That probably sounds quite mysterious right now, but let's take a broader view with a closely related question with which you have had some experience. How can we detect structure exhibited by a set of solutions to a given equation or set of equations?

As this is such a broad question, we should expect an equally broad spectrum of answers. For example, when we ask for the solutions to the equation $x^2 + y^2 = 1$, it matters a great deal whether we are looking for solutions in the integers (\mathbb{Z}), in the rationals (\mathbb{Q}), or in the real or complex numbers (\mathbb{R} or \mathbb{C}). Over the real numbers the set of solutions certainly has a geometric structure; over the rationals we may see no apparent structure, but we have completely characterized all the rational points and hence the primitive Pythagorean triples. Over the integers, the equation $x^2 + y^2 = 1$ has only four solutions, and while there is no apparent structure, simply knowing the set of solutions is finite (and, even better, the exact number) is often a huge victory. Let's see how some of these ideas play out with the *Fermat curve*, $x^n + y^n = z^n$, $n > 2$, for which, in the context of Fermat's Last Theorem, we are seeking integer solutions to this equation. We mentioned a few things that were known about the solutions before Wiles's proof, but there is a good deal more we can say now that we have a better perspective.

Just as Pythagorean triples correspond to rational points on the unit circle $x^2 + y^2 = 1$, so too is there a correspondence between nontrivial integer solutions to $x^n + y^n = z^n$ (i.e., where $x, y, z \neq 0$) and rational points on $u^n + v^n = 1$. An integer solution (x, y, z) to $x^n + y^n = z^n$ maps to the rational solution $(u, v) = (x/z, y/z)$ on $u^n + v^n = 1$. And conversely, a rational point $(\frac{a}{b}, \frac{c}{d})$ on $u^n + v^n = 1$, gives a solution to the Fermat equation since

$$\left(\frac{a}{b}\right)^n + \left(\frac{c}{d}\right)^n = 1 \text{ implies } (ad)^n + (bc)^n = (bd)^n.$$

What more can be said? While we will not talk about solutions to curves over complex numbers (except at the end of the text), it is




Figure 2.2

the set of complex points on a curve which actually have a geometric structure. The set of complex points on an elliptic curve is naturally associated to a torus (doughnut); other curves, a sphere. Still others associated to curves have a geometric structure which looks like several tori glued together, sort of like fat links in a chain. There is a geometric invariant associated to such surfaces called the *genus* of the surface, and the genus refers to the number of holes in the surface, the sphere having genus 0, an elliptic curve having genus 1, and so on (see Figure 2.2).

In 1922 Mordell made a conjecture that, for any curve of genus at least 2, the set of rational points on the curve is finite, and in 1983 Gerd Faltings proved Mordell's conjecture. So given our knowledge that there are infinitely many rational points on $x^2 + y^2 = 1$, we know for sure it has genus 0 or 1. But the real import of Falting's theorem is that for $n > 4$, the curve $u^n + v^n = 1$ has genus ≥ 2 . As a result, there are only finitely many rational points and, by our observation above, only a finite number of primitive integer solutions to $x^n + y^n = z^n$. So if there were counterexamples to the Fermat conjecture, as of 1983 there could only be finitely many for a given value of n .

Exercise. Find all of the rational points on the curve $x^n + y^n = 1$ where n is an integer, $n > 2$.

Now let's broaden our scope even further, and move from geometric considerations to algebraic ones. Let's first look at a set of solutions with both a geometric and algebraic structure. Let

$$S = \{(x, y, z) \in \mathbb{R}^3 \mid x + y + z = 0\}.$$

Depending upon your background, you may recognize this solution set as a plane in \mathbb{R}^3 passing through the origin with normal vector (parallel to) $(1, 1, 1)$, but the set also has an algebraic structure which in turn gives another interpretation of S as a geometric object. First to the geometry. There are many ways to determine a line: two distinct points, a point and a direction, the intersection of two nonparallel planes in \mathbb{R}^3 , and so on. Analogously, while we have just said a plane in \mathbb{R}^3 can be characterized by a point and a normal vector, it might be even more natural to characterize it by a point in the plane and two independent vectors (directions) which lie the plane. Indeed you can easily convince yourself that if you took a large piece of paper, placed a dot on it, and drew two vectors v_1, v_2 , emanating from that point (so that the two vectors do not lie on a common line), you could navigate to any point on the plane by moving a uniquely determined distance along v_1 followed by a uniquely determined distance in the direction of v_2 (we are also allowed to back up along v_1 or v_2). So with this somewhat imprecise definition of a plane, let's see that our solution set S satisfies it.

First we note that our set S contains a point, namely $(0, 0, 0)$. Also note that since every element (x, y, z) of S satisfies $x + y + z = 0$, the point (x, y, z) has the form $(x, y, -x - y)$. Assuming we know how to add vectors in \mathbb{R}^3 and multiply by scalars (see a formal definition just below), this means that every element of S can be written uniquely as the *linear combination* of the vectors $(1, 0, -1)$ and $(0, 1, -1)$:

$$(x, y, -x - y) = (x, 0, -x) + (0, y, -y) = x(1, 0, -1) + y(0, 1, -1).$$

The vectors $(1, 0, -1)$ and $(0, 1, -1)$ are elements of S and are not scalar multiples of each other, so they represent a choice of two independent directions, discussed above, and every vector in S is uniquely representable in terms of them.

Actually, \mathbb{R}^n is the prototype for a broad class of algebraic structures called *vector spaces*, and it is worthwhile to digress briefly in this direction. The following is an informal definition of a vector space. Note that of course in order to prove theorems about vector spaces, we need a precise definition, but that would take us too far afield, and for now we should be focused on the larger view, not the smaller one.

A vector space is a nonempty set V whose elements are called *vectors* with an associated field F consisting of scalars. Already we have an undefined term, a field. For the discussion in this chapter we may assume $F = \mathbb{Q}$ or \mathbb{R} , or \mathbb{C} . Loosely speaking a field is a set in which you can add, subtract, multiply, and divide (by anything nonzero) and still stay in the set. So while the integers \mathbb{Z} comprise a wonderful set in which you can do three of the four things, you cannot always take an integer and divide by a nonzero integer and have the result be integral; for example, $2/3 \notin \mathbb{Z}$, so \mathbb{Z} is not a field, but the others we mentioned are.

We start again. A *vector space* V over a field F is a nonempty set V (of vectors) with some notion of addition which is compatible with multiplication by scalars from F in which some basic properties hold. Let's make this a bit more precise for $V = \mathbb{R}^n$. Given $v_1 = (a_1, \dots, a_n)$, $v_2 = (b_1, \dots, b_n) \in V = \mathbb{R}^n$, and $c \in F = \mathbb{R}$, we define the operations of vector addition and scalar multiplication by acting component-wise:

$$\begin{aligned} v_1 + v_2 &= (a_1 + b_1, \dots, a_n + b_n) && \text{(vector addition),} \\ c \cdot v_1 &= (ca_1, \dots, ca_n) && \text{(scalar multiplication).} \end{aligned}$$

As to the basic properties, we want there to be an *identity* for the operation of addition—in our case the vector $\mathbf{0} = (0, \dots, 0) \in V$, which has the property $v + \mathbf{0} = v = \mathbf{0} + v$ for every $v \in V$ (clear from the definition). Also for each $v \in V$, we want a notion of its *additive inverse*, $-v$, which has the property that $v + (-v) = \mathbf{0}$ for every $v \in V$. So it seems pretty clear from the definitions above that if $v = (a_1, \dots, a_n)$, then $-v = (-a_1, \dots, -a_n)$. There are a few more properties, but they all feel completely natural, so we choose not to pull them into this discussion.

What we do want to bring into this discussion is that an amazing variety of sets have the structure of a vector space. For our examples we shall just consider vector spaces over \mathbb{R} , that is, with scalar multiplication coming from \mathbb{R} . Here are just a few examples.

- (1) $V = \mathbb{R}^n$ (of course).
- (2) V is the set of all polynomials with real coefficients. Vector addition is just the addition of polynomials, and multiplying a polynomial by a scalar just multiplies all the coefficients by the same scalar. The zero polynomial is the identity, and for any polynomial $p(x) = a_n x^n + \cdots + a_1 x + a_0$, we have the inverse $-p(x) = -a_n x^n - \cdots - a_1 x - a_0$.
- (3) $V = \{f : \mathbb{R} \rightarrow \mathbb{R} \mid f \text{ is continuous}\}$ is a vector space over \mathbb{R} . What does that mean? Maybe an example is easier to swallow than a definition. Suppose $f = \sin x$ and $g = 3x^2 + 2$. Both f and g are continuous functions from $\mathbb{R} \rightarrow \mathbb{R}$, so they are vectors in V . Does $f + g$ have a meaning? Sure it does. Somewhat innocently, we would just write $f + g = \sin x + 3x^2 + 2$ and not give it a second thought. But what we are really saying is that there is a new function $h = f + g$, and we compute its values by the rule $h(x) = f(x) + g(x)$. Also something like $\sqrt{2} \cdot f$ makes perfect sense to us: $\sqrt{2} \cdot f = \sqrt{2} \sin x$.

Now there is something deeper going on: How do we know that if $f, g \in V$, then $f + g$ is also? What the question is really asking is how do we know that the sum of two continuous functions is continuous? The answer is provided by an early theorem from calculus. The same is true for scalar multiplication of a continuous function.

- (4) $V = \{f : \mathbb{R} \rightarrow \mathbb{R} \mid f \text{ is differentiable}\}$ is also a vector space, with the definitions as in the previous example, and again the only serious question is whether V is closed under the operations of addition and scalar multiplication, meaning if $f, g \in V$ and $c \in F = \mathbb{R}$, are $f + g, cf \in V$? But again it is a theorem from calculus which tells us this is true.

Now in some vector spaces, we have a notion that a certain finite number of vectors will allow us to determine all others as a *linear combination* of them. For example, every vector in \mathbb{R}^n can be written uniquely as a linear combination of n vectors:

$$(a_1, \dots, a_n) = a_1(1, 0, \dots, 0) + a_2(0, 1, 0, \dots, 0) + \dots + a_n(0, \dots, 0, 1).$$

Any such set \mathcal{B} (such as $\{(1, 0, \dots, 0), (0, 1, 0, \dots, 0), \dots, (0, \dots, 0, 1)\}$ for \mathbb{R}^n) with which one can express every element of the vector space uniquely as a linear combination of the elements of \mathcal{B} is called a *basis* for the vector space, and the number of elements in \mathcal{B} is called the *dimension* of the vector space.

So \mathbb{R}^n is a vector space (over \mathbb{R}) having dimension n , and we have showed the set

$$S = \{(x, y, z) \in \mathbb{R}^3 \mid x + y + z = 0\}$$

is a vector space of dimension 2. The notion of dimension and of a basis are of enormous importance in mathematics.

Now let's come full circle and look at a few more sets with an algebraic structure, namely that of a vector space, which arise as solution sets to equations. We begin with a differential equation that every first-term calculus student can solve: Find the general solution to the differential equation $y' - 2y = 0$. The answer is $y = Ce^{2x}$ for an arbitrary constant C .

Let's approach this in another way. Let V be the set of all possible solutions to the differential equation $y' - 2y = 0$. We observe that V is not the empty set, since it is obvious that $y = 0$ is a solution to the equation. Now if f and g are two solutions to $y' - 2y = 0$, then so are $f + g$ and cf for any scalar $c \in \mathbb{R}$. To see this is easy. To say that $f, g \in V$ is to say that $f' - 2f = 0 = g' - 2g$, so

$$\begin{aligned} (f + g)' - 2(f + g) &= f' + g' - 2f - 2g \\ &= (f' - 2f) + (g' - 2g) = \mathbf{0} + \mathbf{0} = \mathbf{0} \text{ and} \\ (cf)' - 2(cf) &= cf' - 2cf = c(f' - 2f) = c \cdot \mathbf{0} = \mathbf{0}, \end{aligned}$$

where we have used that old chestnut from first-term calculus, that “the derivative of a sum is the sum of the derivatives”, and “the derivative of a constant times a function is the constant times the derivative of the function”. You see—those theorems do have a purpose.

As we have seen in the examples above, this makes the set V of solutions into a vector space. And since we know that every element of V (that is every solution to $y' - 2y = 0$) has the form $y = Ce^{2x}$, we have simply inferred that V is a one-dimensional vector space over \mathbb{R} with a basis consisting of one element $\{e^{2x}\}$.

Now let V be the set of solutions to the homogeneous differential equation $y'' + 9y = 0$. Depending upon your calculus (or physics) background, you may know that this differential equation describes a simple harmonic oscillator, and while you may not know the general solution, everyone can easily check that both $y = \cos 3x$ and $y = \sin 3x$ are solutions. So again, let’s examine the structure of V . By exhibiting solutions, we have shown that the set V is nonempty. Now suppose that $f, g \in V$. Is $f + g \in V$? Well to check, you would need to know that $(f+g)'' + 9(f+g) = 0$, but $f, g \in V$ so $f'' + 9f = 0$ and $g'' + 9g = 0$, and since “the derivative of a sum is the sum of the derivatives” (twice), we see that

$$(f+g)'' + 9(f+g) = (f'' + g'') + 9(f+g) = f'' + 9f + g'' + 9g = \mathbf{0} + \mathbf{0} = \mathbf{0},$$

so $f + g$ is a solution, hence in V . Similarly, any constant times a solution is a solution. So immediately we see that the set

$$S = \{a \sin 3x + b \cos 3x \mid a, b \in \mathbb{R}\}$$

are all solutions, so

$$S = \{a \sin 3x + b \cos 3x \mid a, b \in \mathbb{R}\} \subseteq V.$$

It is clear that $\sin 3x$ and $\cos 3x$ are not scalar multiples of each other, so we really need both of them in S , and indeed S is a vector space of dimension 2. Now when you take a course in differential equations, you will learn that V is a two-dimensional vector space, so it follows that $S = V$, and you have found a basis for V . In the language of a differential equations course, you would say instead that the general solution to $y'' + 9y = 0$ has the form $y = a \sin 3x + b \cos 3x$.

Let's consider a final example, the set V of solutions to $y''' - 2y'' + y' - 2y = 0$. Yes, V is again a vector space, and it follows from the general theory that its dimension is 3. We note that $\cos x$, $\sin x$ are e^{2x} are each solutions, hence $ae^{2x} + b\cos x + c\sin x$ is also, and this is the general solution meaning that $\{e^{2x}, \cos x, \sin x\}$ is a basis for this vector space.

Exercise. Let V be the set of functions $f : \mathbb{R} \rightarrow \mathbb{R}$ which satisfy the differential equation $f'' - f = 0$. Show that V is a vector space over \mathbb{R} and, assuming its dimension is 2, find a basis for V .

2.5. More about Points on Curves

We want to begin to extend our discussion of the structure of solutions to the set of points on curves, which is a good deal more delicate, so we shall ease into things.

We have only looked at a few curves so far. We had a great deal of success characterizing the rational points on $x^2 + y^2 = 1$, and with the help of Wiles and Taylor, we know the story for the curves $x^n + y^n = 1$ for $n \geq 3$. But that's a pretty small sample, and for all but a few of those curves, it took more than 350 years to get the answer.

Perhaps something more middle-of-the-road would be nice. Our interest is in understanding elliptic curves, so let's take a simple case. We will look at the family of curves of the form

$$y^2 = x^3 + k, \quad k \neq 0.$$

The value $k = 0$ is excluded because for a curve to be an elliptic curve the cubic $x^3 + k$ must have distinct roots (in \mathbb{C}). We consider a typical example $y^2 = x^3 + 17$. Figure 2.3 is a plot depicting the real points on the curve.

Notice a few obvious things. The curve is symmetric about the x -axis as it has the form $y^2 = f(x)$, so if (x, y) is a solution, so is $(x, -y)$. Second, the real points are restricted to values of x for which $x^3 + 17 \geq 0$ (since that value must be a square of a real number), so $x \geq -\sqrt[3]{17} \approx -2.57128\dots$. We can calculate points somewhat randomly: $(-\sqrt[3]{17}, 0)$, $(0, \pm\sqrt{17})$, or more generally (for $r \geq -\sqrt[3]{17}$)




Figure 2.3

the point $(r, \pm\sqrt{r^3 + 17})$, but if we asked for which values of r would the resulting point be a rational point, this would prove more challenging. Now for this particular curve we easily see that $(2, \pm 5)$ are rational points (even integer points) on the curve. But did you know that

$$\left(-\frac{64}{25}, -\frac{59}{125}\right), \left(\frac{38194304}{87025}, -\frac{236046706033}{25672375}\right) \text{ or}$$

$$\left(\frac{532027047589930897040873195264}{4848863077511293855911670225}, \frac{388064005784387552318916270407513322740532287}{337644656448214941842939018840311120390375}\right)$$

are rational points as well? The first is easy enough to check. For the others we might be happier with an explanation.

While we may not have a clue where these other points came from, Claude Gaspar Bachet did. He proved quite a remarkable formula in 1621 now known as Bachet's duplication formula. It said if you know

the coordinates (x, y) of some point on the elliptic curve $y^2 = x^3 + k$ ($k \neq 0$), then assuming that $y \neq 0$,

$$(2.3) \quad \left(\frac{x^4 - 8kx}{4y^2}, \frac{-x^6 - 20kx^3 + 8k^2}{8y^3} \right)$$

is another point on the curve. And yes; the points listed above were obtained via his formula, starting with the point $(2, 5)$ and iterating the process.

While it is tedious, it is completely straightforward to check that Bachet's formula works, but where in the world did it come from? Below we outline a method that uses both our familiar Cartesian coordinate system (invented by Descartes a bit later in the century) and the calculus invented by Newton and Leibniz which followed Descartes's work.

If we had started with the point $(2, -5)$ instead of $(2, 5)$, the first two points out would be

$$\left(-\frac{64}{25}, \frac{59}{125} \right), \left(\frac{38194304}{87025}, \frac{236046706033}{25672375} \right).$$

We can see this respects the symmetry, so there are some interesting properties to this formula. Now let's give another clue, such as




Figure 2.4

plotting the points in this progression and the line segments between them (see Figure 2.4). Hopefully, it seems clear (or at least reasonable) that if we start with a point (a, b) on the curve, then the formula gives, as the next point, the point of intersection of the curve with the tangent line to the curve $y^2 = x^3 + 17$ at (a, b) . Of course that presumes there is another point of intersection, and indeed only one other. We shall pursue some of these ideas below.

Remark 2.5. It is certainly worth noting that if the curve $y^2 = x^3 + k$ is chosen with k a rational number, then if we start with a rational point, the formula in equation (2.3) will yield another rational point on the curve. So while we have not discovered a way to list all the rational points, we seem to have found an algorithm by which to produce a sequence of them.

There are a number of interesting questions to ponder:

- (1) Do you think that this sequence of points ever cycles back to the start? Do you think this might be a common occurrence or a rare one?
- (2) Are we always guaranteed that the tangent line to the curve always intersects the curve in a second point? Could it intersect in more than one point?
- (3) Certainly this method “fails” if we start at the point where $y = 0$, but at least from the point of view of the formula, that case was precluded. But is there something deeper going on?

Derive Bachet’s formula using the sequence of steps in the exercise below.

Exercise. Using the ideas above, prove the Bacht duplication formula for $y^2 = x^3 + k$, $k \neq 0$.

$$(x, y) \mapsto \left(\frac{x^4 - 8kx}{4y^2}, \frac{-x^6 - 20kx^3 + 8k^2}{8y^3} \right).$$

We outline some useful steps.

- (1) Use implicit differentiation to derive a formula for the slope of the tangent line to the curve $y^2 = x^3 + k$ which is valid at any point (x, y) where $y \neq 0$.
- (2) Now write down the equation of the tangent line to the curve at the point (a, b) where we assume $b \neq 0$. It will be convenient if you use m for the slope for the time being until you need to use its actual value.
- (3) Now we want to find the point(s) of intersection of the tangent line with the cubic, and this requires a little work. Substitute the expression for y given by the line into the equation that defines the cubic results in an equation of the form $f(x) = 0$ where f is a polynomial of degree 3. Your job is to factor the polynomial since its roots are the x -coordinates corresponding to the points of intersection. Here we catch a bit of a break. Certainly one of the roots is a , which means $(x - a)$ is a factor. But it should not be too much of a surprise that a is (at least) a double root since the line is tangent to the curve at $x = a$ (much like $y = (x - a)^r$ is tangent to the x -axis at $x = a$ and the root a has *multiplicity* r). After factoring out the first of the $(x - a)$ factors, it would be a good time to put in the real value of m to see what simplifies.

2.6. Gathering Some Insight about Plane Curves

In this last section we want to gain some intuition about a famous theorem in algebraic geometry that concerns the number of points of intersection of two plane curves. In our investigation of the Bacht formula, we seemed to be suggesting that a line and a cubic could intersect in at most three points (counting the point of tangency as two of them). Bézout's theorem is a precise statement of what can

happen, but it depends upon the field ($\mathbb{Q}, \mathbb{R}, \mathbb{C}$, etc.) in which we look for points of intersection, counting the intersection points with *multiplicities* (as necessary), as well as another notion new to us: that of curves in *affine* or *projective* space. We will explore some of these ideas now, leaving the notion of projective space for a later chapter.

We need to define what is meant by a *curve in the plane*. Simply put, we take any polynomial $f(x, y)$ in two variables (this gives the “in the plane” part), and the curve is the set of points (x, y) where $f(x, y) = 0$. Often we refer to this set of points as the *zero set* of the polynomial f , denoted $Z(f)$. If you have had some multivariable calculus, this notion is nothing other than a particular level set of the function $f(x, y)$.

Coming from a background focused on functions of a single variable, it is easy to mistakenly conflate things that are quite different. For example, in calculus or high-school algebra, one might casually refer to x^2 as a parabola, but it is not. More appropriately, one would have said the graph of the function $f(x) = x^2$ is a parabola, and that notion is consistent with our characterization that a parabola is the zero set of $f(x, y) = y - x^2$. For the record, the zero set of x^2 is the y -axis, since that is the set of points (x, y) where $f(x, y) = x^2 = 0$. And of course the notion of the zero set changes as we change the field. If the field is \mathbb{Q} , the zero set is the set of rational points; if the field is \mathbb{R} , the zero set is the set of real points; and so on.

To start thinking correctly about curves, we need to move away from only thinking about the graphs of functions. Of course, graphs of functions are special cases of curves, but certainly not all curves. For example, if $g(x)$ is a polynomial in one variable, then its graph is the zero set of $f(x, y) = y - g(x)$, a special kind of curve. But curves like the zero set of $x^2 + y^2 - 1$, the unit circle, are not the graphs of functions.

And to stretch our terminology a bit further, depending on the field in question, the notion of a curve as a zero set can be a bit unintuitive: the zero set of $f(x, y) = x^2 + y^2 - 1$ in the plane \mathbb{R}^2 is the unit circle, but the zero set of $f(x, y) = x^2 + y^2$ in the plane \mathbb{R}^2 is just the point $\{(0, 0)\}$, while the zero set of $f(x, y) = x^2 + y^2 + 1$ in \mathbb{R}^2 is the empty set, since there are no real solutions to $x^2 + y^2 = -1$.

We need just a bit more terminology, the degree of a polynomial in two variables. A polynomial, $f(x, y)$, in two variables is a sum of terms (monomials), each of the form $cx^m y^n$ where c is a nonzero constant and m, n are nonnegative integers. We say that the *degree of the monomial* $cx^m y^n$ is $m + n$, and we say that the (*total*) *degree of* $f(x, y)$ is the maximum of the degrees of monomials whose sum gives f . So we would say that

$$f(x, y) = x^4 - 3x^3y^2 + y^4 - 3x^2y + 2$$

has degree 5. We say that a *line* is a plane curve $Z(f)$, where $f = f(x, y)$ has degree 1; that is, a line is the zero set of $f(x, y) = ax + by + c$ or, said most simply, a line is the set of points in the plane which satisfies $ax + by + c = 0$. Similarly we call the zero set of a polynomial $f(x, y)$ having degree 2 a *conic*; so a conic is the set of points where $f(x, y) = ax^2 + bxy + cy^2 + dx + ey + f = 0$. *Cubics* are the zero sets of polynomials of degree 3, and *quartics* are of degree 4. We don't fuss about terminology of curves of higher degree, referring only to a plane curve of degree n as the zero set of a polynomial $f(x, y)$ of total degree n . Finally, we refer to a plane curve as a *rational curve* if the polynomial used to define it has all rational coefficients, so $Z(y^2 - x^3 + x)$ is a rational cubic.

Now we are equipped with a bit of vocabulary which will allow us to gain insight into Bézout's theorem. We begin with some simple examples and questions to tease out where the subtleties lie.

Exercise. Properties of rational lines in the plane.

- (1) Is every point on a rational line a rational point?
- (2) If a line passes through at least two rational points, is it a rational line? What about lines if we only know one rational point through which they pass?
- (3) Consider two distinct rational lines which intersect. Do they intersect in a rational point?

Exercise. Characterizing the intersection of lines and conics.

- (1) In how many points can two arbitrary lines (in the plane) intersect?
- (2) In how many points can a line and a conic intersect?

In considering the exercises above, you may have noted that the line $Z(y - x)$ and the conic $Z(y^2 - x^2)$ did not behave in an expected manner, intersecting in an infinite number of points. There are a couple of ways in which to avoid or at least anticipate this kind of degeneracy, but both involve the notion of factoring a polynomial in two variables. As a topic of study, this will appear in an abstract algebra course when you study polynomial rings and unique factorization. What is remarkable is that in many ways the set of polynomials with coefficients in a field have factorization properties similar to those of the integers. Over \mathbb{Z} , the Fundamental Theorem of Arithmetic says that every integer 2 or greater can be factored uniquely into a product of primes. Just as with primes in the integers, there is the notion of polynomials that have no nontrivial factorization; we call such polynomials irreducible. More formally, we say that a polynomial $f(x, y)$, with coefficients in a field and of degree ≥ 1 (we exclude the constant polynomials), is *irreducible* if whenever we factor $f(x, y) = g(x, y)h(x, y)$, one of g or h is just a constant. When a polynomial is not irreducible, we refer to it as *reducible*, such as $f(x, y) = y^2 - x^2$. And once again the field of coefficients becomes relevant. For example $f(x, y) = x^2 + y^2$ is irreducible if the coefficients are \mathbb{Q} or \mathbb{R} , but f is reducible over \mathbb{C} , since $f(x, y) = x^2 + y^2 = (x + iy)(x - iy)$.

The problem we encountered above with a line and a conic intersecting in an infinite number of points is precisely the issue that the two polynomials $f(x, y) = y - x$ and $g(x, y) = y^2 - x^2$ are not relatively prime (i.e., they share the common factor $y - x$), so generally we shall preclude that case. Even more of a constraint is that we consider curves $Z(f)$ where $f(x, y)$ is irreducible. Then if we have curves $Z(f)$ and $Z(g)$, where f and g are irreducible but not relatively prime, it will be the case that $f(x, y)$ is a constant multiple of $g(x, y)$ which will mean the curves $Z(f)$ and $Z(g)$ are identical.

Exercises. In the questions below, we assume all the plane curves are irreducible, meaning they are the zero sets of polynomials $f(x, y)$ where $f(x, y)$ is irreducible. It follows (from abstract algebra) that two distinct irreducible plane curves can only intersect in a finite number of points. The questions below try to get at discovering what that number might be.

For all the problems below, consider your curves in \mathbb{R}^2 . Can you come up with examples that suggest answers to these questions? Can you prove any of your assertions?

- (1) In how many points can two (distinct) conics intersect?
- (2) In how many points can a conic and a cubic intersect?
- (3) In how many points can two (distinct) cubics intersect?
- (4) What would be your guess for a generalization?
- (5) Consider the intersection of a rational line with a rational conic.
 - (a) Are the point(s) of intersection necessarily rational? Give a proof or provide a counterexample.
 - (b) Now let's suppose that the line intersects the conic in two points, one of which is rational. Is the second point necessarily rational? Give a proof or a counterexample.

To assist with your intuition, Figure 2.5 illustrates a few curves to consider. The first set fixes a parabola and slides the circle up the y -axis. The second set is a cubic and quartic, and a cubic and conic.

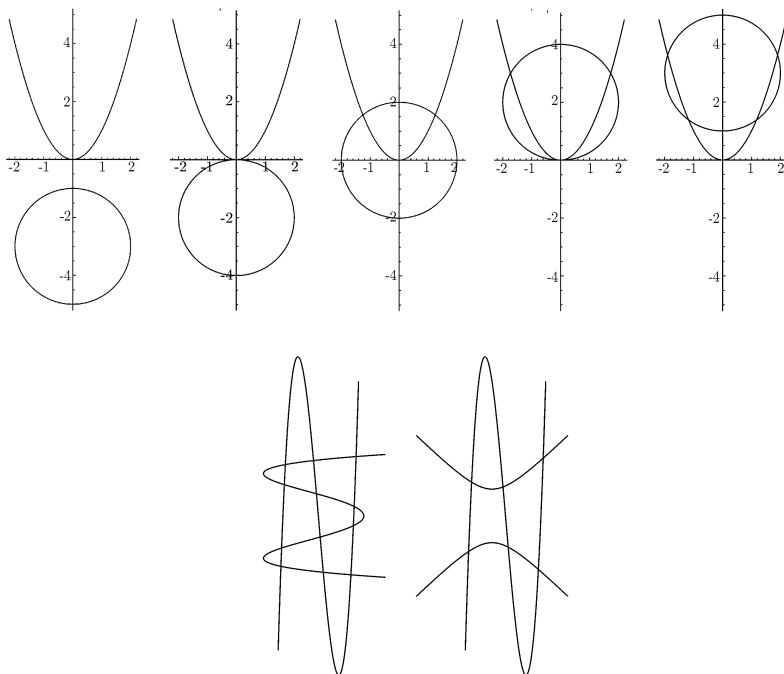


Figure 2.5

2.7. Additional Exercises

In Bachet's duplication formula, we made the assumption that the tangent line intersects the cubic with multiplicity 2 at the point of tangency; actually, you probably proved that it did in working out the duplication formula as an exercise. Let's investigate this concept more carefully.

Suppose we have a plane curve given as the zero set of a polynomial $f(x, y)$ with coefficients in a field. For example, for elliptic curves we are looking at $f(x, y) = y^2 - g(x) = 0$ where g is a nonsingular cubic polynomial. Suppose that a (nonvertical) line $y = mx + b$ intersects the curve. Then the roots of $f(x, mx + b)$ are precisely the x -coordinates of the points of intersection of the line and the curve.

We want to define what is meant by the line $y = mx + b$ intersecting the curve $f(x, y)$ at $x = a$ with multiplicity k .

Let $h(x)$ be a polynomial with coefficients in a field and having degree n , and let a be any element of the field. In our example above, we would have $h(x) = f(x, mx + b)$. We say that h has a *zero of order k* at $x = a$ (or the *order of vanishing of h at $x = a$* is k) if $h(x)$ can be written as the product $h(x) = (x - a)^k q(x)$ with q a polynomial satisfying $q(a) \neq 0$. So we shall say that the line $y = mx + b$ intersects the curve $f(x, y) = 0$ with multiplicity k at $x = a$ if $h(x) = f(x, mx + b)$ has a zero of order k at $x = a$.

Exercise. As a simple example show that the curve $y = (x - a)^k$ intersects the x -axis with multiplicity k at $x = a$ and with multiplicity 0 at all other points $x = b$.

Exercise. Next, let's gain a little more insight by examining the case of zeroes of order 1 and 2. Let $h(x)$ be a polynomial of degree $n \geq 2$ with coefficients in a field F , and let $a \in F$. Then prove the following.

- (1) $h(x) = (x - a)q(x) + h(a)$ for some polynomial q having coefficients in F .
- (2) $h(a) = 0$ if and only if $h(x) = (x - a)q(x)$.
- (3) h has a double root at a if and only if $h(a) = h'(a) = 0$, where $h'(x)$ is the first derivative of $h(x)$.

Exercise. Establish the following generalization of the work we have started above. Show that h has a zero of order k at $x = a$ if and only if $h(a) = h'(a) = \dots = h^{(k-1)}(a) = 0$ and $h^{(k)}(a) \neq 0$, where $h^{(i)}$ is the i th derivative of h . *Hint:* Taylor polynomials are your friend.

Exercise. Now consider $y^2 = g(x)$ where g is a cubic, that is the zero set of $f(x, y) = y^2 - g(x)$. We want to see that a nonvertical tangent has multiplicity at least 2 at the point of tangency.

Chapter 3

Some Elementary Number Theory

In the last chapter, we talked at length about curves, points on curves, and even rational points on curves. We considered the impressive Bachet duplication formula, which, given one rational point, produced another. By the end of this book we want to be doing some very sophisticated arithmetic using rational points on elliptic curves. We shall consider the set of all rational points on an elliptic curve and show that the set has an algebraic structure: given two rational points, we can produce a third. We may even have gained some insight into how this might happen given our excursions in the last chapter, but the actual procedure will turn out to be somewhat more complicated than our initial impressions might lead us to believe. We could attempt to describe the complications, but it is a bit more natural to bump into them and use them as motivation for making new definitions and developing new tools and perspectives.

Still, we should give a few hints. The algebraic structure we referred to is called a *group structure*. Before trying to come to grips with it, we should understand that it is a very natural structure that many sets have; indeed sets with a group structure are even more ubiquitous than sets with a vector space structure. So we will start with a set with which you are familiar, the integers, and progress to

the notion of modular arithmetic. Then in the next chapter, we shall give a high-level description of modular arithmetic by thinking of it not as an operation on the integers but by defining a group structure on a set of equivalence classes of integers. Hmm; it seems we need to know about equivalence classes, so we'll have to develop that too. We seem to have a full agenda, so let's start with developing some basic properties of the integers, and later give a much more revealing description of what you were doing in third grade with arithmetic.

3.1. The Integers

Let \mathbb{Z} denote the set of all integers, and let $\mathbb{N} = \{0, 1, 2, \dots\}$ be the subset of nonnegative integers, typically called the *natural numbers*. One might think that there is little which you do not already know about the integers, but actually that is far from the truth. The systematic study of the integers is called *number theory* (more precisely *elementary number theory*), and many a mathematician has spent his or her career teasing out some of its interesting properties. Yet to the uninitiated, it probably seems “inconceivable”¹ that one could spend a lifetime studying things that at first blush seem to hold no intrinsic interest. I mean how exciting is the number 31, or $2^{31} - 1$? Actually, they are interesting in that both of them are prime numbers (as are 13 and $2^{13} - 1$, 17 and $2^{17} - 1$, and 19 and $2^{19} - 1$), but we are getting a bit ahead of ourselves.

For some people, collections of integers with a specific property are fascinating: numbers called primes, perfect numbers, polygonal numbers, and so on; for others, it is the algebraic structure possessed by the set of all integers that is the real draw, mainly since that structure serves as probably the single most important prototype of the objects of study in abstract algebra and of all of algebraic number theory, called groups and rings. But what do we mean by structure? This is a question we first broached in Chapter 2, and one which will command a great deal of our attention, so let's start with some basic properties of the integers, some of which you have probably seen, but perhaps not at this level of detail.

¹As said by the character Vizzini from *The Princess Bride* (http://www.imdb.com/title/tt0093779/?ref_=fn_al_tt_1)

3.2. Some Basic Properties of the Integers

Most of this section is very standard; see e.g., [JJ98] or [Ros05] for expanded discussions.

We know that if we take any two integers, we can add, subtract, or multiply them together and produce a third integer. This is not true for division unless the one integer “divides” another. Yet from early arithmetic lessons, it still is often the case that you want to divide integers; the result was just slightly more complicated to describe. There was something about quotients and remainders that we need to think about more carefully.

We characterize the pairs of integers (a, b) , whose quotient a/b is an integer c , by saying b *divides* a , denoted $b \mid a$. That is, $b \mid a$ if and only if there is an integer c so that $a = bc$. Note this says $a/b = c \in \mathbb{Z}$. But of course, not all quotients of integers are integral. What do we do?

In elementary school, we learn to do long division of integers, and we can show for example, that 257 divided by 12 has quotient 21 and remainder 5. Put another way, $257 = 12(21) + 5$. There are many things we assumed about this process. One was that the remainder was nonnegative and smaller than the number by which we were dividing, and the second is that these numbers are unique. After all,

$$257 = 12(21) + 5 = 12(22) - 7 = 12(20) + 17.$$

So let’s try to get back to where we were in elementary school. Given two integers a, b with $b \neq 0$, we want to divide b into a and obtain a quotient and remainder where the remainder is “smaller” than the integer by which we were dividing. We see the need to be a bit more precise here since, in the example above, both remainders of 5 and -7 are smaller than 12, even smaller in absolute value.

We summarize this procedure in the following theorem, called the *division algorithm*.

Theorem 3.1 (Division algorithm). *Let $a, b \in \mathbb{Z}$, $b \neq 0$. Then there exist unique integers q and r with $a = bq + r$ and $0 \leq r < |b|$, where $|b|$ is the absolute value of b .*

Remark 3.2. Often, some intuition is useful in order to construct a formal proof. For convenience in this motivation, assume that $a, b > 0$, so a lies in the interval $[0, \infty)$. We want to partition that large interval into smaller ones of length b :

$$[0, \infty) = [0, b) \cup [b, 2b) \cup [2b, 3b) \cup \dots,$$

or more succinctly, $[qb, (q+1)b) = [qb, qb+b)$ for $q = 0, 1, 2, \dots$. Now since those intervals are disjoint, a lies in exactly one of those subintervals $[qb, qb+b)$, so $a = qb + r$ where $0 \leq r < b$. \square

In part, one of the goals of this text is to introduce proof-writing to the reader, and number theory is a wonderful subject through which to achieve this goal. Writing a proof demands great attention to detail; each assertion must be justified. As one's exposure to mathematics increases, certain statements can be taken more easily on faith (meaning you're confident you could produce a proof of the assertion), but in the beginning we strive to be pedantic so as to reveal all the assertions which need justification, though admittedly at the expense that the proof below is somewhat long winded. As we move forward through the text, more sophistication will gradually be assumed, and proofs will become more streamlined.

Another important point is that there is not necessarily one correct proof of a result. Different proofs often contain different, but equally interesting, ideas.

Exercise. In contrast to the proof we give below, which has certain pedagogical motivations, construct a proof using the idea in the remark above. For example, consider the intervals of the form $[q|b|, (q+1)|b|)$ where $|b|$ is the absolute value of b and q ranges over all the integers.

To move from our intuition to a rigorous proof, we rely on a fact of enormous importance, that by construction, the natural numbers are *well-ordered*, meaning that every nonempty subset of \mathbb{N} contains a least element. This is a fundamental fact in number theory, equivalent to the notion of mathematical induction, and it is essential to the proof below.

Proof. In proving a statement like “there exist unique integers ...”, there are actually two statements to prove: integers exist with the stated properties and only those integers have the stated property.

First we show *existence*, that is, there exist a quotient q and a remainder r with $a = bq + r$ and $0 \leq r < |b|$. Note that if we solve the desired equation $a = bq + r$ for r , we have an expression of the form $r = a - bq$. Since we are looking for a value of r with a particular property, this suggests we consider the set of all possible values of r : $S = \{a - nb \mid n \in \mathbb{Z}\} = \{a, a \pm b, a \pm 2b, \dots\}$. Now our goal is to find the smallest nonnegative element of S and show that it has the requisite properties. To force the issue of nonnegativity, we consider the subset $T = S \cap \mathbb{N}$. Certainly T is a subset of natural numbers, so if it is a nonempty set, well-ordering demands it have a smallest element. That will be our remainder.

We check that $a - nb \geq 0$ if and only if $n \leq a/b$ (if $b > 0$) or $n \geq a/b$ (if $b < 0$), and since there are infinitely many integers that can satisfy either inequality, we have that T is nonempty. By well-ordering, T has a least element $r = a - qb$ for some integer q . All we need to verify is that $r < |b|$.

We prove this *by contradiction*. If r is not less than $|b|$, then $r \geq |b|$, so consider $r' = r - |b| \geq 0$. Since $|b| > 0$, we see $r' = r - |b| < r$, so putting the conclusions together, we have $0 \leq r' < r$. But

$$r' = r - |b| = a - bq - |b| = \begin{cases} a - b(q + 1) & \text{if } b > 0 \\ a - b(q - 1) & \text{if } b < 0 \end{cases} \in S \cap \mathbb{N} = T.$$

Since this shows r' is strictly smaller than r and is an element of T , it contradicts that r was chosen to be the smallest element of T . Thus our assumption ($r \geq |b|$) was false, so we have $a = bq + r$ with $0 \leq r < |b|$.

As to *uniqueness*, suppose that $a = bq + r = bq' + r'$ with $0 \leq r, r' < |b|$, that is, we have potentially two different quotients and/or two different remainders. Subtracting and rearranging the expressions for a , we derive that $b(q - q') = r' - r$. Now without loss of generality we may assume $r \leq r'$ (otherwise, write the equation as $b(q' - q) = r - r'$), so $0 \leq r \leq r' < |b|$, which means their difference $r' - r$ satisfies $0 \leq r' - r < |b|$, but it is also an integer multiple of b . The

only possibility is the multiple is zero. Thus $r = r'$, and then since $b(q - q') = 0$ (and $b \neq 0$), we have $q = q'$ as well. \square

Remark 3.3. We now comment about an alternate way in which to think about divisibility. Above we said that $b \mid a$ if and only if $a/b = q \in \mathbb{Z}$. The division algorithm gives that $a = bq + r$ with $0 \leq r < |b|$, so $a/b = q + r/b \in \mathbb{Z}$ if and only if $r = 0$.

Exercise.

- Show that $3 \mid 0$, but $0 \nmid 3$.
- Show that if $a \mid b$ and $b \mid c$, then $a \mid c$.
- Show that if $a \mid b$ and $c \mid d$, then $ac \mid bd$.
- Show that if $m \neq 0$, then $a \mid b$ if and only if $am \mid bm$.

Next we state a couple of extremely useful properties of divisibility.

Proposition 3.4. *Let $a, b \in \mathbb{Z}$, $b \neq 0$.*

- (1) *If a_1, \dots, a_r are integers each of which is divisible by b (i.e., $b \mid a_i$ for all i), then $b \mid (m_1a_1 + \dots + m_ra_r)$ for any integers m_i .*
- (2) *Assume also that $a \neq 0$. If $a \mid b$ and $b \mid a$, then $a = \pm b$.*

Proof. For the first statement, if $b \mid a_i$ for all $i = 1, \dots, r$, then there exist $c_i \in \mathbb{Z}$ with $a_i = bc_i$, $i = 1, \dots, r$. Thus,

$$m_1a_1 + \dots + m_ra_r = m_1bc_1 + \dots + m_rc_r = b(m_1c_1 + \dots + m_rc_r),$$

which shows $b \mid (m_1a_1 + \dots + m_ra_r)$.

For the second statement, if $a \mid b$, then $b = ac$ for some integer c . If $b \mid a$, then $a = bd$ for an integer d . Thus,

$$b = ac = bdc, \quad \text{or} \quad b - bdc = 0 = b(1 - dc).$$

Now in the integers, the product of two integers $mn = 0$ if and only if (at least) one of m or $n = 0$. We have $b(1 - dc) = 0$ and $b \neq 0$, so we must have $1 - dc = 0$, that is $dc = 1$, the only solution with $c, d \in \mathbb{Z}$ is $c = d = \pm 1$. Returning to our expressions for a (or b), we see $a = \pm b$. \square

Definition 3.5. If $d \mid a$ and $d \mid b$, we say d is a *common divisor* of a and b . If $d = \pm 1$ are the only common divisors of a and b , we say that a, b are *relatively prime* or *coprime*.

Remark 3.6. Often in elementary number theory we focus more on the positive integers, so we would say 1, 2, 3, 6, and 12 are the common divisors of 24, 36, when indeed the more correct statement would be that $\pm 1, \pm 2, \pm 3, \pm 4, \pm 6$, and ± 12 are the common divisors of 24 and 36; the integers 24, 25 are relatively prime.

Definition 3.7. Let $a, b \in \mathbb{Z}$, not both zero. The *greatest common divisor* of a, b is the unique positive integer d so that

- (1) $d \mid a$ and $d \mid b$ (i.e., d is a common divisor), and
- (2) if $c \mid a$ and $c \mid b$, then $c \mid d$ (so d is greatest among positive divisors).

We denote this as $d = \gcd(a, b)$ or simply $d = (a, b)$ if the context is clear.

Remark 3.8. Two comments are in order. First, in many textbooks on elementary number theory, the second condition is replaced with

- (2') if $c \mid a$ and $c \mid b$, then $c \leq d$.

It is easy to show that in \mathbb{Z} , these two conditions are equivalent, but the one we have chosen to use as part of our definition generalizes more naturally to other settings such as when polynomials have a common divisor.

The second comment is that we should perhaps justify the definition. After all, it asserts that there exists such a unique positive integer. Of course a and b have common positive divisors, namely $d = 1$, and if d is positive and $d \mid a$, then $d \leq |a|$, where $|a|$ is the absolute value of a . Since there are only a finite number of integers between 1 and $|a|$, there clearly must be a greatest common divisor, which is necessarily uniquely determined by being the largest integer in a finite list.

Below we make a few easy observations about gcds, the proofs of which we leave as an exercise.

Proposition 3.9. *Suppose that $a, b \in \mathbb{Z}$, not both zero. Then*

- $\gcd(\pm a, \pm b) = \gcd(|a|, |b|)$.
- If $a \neq 0$, the $\gcd(a, 0) = |a|$.
- If $a \neq 0$, then $\gcd(a, a) = |a|$.
- If $b \mid a$, then $\gcd(a, b) = |b|$.

Example. $\gcd(24, -36) = 12$; $\gcd(24, 25) = 1$; $\gcd(0, -15) = 15$.

3.3. Euclid's Algorithm

For small integers, we tend to rely on our ability to factor integers in order to compute a greatest common divisor. For large numbers, factoring is impractical (indeed we shall see that the security of RSA encryption depends upon that assumption), so we need to rely on a more computationally feasible means of extracting the gcd. The method is called Euclid's algorithm, and it is based on the division algorithm we have already established. We begin with an easy but pivotal lemma.

Lemma 3.10. *If a, b are integers with $b \neq 0$, and we write $a = bq + r$, then $\gcd(a, b) = \gcd(b, r)$.*

Proof. We claim that the set of common divisors of a and b are the same as the set of common divisors of b and r . Given this claim, the greatest among these is necessarily the same. To establish the claim we need only show that every common divisor of a, b is a common divisor of b, r , and conversely.

If $d \mid a$ and $d \mid b$, then by Proposition 3.4, d divides any linear combination of a and b , namely $d \mid r = a(1) - bq$. Conversely, if $d \mid b$ and r , then by the same reasoning $d \mid a = bq + r$. \square

Our goal is to compute the gcd of two integers a, b , where at least one of them is nonzero. Proposition 3.9 handles the degenerate case where one of the integers is zero and gives also $\gcd(\pm a, \pm b) = \gcd(|a|, |b|)$, so we may assume without loss of generality that

$a \geq b > 0$. To find their gcd, we iterate the division algorithm as follows:

$$\begin{array}{ll}
 a = bq_1 + r_1, & \text{with } 0 \leq r_1 < b, \\
 b = r_1q_2 + r_2, & \text{with } 0 \leq r_2 < r_1, \\
 r_1 = r_2q_3 + r_3, & \text{with } 0 \leq r_3 < r_2, \\
 \vdots & \vdots \\
 r_{n-3} = r_{n-2}q_{n-1} + r_{n-1}, & \text{with } 0 \leq r_{n-1} < r_{n-2}, \\
 r_{n-2} = r_{n-1}q_n + r_n, & \text{with } r_n = 0.
 \end{array}$$

Note that $0 \leq r_n < r_{n-1} < \cdots < r_1 < b$ is a strictly decreasing sequence of nonnegative integers, so the algorithm must terminate in fewer than b steps. Actually, by an 1844 result of Gabriel Lamé, the algorithm will terminate in no more than $5 \log_{10}(b)$ steps, which in the language of computational complexity is linear in the length of the smaller input.

The point of Euclid's algorithm is that

Theorem 3.11. *With the notation as above, $\gcd(a, b) = r_{n-1}$, that is the last nonzero remainder in Euclid's algorithm.*

Proof. By the lemma, $\gcd(a, b) = \gcd(b, r_1) = \gcd(r_1, r_2) = \cdots = \gcd(r_{n-2}, r_{n-1}) = \gcd(r_{n-1}, r_n) = \gcd(r_{n-1}, 0) = r_{n-1}$. \square

Example 3.12. Let's compute the gcd of 252 and 198:

$$\begin{aligned}
 252 &= 198(1) + 54, \\
 198 &= 54(3) + 36, \\
 54 &= 36(1) + 18, \\
 36 &= 18(2) + 0.
 \end{aligned}$$

So the $\gcd(252, 198) = 18$, the last nonzero remainder.

Now we come to a major application of Euclid's algorithm, called Bézout's identity (not to be confused with Bézout's theorem discussed in the previous chapter).

Theorem 3.13 (Bézout's identity). *Let a, b be integers, not both zero. Then there exist $u, v \in \mathbb{Z}$ so that $\gcd(a, b) = au + bv$.*

Proof. Note that there is no loss of generality assuming both a and b are nonnegative, for say $a < 0$ and $b \geq 0$. We have $\gcd(a, b) = \gcd(|a|, b) = \gcd(-a, b) = -a(u) + bv = a(-u) + bv$, as desired. Other cases are similar.

The proof is simply to realize that we can run Euclid's algorithm backwards starting with the gcd which equals r_{n-1} and back substitute until we have a combination of a and b . \square

We illustrate this with the example computed above: $\gcd(252, 198) = 18$:

$$\begin{aligned} 18 &= 54 - 36(1) \\ &= 54 - (1)(198 - 54(3)) = 198(-1) + 54(4) \\ &= 198(-1) + 4(252 - 198(1)) = 198(-5) + 252(4). \end{aligned}$$

Via Bézout's identity, we obtain two hugely useful corollaries.

Corollary 3.14. *Let a, b be nonzero integers. Then*

- (1) *If $a \mid c$, $b \mid c$, and $\gcd(a, b) = 1$, then $ab \mid c$.*
- (2) *If $a \mid bc$ and $\gcd(a, b) = 1$, then $a \mid c$.*

Proof. It is always a good idea to see whether hypotheses are necessary. For each statement, we see that the condition $\gcd(a, b) = 1$ is essential. For example, in the first statement if $a = 4$, $b = 6$, and $c = 12$, we see $a \mid c$ and $b \mid c$, but $ab \nmid c$. Similarly, for the second, if $a = 6$, $b = 3$, and $c = 4$, we have that $a \mid bc$, but clearly $a \nmid b$ and $a \nmid c$.

So we begin with the necessary assumption that $\gcd(a, b) = 1$. Bézout's identity says there are integers u, v so that $au + bv = 1$. For the first assertion, we can write $c = am = bn$ for some integers m, n . Then

$$c = c \cdot 1 = c(au + bv) = cau + cbv = (bn)au + (am)bv = ab(nu + mv),$$

showing that $ab \mid c$.

Similarly, for the second assertion, we write

$$c = c \cdot 1 = c(au + bv) = cau + cbv.$$

We note that $a \mid a$ and, by hypothesis, $a \mid bc$, so, by Proposition 3.4, a divides the linear combination $cau + cbv = c$. \square

With these results in hand, we are in a position to give a proof of the Fundamental Theorem of Arithmetic. While we have used the word prime before, it is time to be precise in its definition. An integer $n > 1$ is either prime or composite: It is *prime* if its only positive divisors are 1 and n . It is *composite* if it is not prime, and therefore we infer that an integer $n > 1$ is composite means that it can be factored, $n = ab$ where $1 < a, b < n$. First a useful exercise.

Exercise. Let $p > 1$ be a prime. Show that

- For any integer n , $\gcd(p, n) = 1$ or p .
- For integers m, n if $p \mid mn$, then either $p \mid m$ or $p \mid n$.

Theorem 3.15 (Fundamental Theorem of Arithmetic). *Every integer $n > 1$ can be factored uniquely as a product of primes.*

Proof. First we prove the existence part of the statement, that every integer $n > 1$ can be written as the product of primes. The intent of this statement is that a prime is to be thought of as the product of one prime. We prove this by induction on n . For the base case, we observe that $n = 2$ is prime. So now we assume $n > 2$ and, for all integers m with $1 < m < n$, that m can be written as the product of primes. Consider n . If n is prime, we are done. If not, it is composite, and so it can be written as $n = ab$ with $1 < a, b < n$. By induction, each of a, b can be written as a product of primes, so concatenating the two products gives us a representation of n as the product of primes. So induction gives us existence.

Now for uniqueness, and this takes some explanation. After all, $6 = 2 \cdot 3 = 3 \cdot 2$, but we don't want to call these two different factorizations since multiplication in the integers is commutative. So the statement we shall prove is that whenever an integer $n > 1$ is written in the product of primes (in two possibly different ways) as $n = p_1 p_2 \cdots p_r = q_1 q_2 \cdots q_s$, with p_i and q_j all primes (not necessarily

distinct), that $r = s$ (the number of primes is the same) and, that after some reordering, $p_i = q_i$ for $i = 1, \dots, r$.

In any such factorizations, we may assume without loss of generality that $r \leq s$, and the proof of uniqueness will be by induction on r . The base case is $r = 1$, that is $n = p_1 = q_1 \cdots q_s$. This says $n = p_1$ is prime, so by definition the only way to factor it is $p_1 \cdot 1$, so we may conclude $s = 1$ and $q_1 = p_1$. Now for the induction hypothesis: Assume that $r > 1$ and, for any integer m which admits a factorization into t primes with $t < r$, that such a factorization is unique in the sense specified above. So we consider $n = p_1 \cdots p_r = q_1 \cdots q_s$. Using the exercise above, we see that since $p_1 \mid n = q_1 \cdots q_s$, we must have that $p_1 \mid q_i$ for some prime q_i . But that says that $q_i = p_1 \cdot a$ for some integer a . But as q_i is a prime, we may conclude that $a = 1$, so $p_1 = q_i$. Now reorder the primes q_j making $q_1 = p_1$, so $n = p_1 \cdots p_r = p_1 q_2 \cdots q_s$. Canceling the p_1 's from both factorizations leaves us with $n' = n/p_1 = p_2 \cdots p_r = q_2 \cdots q_s$. Visibly, n' can be written as the product of $r - 1$ primes, so by induction, we know that factorization is unique, meaning $r - 1 = s - 1$ (so $r = s$), and after reordering $q_j = p_j$ for $j = 2, \dots, r$. \square

3.4. A First Pass at Modular Arithmetic

As another application of divisibility, we define the notion of modular arithmetic as an operation on the integers. Later we shall reinterpret modular arithmetic as algebraic operations on a finite set, denoted \mathbb{Z}_n , which we shall see gives the set the structure of what is called a ring.

Let n be a positive integer, and let a, b be arbitrary integers. We want to define what it means to say that a and b are *congruent modulo* n . First we shall give a definition that is fairly intuitive, and then we shall show this intuitive notion is equivalent to a divisibility condition which is easier to check in practice.

Loosely speaking, we will say that a and b are congruent modulo n , if they have the same remainder when divided by n . More precisely, dividing a and b by n via the division algorithm gives expressions $a = nq + r$ and $b = nq' + r'$ with unique remainders $0 \leq r, r' < n$.

Definition 3.16. We say that a is *congruent to b modulo n* , written $a \equiv b \pmod{n}$ or $a \equiv b(n)$ if and only if $r = r'$.

Let's make a couple of simple observations about our definition. Fix the positive integer n . Let $a \in \mathbb{Z}$, and write $a = nq + r$ with $0 \leq r < n$.

First observe that $a \equiv r \pmod{n}$. To see this, we note that $r = n \cdot 0 + r$ with $0 \leq r < n$ so, by the proof of the division algorithm, r is the unique remainder, so a and r have the same remainder. What does this mean?

Proposition 3.17. *Every integer is congruent modulo n to exactly one of the integers $0, 1, 2, \dots, n - 1$.*

Proof. The argument above shows that an integer is congruent modulo n to its unique remainder given via the division algorithm. This says that an integer is congruent to one and to only one of $0, 1, \dots, n - 1$. \square

To state the last sentences of the proof in another way, this also says that given two remainders $0 \leq r, r' < n$, we have $r \equiv r' \pmod{n}$ if and only if $r = r'$.

Example 3.18. Every integer is congruent modulo 2 to either 0 or 1, that is dividing by 2 leaves a remainder of 0 or 1. Those with remainder 0 are called *even*, and those with remainder 1 are called *odd*. So when we say $a \equiv b \pmod{2}$, we know that a and b are either both even or odd, and are said to have the same *parity*.

The next proposition gives a number of equivalent ways of thinking about congruence.

Proposition 3.19. *Let $a, b, n \in \mathbb{Z}$ with $n > 0$. The following conditions are equivalent:*

- (1) $a \equiv b \pmod{n}$.
- (2) $n \mid (a - b)$.
- (3) $a = b + kn$ for some integer k .

Proof. To show these statements are equivalent, one must show each implies the other, and the converse. Technically, that is six statements to prove, but one can do it efficiently with three by proving that $(1) \implies (2) \implies (3) \implies (1)$, which is what we shall do.

(1) \implies (2): If $a \equiv b \pmod{n}$, then via the division algorithm, $a = nq + r$ and $b = nq' + r$ with the same remainder r . But then $a - b = n(q - q')$, which implies $n \mid (a - b)$.

(2) \implies (3): If $n \mid (a - b)$, then $a - b = kn$ for some integer k , which in turn means $a = b + kn$.

(3) \implies (1): Let $a = b + kn$. Via the division algorithm write $b = nq + r$ with $0 \leq r < n$. Then $a = b + kn = n(q + k) + r$. Since $0 \leq r < n$, this is the unique representation of a via the division algorithm. We see both expressions have the same remainder r , so $a \equiv b \pmod{n}$. \square

Below, we give a few examples of the uses of congruences, but our lives will be a great deal easier if we first understand a few elementary properties concerning congruence. In the next chapter, these properties will be interpreted in a much broader context. For congruences, not only are these properties elementary, they are so elementary, they probably would not occur to you that you are using them, so we will try to point out why these properties are crucial to supporting computations with congruences.

The first of these properties is that congruence modulo n satisfies three properties which will characterize it as an equivalence relation on \mathbb{Z} , a notion we will explore more in the next chapter.

Proposition 3.20. *Let n be a positive integer, and let $a, b, c \in \mathbb{Z}$. Then*

- (1) *Reflexive.* $a \equiv a \pmod{n}$.
- (2) *Symmetric.* $a \equiv b \pmod{n}$ if and only if $b \equiv a \pmod{n}$.
- (3) *Transitive.* If $a \equiv b \pmod{n}$ and $b \equiv c \pmod{n}$, then $a \equiv c \pmod{n}$.

The reflexive property seems almost silly, but it is not. If instead of congruence we were talking about the relation “less than” ($<$),

then it would not be true that $a < a$. The symmetric property allows us not to worry about whether we write $a \equiv b \pmod{n}$ or $b \equiv a \pmod{n}$. It too is not silly, for if the relation were the subset relation (\subseteq) and we had two subsets for which $A \subseteq B$, it would not necessarily be the case that $B \subseteq A$: the integers are a subset of the real numbers, but not conversely. The transitive law is simply handy. What it says is that you can do work in stages, first reducing a to b , then perhaps by a different observation from b to c , and in the end you know that it was possible to go directly from a to c . We shall see this in play in the examples below.

There are many proofs one can give of the above using the equivalences in Proposition 3.19.

Proof. We will use $c \equiv d \pmod{n}$ if and only if $c = d + kn$ for some integer k .

- (1) $a \equiv a \pmod{n}$ since $a = a + 0 \cdot n$.
- (2) If $a \equiv b \pmod{n}$, then $a = b + kn$. But then $b = a + (-k)n$, so $b \equiv a \pmod{n}$, and the converse by symmetry.
- (3) Given $a \equiv b \pmod{n}$ and $b \equiv c \pmod{n}$, we can write $a = b + kn$ and $b = c + \ell n$, so that

$$a = b + kn = c + \ell n + kn = c + (k + \ell)n,$$

hence $a \equiv c \pmod{n}$. □

We also need the following properties. These too may appear at first to be “symbol pushing”, but they are actually key to being able to do computations efficiently and underlie important applications later on.

Proposition 3.21. *Let $a, b, a', b' \in \mathbb{Z}$, let n be a positive integer with $a \equiv a' \pmod{n}$, and let $b \equiv b' \pmod{n}$. Then*

- (1) $a \pm b \equiv a' \pm b' \pmod{n}$,
- (2) $ab \equiv a'b' \pmod{n}$.

Proof. $a \equiv a' \pmod{n}$ and $b \equiv b' \pmod{n}$ means that $a = a' + kn$ and $b = b' + \ell n$ for some integers k and ℓ . Thus,

$$\begin{aligned} a \pm b &= a' \pm b' + n(k \pm \ell), \\ ab &= a'b' + n(kb' + \ell a' + k\ell n). \end{aligned}$$

Rewriting these equalities as congruences yields the result. \square

Let's put these properties to work.

Example 3.22. Find the last decimal digit of $1! + 2! + \cdots + 100!$. Oh my! This seems like a daunting amount of work. What if you knew that the last decimal digit of $1! + 2! + \cdots + 100!$ is the same as the last decimal digit of $1! + 2! + \cdots + 100! + \cdots + n!$ for any $n > 100$? That would actually be a big clue. Why? Well, it would say that adding those larger numbers is not changing the last digit of our number.

So how do we recognize the last digit of a number? It seems like a silly question—we just look at the number. But that's not terribly helpful. So let's ask again: how do we know the last decimal digit of 12345 is a 5, and what does that mean? The meaning is hidden in the word “decimal”, meaning that 12345 is the base 10 expansion of a number. Said another way, it tells us that

$$\begin{aligned} 12345 &= 1 \cdot 10^4 + 2 \cdot 10^3 + 3 \cdot 10^2 + 4 \cdot 10^1 + 5 \cdot 10^0 \\ &= 1234 \cdot 10 + 5, \end{aligned}$$

but the later expression tells us that $12345 \equiv 5 \pmod{10}$. To be even more precise, given the base 10 expansion of an integer, the last decimal digit is the remainder given by the division algorithm upon dividing by 10.

Now let's return to the problem of finding the last decimal digit of

$$1! + 2! + \cdots + 100!.$$

We note that for $n \geq 5$, $n!$ is a multiple of 10 (since both 2 and 5 divide $n!$ and they are coprime), that is $n! \equiv 0 \pmod{10}$ for each $n \geq 5$. Thus

$$1! + 2! + \cdots + 100! \equiv 1! + 2! + 3! + 4! + 0 + \cdots + 0 \pmod{10},$$

where we have just used Proposition 3.21(1). Thus,

$$1! + 2! + \cdots + 100! \equiv 1! + 2! + 3! + 4! \equiv 33 \pmod{10}$$

$$\text{and } 33 \equiv 3 \pmod{10}, \text{ so}$$

$$1! + 2! + \cdots + 100! \equiv 3 \pmod{10},$$

where we have used the transitive property of congruence. It follows that the last digit is 3. \square

In the next example, we use congruences in a different way. Sometimes we are not sure whether a statement is true or false. Of course if you can prove it, we know it is true. On the other hand, if you can find no proof, there is little you can conclude. But in order to show that some statement is false, it is often enough to show that if it were true, an impossible situation would ensue.

Example 3.23. Is 12345678 a square in \mathbb{Z} ? We could settle this question quickly with a calculator, but where's the fun in that? Besides, we could as easily have asked if

825949768513252587956123427457683521546788956431258612345678

is a square, but that's too long to type. The technique we have in mind works as easily with one of those integers as the other, and indeed we would have the answer before you even typed that number into your favorite device.

So let's work with 12345678. If it is a square in \mathbb{Z} , then $12345678 = m^2$ for an integer m . We would like to see that this is impossible without trying to find m . What comes next may seem unintuitive, but new tools require learning how to use them, so let's take a small excursion.

We know that integers are either odd or even, that is of the form $m = 2n$ or $m = 2n + 1$, so their squares $m^2 = 4n^2$ or $4n^2 + 4n + 1$ satisfy $m^2 \equiv 0 \pmod{4}$ or $m^2 \equiv 1 \pmod{4}$. Did you notice that we used Proposition 3.21 once again?

Now by the division algorithm, we know an arbitrary integer is congruent to exactly one of 0, 1, 2, or 3 $\pmod{4}$, so if n is an integer which is congruent to 2 or 3 $\pmod{4}$, then it is impossible for it to be a square in \mathbb{Z} . So all we need to do is compute $12345678 \pmod{4}$.

Can we do this easily? Sure, and we won't even use the division algorithm.

The base 10 expansion of our number tells us that

$$12345678 = 123456(100) + 78 = 123456(25)(4) + 78 \equiv 78 \pmod{4}.$$

And direct computation tells us that $78 = 4(19) + 2$, so $78 \equiv 2 \pmod{4}$, and by transitivity, $12345678 \equiv 2 \pmod{4}$, so it cannot be a square in \mathbb{Z} . Note that the larger number above is also (trivially) congruent to $2 \pmod{4}$ (why?), so it cannot be a square either.

Finally, note that this observation only works in one direction. For example, $5 \equiv 1 \pmod{4}$ but 5 is not a square in \mathbb{Z} . \square

For a final example, we return to a matter from Chapter 2.

Example 3.24. When we talked about Pythagorean triples (positive integers A, B, C with $A^2 + B^2 = C^2$), we asserted and rather awkwardly justified that A and B could not both be odd. Now we easily see why. A square integer is congruent to 0 or 1 modulo 4, so $C^2 \equiv 0, 1 \pmod{4}$. If A and B are both odd, then $A^2 + B^2 \equiv 2 \pmod{4}$, so there can be no equality $A^2 + B^2 = C^2$ since $A^2 + B^2 \not\equiv C^2 \pmod{4}$.

Exercise. Can you find integers x, y, z so that $987654319 = x^2 + y^2 + z^2$? *Hint:* Determine the possible values of $x^2 + y^2 + z^2 \pmod{8}$.

Exercise (A precursor to the Chinese Remainder Theorem). Find the smallest number of marbles in a jar so that one remains if taken out 2, 3, 5 at a time, but none remain if taken out 11 at a time.

Exercise. To get more of a feel for congruences and how to move between congruences and equalities, consider the following exercises:

- Show (by example) that the congruence $ax \equiv ay \pmod{n}$ does not necessarily imply that $x \equiv y \pmod{n}$.
- On the other hand, show that if $x \equiv y \pmod{n}$, then $ax \equiv ay \pmod{n}$ for any integer a .
- Show that there exist integers u, v so that $au + nv = b$ if and only if $ax \equiv b \pmod{n}$ is solvable.

3.5. Elementary Cryptography: Caesar Cipher

An example of congruences having historical significance is a mathematical characterization of what is called the Caesar cipher. We also discuss this example as a way of introducing some basic terminology in cryptography, which we shall take up in more depth in a later chapter.

In brief, to send messages to his commanders on the battlefield, Caesar needed a means to write a message and have a courier carry it to his commanders so that only the commander could understand the message. In particular, if the courier was intercepted and the message read by a third party, they would be unable to understand the message. To accomplish this, Caesar encrypted them. He would generate his intended message; this is called a *plaintext* message. He would then encrypt it (as described below), producing what is called a *ciphertext* message. The ciphertext message was then sent by courier to his commanders and then decrypted. The intent, of course, was that even if the courier was intercepted and the message read, the real (plaintext) message could not be recovered.

The process of encryption was to write out the message, and then shift each letter in the message forward by three, that is $A \mapsto D$, $B \mapsto E$, and so on. The complete lexicon is given by:

Plaintext:	<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>E</i>	<i>F</i>	<i>G</i>	<i>H</i>	<i>I</i>	<i>J</i>	<i>K</i>	<i>L</i>	<i>M</i>
	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
Ciphertext:	<i>D</i>	<i>E</i>	<i>F</i>	<i>G</i>	<i>H</i>	<i>I</i>	<i>J</i>	<i>K</i>	<i>L</i>	<i>M</i>	<i>N</i>	<i>O</i>	<i>P</i>

Plaintext:	<i>N</i>	<i>O</i>	<i>P</i>	<i>Q</i>	<i>R</i>	<i>S</i>	<i>T</i>	<i>U</i>	<i>V</i>	<i>W</i>	<i>X</i>	<i>Y</i>	<i>Z</i>
	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
Ciphertext:	<i>Q</i>	<i>R</i>	<i>S</i>	<i>T</i>	<i>U</i>	<i>V</i>	<i>W</i>	<i>X</i>	<i>Y</i>	<i>Z</i>	<i>A</i>	<i>B</i>	<i>C</i>

Thus the word CAT in plaintext would translate to FDW in ciphertext. Upon receipt of the ciphertext, the commanders would decrypt it by shifting each letter back by three.

Mathematically, we achieve this by converting each letter of the alphabet into a numerical equivalent:

<i>A</i>	<i>B</i>	<i>C</i>	<i>D</i>	<i>E</i>	<i>F</i>	<i>G</i>	<i>H</i>	<i>I</i>	<i>J</i>	<i>K</i>	<i>L</i>	<i>M</i>
↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
0	1	2	3	4	5	6	7	8	9	10	11	12
<i>N</i>	<i>O</i>	<i>P</i>	<i>Q</i>	<i>R</i>	<i>S</i>	<i>T</i>	<i>U</i>	<i>V</i>	<i>W</i>	<i>X</i>	<i>Y</i>	<i>Z</i>
↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓	↓
13	14	15	16	17	18	19	20	21	22	23	24	25

So C-A-T would be encoded 2-0-19; the ciphertext F-D-W would be 5-3-22. So if P is the numeric equivalent of a plaintext letter, the corresponding ciphertext letter would be determined by the congruence

$$C \equiv P + 3 \pmod{26}.$$

The plaintext would be recovered from the ciphertext by

$$P \equiv C - 3 \equiv C + 23 \pmod{26}$$

where in all cases one would take the values for P and C from among 0 – 25.

The Caesar cipher is an example of a *shift cipher*, that is any cipher where each character of plaintext is converted to ciphertext by a procedure of the form $C \equiv P + k \pmod{n}$. The set of values of k which produce distinct (nontrivial) ciphers is called the *keyspace* associated to the encryption scheme; clearly, it has size $n - 1$ ($k = 0$ is a trivial encryption). The Caesar cipher is a shift cipher where $n = 26$, but n could be as large as needed to accommodate a larger alphabet, upper and lower case letters, as well as special symbols.

While this primitive cryptographic scheme may have worked adequately for Caesar, any person trying to break this cryptosystem would have little challenge, especially if they suspected the nature of the encryption scheme. And this, perhaps surprisingly, is actually a very important issue.

In designing a strong cryptographic system, how important do you think it is to keep the nature of the method of encryption and decryption a secret? After all, if you knew someone was using a shift cipher, you could try all $n - 1$ nontrivial keys in a matter of

seconds on a computer, so it might seem if you did not know the nature of the cipher, it might be more secure. Indeed that was the prevailing wisdom—even for the government and military—for a long time. Here’s the rub: as long as an adversary does not know the nature of the method, there actually may be some additional security, but if the adversary learns the nature of the encryption and this knowledge is unknown to those sending encrypted messages, security can be dramatically reduced. A prominent example was the cracking of the Enigma machine in World War II, something to which entire books have been devoted. While a fascinating story, it would take us too far afield, so we shall leave it to the interested party to explore on his or her own.

The danger of hiding the method of encryption and decryption is a lesson now well learned, and for new cryptographic schemes it is absolutely necessary that there be complete transparency in terms of how the encryption and decryption algorithms work. For government and commercial use, this is not only necessary for the reason suggested above, but also to ensure the algorithms do not contain a “back door”, which would allow the designers an easy way to break the cipher. This particular issue is one which confronted the Apple corporation in February 2016 with regard to encryption on its iPhones. We shall learn more about these issues in later chapters.

To advance a broader discussion of cryptography, we introduce a little notation and terminology. A basic cryptographic setup has two functions, E and D , representing an encryption and decryption scheme. If P denotes a plaintext message and C the corresponding ciphertext (encrypted message), then the requirements for a cryptosystem are pretty basic.

We take plaintext P and use E to encrypt the message, producing ciphertext $C = E(P)$. Ideally, it is very difficult to discover P from C . The other essential feature is that the decryption scheme must work, that is $P = D(C) = D(E(P))$, that is D is a left-hand inverse to E .

For the Caesar cipher we had E and D defined as

$$\begin{aligned} C &= E(P) \equiv P + 3 \pmod{26}, \\ P &= D(C) \equiv C - 3 \pmod{26} \equiv C + 23 \pmod{26}. \end{aligned}$$

Exercise. Explore an encryption scheme known as ROT13; it is a shift cipher. What can you say about the encryption and decryption functions E and D ?

3.6. Affine Ciphers and Linear Congruences

Having dismissed the shift cipher as too simple, perhaps we can tweak it and produce a slightly more complicated scheme, called an *affine* cipher. Here we take $C \equiv aP + b \pmod{26}$ for integers a and b . Clearly, $a = 1$ recovers our shift cipher, but when $a \neq 1$, we must ask when is the congruence $C \equiv aP + b \pmod{26}$ uniquely solvable for $P \pmod{26}$, that is, when can we find a decryption algorithm to accompany this encryption scheme?

Exercise. Before reading the propositions that follow, can you determine the values of a for which a decryption algorithm to $C \equiv aP + b \pmod{26}$ be produced?

We summarize what we have learned from the exercise. We begin with a simple case.

Proposition 3.25. *The congruence $ax \equiv 1 \pmod{n}$ is solvable if and only if $\gcd(a, n) = 1$ and, when solvable, there is a unique solution modulo n .*

Proof. If $ax \equiv 1 \pmod{n}$ is solvable, then there exists a $y \in \mathbb{Z}$ with $ax + ny = 1$. Let d be the gcd of a and n . Then, by definition of a gcd, $d \mid a$ and $d \mid n$, so we know that d divides any combination of a and n ; in particular $d \mid (ax + ny) = 1$, so d is clearly 1.

Conversely, if $d = \gcd(a, n) = 1$, then Bézout's identity tells us that there exists $u, v \in \mathbb{Z}$ with $au + nv = d = 1$, but this means $au \equiv 1 \pmod{n}$, so the congruence $ax \equiv 1 \pmod{n}$ is solvable with solution $x \equiv u \pmod{n}$.

To see that there is a unique solution modulo n , suppose there were two, x, y . Then $ax \equiv ay \equiv 1 \pmod{n}$. Knowing $au + nv = 1$, we multiply both sides of the congruence by u (see the exercise before §3.5), and we see $uax \equiv uay \pmod{n}$. But $au \equiv 1 \pmod{n}$, so we deduce $x \equiv y \pmod{n}$, as required. \square

Implicit in the proof above is the following result, allowing us to recognize when two integers are relatively prime.

Corollary 3.26. *Let $a, b \in \mathbb{Z}$. Then $\gcd(a, b) = 1$ if and only if there exist $u, v \in \mathbb{Z}$ with $au + bv = 1$.*

Now we proceed to handle general linear congruences.

Proposition 3.27. *If the congruence $ax \equiv b \pmod{n}$ is solvable, then $d = \gcd(a, n) \mid b$.*

Note that the contrapositive is more instructive: If $d \nmid b$, then the congruence is not solvable.

Proof. The proof is similar to the one above. If $ax \equiv b \pmod{n}$ is solvable, then there exist integers x, y so that $ax + ny = b$. Since $d = \gcd(a, n)$, we know $d \mid a$ and $d \mid n$, so d divides any combination of a and n , in particular, $d \mid b$. \square

The converse is where the substance lies.

Theorem 3.28. *Let $d = \gcd(a, n)$. If $d \mid b$, then the congruence $ax \equiv b \pmod{n}$ is solvable, and there are precisely d incongruent solutions modulo n . Indeed they are all of the form $x_0 + \frac{n}{d}t \pmod{n}$ for $t = 0, 1, \dots, d - 1$ where x_0 is any particular solution.*

Before proving the theorem, we sum up the results so far:

Corollary 3.29. *The congruence $ax \equiv b \pmod{n}$ is solvable if and only if $d = \gcd(a, n) \mid b$. When solvable, there are precisely d incongruent solutions modulo n .*

Proof of the theorem. Let $d = \gcd(a, n)$ and assume that $d \mid b$. By Bézout's identity, we know there exist integers u, v so that

$$\begin{aligned} au + nv &= d, & \text{so} \\ \frac{a}{d}u + \frac{n}{d}v &= 1 & \text{(divide by } d\text{). Thus} \\ \frac{a}{d}ub + \frac{n}{d}vb &= b & \text{(multiply by } b\text{), and finally,} \\ a\left(u\frac{b}{d}\right) + n\left(v\frac{b}{d}\right) &= b & \text{(redistribute since } d \mid b\text{).} \end{aligned}$$

This says that $x_0 = u\frac{b}{d}$ is a solution to $ax \equiv b \pmod{n}$, so solutions exist.

Now, how many are there? Suppose that y_0 is another solution to the congruence. Then

$$\begin{aligned} ay_0 \equiv ax_0 \equiv b \pmod{n} &\implies a(y_0 - x_0) \equiv 0 \pmod{n} \\ &\iff a(y_0 - x_0) = nk \text{ for some integer } k \\ &\iff \frac{a}{d}(y_0 - x_0) = \frac{n}{d}k \\ &\iff \frac{a}{d}(y_0 - x_0) \equiv 0 \pmod{\frac{n}{d}} \\ &\implies u\frac{a}{d}(y_0 - x_0) \equiv 0 \pmod{\frac{n}{d}}, \end{aligned}$$

and since from above, $\frac{a}{d}u + \frac{n}{d}v = 1$, we have $\frac{a}{d}u \equiv 1 \pmod{\frac{n}{d}}$, so

$$\begin{aligned} u\frac{a}{d}(y_0 - x_0) \equiv 0 \pmod{\frac{n}{d}} &\implies (y_0 - x_0) \equiv 0 \pmod{\frac{n}{d}} \\ &\implies y_0 \equiv x_0 \pmod{\frac{n}{d}} \\ &\implies y_0 = x_0 + \frac{n}{d}t \text{ for } t = 0, 1, \dots, d-1, \end{aligned}$$

all of which are distinct modulo n . \square

Remark 3.30. Note that $ax \equiv b \pmod{n}$ is solvable if and only if $\frac{a}{d}x \equiv \frac{b}{d} \pmod{\frac{n}{d}}$ is, and modulo $\frac{n}{d}$ there is a unique solution which propagates back to d solutions modulo n .

We consider some examples.

Example 3.31. Consider the congruence $6x \equiv 9 \pmod{15}$. We see that $d = \gcd(6, 15) = 3$ divides 9, so there will be three incongruent

solutions modulo 15. As above we see that $6x \equiv 9 \pmod{15} \iff 2x \equiv 3 \pmod{5}$, to which there is a unique solution $x \equiv 4 \pmod{5}$. Thus modulo 15, the solutions are $x \equiv 4 + 5t \pmod{15}$, $t = 0, 1, 2$, that is $x \equiv 4, 9, 14 \pmod{15}$.

Next we present a more complicated example that mimics the proof and uses Bézout's identity.

Example 3.32. We consider $198x \equiv 90 \pmod{252}$. We begin by computing the gcd of 252 and 198 (which we did previously in Example 3.12).

$$252 = 198(1) + 54$$

$$198 = 54(3) + 36$$

$$54 = 36(1) + 18$$

$$36 = 18(2) + 0$$

So $d = \gcd(252, 198) = 18$, and we see that $18 \mid 90$, so the congruence is solvable and has 18 incongruent solutions modulo 252.

To gain a particular solution, we work Euclid's algorithm backward, solving for 18 as follows:

$$\begin{aligned} 18 &= 54 - 36(1) \\ &= 54 - (1)(198 - 54(3)) = 198(-1) + 54(4) \\ &= 198(-1) + 4(252 - 198(1)) = 198(-5) + 252(4). \end{aligned}$$

Now

$$\begin{aligned} 198(-5) + 252(4) &= 18 \quad \text{means} \\ 198(-25) + 252(20) &= 90 \quad (\text{multiply by } 5), \end{aligned}$$

so $x_0 = -25$ is one solution to the original congruence $198x \equiv 90 \pmod{252}$. All solutions are of the form $x \equiv x_0 + \frac{252}{18}t = -25 + 14t \pmod{252}$, $t = 0, 1, \dots, 17$.

Note that alternatively, we might have simplified the congruence $198x \equiv 90 \pmod{252}$ to the equivalent $11x \equiv 5 \pmod{14}$, a quick inspection of which suggests $x \equiv 3 \pmod{14}$ as a solution. Thus the solutions to the original congruence are of the form $x \equiv 3 + 14t \pmod{252}$, $t = 0, 1, \dots, 17$.

While these may look rather different, they are not. The first set of solutions produces the congruences

$$x \equiv -25, -11, 3, 17, \dots, 213 \pmod{252},$$

while the second set produces

$$3, 17, \dots, 213, 227, 241 \pmod{252}.$$

The classes $3, 17, \dots, 213$ are obviously the same, and we see $-25 \equiv 227 \pmod{252}$ and $-11 \equiv 241 \pmod{252}$.

As a final example, we consider a simple congruence with solutions obtained in three different ways.

Example 3.33. The congruence $10x \equiv 15 \pmod{35}$ is solvable since $d = \gcd(10, 35) = 5 \mid 15$, and there will be five solutions modulo 35.

We can always reduce to the congruence $2x \equiv 3 \pmod{7}$.

The first method of solution is by inspection: $x \equiv 5 \pmod{7}$ works and is the unique solution mod 7, so we get $x \equiv 5+7t \pmod{35}$, $t = 0, 1, 2, 3, 4$ for the five solutions modulo 35.

The second method is to note $2x \equiv 3 \pmod{7}$ is solvable since $\gcd(2, 7) = 1$. Use Bézout's identity to write $1 = 2(-3) + 7(1)$. Multiplying by 3 gives $2(-9) + 7(3) = 3$, that is $x \equiv -9 \pmod{35}$ is a solution to the original congruence, so solutions are $-9+7t \pmod{35}$, $t = 0, 1, 2, 3, 4$.

The third method suggests that perhaps you can guess a solution to $2x + 7y = 3$, say $2(-2) + 7(1) = 3$. This would give $x \equiv -2 + 7t \pmod{35}$, $t = 0, 1, 2, 3, 4$, for a complete set of solutions.

3.7. Systems of Congruences

Just as a course in linear algebra often begins with a discussion of how to solve systems of linear equations, we frequently find ourselves in a situation where we wish to solve multiple congruences simultaneously. An extremely useful theorem that tells us how to do this in certain cases is called the *Chinese Remainder Theorem* (CRT). We state a general version, but leave the proof to a series of well-hinted exercises.

Theorem 3.34 (CRT). *Let m_1, m_2, \dots, m_r be integers with $m_i \geq 2$, and assume that for each $i \neq j$, $\gcd(m_i, m_j) = 1$. (We say the moduli m_i are coprime in pairs.) Let $a_1, \dots, a_r \in \mathbb{Z}$ be arbitrary. Then the system of congruences*

$$\begin{aligned} x &\equiv a_1 \pmod{m_1} \\ x &\equiv a_2 \pmod{m_2} \\ &\vdots \\ x &\equiv a_r \pmod{m_r} \end{aligned}$$

is solvable, and any two solutions x, x' satisfy $x \equiv x' \pmod{m}$ where $m = m_1 m_2 \cdots m_r$.

We note that this theorem can be proven by induction on r , and the case of $r = 1$ is trivial, so we start with the case of $r = 2$.

Exercise. Let $m, n > 1$ be coprime integers, and let a, b be arbitrary integers. Then the system of congruences

$$\begin{aligned} x &\equiv a \pmod{m}, \\ x &\equiv b \pmod{n} \end{aligned}$$

has a unique solution modulo mn . *A generous hint:* Note that since $\gcd(m, n) = 1$, Bézout's identity says there exists $u, v \in \mathbb{Z}$ so that $mu + nv = 1$. Show that the number $bmu + anv$ is a solution to the system, and then prove it is unique modulo mn .

Exercise. Explain how to use the above version of the CRT to solve a system

$$\begin{aligned} x &\equiv a \pmod{\ell}, \\ x &\equiv b \pmod{m}, \\ x &\equiv c \pmod{n}, \end{aligned}$$

where $\ell, m, n > 1$ are integers which are coprime in pairs.

Chapter 4

A Second View of Modular Arithmetic: \mathbb{Z}_n and U_n

4.1. Groups and Rings

Given two integers, we can add, subtract, or multiply them and always get another integer. We have seen that division is more problematic but led to many interesting ideas like the division algorithm. For now we focus just on addition, subtraction, and multiplication on the integers. Each of these operations is what is called a *binary operation* on the integers, meaning a rule that takes two inputs from the set and produces a third, in our case the sum, difference, or product of two integers. It may seem that these operations should all have similar properties, but that is not true, and understanding the distinctions leads us to a formal definition of a group (and a ring) which we need and want to leverage.

Let's slow down and take a closer look. As a concrete example, let's examine the operation of addition on the integers. Addition is the binary operation that takes an ordered pair of integers (m, n) and produces their sum, $m + n$. You might wonder what there can possibly be to understanding addition. But an operation all by itself

is just a function, and unless that function has interesting properties, it may not have much utility.

For example, while our binary operation gives the rule for adding two things together, how do we add three, $a + b + c$? In the past, you would never have hesitated in finding the sum of those three integers. You would grab two, add them, and then add the result to the third, using the binary operation twice. But how do we choose the first two? Is there a difference if we first add a and b together and add that result to c , or if we first add b and c , and add that result to a . Symbolically, how can we be sure that

$$(a + b) + c = a + (b + c)?$$

Probably somewhere in some distant corner of your brain, this relation seems familiar, and indeed you pull from the depths its name: the *associative* law. It seems so obvious that an operation should be associative. Isn't every operation associative? Actually, no. An easy counterexample is that subtraction on the integers is not an associative operation. For example,

$$-7 = (2 - 3) - 6 \neq 2 - (3 - 6) = 5.$$

So already we see that the operations of addition and subtraction on the integers are not on the same footing. Fortunately, this won't give us trouble, and even better, we have a way of tying subtraction to addition in a useful manner. For the record, multiplication on the integers is also associative, so we will say that \mathbb{Z} comes equipped with two associative binary operations, addition and multiplication.

In exploring the properties of addition and multiplication, we recall other facts from elementary school. The integers 0 and 1 are very special. For every integer m ,

$$m = m + 0 = 0 + m \quad \text{and} \quad m = 1 \cdot m = m \cdot 1.$$

These elements are called the *identities* associated to the operations of addition and multiplication. These operations have one more common property: they are both *commutative* operations, that is

$$m + n = n + m \quad \text{and} \quad mn = nm,$$

for all integers m and n .

There is one more property enjoyed by addition that is not enjoyed by multiplication on \mathbb{Z} , and that is the notion of an *inverse*. Formally, the property says that for every integer m there exists an integer n , so that $m + n = 0$. We are very comfortable with this notion, and typically we denote the (unique) element n by $n = -m$, and it is here that subtraction comes in.

If we accept subtraction as an operation, then

$$-m := 0 - m,$$

while if we accept the existence of inverses for all elements, then we can define subtraction as

$$m - n := m + (-n),$$

where the last expression is adding the inverse of n to m .

Perhaps you think all this formality is too much fussing for something you already understand very well, but we are counting on that implicit understanding since we want to extend these notions to sets with unfamiliar binary operations. Indeed, one very large goal of this text is to consider the set of points (x, y) on an elliptic curve, both of whose coordinates are rational numbers. We shall show how to define a binary operation on this set of points, which in many ways acts like the integers under addition, so we are relying heavily on your deep understanding of operations on the integers to allow us to extend these notions to new and important settings.

The subject of abstract algebra is the study of sets with binary operations which carry algebraic structure. Most algebra courses (and most textbooks) introduce these algebraic objects in increasing order of complexity. At the simple end of the spectrum is a set having a single binary operation, called a group.

Definition 4.1. A nonempty set G is called a *group* if it has a binary operation $*$ (that is, a map $G \times G \rightarrow G$ written $(g, h) \mapsto g * h$) satisfying the following:

- Identity. There is an element $e \in G$ so that $g * e = e * g = g$ for all elements $g \in G$.
- Inverses. For every $g \in G$ there is an $h \in G$ so that $g * h = h * g = e$.

- Associative. For all $g, h, k \in G$, $g * (h * k) = (g * h) * k$.

The set \mathbb{Z} of integers has two associative binary operations, addition and multiplication, so we can ask whether \mathbb{Z} is a group under each of those operations, and we quickly see the answer is no. The set \mathbb{Z} is a group under addition (with $* = +$): it has identity 0, it has inverses $-m$ for each m , and it is associative. Moreover, we know that addition is commutative as well on \mathbb{Z} , so \mathbb{Z} is a commutative group under addition. For historical reasons (in honor of Abel's work in group theory), commutative groups are instead referred to as *abelian* groups.

We now verify that \mathbb{Z} is not a group under multiplication. While it has an identity (1), and the operation is associative, inverses in general do not exist. For example, given the integer 2, there is no integer n so that $2n = 1$.

Nonetheless, the operations of addition and multiplication on \mathbb{Z} are quite compatible. The properties which characterize this compatibility are those that define an algebraic object called a ring.

Definition 4.2. A nonempty set R is called a *ring* (with identity) if it has two binary operations, one called addition, $+$, and one called multiplication, \cdot , so that R is an abelian group under addition, and it has the following additional properties:

- Multiplicative identity. There is an element $1 \in R$ so that $r \cdot 1 = 1 \cdot r = r$ for all elements $r \in R$.
- Associative. For all $r, s, t \in R$, $(r \cdot s) \cdot t = r \cdot (s \cdot t)$.
- Distributive. For $r, s, t \in R$, we have both
 $(r + s) \cdot t = r \cdot t + s \cdot t$ and $r \cdot (s + t) = r \cdot s + r \cdot t$.

The distributive property simply says that addition and multiplication play together well (i.e., interact compatibly). If in addition, the operation of multiplication is commutative, then R is called a commutative ring (with identity). Recall that the operation of addition is assumed to be commutative as part of the definition of a ring.

Finally, it does make sense to talk about rings that do not have an identity, but all the rings we shall meet in this text will have one.

Exercise. Show that $2\mathbb{Z}$ (the even integers) is a ring (without a multiplicative identity) under the usual operations of addition and multiplication.

4.2. Fractions and the Notion of an Equivalence Relation

The goal of the next two sections is to revisit the notion of modular arithmetic, but from a sophisticated point of view, using the notion of an equivalence relation to define a new set that we endow with the algebraic structure of a ring. Don't be concerned if that sounds rather daunting; we shall take it one step at a time.

Before making abstract definitions, let's examine an example with which you are very comfortable. The object we want to look at is the set of rational numbers, \mathbb{Q} . We can write down a definition of \mathbb{Q} without much fuss:

$$\mathbb{Q} = \{a/b \mid a, b \in \mathbb{Z}, b \neq 0\}.$$

We have known this characterization for so long that we have probably forgotten what the symbols actually mean. For example, surely we could have written the symbol a/b as $\frac{a}{b}$, or perhaps even $a \div b$. Perhaps you are inclined to ask, what significance is in the symbol? The symbol seems only intended to distinguish the numerator from the denominator. So to float an idea, why don't we just write $a/b = (a, b)$, that is, as an ordered pair where the first coordinate is the numerator, and the second is the denominator? This example will show why mathematicians are so careful with notation. Perhaps we are suggesting that we could write $\mathbb{Q} = \{(a, b) \mid a, b \in \mathbb{Z}, b \neq 0\}$?

No, not really—the notation (a, b) has its own separate meaning. If we think about ordered pairs, we are quite sure that $(1, 2) \neq (2, 4)$, say, for example, as points in the plane. So if $a/b = (a, b)$, we would have $1/2 \neq 2/4$, but those rational numbers are equal. What is the idea we are trying to get at? After all, we seem perfectly happy with the notation a/b . Why mess with it?

We are pretty sure now that the forward slash symbol in a/b has more significance than just being a separator for the numerator

and denominator. Consider a problem from elementary school: $\frac{1}{2} + \frac{1}{3} = \square$? How do we find the answer?

$$\frac{1}{2} + \frac{1}{3} = \frac{3}{6} + \frac{2}{6} = \frac{5}{6}.$$

Using some long-forgotten rules, we write $\frac{1}{2} = \frac{3}{6}$ and $\frac{1}{3} = \frac{2}{6}$ and then follow a procedure for addition of fractions having the same denominator. How in the world did we know to do that, and more to the point, what does it mean to write $1/2 = 3/6$?

The point here is that the symbol a/b means many things to us. For sure, it at least represents any fraction of the form at/bt where t is any nonzero integer. So how do we define the rational numbers? We start with a set S of ordered pairs representing all possible numerators and denominators, so we write

$$S = \{(a, b) \mid a, b \in \mathbb{Z}, b \neq 0\}.$$

Now to make the rational numbers, we want to identify many of these pairs. You learned that $a/b = c/d$ if and only if $ad - bc = 0$, so let us say that two ordered pairs $(a, b), (c, d) \in S$ are *related*, written $(a, b) \sim (c, d)$, if and only if $ad - bc = 0$. This relation is quite special. It satisfies the *reflexive*, *symmetric*, and *transitive* properties that we first met when we defined the notion of congruence. In the context of our new relation, those terms mean the following:

- Reflexive. $(a, b) \sim (a, b)$ for all $(a, b) \in S$. This is clear since the rule says they are related if $ab - ba = 0$, which is true in the integers.
- Symmetric. If $(a, b), (c, d) \in S$, then $(a, b) \sim (c, d)$ implies that $(c, d) \sim (a, b)$. By our rule, $(a, b) \sim (c, d)$ implies $ad - bc = 0$. For $(c, d) \sim (a, b)$, we would need that $cb - da = 0$, but this is immediate from the given expression and properties of the integers.
- Transitive. If $(a, b), (c, d), (e, f) \in S$, and $(a, b) \sim (c, d)$, and $(c, d) \sim (e, f)$, then $(a, b) \sim (e, f)$. To see this, we observe that $(a, b) \sim (c, d)$ implies $ad - bc = 0$, and $(c, d) \sim (e, f)$ implies that $cf - ed = 0$. Multiplying the first equality by f and the second by b , we have $adf - bcf = 0 = bcf - bed$, and hence by adding, $adf - bed = d(af - be) = 0$. Since $d \neq 0$,

we conclude that $af - be = 0$, which means $(a, b) \sim (e, f)$, as required.

We recall that a relation that is reflexive, symmetric, and transitive is called an *equivalence relation*, and by the *equivalence class* of an element (a, b) , we mean the subset of S consisting of all elements that are related (equivalent) to (a, b) . So the equivalence class of $(1, 2)$ (which we usually denote $1/2$) is the set of all pairs (c, d) which give the same fraction. It is this notion we are using when we write $1/2 = 3/6$ in order to solve our addition problem above. Indeed the statement $1/2 = 3/6$ is the statement that the subset of elements of S equivalent to $(1, 2)$ is the same as the subset of elements of S equivalent to $(3, 6)$.

So the bottom line is while we seem to have learned many new definitions, they are just clarifying the concept of a fraction which we have understood implicitly for a very long time. Before leaving fractions for a new look at modular arithmetic, let's give an alternate characteristic of our equivalence relation which defines fractions.

It is clear that for any $(a, b) \in S$ and $t \in \mathbb{Z}$, $t \neq 0$, that $(a, b) \sim (at, bt)$ since $abt - bat = 0$. So given an element $(a, b) \in S$, we can let $t = \gcd(a, b)$ and write $(a, b) = (a_0t, b_0t)$ where $a_0, b_0 \in \mathbb{Z}$ and $\gcd(a_0, b_0) = 1$, so $(a, b) \sim (a_0, b_0)$. In more familiar terms, $a/b = a_0/b_0$ where a_0/b_0 has been reduced to lowest terms. Now it is an easy exercise to show that given $(a, b) \in S$,

$$\{(c, d) \in S \mid (c, d) \sim (a, b)\} = \{(a_0t, b_0t) \mid t \in \mathbb{Z}, t \neq 0\}.$$

We shall see that this alternate characterization is natural when we define projective space in Chapter 6.

4.3. Modular Arithmetic

In the previous section we talked about an equivalence relation on the set $S = \{(a, b) \mid a, b \in \mathbb{Z}, b \neq 0\}$, and using that equivalence relation, we considered the set of equivalence classes, the rational numbers, and algebraic operations on that set.

In this section we do something analogous. We define an equivalence relation on the integers and then consider the set of equivalence

classes and work to understand the algebraic structure they inherit. Actually, for each positive integer n , we shall define an equivalence relation called *congruence modulo n* , and we will see that the equivalence relation *partitions* the integers into exactly n equivalence classes, terms we shall make precise below.

We recall the definition of congruence from the previous chapter. Let n be a positive integer, and let a, b be arbitrary integers. We divide a and b by n via the division algorithm, so we write $a = nq + r$ and $b = nq' + r'$ with $0 \leq r, r' < n$.

Definition 4.3. We say that a is *congruent to b modulo n* , written $a \equiv b \pmod{n}$ or $a \equiv b(n)$ if and only if $r = r'$. Equivalently, $a \equiv b \pmod{n}$ if and only if $n \mid (b - a)$, which is also equivalent to saying that $a = b + kn$ for some integer k .

As interim notation, to distinguish clearly between equivalence classes of integers and the integers which compose the equivalence class, we shall write

$$[a]_n = \{k \in \mathbb{Z} \mid k \equiv a \pmod{n}\};$$

in words, $[a]_n$ is the set of integers congruent modulo n to the integer a .

So for example,

$$\begin{aligned} [0]_2 &= \{0, \pm 2, \pm 4, \pm 6, \dots\} \text{ (the even integers),} \\ [1]_2 &= \{\pm 1, \pm 3, \pm 5, \dots\} \text{ (the odd integers).} \end{aligned}$$

We note a couple of simple, but very useful properties concerning the sets $[a]_n$.

Proposition 4.4. *Let $a, n \in \mathbb{Z}$, and let $n \geq 1$.*

- (1) $[a]_n = \{k \in \mathbb{Z} \mid k \equiv a \pmod{n}\} = \{a + n\ell \mid \ell \in \mathbb{Z}\}$.
- (2) $b \in [a]_n$ if and only if $[a]_n = [b]_n$.

Remark 4.5. We see that the first statement justifies the example above:

$$\begin{aligned} [0]_2 &= \{0 + 2\ell \mid \ell \in \mathbb{Z}\} = \{0, \pm 2, \pm 4, \pm 6, \dots\}, \\ [1]_2 &= \{1 + 2\ell \mid \ell \in \mathbb{Z}\} = \{\pm 1, \pm 3, \pm 5, \dots\}. \end{aligned}$$

The second statement says that an equivalence class can be named by any element in the class, so that

$$[0]_2 = [-4]_2 = [123456]_2 \quad \text{and} \quad [1]_2 = [-5]_2 = [12345]_2.$$

Proof. Both statements require us to show that two sets are equal. To show that sets $A = B$, we typically show $A \subseteq B$ and $B \subseteq A$. To show $A \subseteq B$, it is enough to take an arbitrary element $a \in A$ and show that $a \in B$. Now that we have a plan, we begin.

For the first statement, we want to show that $\{k \in \mathbb{Z} \mid k \equiv a \pmod{n}\} = \{a + n\ell \mid \ell \in \mathbb{Z}\}$. So first take an element $k \equiv a \pmod{n}$. By definition, $n \mid (k - a)$ which means $k - a = n\ell$ or $k = a + n\ell$, which tells us that $\{k \in \mathbb{Z} \mid k \equiv a \pmod{n}\} \subseteq \{a + n\ell \mid \ell \in \mathbb{Z}\}$. Now let $b = a + n\ell$. We see immediately that $b \equiv a \pmod{n}$ since $b - a = n\ell$, i.e., $n \mid (b - a)$, so $\{a + n\ell \mid \ell \in \mathbb{Z}\} \subseteq \{k \in \mathbb{Z} \mid k \equiv a \pmod{n}\}$. Having established both containments, we have equality of the two sets.

For the second statement, first suppose that $[a]_n = [b]_n$. Note that $b \in [b]_n$ since $b \equiv b \pmod{n}$ (congruence is a reflexive relation), so $b \in [a]_n = [b]_n$. To show the converse, suppose that $b \in [a]_n$, so $b \equiv a \pmod{n}$. We need to show that $[a]_n = [b]_n$, so we begin by showing $[a]_n \subseteq [b]_n$. To that end, let $c \in [a]_n$; we show that $c \in [b]_n$. Since $c \in [a]_n$, $c \equiv a \pmod{n}$, but $b \equiv a \pmod{n}$, so by transitivity, $c \equiv b \pmod{n}$, so $c \in [b]_n$, thus $[a]_n \subseteq [b]_n$. Since the argument is symmetric (that is, we could exchange a and b in all our statements), we can conclude $[b]_n \subseteq [a]_n$, so we have $[a]_n = [b]_n$, as desired. \square

For the rational numbers, the analog of the second statement is precisely what allows us to say that $1/2 = 3/6$, which we have found useful in computations. Nonetheless, even though we know that

$$1/2 = 3/6 = (-2)/(-4) \dots,$$

we typically settle on $1/2$ as the preferred representative of the equivalence class. We can do the same with congruence classes.

As we shall prove below, Proposition 4.4 also tells us that the set of congruence classes modulo n form what is called a *partition* of \mathbb{Z} . The notion of a partition of a set is a very simple one. It is simply a set of subsets so that every element of the original set is in precisely

one subset. Said another way, it is a set of subsets whose union is the whole set and any two subsets that are not equal are actually disjoint. Perhaps this is still too abstract. Take a bowl and pour a bag of plain and a bag of peanut M&M's into the bowl, mix thoroughly, and pour the contents onto a table. The set we shall partition is the set of M&M's on the table. What are some partitions of the set?

We could divide the set of M&M's into piles by color. These piles would be a partition. We could divide the M&M's into piles depending upon whether they were plain or peanut. These piles would be a different partition. We could even have kids gather around the table and grab as many M&M's as they can until the pool was gone. This too is a partition, though one in which it is more difficult to describe to which subset a given M&M belongs.

Now let's prove our claim: the set of congruence classes modulo n forms a partition of \mathbb{Z} —every integer belongs to one and only one congruence class. This is easy to see.

Fix a positive integer n , and let $a \in \mathbb{Z}$ be arbitrary. By Proposition 4.4, we trivially have that $a \in [a]_n$, so every integer is in at least one congruence class. Can the classes overlap? If $c \in [a]_n \cap [b]_n$, then by Proposition 4.4 $[a]_n = [c]_n = [b]_n$, so they lie in exactly one congruence class. Are there nice representatives for the classes? Actually, there is more than one set of nice representatives, but for now, we settle on one very natural set.

Write $a = nq + r$ with $0 \leq r < n$ by the division algorithm. Then since $a \equiv r \pmod{n}$, we have that $[a]_n = [r]_n$. So every element of \mathbb{Z} lies in $[0]_n \cup [1]_n \cup \cdots \cup [n-1]_n$. Moreover, given $0 \leq r, r' < n$, $[r]_n = [r']_n$ if and only if $r \equiv r' \pmod{n}$, but since $0 \leq |r - r'| < n$, the only way to have $n \mid (r - r')$ is for $r = r'$, so those n congruence classes are disjoint.

Definition 4.6. For a positive integer n , we let $\mathbb{Z}_n = \{[a]_n \mid a \in \mathbb{Z}\}$; that is, \mathbb{Z}_n is the set of congruence classes modulo n . So for each n , we have a different partition \mathbb{Z}_n of \mathbb{Z} .

For example,

$$\mathbb{Z}_2 = \{[0]_2, [1]_2\},$$

$$\mathbb{Z}_5 = \{[0]_5, [1]_5, [2]_5, [3]_5, [4]_5\}, \text{ and in general,}$$

$$\mathbb{Z}_n = \{[0]_n, [1]_n, \dots, [n-1]_n\}.$$

Exercise. It is actually not difficult to show that there is a one-to-one correspondence between partitions of a set and equivalence relations on the set. While we have seen the example that the equivalence relation of congruence modulo n gives rise to the partition of the integers into congruence classes, consider the equivalence relation associated to the partitions of M&M's we gave above. Then see if you can prove the general statement.

Now we want to translate arithmetic with congruences into algebraic operations on the congruence classes in \mathbb{Z}_n . In particular, we are going to give the set \mathbb{Z}_n the structure of a ring. The ring operations on \mathbb{Z}_n will be analogous to arithmetic with fractions. Recall that finding the sum $1/2 + 1/3$ involves understanding the equivalence classes represented by these fractions.

We define two binary operations on \mathbb{Z}_n in what should seem a very natural manner:

$$[a]_n + [b]_n = [a + b]_n,$$

$$[a]_n \cdot [b]_n = [ab]_n.$$

Before declaring them to be binary operations on the set, we have to check that the definition actually makes sense. After all, we are defining the addition and multiplication of two sets, and while symbolically the definition seems natural, there really is something to check, and that is that our definition is *well-defined*. To see the essence of the problem, we first do an example.

Example 4.7. Let $n = 10$, so $\mathbb{Z}_{10} = \{[0]_{10}, [1]_{10}, \dots, [9]_{10}\}$. Our rules say things like $[2]_{10} + [3]_{10} = [2 + 3]_{10} = [5]_{10}$, and $[2]_{10} \cdot [4]_{10} = [2(4)]_{10} = [8]_{10}$, which seem fine. So too does $[6]_{10} + [7]_{10} = [13]_{10} = [3]_{10}$. But here's the rub. These classes are sets, so the definition we give can depend only on the set and not the representative of

the set we choose. Concretely, we know that $[4]_{10} = [1234]_{10}$ and $[7]_{10} = [-243]_{10}$, so to be well-defined it must be the case that $[4]_{10} + [7]_{10} = [1234]_{10} + [-243]_{10}$. Our rules say $[4]_{10} + [7]_{10} = [11]_{10}$ and $[1234]_{10} + [-243]_{10} = [991]_{10}$, so it better be true that $[11]_{10} = [991]_{10}$. Well it is, but being well-defined means that this will always be the case no matter how we choose our representatives.

To settle this issue, we take two classes each with different elements naming the class. Suppose $[a]_n = [a']_n$ and $[b]_n = [b']_n$. We need to show that

$$[a]_n + [b]_n = [a + b]_n = [a' + b']_n = [a']_n + [b']_n$$

and, similarly,

$$[a]_n \cdot [b]_n = [ab]_n = [a'b']_n = [a']_n \cdot [b']_n.$$

But we have already checked this via congruences in Proposition 3.21. The assumption $[a]_n = [a']_n$ and $[b]_n = [b']_n$ says $a \equiv a' \pmod{n}$ and $b \equiv b' \pmod{n}$, so by Proposition 3.21 we know that $a + b \equiv a' + b' \pmod{n}$ and $ab \equiv a'b' \pmod{n}$, which in turn tells us that $[a + b]_n = [a' + b']_n$ and $[ab]_n = [a'b']_n$.

Now that we have well-defined operations on the set \mathbb{Z}_n , we establish the following.

Proposition 4.8. *The operations defined above make \mathbb{Z}_n into a commutative ring with identity.*

Proof. First we show that \mathbb{Z}_n is an abelian group under addition. The identity is $[0]_n$ and the inverse of $[a]_n$ is $[-a]_n$. The operation is associative because it *inherits* this property from \mathbb{Z} , that is,

$$\begin{aligned} ([a]_n + [b]_n) + [c]_n &= [(a + b)]_n + [c]_n = [(a + b) + c]_n \\ &= [a + (b + c)]_n = [a]_n + [b + c]_n \\ &= [a]_n + ([b]_n + [c]_n), \end{aligned}$$

where the key step $[(a + b) + c]_n = [a + (b + c)]_n$ uses associativity of addition in \mathbb{Z} , which is why we say the property is inherited from \mathbb{Z} .

The properties of multiplication are easily checked: $[1]_n$ is the multiplicative identity, and associativity and the distributive laws are

inherited from the corresponding ones in \mathbb{Z} , just as we have shown above with associativity under addition. \square

We want to investigate the multiplicative structure of \mathbb{Z}_n in a bit more detail.

Definition 4.9. For a positive integer n , let

$$U_n = \{[a]_n \mid \gcd(a, n) = 1\}.$$

Once again, we need to be a little careful to make sure our definition makes sense. Let's suppose that $a \in \mathbb{Z}$ and $\gcd(a, n) = 1$. Is it true if $[a]_n = [a']_n$ that $\gcd(a', n) = 1$ as well? Fortunately, the answer is yes, and it is easy to see. If $[a]_n = [a']_n$, then $a = a' + kn$ for some integer k . We want to see that $\gcd(a, n) = 1$ implies that $\gcd(a', n) = 1$, so let d be any common divisor of a' and n . Then we know that d divides any combination of a' and n . In particular $d \mid a' + kn = a$, so d is a common divisor of a and n , so must divide the $\gcd(a, n) = 1$, so $d = \pm 1$ which means $\gcd(a', n) = 1$, so there is no ambiguity.

Proposition 4.10. For a positive integer n , U_n is an abelian group under multiplication.

Proof. You might think this should be trivial, but there are important properties to verify. First, while we know that \mathbb{Z}_n is closed under multiplication, we need to verify that U_n is as well; that is if $[a]_n, [b]_n \in U_n$, then so is $[a]_n[b]_n = [ab]_n$. This boils down to saying that if $\gcd(a, n) = 1 = \gcd(b, n)$, then $\gcd(ab, n) = 1$. If $\gcd(a, n) = 1 = \gcd(b, n)$, then by Bézout's identity, there exist integers u, v, u', v' so that $au + nv = 1 = bu' + nv'$. Multiplying them together shows that $ab(uu') + n(bu'v + auv' + nvv')$ which implies ab and n are coprime, so multiplication is actually a binary operation on U_n . Now things get easier.

Given closure, the operation is associative since multiplication is associative on \mathbb{Z}_n , and as a set $U_n \subset \mathbb{Z}_n$. The identity is $[1]_n$, so we need only show that every element has an inverse; that is, given an element $[a]_n \in U_n$, we need to show there is a $[b]_n \in U_n$ with $[a]_n[b]_n = [1]_n = [b]_n[a]_n$.

We have already done the work via linear congruences. If $[a]_n \in U_n$, then $\gcd(a, n) = 1$, so by Corollary 3.29, the congruence $ax \equiv 1 \pmod{n}$ has a unique solution modulo n , say b . This means that $ab \equiv 1 \pmod{n}$. Translating this back to congruence classes, we have

$$[1]_n = [ab]_n = [a]_n[b]_n = [b]_n[a]_n,$$

using that multiplication is commutative in \mathbb{Z}_n . \square

As a prelude to RSA cryptography, we need to define a rather well-known arithmetic function called the *Euler totient function* (ϕ) and prove a simple result of Euler's (which actually generalizes one of Fermat's—no, not that one).

Perhaps we should begin with the usual definition of the Euler's function. Let \mathbb{Z}_+ denote the positive integers. We define the function $\phi : \mathbb{Z}_+ \rightarrow \mathbb{Z}_+$ by

$$\phi(n) = \#\{k \in \mathbb{Z} \mid 1 \leq k \leq n \text{ and } \gcd(k, n) = 1\},$$

that is $\phi(n)$ is the number of integers between 1 and n which are relatively prime to n .

One easily checks a few values: $\phi(1) = 1$, $\phi(2) = 1$, $\phi(3) = 2$, $\phi(4) = 2$, $\phi(5) = 4$, but it is somewhat nicer (and more useful) to realize that the Euler function counts the size of a group.

Proposition 4.11. *The Euler totient function is given by $\phi(n) = \#U_n$.*

Proof. We have seen that $\mathbb{Z}_n = \{[a]_n \mid a \in \mathbb{Z}\} = \{[0]_n, \dots, [n-1]_n\}$. For this proof it is more convenient to choose a slightly different set of representatives for the classes in \mathbb{Z}_n , so we write $\mathbb{Z}_n = \{[1]_n, [2]_n, \dots, [n]_n\}$ where we have written $[n]_n$ instead of $[0]_n$. Since $U_n = \{[a]_n \in \mathbb{Z}_n \mid \gcd(a, n) = 1\}$, we see that

$$U_n = \{[k]_n \mid 1 \leq k \leq n \text{ and } \gcd(k, n) = 1\},$$

which gives $\phi(n) = \#U_n$. \square

We are going to begin to lighten our notation just a bit. When n is understood, we shall write $\mathbb{Z}_n = \{[0]_n, \dots, [n-1]_n\}$ as $\mathbb{Z}_n = \{\overline{0}, \overline{1}, \dots, \overline{n-1}\}$.

Example 4.12. A few examples of \mathbb{Z}_n and U_n .

$$\begin{array}{lll} \mathbb{Z}_2 = \{\overline{0}, \overline{1}\}, & U_2 = \{\overline{1}\}, & \phi(2) = 1, \\ \mathbb{Z}_3 = \{\overline{0}, \overline{1}, \overline{2}\}, & U_3 = \{\overline{1}, \overline{2}\}, & \phi(3) = 2, \\ \mathbb{Z}_4 = \{\overline{0}, \overline{1}, \overline{2}, \overline{3}\}, & U_4 = \{\overline{1}, \overline{3}\}, & \phi(4) = 2, \\ \mathbb{Z}_5 = \{\overline{0}, \overline{1}, \overline{2}, \overline{3}, \overline{4}\}, & U_5 = \{\overline{1}, \overline{2}, \overline{3}, \overline{4}\}, & \phi(5) = 4, \\ \mathbb{Z}_6 = \{\overline{0}, \overline{1}, \overline{2}, \overline{3}, \overline{4}, \overline{5}\}, & U_6 = \{\overline{1}, \overline{5}\}, & \phi(6) = 2, \\ \mathbb{Z}_p = \{\overline{0}, \overline{1}, \dots, \overline{(p-1)}\}, & U_p = \{\overline{1}, \dots, \overline{(p-1)}\}, & \phi(p) = p-1 \text{ (} p \text{ prime)}. \end{array}$$

The case for p a prime is immediate from the observation that for $1 \leq k < p$, $\gcd(k, p) = 1$ since p a prime means its only positive divisors are 1 and p . \square

Below we shall prove Fermat's little theorem and Euler's theorem, and while they both have proofs whose tools lie wholly within elementary number theory, we give a proof using some basic properties from groups, since it is this argument we shall generalize later when talking about the group of points on an elliptic curve.

We begin with a leisurely stroll. Let G be a group, and let $g \in G$. The axioms for a group guarantee the existence of an inverse, that is an element h so that $g * h = h * g = e$. While the axioms don't say it, the inverse is unique, for if k is another (possibly different) inverse (i.e., $g * k = k * g = e$), then

$$h = h * e = h * (g * k) = (h * g) * k = e * k = k,$$

showing they must be equal.

Denote the unique inverse of g by g^{-1} ; we now extend this notation. For $k > 0$, denote by g^k the element $\underbrace{g * g * \dots * g}_{k \text{ times}}$, by $g^0 = e$,

the identity, and for $k < 0$, let g^k be the inverse of $g^{|k|}$. With this shorthand convention, g^k satisfies the usual rules for exponents: $g^k g^\ell = g^{k+\ell}$ and $(g^k)^\ell = g^{k\ell}$ for all integers k, ℓ .

Now suppose that G is a finite group, and consider the elements

$$g, g^2, g^3, \dots, g^{r+1},$$

where $r = \#G = |G|$ is the *order* (cardinality) of the group G . Since there are $r + 1$ elements in the list, all of which are in G , and there are only r distinct elements in G , two elements in the list must be the same, that is

$$g^i = g^j \quad \text{for some } 1 \leq i < j \leq r + 1.$$

Multiplying both sides by g^{-i} , we see that $g^{j-i} = g^i g^{-i} = g^0 = e$, and note that $1 \leq j-i \leq r$. We summarize this as a little proposition.

Proposition 4.13. *Let G be a finite group. Then for any element $g \in G$, there is a positive integer $n \leq |G|$ with $g^n = e$.*

Definition 4.14. We define the *order* of an element g , denoted $|g|$, as the smallest positive integer m so that $g^m = e$, if one exists. If no such positive integer exists, we say g has infinite order. We have just shown that in a finite group, every element has a finite order.

For example consider the group $G = U_9$ which has order $\phi(9) = 6$: $U_9 = \{\bar{1}, \bar{2}, \bar{4}, \bar{5}, \bar{7}, \bar{8}\}$, with $e = \bar{1}$.

$$\bar{1}^1 = \bar{1}, \text{ so } |\bar{1}| = 1,$$

$$\bar{2}^1 = \bar{2}, \bar{2}^2 = \bar{4}, \bar{2}^3 = \bar{8}, \bar{2}^4 = \bar{7}, \bar{2}^5 = \bar{5}, \bar{2}^6 = \bar{1}, \text{ so } |\bar{2}| = 6,$$

$$\bar{4}^1 = \bar{4}, \bar{4}^2 = \bar{7}, \bar{4}^3 = \bar{1}, \text{ so } |\bar{4}| = 3,$$

$$\bar{5}^1 = \bar{5}, \bar{5}^2 = \bar{7}, \bar{5}^3 = \bar{8}, \bar{5}^4 = \bar{4}, \bar{5}^5 = \bar{2}, \bar{5}^6 = \bar{1}, \text{ so } |\bar{5}| = 6,$$

$$\bar{7}^1 = \bar{7}, \bar{7}^2 = \bar{4}, \bar{7}^3 = \bar{1}, \text{ so } |\bar{7}| = 3,$$

$$\bar{8}^1 = \bar{8}, \bar{8}^2 = \bar{1}, \text{ so } |\bar{8}| = 2.$$

An important theorem in group theory that we want to leverage here but whose proof would take us a bit too far afield, is due to Lagrange, a special case of which we state here.

Theorem 4.15 (Lagrange). *Let G be a finite group, and let $g \in G$. Then $g^{|G|} = e$.*

Above we had shown that there was an integer $n \leq |G|$ for which $g^n = e$, but this result says that $n = |G|$ is a universal exponent for every element in the group. So any element when raised to a power equal to the order of the group is the identity of the group. We saw

this in our example above since the orders of all the elements divide 6, the order of U_6 . In fact, this later observation is a fact we can prove.

Proposition 4.16. *Let G be a group, $g \in G$, and suppose that for some positive integer m , $g^m = e$. Then $|g| \mid m$. In particular, in a finite group, we have that $|g|$ divides $|G|$; the order of an element divides the order of the group.*

Proof. Since there is a positive integer m with $g^m = e$, we know g has finite order, so let $d = |g|$. Then $g^d = g^m = e$ and d is by definition the smallest positive integer with that property. Use the division algorithm and write $m = dq + r$ with $0 \leq r < d$. Observe that $e = g^m = (g^d)^q g^r = g^r$, so if $0 < r < d$, it would violate that d is the order of g , thus $r = 0$ and we have $d \mid m$.

The second statement now follows since, by Lagrange's theorem, $g^{|G|} = e$. \square

Indeed the theorem above gives one statement of Euler's theorem:

Theorem 4.17 (Euler's theorem). *Let $a \in \mathbb{Z}$ with $\gcd(a, n) = 1$. Then $a^{\phi(n)} \equiv 1 \pmod{n}$.*

Note that Euler's theorem is simply the statement that if $\bar{a} \in U_n$, then $\bar{a}^{|U_n|} = \bar{1}$, that is a special case of the general result for finite groups. Since we have not proven the general result, we give a separate proof of Euler's theorem but make the point that this is a good example of the power of abstraction.

Indeed Euler's theorem is itself a generalization of Fermat's little theorem, since for a prime p , we know $|U_p| = p - 1$.

Theorem 4.18 (Fermat's little theorem). *Let p be a prime and a an integer with $p \nmid a$; then $a^{p-1} \equiv 1 \pmod{p}$. Equivalently, for any integer a , we have $a^p \equiv a \pmod{p}$.*

To prove Euler's theorem, we need to produce two different but related sets of representatives for U_n . We start with

Proposition 4.19. *Suppose that $U_n = \{\bar{b}_1, \dots, \bar{b}_{\phi(n)}\}$ and that a is an integer with $\gcd(a, n) = 1$. Then $U_n = \{\overline{ab_1}, \dots, \overline{ab_{\phi(n)}}\}$. Equivalently, for each i with $1 \leq i \leq \phi(n)$, there is a unique j with $1 \leq j \leq \phi(n)$ so that $ab_i \equiv b_j \pmod{n}$.*

Before proving this, we give an example:

Example 4.20. $U_{10} = \{\bar{1}, \bar{3}, \bar{7}, \bar{9}\}$. Let $a = 13$ which is clearly coprime to 10. The proposition says that $U_{10} = \{\overline{13}, \overline{39}, \overline{91}, \overline{117}\} = \{\bar{3}, \bar{9}, \bar{1}, \bar{7}\}$. So we get the same classes, just in a different order.

Proof. First observe that each class $\overline{ab_i} \in U_n$. We could *reprove* this with congruences, but we have already established that U_n is a group under multiplication, and so is closed under multiplication. Since $\gcd(a, n) = 1$, we have $\bar{a} \in U_n$, so $\bar{a}\bar{b}_i = \overline{ab_i} \in U_n$ by closure.

Next we observe that all the elements are distinct. Again, we do this using group theory.

Suppose that $\overline{ab_i} = \overline{ab_j}$, then $\bar{a}\bar{b}_i = \bar{a}\bar{b}_j$. Multiplying both sides by the inverse of \bar{a} (U_n is a group!), we see that $\bar{b}_i = \bar{b}_j$, so all the representatives are distinct (i.e., different congruence classes), and since there are the correct number, they fill out all of U_n . \square

Theorem 4.21 (Euler's theorem). *Let $a, n \in \mathbb{Z}$ with $n \geq 1$ and $\gcd(a, n) = 1$. Then $a^{\phi(n)} \equiv 1 \pmod{n}$.*

Proof. Write $U_n = \{\bar{b}_1, \dots, \bar{b}_{\phi(n)}\}$. Then as in the proposition, $U_n = \{\overline{ab_1}, \dots, \overline{ab_{\phi(n)}}\}$. Since these are the same elements (possibly in a different order) and U_n is abelian, their products are equal, that is

$$\bar{b}_1\bar{b}_2 \cdots \bar{b}_{\phi(n)} = \overline{ab_1}\overline{ab_2} \cdots \overline{ab_{\phi(n)}} = \bar{a}^{\phi(n)}\bar{b}_1\bar{b}_2 \cdots \bar{b}_{\phi(n)}.$$

Since U_n is a group, the element $\bar{b}_1\bar{b}_2 \cdots \bar{b}_{\phi(n)}$ has an inverse in U_n . Multiplying both sides by the inverse produces $\bar{a}^{\phi(n)} = \bar{1}$ or, equivalently, $a^{\phi(n)} \equiv 1 \pmod{n}$. \square

As an amusing corollary, we know that the congruence $ax \equiv b \pmod{n}$ has a unique solution mod n if and only if $\gcd(a, n) = 1$. Well, if it's unique, what is it? Since $\gcd(a, n) = 1$, we know $a^{\phi(n)} \equiv 1 \pmod{n}$, so multiplying both sides of the congruence by $a^{\phi(n)-1}$ yields $x \equiv a^{\phi(n)-1}b \pmod{n}$. Of course computing this solution may

take time if $\phi(n)$ is large, which leads us to consider fast modular exponentiation.

Example 4.22. Let's compute $5^{123} \pmod{13}$. We consider Euler's (or Fermat's little) theorem and realize that $5^{12} \equiv 1 \pmod{13}$, so $5^{123} = (5^{12})^{10} \cdot 5^3 \equiv 5^3 \equiv 8 \pmod{13}$.

Example 4.23. Next observe that $341 = 11 \cdot 31$. We want to show that $2^{341} \equiv 2 \pmod{341}$. Comparing this to Fermat's little theorem, we might wonder whether 341 is prime since it seems to behave like one when exponentiating with respect to the base 2. Indeed we know that 341 is not prime, but this behavior earns it the name *pseudoprime to the base 2*, an important notion in the realm of primality testing.

Since we fortuitously know the factorization of 341, we claim that $2^{341} \equiv 2 \pmod{341}$ if and only if $2^{341} \equiv 2 \pmod{11}$ and $2^{341} \equiv 2 \pmod{31}$. One direction is obvious and always true: If $a \equiv b \pmod{n}$, then $a \equiv b \pmod{d}$ for any $d \mid n$. But the converse depends on the moduli 11 and 31 being coprime. Suppose that $a \equiv b \pmod{11}$ and $a \equiv b \pmod{31}$. We see that $11 \mid (a - b)$ says that $a - b = 11k$ for some integer k . Since $31 \mid (a - b)$, we see $31 \mid 11k$, and since $\gcd(11, 31) = 1$, we have that $31 \mid k$, so that $k = 31k'$. Thus $a - b = 11k = 11 \cdot 31k' = 341k'$ or $a \equiv b \pmod{341}$.

Now we establish the simpler congruences. By Euler's (or Fermat's little) theorem, we know that $2^{10} \equiv 1 \pmod{11}$ and $2^{30} \equiv 1 \pmod{31}$. Thus

$$2^{341} \equiv (2^{10})^{34} \cdot 2^1 \equiv 2 \pmod{11}$$

and

$$2^{341} \equiv (2^{30})^{11} \cdot 2^{11} \equiv 2^{11} \pmod{31}.$$

To finish, we notice serendipitously that $2^5 = 32 \equiv 1 \pmod{31}$ so that $2^{11} \equiv (2^5)^2 \cdot 2^1 \equiv 2 \pmod{31}$.

Example 4.24. We do the last example again, but this time assuming we know nothing about the factorization. Being able quickly to perform modular exponentiation is critical to the implementation of the RSA cryptosystem.

Let's discuss this in general. We want to find $a^{341} \pmod{n}$ for some a and n . Their values are not really too important to the general plan. The key is to write the exponent in its base 2 expansion: $341 = 2^8 + 2^6 + 2^4 + 2^2 + 2^0$, so that

$$a^{341} \equiv a^{2^8+2^6+2^4+2^2+2^0} \equiv a^{2^8} a^{2^6} a^{2^4} a^{2^2} a^{2^0} \pmod{n}.$$

Now we observe that the terms a^{2^k} can be obtained by successive squaring, that is

$$a^{2^{k+1}} = a^{2^k \cdot 2} = (a^{2^k})^2.$$

So we compute:

$$\begin{aligned} 2^{2^0} &\equiv 2 \pmod{341}, \\ 2^{2^1} &\equiv (2^1)^2 \equiv 4 \pmod{341}, \\ 2^{2^2} &\equiv (2^{2^1})^2 \equiv 16 \pmod{341}, \\ 2^{2^3} &\equiv (2^{2^2})^2 \equiv 256 \pmod{341}, \\ 2^{2^4} &\equiv (2^{2^3})^2 \equiv 256^2 \equiv 64 \pmod{341}, \\ 2^{2^5} &\equiv (2^{2^4})^2 \equiv 64^2 \equiv 4 \pmod{341}, \\ 2^{2^6} &\equiv (2^{2^5})^2 \equiv 4^2 \equiv 16 \pmod{341}, \\ 2^{2^7} &\equiv (2^{2^6})^2 \equiv 16^2 \equiv 256 \pmod{341}, \\ 2^{2^8} &\equiv (2^{2^7})^2 \equiv 256^2 \equiv 64 \pmod{341}. \end{aligned}$$

Thus

$$2^{341} \equiv 2^{2^8} 2^{2^6} 2^{2^4} 2^{2^2} 2^{2^0} \equiv 64 \cdot 16 \cdot 64 \cdot 16 \cdot 2 \equiv 2 \pmod{341}.$$

4.4. A Few More Comments on the Euler Totient Function

By the Fundamental Theorem of Arithmetic (Theorem 3.15), we know that every integer $n > 1$ factors uniquely as a product of primes, so in number theory we are often interested in functions that respect factorizations. That means that the ideal situation is one in which we have a function f so that $f(mn) = f(m)f(n)$ for any integers m, n ; for example, $f(n) = n^k$ for any integer k is such a function. It is easy to show that the Euler totient function is not quite this nice, but it still has many important properties (see the exercises below). One of them is that for relatively prime inputs m, n , $\phi(mn) = \phi(m)\phi(n)$. This becomes very important as we discuss RSA. At the heart of RSA is a modulus n which is the product of two primes, $n = pq$. As part of the encryption scheme, we must compute $\phi(n)$. If we know the factorization, this becomes trivial, $\phi(n) = \phi(pq) = \phi(p)\phi(q) = (p-1)(q-1)$, but it will turn out that calculating the value of $\phi(n)$ without knowing the factorization $n = pq$ is as hard as factoring the integer n .

To develop a few more properties of the Euler totient function, we may take its definition as $\phi(n) = \#U_n$, so we have seen $\phi(p) = p-1$ when p is a prime. Consider the exercises below.

Exercise. Let p be a prime. Determine the value of $\phi(p^r)$ for any positive integer r . *Hint:* It may be easier to count the number of elements of $a \in \mathbb{Z}_{p^r}$ which are not relatively prime to p^r and use that to determine the value of the function. Of course be sure to check your answer against a few examples you can compute by hand.

Exercise. It is easy to show that in general $\phi(mn) \neq \phi(m)\phi(n)$ for general m, n , but what is remarkable is that when $\gcd(m, n) = 1$, $\phi(mn) = \phi(m)\phi(n)$. The function ϕ is an example of a *multiplicative* function in number theory. Perhaps more surprising is that this is a direct consequence of the Chinese Remainder Theorem. Give a proof that ϕ is multiplicative using the following idea. Suppose that $m, n \geq 2$ and $\gcd(m, n) = 1$. Show that there is a bijection between the sets U_{mn} and $U_m \times U_n$ (ordered pairs (a, b) with $a \in U_m$, $b \in U_n$). Note that U_{mn} has cardinality $\phi(mn)$ and $U_m \times U_n$ has cardinality $\phi(m) \cdot \phi(n)$. To establish the bijection, define a map $F : U_{mn} \rightarrow U_m \times U_n$ by $F([a]_{mn}) = ([a]_m, [a]_n)$. You need to show this map is well-defined, one-to-one, and onto. Then deduce the result.

Some of these words may be new to you, so here are some definitions.

- We have encountered the term “well-defined” before. In this context it means that if $[a]_{mn} = [b]_{mn}$, then $F([a]) = F([b])$.
- The map F is one-to-one (injective) if $F([a]) = F([b])$ implies $[a]_{mn} = [b]_{mn}$.
- The map F is onto (surjective) if given $([b]_m, [c]_n) \in U_m \times U_n$, there exists $[a]_{mn} \in U_{mn}$ so that $F([a]) = ([b], [c])$.
- A map is bijective if it is one-to-one and onto.
- If $f : S \rightarrow T$ is a bijection, then S and T are said to have the same cardinality (size), and the result you are to prove is simply that when $\gcd(m, n) = 1$, the size of U_{mn} and $U_m \times U_n$ is the same.

Putting together the two exercises we see that we can write down an explicit formula for $\phi(n)$ for any $n \geq 2$. Factor $n = p_1^{e_1} \cdots p_r^{e_r}$ into powers of distinct primes. Because ϕ is multiplicative, $\phi(n) = \phi(p_1^{e_1}) \cdots \phi(p_r^{e_r})$, and by the first exercise we have $\phi(n) = p_1^{e_1-1}(p_1 - 1) \cdots p_r^{e_r-1}(p_r - 1) = p_1^{e_1}(1 - 1/p_1) \cdots p_r^{e_r}(1 - 1/p_r)$. A fancy way to write this is

$$\phi(n) = n \prod_{p|n} \left(1 - \frac{1}{p}\right),$$

where the product is over all primes $p \mid n$.

4.5. An Application to Factoring

In Theorem 3.15, we gave a proof that every integer $n > 1$ can be factored as a product of primes in an essentially unique way. On the other hand, there is often a difference between what can be done theoretically and how efficient it is to do practically. An excellent example of this is the process of factoring an integer. As we have alluded, the difficulty of factoring an integer—even an integer which is the product of only two primes—is the basis for the security of the public key cryptosystem known as RSA which we investigate in the next chapter. So this seems a good time to take a first look at factoring, first with a rather brute force mentality, and second with a good deal more sophistication using some of the elementary number theoretic results we have established so far.

Factoring is inherently a recursive process, so we will always be content to start with a given $n > 1$ and write $n = ab$ with $1 < a, b < n$. If this is not possible, then of course n is a prime. If n can be written as $n = ab$ with $1 < a, b < n$, then one of the factors a or b must be $\leq \sqrt{n}$. So as the first method of factoring, we consider trial division.

4.5.1. Trial Division. A composite number $n > 1$ must have a divisor whose value is $\leq \sqrt{n}$, so given an integer $n > 1$, we can try dividing by $2, 3, 4, 5, 6, \dots, \lfloor \sqrt{n} \rfloor$. If none of these divides n , we know that n is prime. This is certainly perfectly reasonable in theory, but in practice, it is a terrible method except for very small integers n .

For example, if n is an integer with approximately 100 decimal digits and our computer can check 1 million trial divisions per second, it could take as long as 3.2×10^{37} years to check all the divisors. If we used a computer a million times faster (i.e., 1 trillion trial divisions per second), it could take up to 3.2×10^{31} years. Given that the estimated age of the universe is approximately 13.5 billion years (1.35×10^{10}) we are talking about a process whose length would be the cube of the age of the universe. Most would consider this less than optimal.

But how else can we prove that a number is composite, without exhibiting a nontrivial divisor? It turns out that factoring is quite an art in which certain types of integers are easier to factor than others. Below we look at one method that is effective on a certain class of

integers, a method which will be greatly generalized once we have some knowledge of elliptic curves under our belts.

Perhaps the first thing one should know before attempting to factor an integer n is that n is composite, i.e., it is not a prime. Fortunately the task of determining whether a number is prime is much simpler, meaning there is a deterministic polynomial time algorithm that can answer whether an integer is prime. The subject of primality testing is one worth its own set of chapters, but time constrains us to assume that our given integer is composite (such as an RSA modulus which is the product of two distinct primes) and go from there.

4.5.2. Pollard's $p - 1$ Method. Since the ability to choose an integer that is hard to factor is the basis for the security of RSA, we should think of factoring as an adversarial game, one side trying to pick composite integers difficult to factor and the other side trying to classify those integers that are easy to factor. What can that possibly mean? One simple observation is that while the size of an integer can influence the degree of difficulty in factoring it, there are exceptions. Nobody would have much difficulty in factoring 100000000000000000000000000000000. And while a complete factorization of $n = 1234567890123456$ might be challenging, it takes little effort to infer that it factors as $n = 2^6 \cdot 19290123283179$, that is, even numbers make a poor choice if you are challenging someone to factor your integer. Similarly, an integer divisible by mostly small primes would also be easy to factor.

The method of Pollard, his $p - 1$ method, is generally effective on those integers n which possess a prime divisor p , so that $p - 1$ is itself the product of small primes. You probably would never have thought to consider a class of integers defined by that property, but it gives you a sense of how complicated this adversarial game of factoring really is.

Of course if your integer does not have that property, the algorithm we describe will run a very long time and reveal nothing. This is part of why factoring is still as much an art as science. So we shall assume our n has such a prime divisor p , though it will take a few moments for us see where this assumption can be used.

Even if we make this assumption, how does it help? After all, we don't know p . Let's keep our eye on the fact that what we are after is a nontrivial divisor of n , that is, we want to write $n = ab$ with $1 < a, b < n$. One observation is that if you have such a factorization, then in particular, you have found an integer a for which $1 < a = \gcd(a, n) < n$, so finding such an integer a with $1 < \gcd(a, n) < n$ is equivalent to finding nontrivial divisor of n . Randomly guessing values for a seems pretty much the same as trial division, so we need to leverage things in our favor, and here we must (eventually) use our assumption about the structure of the integer n . Still, we can take a few more steps without the assumption.

So to begin, we choose an integer $1 < a < n$ at random. If $1 < \gcd(a, n) < n$, we were very lucky and have a nontrivial factorization. If not, then $\gcd(a, n) = 1$. Now what?

If p is a prime with $p \mid n$ (remember we are assuming n is composite), then since $\gcd(a, n) = 1$ we know $p \nmid a$, so by Fermat's little theorem or Euler's theorem (Theorems 4.18 and 4.21, respectively), we have that $a^{p-1} \equiv 1 \pmod{p}$. Said another way, $p \mid (a^{p-1} - 1)$, so since $p \mid n$, we have $p \mid \gcd(a^{p-1} - 1, n)$; in particular $\gcd(a^{p-1} - 1, n) > 1$. Great, but we still don't know p . True, but if M is any integer so that $(p-1) \mid M$, we would also have $a^M = (a^{p-1})^{M/(p-1)} \equiv 1 \pmod{p}$, so that $p \mid \gcd(a^M - 1, n)$. So what is the punch line?

We hope there is a prime $p \mid n$ so that $p-1$ is the product of small primes to small powers. Choose a bound B and compute (a table of) prime powers less than or equal to B :

$$2^{e_2} \leq B, \quad 3^{e_3} \leq B, \quad 5^{e_5} \leq B, \quad \dots, \quad p_r^{e_r} \leq B.$$

For example with $B = 11$, we would have computed $2^3, 3^2, 5^1, 7^1, 11^1$. Put $M = 2^{e_2} 3^{e_3} \dots p_r^{e_r}$, and compute $\gcd(a^M - 1, n)$. If $p-1$ is the product of small primes to small powers, then $p-1$ will divide such an M for B large enough, forcing $\gcd(a^M - 1, n)$ to be divisible by p . So the key to this method is to generate M and to compute $\gcd(a^M - 1, n)$.

At first blush, this looks like it might be computationally intense, but it is not. We know that if $a \equiv b \pmod{n}$, then $\gcd(a, n) = \gcd(b, n)$, so we need not compute $a^M - 1$ exactly; we simply compute

it modulo n , for which we already know of fast methods (modular exponentiation) by expanding the exponent M in binary and then performing successive squarings.

The General Pollard Algorithm. Given an odd, composite number n and an initial bound B , proceed with the following.

- (1) Find primes and prime powers $p_i^{a_i} \leq B$, $i = 1, 2, \dots, r$. This can be done via known tables of primes, sieving, and other fast methods.
- (2) Choose an integer $1 < a < n$ at random. If the $\gcd(a, n) > 1$, we have found a nontrivial factor; otherwise, $\gcd(a, n) = 1$ and Euler's theorem applies.
- (3) For $M = p_1^{a_1} p_2^{a_2} \cdots p_r^{a_r}$, compute $a^M - 1 \pmod{n}$.
- (4) Test $d = \gcd(a^M - 1, n)$. If $1 < d < n$, we have succeeded; otherwise, this iteration has failed.

After a couple of examples, we give options for how to proceed in the event of a failure.

Example 4.25. Let $n = 246082373$. Choose $a = 2$. Since n is odd, we know that $\gcd(a, n) = 1$. Choose an initial bound $B = 8$. Then $M = 2^3 \cdot 3 \cdot 5 \cdot 7$. We compute $\gcd(a^M - 1, n) = 1$ (using Euclid's algorithm). We have failed to find a nontrivial divisor, so we increase B to 10. Now $M = 2^3 3^2 \cdot 5 \cdot 7 = 2520$. We find that $\gcd(a^M - 1, n) = 2521$, a prime! That is, $a^M = a^{p-1} \equiv 1 \pmod{p}$ (by Fermat) and so $p = 2521$ divides both n and $a^M - 1$. Using this divisor, we factor $n = 2521 \cdot 97613$, both of which are primes.

Example 4.26. Let's take our earlier example from above, $n = 1234567890123456$. We first peel off any powers of 2, leaving us with $n = 2^6 \cdot 19290123283179$. Again we can choose $a = 2$ and set a modest bound of $B = 8$, so $M = 2^2 3^2 \cdot 5 \cdot 7 = 2520$ as above. We compute $\gcd(a^M - 1, 19290123283179) = 147 = 3 \cdot 7^2$, so we are left to factor $19290123283179/147 = 131225328457$. We know without trying that $\gcd(a^M - 1, 131225328457) = 1$, so we need to increase B , but by how much? That is the game: when to keep going and when to give up? Indeed for $B \leq 112$, the gcd will equal 1. Fi-

nally, for $B = 113$ and $M = 2^6 \cdot 3^4 \cdot 5^2 \cdot 7^2 \cdot 11 \cdot 13 \cdot 17 \cdot 19 \cdot 23 \cdot 29 \cdot 31 \cdot 37 \cdot 41 \cdot 43 \cdot 47 \cdot 53 \cdot 59 \cdot 61 \cdot 67 \cdot 71 \cdot 73 \cdot 79 \cdot 83 \cdot 89 \cdot 97 \cdot 101 \cdot 103 \cdot 107 \cdot 109 \cdot 113 = 955888052326228459513511038256280353796626534577600$, we find $\gcd(2^M - 1, 131225328457) = 435503$ a prime, and $131225328457/435503 = 301319$, a prime as well, which finally gives

$$n = 1234567890123456 = 2^6 \cdot 3 \cdot 7^2 \cdot 301319 \cdot 435503.$$

The key here (which we cannot know beforehand) is that the prime $p = 435503$ has the property that $p - 1 = 2 \cdot 41 \cdot 47 \cdot 113$ (the product of relatively small primes and prime powers), but which still pushed the bound B to 113. Of course that is better than the other prime $p = 301319$ for which $p - 1 = 2 \cdot 150659$ which would have pushed B to 150659(!).

Finally, we give an example where we cannot fix things by increasing the bound.

Example 4.27. $n = 2047 = 2^{11} - 1$. We choose $a = 2$ and $B = 10$, so $M = 2^3 3^2 \cdot 5 \cdot 7 = 2520$, and we find $\gcd(a^M - 1, n) = 1$, so we increase B to 11. Now $M = 2^3 3^2 \cdot 5 \cdot 7 \cdot 11$, and we find $\gcd(a^M - 1, n) = 2047 = n$, and increasing B can't help. As we increase B , it will always be the case that M is divisible by 11, so write $M = 11k$. Now we observe that $\gcd(2^{11k} - 1, 2^{11} - 1) = \gcd(2^M - 1, n) = n$ for any k by using the polynomial identity $x^k - 1 = (x - 1)(x^{k-1} + x^{k-2} + \dots + x + 1)$, and substituting $x = 2^{11}$. We see $2^{11k} - 1 = (2^{11} - 1)(2^{11(k-1)} + \dots + 1) \equiv 0 \pmod{2^{11} - 1}$, so our alternative is now to change the choice of a to, say, $a = 3$. This type of failure occurs infrequently. Indeed, if n is divisible by at least two odd primes, the probability that a randomly chosen value for a will fail is less than one-half; see [CP05].

In summary, Pollard's $p-1$ test works if there is a prime p dividing our composite number n so that $p - 1$ is the product of small primes to small powers, and it depends on the Fermat's little theorem, a special case of Lagrange's theorem, which in this case said that for any element $a \in U_p$, $a^{|U_p|} = 1$ in U_p . Our eventual goal is to replace U_p with the group of points of an elliptic curve which will afford us much greater flexibility in attacking the problem of factorization.

Chapter 5

Public-Key Cryptography and RSA

“We stand today on the brink of a revolution in cryptography” is the opening line to a seminal paper in cryptography [DH76] written by Whitfield Diffie and Martin Hellman. The genesis of public-key cryptography is generally attributed to Diffie, Hellman, and Ralph Merkle [Mer82], and the first practical method of implementation of a public-key cryptosystem is credited to Ronald Rivest, Adi Shamir, and Leonard Adleman [RSA78].¹ And while that line was written 40 years ago, it remains timely today. It seems each day brings us to a new brink in terms of the need for ever higher levels of digital security and secure communication. One interesting reality is that perfect security is possible with the use of so-called *one-time pads*; unfortunately, their use is simply impractical in the context and scale with which we must now communicate. While cryptography has a long and rich history well worth exploring, we shall restrict ourselves to this more recent era ushered in by Diffie, Hellman, et al.

¹The exact history is more complicated with some of these notions originating in earlier research by British cryptographers that was only declassified in 1997.

5.1. A Brief Overview of Cryptographic Systems

For better or worse, everywhere we go today, we are often in search of WiFi hotspots. A large segment of the population needs to check everything, seemingly all the time. Perhaps that's a bit hyperbolic, but do you recall how easy it was to get a seat in a coffee shop before WiFi was ubiquitous?

Perhaps we seem a bit off track. But as we sit in the coffee shop with our phone, tablet, or laptop, many or most of us simply connect to the local wireless network and go about our business. We are pre-occupied with the data that moves back and forth across that network, but not so much about how that actually takes place. Perhaps you are in the midst of a delicate online chat or are writing a sensitive letter to someone. How would you feel if every word you typed was said aloud over a public address system? It might give you pause. But of course the reality is that is exactly what's happening, just in a way that is not easily perceived by those in the room—at least those listening just with their ears. But your electronic device is very happily broadcasting everything you type across the airways to any willing receiver, in particular to that wireless router in the coffee shop, but to everyone else as well. You are effectively the radio tower, and anyone with a properly tuned radio can listen to what you are broadcasting.

Now maybe what you are doing on your device is checking your credit card account or making an online purchase, and you are happily entering your credit card number, passwords, PINs, and so on. Remember, you are the radio tower and any interested party can listen in with very little trouble. Now it is not our intention here to proselytize the concept of internet security, but issues concerning the lack of it certainly do pop up in the news with alarming frequency. So what's the point? We tend to think of cryptography as the stuff of spies and military communication, and of course in part it is, but its use has become as ubiquitous as the WiFi hotspot. It affects our daily lives and will continue to be an essential part of a modern world, so it seems a topic well worth understanding.

As a start, let's consider a simplified version of how cryptography is applied to internet transactions. There you are sipping that latte, and staring at a page on Amazon. No, not what you were looking for; search again. Still not right. Finally you have found what you want. Has your web browsing been secure? No, but then it did not really need to be unless you were concerned about others knowing your browsing habits. But now it is time to order the item, and your vendor asks if you would like to sign on to their secure server. You click the appropriate button without thinking, almost as easily as the "agree" button on all those end user licensing agreements (EULAs) you see when installing a new app on your phone or laptop. What happens when you click the option to do a secure sign on? Well, it's secure, right? And that means what exactly? Well, it's encrypted so you can enter your credit card information safely. Let's face it—what happens is magic.² Just how does Amazon know who you are? More to the point, how do you know you are really talking to Amazon—and how does it set up this encrypted tunnel through which you are going to conduct your secure transaction? There are more details than we really want or need at this point, so let's take a light pass at understanding some of the pieces of what is known as Transport Layer Security (TLS).

How is this encrypted tunnel going to be established between the server (Amazon) and the client (your web browser) when neither of you knows each other? And what is an encrypted tunnel? The first fact to acknowledge is while public-key cryptography is an indispensable piece of this complicated process of secure internet communication, its role is actually quite small in the sense that the bulk of information which is encrypted and transferred over the internet is encrypted using a secret-key cryptographic procedure, and not a public-key cryptographic system. In current use in the United States (as well as to a large extent internationally) is the algorithm Rijndael, chosen in 2002 as the algorithm to be used to implement the Advanced Encryption Standard (AES) by the National Institute of Standards and Technology and approved by the U.S. government for

²Arthur C. Clarke: "Any sufficiently advanced technology is indistinguishable from magic."

official use in commerce. Algorithms such as Rijndael are symmetric-key algorithms, meaning that the same key is used by sender and recipient to encrypt and decrypt data.

The first major obstacle which presents itself in using any secret-key cryptosystem is key distribution. In the past, when the number of people who needed to communicate securely was small and, in particular, when they were known to each other, conveying/exchanging keys was a manageable task. But now Amazon must exchange keys with every single patron who wants to place an order, and whether you have done business before with Amazon, each session gets a new key for you and Amazon to share for this transaction.

So when your web browser (the client) knocks on Amazon's door (the server), and asks to transact business, the first question the client asks is are you really who you claim to be (and server's credentials are presented to and verified by the client). Then the question is what kind of secret-key system would your client like to use? The server offers a number of options. Your web browser replies, and at this point a secret key is created jointly by both client and server. The details of how this key exchange (key creation) is accomplished will be covered in the next chapter when we have a bit more background, but this is one large part of the use of public-key cryptography as outlined by Diffie and Hellman. Once both the server and client have agreed on a key and a secret-key encryption method, secure communication between the client and server begins.

To give a brief introduction to public-key cryptography, we return to the problems that Diffie, Hellman, and Merkle sought to solve: how do we establish private communication between parties previously unknown to each other (which we have discussed in broad strokes above), and how do we establish the authenticity and origin of a digital document. In this chapter, we shall focus on RSA as an implementation of public-key cryptography since the algorithm relies only on the small bit of modular arithmetic we have already established. We shall leave the details of the Diffie–Hellman key exchange to the next chapter when we have discussed the notion of primitive roots and discrete logarithms.

Let's set a bit of notation. We let P denote a plaintext (i.e., an unencrypted) message, let C denote the encrypted ciphertext, let E be an encryption procedure, and let D be the decryption procedure paired with E .

Any standard cryptographic system embodies the following properties:

- For any plaintext message P , $P = D(E(P))$, that is, if you encrypt the message P and then decrypt it, you return to the original message.
- Ideally, both $C = E(P)$ and $P = D(C)$ are very fast to compute, and of course P should be hard to derive from C .

For a public-key cryptosystem, we have some additional requirements:

- A significant portion of the “public” part is that revealing E does not reveal an easy way to deduce D , meaning both that only the owner of D can decrypt messages encrypted by her procedure E , and that only the owner has an efficient way of producing D , given E . So this is very different than the symmetric-key system we discussed in establishing an encrypted tunnel for web transactions; here the encryption key is made public, and it is presumed that the (different) decryption key is very hard to deduce.
- The other feature, which is essential for authentication of both the message and sender, is that for every message M , $E(D(M)) = M$; that is, if you first apply the decryption procedure to a message and then the encryption procedure, you return to the original message. Combined with the properties of E and D under a general cryptosystem, we see that E and D are (mathematically) two-sided inverse functions.

In what follows we present successively more honest representations of how any public-key cryptosystem (PKCS)—in particular RSA—is used. We start with overly simplified models to illustrate the essential features of a PKCS.

As is standard in the literature concerning secure communication, our two players are always Alice and Bob. Each individual generates a pair of encryption and decryption algorithms, E_A, D_A, E_B, D_B . Each “publishes” their encryption algorithms in a publicly accessible repository. Alice wants to send an encrypted message to Bob. She generates her plaintext message P_A , encrypts it with Bob’s public encryption algorithm E_B , and sends $E_B(P_A)$ to Bob over whatever insecure channel she likes. Bob receives the message and extracts $P_A = D_B(E_B(P_A))$ using his decryption procedure D_B . So this is just like a secret-key cryptosystem, except for the new ability to send the message to anyone with a published encryption key. No prior contact is required.

The issue of authentication is another critical issue in an age when banking and legal transactions must take place over the internet. It is here that we see how easily a PKCS facilitates the signing of electronic documents. Once again, we first take a simplistic approach. Alice wishes to send a message to Bob, which Bob can then prove to a third party that the message is indeed from Alice and also that the message has not been altered.

Alice generates her message P_A and (effectively) signs it electronically by creating $S_A = D_A(P_A)$, that is, she uses her private decryption procedure and applies it to her plaintext message. She then sends this signed packet to Bob, as $E_B(S_A)$, using Bob’s publicly available encryption algorithm. Bob receives the message from someone whom he believes is Alice and recovers $S_A = D_B(E_B(S_A))$ by applying his private decryption algorithm to it. Now Bob uses Alice’s public encryption algorithm E_A to recover $P_A = E_A(S_A) = E_A(D_A(P_A))$. So at this point, Bob has the message in plaintext that Alice intended for him to have. Bob can now hand P_A and S_A to a judge who verifies the message is from Alice since only Alice’s public E_A can undo the signed document $S_A = D_A(P_A)$. The judge is also confident that the message P_A has not been tampered with, since to substitute a new message P'_A for P_A , the forger would have to produce $D_A(P'_A)$, which he cannot. Of course this points out how crucial it is to keep the decryption procedure D_A private since having D_A allows anyone to forge documents appearing to come from Alice.

In theory, the procedure above works perfectly fine for authentication purposes. In reality, the process of signing an entire document by using public-key decryption procedure has certain practical drawbacks. While there are issues concerning the amount of computation time and memory a PKCS requires, perhaps a simpler constraint comes from the fact that, in general, public-key cryptosystems are block ciphers, meaning that in order to encrypt a large document, the document must be broken into blocks and each block encrypted successively. So a signature of a long document could consist of perhaps thousands of signed pieces each of which would have to be verified.

Instead, what is typically done is for Alice to take her message and pass it through a so-called *hash* function (which we shall discuss more in depth below), but for now suffice it to say that two of the properties of a hash function are that they take documents of arbitrary size and reduce them to a string of characters of a short, fixed length, and that changing even a single character of the original long document produces a radically different hash value. So to attach a signature of her document (which will be verified), Alice first computes its hash: $H_A = \text{hash}(P_A)$, to serve as a digital fingerprint. She then signs the hash as she did before with the message $S_A = D_A(H_A)$, sending both the plaintext message P_A and signature S_A to Bob. Bob then verifies the authenticity of P_A by taking the publicly available hash function, forming a hash of his received copy of P_A , and then comparing it to $E_A(S_A) = E_A(D_A(H_A)) = H_A$.

To maintain the thread of our discussion, we leave hash functions temporarily and return to public-key cryptosystems and their implementation by RSA.

5.2. RSA

One of the most important questions about any cryptographic system regards its inherent security. How hard is it to break? The security of every public-key cryptosystem depends upon a task that is easy to do if one has privileged information (a trap door) but is difficult to do otherwise. For RSA, the easy task is finding two large primes and multiplying them together; the hard task is factoring their product. Slightly more to the point, RSA has as its basis the following idea.

Take two large primes $p \neq q$ and form $n = pq$. What is $\phi(n)$, where ϕ is the Euler totient function from the previous chapter? If we know p and q , the answer is trivial: $\phi(n) = \phi(p)\phi(q) = (p-1)(q-1)$. On the other hand, if we do not know how to factor n , this is very hard. Indeed, we shall show that being able to compute $\phi(n)$ is equivalent to knowing how to factor n .

Before continuing, it should come as no surprise that the appearance of the RSA algorithm was the impetus for a huge amount of research into the problems of how to compute very large primes (primality testing) and how to factor integers efficiently. So, in particular, if one had a fast way to factor integers, one could easily defeat RSA. While we have chosen in this book to talk more about cryptography, which, by definition, is the study of mathematical solutions to privacy and authentication problems, it is only one half of the subject known as cryptology; the other half of cryptology is cryptanalysis, the art of breaking cryptographic systems. This is an equally fascinating area to explore, but we shall discuss only general themes that characterize it.

5.2.1. Implementation of RSA. Let us now describe precisely how to implement RSA. To be somewhat more accurate, we shall explain how to implement what some would call *textbook RSA*, meaning that RSA is not really deployed as indicated below for security reasons. We shall say a bit more about this in the section on security of RSA, but for now the textbook implementation fits nicely with the mathematical concepts we have developed so far.

As above, choose two large distinct primes p and q , form their product $n = pq$, and compute $\phi(n) = (p-1)(q-1)$. The integer n is referred to as the RSA modulus. Assume that your plaintext message P has been converted to an integer between 0 and $n-1$. This can be done in many ways, starting perhaps first with some translation of your natural character set into integers, as we primitively did to use the Caesar cipher. Of course, the complication of accurately rendering messages with many special characters and in different languages is now handled by Unicode character encodings (such as the one most prevalent on the web, utf-8), but we sweep these complications under the rug, since they are not our main focus. We assume the message

we wish to encrypt has been converted to an integer P between 0 and $n - 1$.

Next, choose (more or less at random) a positive integer e with $\gcd(e, \phi(n)) = 1$. This turns out to be quite fast since one can use Euclid's algorithm to test whether a randomly chosen integer is relatively prime to $\phi(n)$. If not, choose another and test again. For example, any prime number not dividing $\phi(n) = (p - 1)(q - 1)$ will do. There are some practical considerations that slightly constrain the choice of e , but we shall say more about them in the section on the security of RSA. Once the number e is chosen, we know from Proposition 3.25 that there is a unique $d \pmod{\phi(n)}$ so that $ed \equiv 1 \pmod{\phi(n)}$; moreover, it is Euclid's algorithm (§3.3) which finds this value quickly.

So let's summarize what parameters we have so far.

- We have chosen primes $p \neq q$, and we set $n = pq$.
- We have computed $\phi(n) = \phi(p)\phi(q) = (p - 1)(q - 1)$.
- We have chosen an integer e so that $\gcd(e, \phi(n)) = 1$.
- We have used Euclid's algorithm to find an integer d with $ed \equiv 1 \pmod{\phi(n)}$.

The RSA algorithm is given quite simply by the following.

- Encryption: $C = E(P) \equiv P^e \pmod{n}$ (plaintext to ciphertext)
- Decryption: $P = D(C) \equiv C^d \pmod{n}$ (ciphertext to plaintext)

That is, given a plaintext message P (represented as an integer $0 \leq P < n$), we compute the ciphertext message C by raising P to the e th power and reducing modulo n . The process of exponentiation is made efficient via the process of fast modular exponentiation introduced in Example 4.24. The process for decryption is to exponentiate C and reduce modulo n .

Example 5.1. Let's use a 27-letter alphabet with $A \longleftrightarrow 0$, $B \longleftrightarrow 1$, \dots , $Z \longleftrightarrow 25$, space $\longleftrightarrow 26$. We shall convert between alphabet and numeric plaintext messages using a common scheme: encode as a

base 27 number with a certain block size, in our case 5. We do this as follows: Suppose we want to convert the message “Groups are fun”.

First we break up our plaintext message into blocks of length 5, padding the last block if necessary: “Group” “sare” “funX”. Note we will not distinguish between upper and lower case, but this would easily be done by expanding the size of our alphabet.

“Group” \mapsto 06, 17, 14, 20, 15; “sare” \mapsto 18, 26, 00, 17, 04; “funX” \mapsto 26, 05, 20, 13, 23, where the numbers represent the base 27 digits. We now encode these as base 27 numbers:

$$\begin{aligned} \text{Group} &\mapsto 06, 17, 14, 20, 15 \\ &\mapsto 27^4(06) + 27^3(17) + 27^2(14) + 27^1(20) + 27^0(15) \\ &= 3534018, \\ \text{sare} &\mapsto 18, 26, 00, 17, 04 \\ &\mapsto 27^4(18) + 27^3(26) + 27^2(00) + 27^1(17) + 27^0(15) \\ &= 10078170, \\ \text{funX} &\mapsto 26, 05, 20, 13, 23 \\ &\mapsto 27^4(26) + 27^3(05) + 27^2(20) + 27^1(13) + 27^0(23) \\ &= 13930835. \end{aligned}$$

To proceed, we note that all the plaintext messages P will satisfy $0 \leq P < 27^5 = 14,348,907$, so when choosing primes p, q for our RSA modulus $n = pq$, we must make sure $n \geq 27^5$.

Exercise. Suppose we choose primes p and q , so that $n = pq = 59753237$. With the knowledge of those primes, we compute $\phi(n) = (p - 1)(q - 1) = 59737740$ and choose the common encryption exponent $e = 2^{16} + 1 = 65537$ (the last known Fermat prime).

- (1) Find the primes p and q . (This is not necessary to break the code, but it reinforces that knowing $\phi(n)$ is equivalent to factoring n . See the section on security of RSA for a hint.)
- (2) Find the decryption exponent.
- (3) Using the base 27 encoding scheme as above, decrypt the message consisting of two blocks of numerical ciphertext, i.e., given as $C = P^e \pmod{n}$: 10881312 41465338.

5.2.2. Verifying the Algorithm. We need only verify that for

$$0 \leq M < n, \text{ we have } M^{ed} \equiv M \pmod{n}.$$

This will tell both that $D(E(P)) = P$ and $C = E(D(C))$ for all plaintext messages P and ciphertext messages C . With some easy exceptional cases to handle, the key to understanding why RSA works is Euler's theorem (Theorem 4.21). The condition that $ed \equiv 1 \pmod{\phi(n)}$ is equivalent to saying that $ed = 1 + k\phi(n)$ for some integer k . Note that since n is the product of two primes, the only possibility for $\gcd(M, n)$ is $1, p, q$, or $n = pq$, where the last is precluded since we have assumed $M < n$. The cases where the gcd is p or q are completely analogous, so we have two cases to consider.

Case 1: $\gcd(M, n) = 1$. By Euler's theorem, $M^{\phi(n)} \equiv 1 \pmod{n}$, so

$$M^{ed} \equiv M^{1+k\phi(n)} = MM^{\phi(n)k} \equiv M \pmod{n}.$$

Case 2: $\gcd(M, n) > 1$. Then from our comments above we may assume that $\gcd(M, n) = p$. Since p and q are relatively prime, it is easy to see that $M^{ed} \equiv M \pmod{n}$ if and only if $M^{ed} \equiv M \pmod{p}$ and $M^{ed} \equiv M \pmod{q}$, so in particular we need only check these later conditions.

Since $p \mid M$, we have $M^{ed} \equiv 0 \equiv M \pmod{p}$. Now $\gcd(M, n) = p$ and q a prime, $q \neq p$ means that $\gcd(M, q) = 1$, so $M^{(q-1)} \equiv 1 \pmod{q}$ from Fermat's little theorem (Theorem 4.18), or alternatively Euler's theorem (Theorem 4.21). Thus we have

$$M^{ed} \equiv M^1 M^{(q-1)(p-1)k} \equiv M(M^{(q-1)})^{k(p-1)} \equiv M \pmod{q}.$$

So we have established that $M^{ed} \equiv M \pmod{p}$ and $M^{ed} \equiv M \pmod{q}$, which implies $M^{ed} \equiv M \pmod{n}$ as required.

Exercise. Compute the probability that a plaintext message M is not prime to $n = pq$. If we wanted to ensure that our messages were always relatively prime to n , what could be done?

5.2.3. Security of RSA. Now that we have verified that RSA will function as a public-key cryptosystem, we consider the security of RSA. As we said earlier, public-key cryptosystems depend upon problems which are easy to do if given privileged information but are believed to be hard otherwise. For RSA, the task believed to be hard is factoring large integers (hundreds of digits), and so far it has proven to be a hard problem, but it is not provably a hard problem. So research will always continue on new methods for how to factor integers or more generally break RSA. As we first mentioned in the section where we talked about the implementation of RSA, in this section, we first talk about the security of *textbook RSA*, leaving a more technical discussion to the end of this chapter.

In an RSA public-key setup, both integers n and e are a matter of public record; in the general description of a PKCS, this is making the encryption procedure E public. Recall that doing so is precisely what is needed for someone to send you an encrypted message without actually having met you.

How can one defeat RSA? That is, given that you know e and n , if you intercept a piece of ciphertext C , how can you recover the plaintext P which was used to create the ciphertext $C \equiv P^e \pmod{n}$? One way would be to recover the decryption exponent d (satisfying $ed \equiv 1 \pmod{\phi(n)}$), and the other way is to extract e th roots modulo n .

It is generally believed that extracting e th roots modulo n is approximately as difficult as (and often dependent upon) factoring n , so we shall make only a few comments about this approach. First, the belief that extracting e th roots is as hard as factoring does not mean that it is, and efforts continue on precisely this problem independent of efforts to factor large integers. See [JNT07] for some recent work.

On the other hand, one cannot be too cavalier about the choice of e . When we set up RSA, we said you could choose e (more or less at random) to be any integer relatively prime to $\phi(n) = (p-1)(q-1)$. Certainly, $e = 1$ would be a poor choice, since then $C = P$. The value $e = 2$ is precluded since $\phi(n)$ is even. But what about $e = 3$, presuming it is relatively prime to $\phi(n)$? Seems reasonable. We note that if your exponent e (which is public) is very small and n is large compared to your plaintext P , one could simply try taking e th roots in the real numbers to attempt to recover the plaintext. More precisely, if $P^e < n$, then taking “real” e th roots will recover P .

A final remark is that an oft-chosen exponent is $e = 65537 = 2^{2^4} + 1$, the last known Fermat prime, if it doesn’t divide $p-1$ or $q-1$. And if it does, you could choose new p and q ; as it turns out, finding large primes is an “easy” problem, meaning there is a polynomial time deterministic algorithm to verify if a given integer is a prime, but more on that later. Let’s see why that value of e is popular; remember it is public, so popular is alright.

We first examine the objection above. Let’s say n is large, $n = 10^{400}$, which is a four hundred digit number, the product of two large primes. When is $P^{65537} < 10^{400}$? Well, when $P < 10^{400/65537} \approx 1.01$, and since plaintext is encoded as integers, most of our messages will escape discovery by taking real e th roots. So what else is nice? The fact that $e = 65537 = 2^{2^4} + 1 = 2^{16} + 1$ means that computing P^e by fast modular exponentiation is *really* fast: compute

$$P, \quad P^2 = (P^{2^0})^2, \quad P^{2^2} = (P^{2^1})^2, \quad P^{2^3} = (P^{2^2})^2, \\ P^{2^4} = (P^{2^3})^2, \quad \dots, \quad P^{2^{15}} = (P^{2^{14}})^2, \quad P^{2^{16}} = (P^{2^{15}})^2,$$

and $P^e = P \cdot P^{2^{16}}$, requiring only 17 multiplications!

The other attack on RSA is to deduce d , satisfying $ed \equiv 1 \pmod{\phi(n)}$. For that you need to know the value of $\phi(n)$, and we

now show that knowledge of $\phi(n)$ yields knowledge of p and q , so the difficulty of deducing $\phi(n)$ is equivalent to the difficulty of factoring n .

We begin with the simple observation that

$$\phi(n) = (p-1)(q-1) = pq - (p+q) + 1 = n - (p+q) + 1,$$

so knowing n and $\phi(n)$ tells us the value of $n - \phi(n) + 1 = p + q$. Also (assuming without loss that p is greater than q),

$$p - q = \sqrt{(p-q)^2} = \sqrt{(p+q)^2 - 4pq} = \sqrt{(p+q)^2 - 4n}$$

is known. Given both $p + q$ and $p - q$, we immediately deduce p and q , so knowledge of $\phi(n)$ is equivalent to factoring $n = pq$.

Exercise/Project. Explore real implementations of RSA, e.g., PKCS#1 v2.2: RSA Cryptography Standard [Lab12].

Emphasis should be that, in reality, one is not encrypting a generic plaintext message, but an AES key to transmit securely. On the plus side, this is already a number, so encoding is not the issue. The problem is that an AES key is short relative to required modulus size. This means the AES keys are not randomly distributed in the RSA key space, producing a vulnerability and necessitating certain padding schemes.

5.3. Hash Functions

While hash functions are of critical importance in cryptographic applications, their use is actually quite broad. Let's quickly survey some places where they are used.

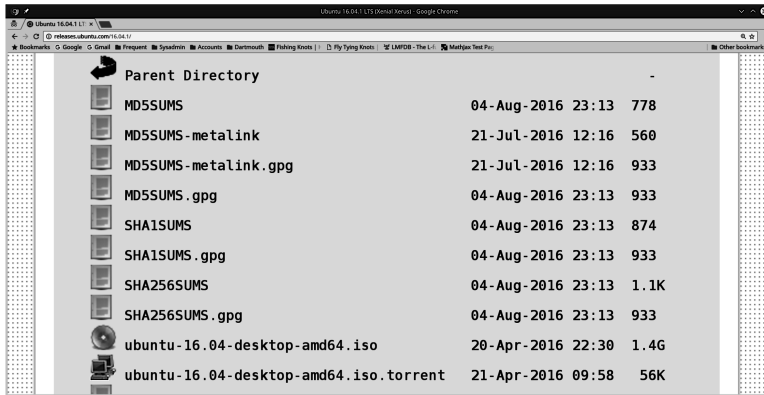
You have probably encountered hashes implicitly or explicitly in your daily online experiences. Often you download some piece of software to install on your phone or computer: a software upgrade, music-streaming software, cloud storage applications, photo-editing software, and so on. The download begins, and when you start to install the software, you will often see a message of the type "Verifying the integrity of the download". Now we sort of understand what that is supposed to mean, but really, how do you verify the integrity of the download? Frankly, it is another miracle of science that all those bits travel through all those routers from their source to your laptop and

arrive intact. As miraculous as the process is, it occasionally does cough, sputter, or stall, and the file on your laptop may not be an exact copy of the one on the server from which the download began. How do you know? Does it really matter? The answer to the first question is hash functions, and the answer to the second might be best given as another question: An on/off switch is often labeled 0/1, that is, with one single bit (binary digit) of data. Does the choice of that bit really make a difference?

While in many cases, such as those above, hashes have been relegated to the background, it is certainly easy to find them, often appearing quite prominently. Indeed you will often see many different types of them. They go by the names MD5sums (MD for message digest), SHA1 (SHA for secure hash algorithm—catchy right?), or SHA256 (no, that’s not version 256 of SHA, but an indicator that the output string is 256 bits long). In the section on digital signatures, we mentioned two features such functions are designed to have: First, they take files of arbitrary length and generate a short string that serves as a digital fingerprint of that file. Second, minor changes in the source file generate large changes in the resulting fingerprint, so the presence of a changed bit somewhere in the midst of an enormous file would be easily detected by its changed hash.

Consider Figure 5.1. It is a bit involved but aims to demonstrate several points. The displayed browser image (<http://releases.ubuntu.com/16.04.1/>) shows the content of a download site for Ubuntu® (a popular Linux distribution),³ and there for interested parties to download are “iso” images of their operating system. An iso image is a set of files in a format meant to be burned to a CD or DVD for installation on another computer. The file `ubuntu-16.04-desktop-amd64.iso` is 1.4G in size, so it is larger than a standard CD, and it contains a complex collection of files with many of them used to install software, drivers to operate the hardware on your computer, and so on. Before you install it, you might like to make sure the image is correct. So Ubuntu® has carefully provided hashes for all the large files you might download from this page, so you can compare the hash of your downloaded file with the original.

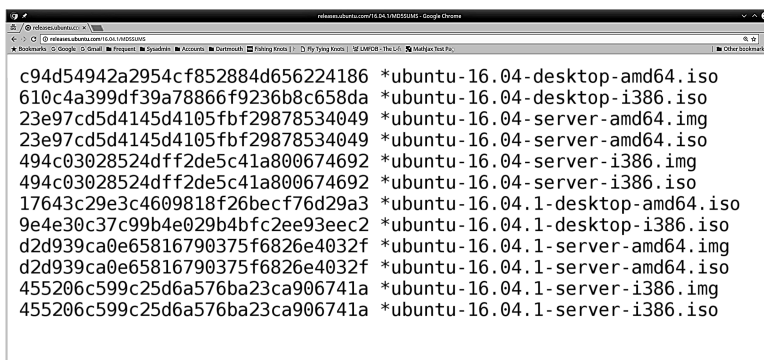
³Ubuntu is a registered trademark of Canonical Ltd.



File Name	Date	Time	Size
Parent Directory	-	-	-
MD5SUMS	04-Aug-2016	23:13	778
MD5SUMS-metalink	21-Jul-2016	12:16	560
MD5SUMS-metalink.gpg	21-Jul-2016	12:16	933
MD5SUMS.gpg	04-Aug-2016	23:13	933
SHA1SUMS	04-Aug-2016	23:13	874
SHA1SUMS.gpg	04-Aug-2016	23:13	933
SHA256SUMS	04-Aug-2016	23:13	1.1K
SHA256SUMS.gpg	04-Aug-2016	23:13	933
ubuntu-16.04-desktop-amd64.iso	20-Apr-2016	22:30	1.4G
ubuntu-16.04-desktop-amd64.iso.torrent	21-Apr-2016	09:58	56K

Figure 5.1

Figure 5.2 is a listing of the file MD5SUMS, which lists the MD5 hashes for various of the iso files.



```

c94d54942a2954cf852884d656224186 *ubuntu-16.04-desktop-amd64.iso
610c4a399df39a78866f9236b8c658da *ubuntu-16.04-desktop-i386.iso
23e97cd5d4145d4105fbf29878534049 *ubuntu-16.04-server-amd64.img
23e97cd5d4145d4105fbf29878534049 *ubuntu-16.04-server-amd64.iso
494c03028524dff2de5c41a800674692 *ubuntu-16.04-server-i386.img
494c03028524dff2de5c41a800674692 *ubuntu-16.04-server-i386.iso
17643c29e3c4609818f26becf76d29a3 *ubuntu-16.04.1-desktop-amd64.iso
9e4e30c37c99b4e029b4bfc2ee93eec2 *ubuntu-16.04.1-desktop-i386.iso
d2d939ca0e65816790375f6826e4032f *ubuntu-16.04.1-server-amd64.img
d2d939ca0e65816790375f6826e4032f *ubuntu-16.04.1-server-amd64.iso
455206c599c25d6a576ba23ca906741a *ubuntu-16.04.1-server-i386.img
455206c599c25d6a576ba23ca906741a *ubuntu-16.04.1-server-i386.iso

```

Figure 5.2

The first line lists the following hash.

```
c94d54942a2954cf852884d656224186 *ubuntu-16.04-desktop-amd64.iso
```

This 32-character string `c94d54942a2954cf852884d656224186` is actually the hexadecimal (base 16) representation of the 128-bit (base 2) output of the MD5 hash algorithm. This is what you actually compare against when verifying the integrity of a download. Admittedly, that is quite a bit to absorb, but before we finish with this example, let's note the files listed on the first image with a gpg suffix, like `MD5SUMS.gpg`. The gpg suffix stands for **Gnu Privacy Guard** (<https://www.gnupg.org/>), means that the file `MD5SUMS` has been digitally signed by the creators of that file. Now why would they do that?

You need to adopt a properly nefarious perspective to appreciate the reason. Would this not be the perfect opportunity to infect millions of computers by replacing the installation software with another which also installs software intended to compromise these systems? The perpetrator would download the original iso, modify select files, and recreate a new iso which they hope to get onto the download site. Now the problem is that the new hashes (MD5, SHA1, SHA256) won't match, so to thwart those who are careful about the software they install on their computers (we all are, right?), they would have to replace the file `MD5SUMS` with a new file by that name containing the hashes of the new files. But there is that gpg file, which is a signed version of the original, and by using the developer's public encryption key (gpg), the original MD5sums can be revealed. So now what shall our nefarious intruder do? Well, they would have to create a new file system with their modifications which would have the original MD5, SHA1, and SHA256 hashes. This leads to our need to understand other important requirements of a hash function. Oh, but to answer the obvious question, beating a single hash is highly improbable; beating three would be a truly impressive feat.

The idea of a hash function is that it should take as input a file (string) of arbitrary length and produce an output string of fixed (short) length. Common hash algorithms (e.g., MD5, SHA-1, SHA-2) have output lengths ranging from 128–512 bits. Their computation should be very fast. Let's understand what is happening mathematically. We have a function h whose domain is a set of strings of arbitrary length and whose codomain is a set of bit strings of bounded

length. This function is definitely not injective, which means there exist many inputs $m \neq m'$ for which $h(m) = h(m')$. Two such values $m \neq m'$ for which $h(m) = h(m')$ is called a *collision*. Imagine a hash function whose job it was to take the names of mathematicians attending a Joint Mathematics Meeting and assign them to rooms in a block of hotels. A collision would be two people showing up to occupy the same room.

Now we make more precise the requirements for a cryptographic hash function. The definitive source for secure hash standards are the NIST documents [NIS12a], [NIS12c]. We offer a short synopsis; see especially §4.1 of [NIS12a].

Properties that a cryptographic hash function should have.

- (1) It should exhibit *collision resistance*. That is, it should be computationally infeasible to find messages $m \neq m'$ with $h(m) = h(m')$. Note that mathematically, this will definitely happen. We are simply asking for the function h to be sufficiently complicated that it takes an unreasonable amount of computation to have high probability of finding such m, m' .
- (2) It should have *preimage resistance*. Given a randomly chosen hash value h_0 , it should be computationally infeasible to find a message m for which $h(m) = h_0$. Mathematically, one is asking for it to be computationally infeasible to find any element of the inverse image $h^{-1}(h_0)$. This is sometimes referred to as the *one-way* nature of the hash function.
- (3) It should have *second preimage resistance*. Given an input m with hash $h(m)$, it should be computationally infeasible to find a different value m' with $h(m) = h(m')$.

Now let's try to understand the significance of and differences among these requirements. Once again, the best perspective to adopt is one in which you are trying to accomplish something criminal.

Let's start in the middle with *preimage resistance*, since it stands somewhat alone. It is simply a part of life that we all have many

online accounts, each with its own peculiarity regarding login names, password construction, and frequency with which passwords must be changed. The most unsafe situation would be if passwords were stored as plaintext on the server that we wanted to access. It is better if (at least) the passwords are hashed. That way the stored passwords all have the same length, look quite random, and do not appear to reveal anything. Now when you log in, your entered password is hashed and compared to the stored hash of your password. If they match, you are in; otherwise, you're not. Now let's suppose that the hash function being used by the site had the property that it ignored the difference between uppercase and lowercase letters, and your password is `MdRimBF4e!` (My dog Riley is my Best Friend for(4) ever—and yes, `BF4e!` is a much better choice cryptographically than `BFF`). If your hash function ignored the case of passwords, then there would be 2^8 passwords which would hash to the same value as `MdRimBF4e!` since each alphabetic character in the password can either be upper or lower case. Said another way, that reduces the size of the keyspace of alphabetic passwords of length n by a factor of 2^n . So if you could easily find a (pass)word whose hash matched an existing one, you could gain access to an account without knowing the real password.

Turning to the first and third conditions, it appears at first blush (maybe even second), that *collision resistance* implies *second preimage resistance*—consider the contrapositive. Given an m and its hash $h(m)$, if it is computationally feasible to produce m' with $h(m') = h(m)$, then it is computationally feasible to produce a collision. The reason both of these requirements are here is that in giving our synopsis of the standards, we have not mentioned that the term “computationally feasible” has different meanings in each of these requirements: if one task is computationally feasible to do with a certain amount of effort, it does not necessarily mean the other is feasible with the same amount of effort. But this technical point is only a small difference between collision resistance and second preimage. The actual distinction is more interesting and is the subject of the so-called birthday paradox in probability.

The *birthday paradox* asks us to distinguish between two questions:

- (1) What is the probability that in a room with n people, at least one other person has my birthday (same month and day)?
- (2) What is the probability that in a room with n people, at least two will have the same birthday (same month and day)?

Second preimage resistance addresses the first of these questions, and *collision resistance* is the subject of the second. So first, what are the answers?

- (1) The probability that at least one other person in the room has my birthday, $P_{my}(n)$, is 1 minus the probability that nobody in the room has my birthday. The probability that my birthday is different than another individual is $\frac{364}{365}$ (assuming 365 days in a year) and, since we will assume that the birthdays among the n people in the room are random (so represent independent events), the probability that n peoples' birthdays differ from mine is $\left(\frac{364}{365}\right)^n$. So the probability that at least one among the n has my birthday is

$$P_{my}(n) = 1 - \left(\frac{364}{365}\right)^n.$$

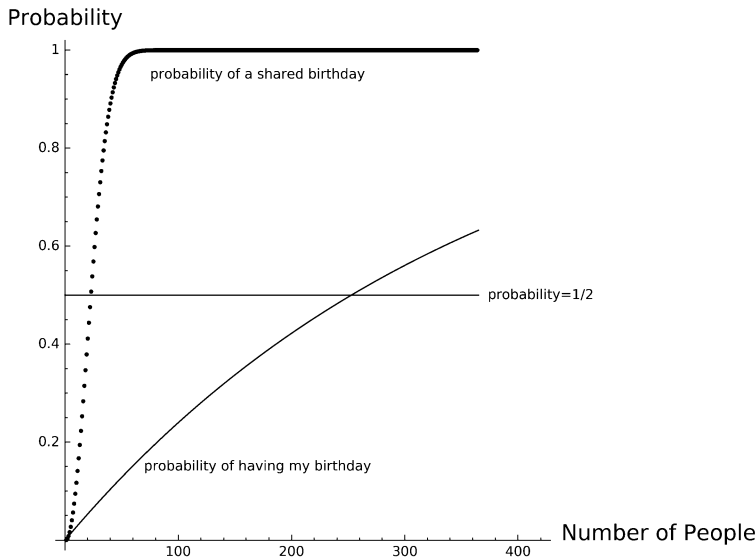
- (2) Let $P_{sh}(n)$ denote the probability that in a room with n people at least two share the same birthday. For sure (even accounting for leap years), we know that if $n > 366$, $P_{sh}(n) = 1$ by the pigeon-hole principle. The quantity $1 - P_{sh}(n)$ is the probability that no two of the n people in the room share the same birthday, so

$$1 - P_{sh}(n) = \frac{365}{365} \cdot \frac{364}{365} \cdot \frac{363}{365} \cdots \frac{366 - n}{365},$$

so

$$P_{sh}(n) = 1 - \frac{365}{365} \cdot \frac{364}{365} \cdot \frac{363}{365} \cdots \frac{366 - n}{365} = 1 - \frac{n! \binom{365}{n}}{365^n}.$$

In Figure 5.3 we graph these two probability functions (as a function of the number of people in the room).

**Figure 5.3**

We see that the probability of two people having a shared birthday rises above $1/2$ as soon as there are at least 23 people in the room, but it takes 253 people in the room before the probability that someone else has my birthday rises above $1/2$. So now let's reconnect these ideas with cryptography.

Let's suppose that Alice and Bob are in business and they create a contract between them, call it M . In the contract Bob agrees to do certain jobs for Alice, and Alice in turn agrees to pay Bob, and they jointly agree upon an amount for the service. Alice lives on the west coast and Bob on the east, so they want to handle the transaction electronically. Bob is happy with the contract M , so he creates a hash of it $h(M)$ (using one of the standard hash algorithms), and he then signs the hash with his private decryption key and sends M , together with $D_B(h(M))$, to Alice. Alice uses Bob's public-encryption key to recover $h(M)$. She has her own copy of the contract; let's call it M_A (even though it is supposed to be the case that $M_A = M$). She

computes $h(M_A)$ and sees that her hash matches Bob's, meaning she has signed a copy of the contract that she holds.

Now let's suppose that Bob is not totally happy with the contract: the compensation is inadequate, the deadlines are too tight, or any number of other things that he might want to change. So Bob wants to replace M with M' and send it to Alice and still have her think he has signed the original contract. What would he need to do? Well, he would have to create a new contract M' , so that $h(M') = h(M)$. This is exactly the context of second preimage resistance, or in the context of the birthday paradox, finding someone else with my birthday, which is generally a hard problem. Of course should Bob succeed, he can later present M' as the valid contract, since the hashes match.

Now creating an M' with the same hash as M is quite hard, so how could one use the birthday paradox? We are still a society in transition from paper transactions and contracts to electronic ones, and even when we don't actually print a document, we use software that "renders" the document in a manner convenient for us to examine. What am I getting at? There are many characters that can appear as part of a file which do not show up when you view or print the document unless you take extra steps. For example, when you tab for an indented line, you simply see the line has been indented, and it is not obvious whether that indentation occurred because of a tab or someone hitting the space bar five times. Of course as characters in the file, these are very different characters (and files). There are many other things that would be difficult to detect visually, like an extra space after a period at the end of a line or an extra comma. Additionally, anywhere there is white space, spaces and tabs and any unprintable characters could be added and make no *visible* difference in the document. So in this game, as Bob and Alice are negotiating the terms of the contract, Bob prepares two contracts M and M' . The contract M' is what Bob wants, and M is what Alice is offering. Bob now manipulates both contracts M and M' looking for a pair whose hashes $h(M)$ and $h(M')$ are equal. This is the situation of finding two people in the room with the same birthday, a significantly easier problem to solve as we have seen.

The use of the birthday paradox in this manner is one of the reasons that the length of hashes have increased over time. At lengths of 256, 384, 512 bits, probability is still on the side of guarding against collisions and preimage computations.

5.4. Breaking Cryptosystems and Practical RSA Security Considerations

In the modern use of cryptography, it is almost always the case that someone trying to break a system knows precisely what algorithms are being used: RSA, AES, ECC (elliptic curve cryptography), etc. The reason is twofold. In the past, people would often try to hide the exact mechanism used for encryption (e.g., rotor machines such as the Enigma from World War II) in the hopes that the lack of that knowledge would add to the security of the system. But, as we have said before, the problem is that if that information becomes known to the cryptanalyst and those using the encryption scheme are unaware of the breach, a great deal of information can often be gleaned leading to the potential collapse of the entire system. This is precisely what happened in the case of the Enigma machine; there are numerous books, articles, and even movies about the enormous efforts brought to bear to break this cryptographic system. There is a second reason algorithms are now almost always known, and that is one of verification. AES (Rijndael) is used for domestic and international commerce. All parties want to know that there is nothing buried in the depth of the algorithm that could be used against them.

There are standards by which the security of a cryptosystem is assessed; most of these terms go back at least to the article of Diffie and Hellman [DH76].

Methods of attacking a cryptosystem:

- (1) Brute-force search of keyspace.

All encryption/decryption algorithms have keys. In the case of Caesar-type ciphers, the key is simply the offset, so the keys are just the integers 0 to 25, and these 26 numbers form the keyspace. So if one knew it was a Caesar cipher

being used, a brute-force search of the keyspace would reveal a correct key almost instantaneously.

AES (Rijndael) allows keys of various lengths 128, 192, or 256 bits. At 256 bits, the keyspace has size 10^{78} , roughly on the order of the number of atoms in the observable universe. The point here is that a brute-force search of this keyspace—that is, trying key after key until the correct one is found—simply takes too much time, and the security of the system resides in that fact, mainly because keys can be changed with each new transaction.

(2) Ciphertext-only attack.

This means the cryptanalyst is given only encrypted messages. Depending upon the system, this may provide little or no information. For example, if your cipher were a simple substitution cipher (meaning using a fixed permutation of the alphabet but always encoding a letter in the same way), ciphertext messages are susceptible to frequency analysis and are often context sensitive making the decryption of a message a bit like the daily jumble.

(3) Known-plaintext attack.

This is more interesting in the sense that the cryptanalyst has copies of both the plaintext message and the corresponding ciphertext message from which to gain insight. Previously intercepted classified documents which are now declassified fall into this category.

(4) Chosen-plaintext attack.

This offers the cryptanalyst the most power. For example, if the encryption is a simple substitution cipher, then all that is needed to reveal the permutation is a single encoding of each letter of the alphabet. Of course this is just the simplest of substitution ciphers; one-time pads, which offer perfect security, are also substitution ciphers; it is just that the key has the same length as the message.

Modern cryptosystems need to be resistant to all of these types of attacks, so let's come back to the security of RSA. First, a simple

question. What are the keys? What is the keyspace? How would one attempt to mount a brute-force attack on RSA?

The information that a person or organization makes public in RSA is the value of modulus n and the encryption exponent e (so that a ciphertext is produced via $C \equiv P^e \pmod{n}$). We have said that to find the decryption key d is a matter of solving $ed \equiv 1 \pmod{\phi(n)}$, and knowing $\phi(n)$ provides as much information as factoring $n = pq$, so the security rests in how hard it is to factor n . Well, just how hard is it to factor a number? The answer to this question is not simply a matter of how large the integer n is, but also of its composition. After all, even though it is very large, few would have trouble factoring $n = 10^{12345}$, since the only prime factors are 2 and 5. So when we publish our pair (e, n) , we want to make sure n is hard to factor. Since our modulus n has the special form of the product of two distinct primes, we want neither prime to be small, so we look for primes whose size are approximately \sqrt{n} . Since it is such a vital matter, many organizations keep track of the state of the art in factoring and issue recommended guidelines for how large the modulus n should be. Recommendations made by the National Institute of Standards and Technology are in sections 5.6.1 and 5.6.2 of [NIS12b] (see also [Gir15]). Up until about 2010, it was felt that a 1024-bit modulus (n having size $2^{1024} \approx 10^{309}$) was safe from the factoring methods currently known, but those standards have been raised now to a modulus of size 2048 bits.

We shall talk more about factoring later in this book, but as we start to broaden our perspective as cryptanalysts, perhaps we should ask different questions. So factoring is hard. And that means its difficulty is measured in exponential time relative to the length of the number to be factored, so think of an n -bit number as requiring 2^n units of time. In contrast, determining whether a randomly generated integer is prime can be done in polynomial time relative to the length of the input, which is to say, this is very fast. So maybe, we can get lucky in searching the keyspace. We know that $n = pq$ where p and q are both primes about the size of \sqrt{n} , so supposing n was a 1024-bit modulus, we would be looking for two primes of size 512 bits, that is, between 2^{512} and 2^{513} . The exercises below help decide whether the

keyspace is large or small and how easy or hard it is to choose two primes from which to construct your RSA modulus.

Project. Suppose you want a 1024-bit RSA modulus, so you want two primes 512 bits long. About how many primes are there of that approximate size? What are the chances that randomly chosen odd integers of that length will be prime?

Exercise. The observations above also provide an answer to another important security question. We know that people often use $e = 65537$ as an encryption exponent. What if they also chose the same value of n ? What would be the security implications?

A technical issue with RSA. If it were the case that the set of plaintext messages to be encrypted by RSA were well distributed throughout the plaintext space, RSA would be as secure as indicated, protected by the difficulty of factoring a 1024-bit (or now a 2048-bit) modulus. But the reality is that RSA is used to encrypt symmetric keys, sometimes as small as a 128-bit AES key or to sign the hash of a document, but in the end it used to encrypt something on the order of 128, 256, 385, 512 bits. This means the plaintexts are not well distributed, which makes RSA more vulnerable.

As it seems in all things, there is the theoretical side and then the practical implementation side. So the practical side of RSA, described in [Lab12], first pads the small plaintext to help the messages fill out the plaintext space and then uses RSA. On the decryption side the padding is removed, leaving the desired message.

Chapter 6

A Little More Algebra

Our exposure to groups and rings has been somewhat narrow so far. We were long ago acquainted with the algebraic properties of the integers \mathbb{Z} but now know those properties give \mathbb{Z} the structure of a commutative ring. In Chapter 4 we introduced the sets \mathbb{Z}_n and U_n , the former a commutative ring (hence an abelian group under addition), and U_n an abelian group under multiplication. In fact U_n is precisely the set of elements in \mathbb{Z}_n that have multiplicative inverses, and to recognize that relation, U_n is often called the *unit group* in \mathbb{Z}_n . In the special case that $n = p$ is prime, we saw that $U_p = \mathbb{Z}_p \setminus \{0\}$. This property makes \mathbb{Z}_p an especially important and distinguished ring called a *field*. So the ring \mathbb{Z}_p (while finite) stands equivalent as an algebraic structure to \mathbb{Q} and \mathbb{R} and \mathbb{C} , all commutative rings in which every nonzero element is a unit (has a multiplicative inverse). We will henceforth recognize this special status by writing \mathbb{F}_p when we want to think of \mathbb{Z}_p as finite field. Finite fields are especially important in number theory and in the area of cryptography in particular. We shall see the pivotal role they play beginning with the next chapter.

On the other hand, sometimes we want to consider \mathbb{Z}_n (even \mathbb{Z}_p) not even as a ring, but just as an abelian group under addition. It can be a bit confusing about how we regard the set \mathbb{Z}_n , but usually the manner in which we choose to view it will be clear from context, and we shall see plenty of examples below.

6.1. Towards a Classification of Groups

An overarching endeavor in mathematics and in the sciences in general is to classify objects according to various criteria. Classification asks us to answer a subtle question: When are two objects the same? If we put two apples on a table and ask if they are the same, the responses could vary: One might say, “Sure; they’re both apples.” Another might say, “Certainly not; one is a Granny Smith and the other a golden delicious.” A third might quip, “How could they be the same? There are two objects on the table.” Ok, so we need some rules for how to classify things. In our case, we want to say when two groups are the same, meaning they act as groups in exactly the same way, they have exactly the same properties, and so on. The precise term is whether the groups are *isomorphic*, meaning not necessarily the same but algebraically indistinguishable (as groups).

Perhaps this is still unclear; maybe it is easier first to decide when two groups are not isomorphic. Let’s gather some thoughts. Certainly if two groups are the same in any sense of the word, they should have the same number of elements, so \mathbb{Z}_7 (as an additive abelian group) and U_5 (as a multiplicative abelian group) are not isomorphic, the first having seven elements, the second having four. Some groups are abelian (like \mathbb{Z}_n and U_n) while others are not. One can check that the set of permutations of three objects, denoted S_3 ,¹ is a non-abelian group with six elements, so even though S_3 and \mathbb{Z}_6 both have order 6 (six elements), they are not isomorphic since multiplication in one doesn’t act like multiplication in the other.

Maybe we are on to something here. After all, what is a group but a nonempty set with a binary operation satisfying a few properties. The orders of elements, whether two elements commute and other properties are all available by looking at a table with all possible products displayed, are called a *Cayley* table.

6.2. Cayley Tables

Consider a Cayley table for a group $G = \{e, a, b, c\}$ with order 4, identity e , and operation $*$. The elements of the group are listed in

¹See the exercise in section 6.3.

the first row and column, and the inner part of the table records their products as shown.

G	e	a	b	c
e	$e * e$	$e * a$	$e * b$	$e * c$
a	$a * e$	$a * a$	$a * b$	$a * c$
b	$b * e$	$b * a$	$b * b$	$b * c$
c	$c * e$	$c * a$	$c * b$	$c * c$

Below are the Cayley tables for two groups U_5 and U_{10} .

U_5	$\bar{1}$	$\bar{2}$	$\bar{3}$	$\bar{4}$
$\bar{1}$	$\bar{1}$	$\bar{2}$	$\bar{3}$	$\bar{4}$
$\bar{2}$	$\bar{2}$	$\bar{4}$	$\bar{1}$	$\bar{3}$
$\bar{3}$	$\bar{3}$	$\bar{1}$	$\bar{4}$	$\bar{2}$
$\bar{4}$	$\bar{4}$	$\bar{3}$	$\bar{2}$	$\bar{1}$

U_{10}	$\bar{1}$	$\bar{3}$	$\bar{7}$	$\bar{9}$
$\bar{1}$	$\bar{1}$	$\bar{3}$	$\bar{7}$	$\bar{9}$
$\bar{3}$	$\bar{3}$	$\bar{9}$	$\bar{1}$	$\bar{7}$
$\bar{7}$	$\bar{7}$	$\bar{1}$	$\bar{9}$	$\bar{3}$
$\bar{9}$	$\bar{9}$	$\bar{7}$	$\bar{3}$	$\bar{1}$

Are the groups the same? Well, no. One has elements $\bar{1}$, $\bar{2}$, $\bar{3}$, and $\bar{4}$, while the other has elements $\bar{1}$, $\bar{3}$, $\bar{7}$, and $\bar{9}$. But if you took a pen and relabeled the Cayley table for U_{10} leaving $\bar{1}$ alone but changed each occurrence of $\bar{3}$ to $\bar{2}$, of $\bar{7}$ to $\bar{3}$, and of $\bar{9}$ to $\bar{4}$, the two tables would be identical. This is what we mean by isomorphic.

More formally, we define a map $\varphi : U_{10} \rightarrow U_5$, by $\varphi(\bar{1}) = \bar{1}$, $\varphi(\bar{3}) = \bar{2}$, $\varphi(\bar{7}) = \bar{3}$, and $\varphi(\bar{9}) = \bar{4}$, which maps one Cayley table to the other. Because the tables (once relabeled) match, the function $\varphi : U_{10} \rightarrow U_5$ defined above necessarily satisfies $\varphi(ab) = \varphi(a)\varphi(b)$ for all $a, b \in U_{10}$, and this property characterizes the notion of what is

called a *homomorphism* between groups; it is a *structure-preserving map*. A bijective homomorphism (one that is also one-to-one and onto) is the formal definition of an *isomorphism*. We write $U_{10} \cong U_5$ (read U_{10} is isomorphic to U_5) to designate this relationship. It not hard to check (but is very important) that the isomorphism of groups is an equivalence relation.

For another example, we compute the Cayley tables of two other groups of order 4 and find the following.

U_8	$\bar{1}$	$\bar{3}$	$\bar{5}$	$\bar{7}$
$\bar{1}$	$\bar{1}$	$\bar{3}$	$\bar{5}$	$\bar{7}$
$\bar{3}$	$\bar{3}$	$\bar{1}$	$\bar{7}$	$\bar{5}$
$\bar{5}$	$\bar{5}$	$\bar{7}$	$\bar{1}$	$\bar{3}$
$\bar{7}$	$\bar{7}$	$\bar{5}$	$\bar{3}$	$\bar{1}$

U_{12}	$\bar{1}$	$\bar{5}$	$\bar{7}$	$\bar{11}$
$\bar{1}$	$\bar{1}$	$\bar{5}$	$\bar{7}$	$\bar{11}$
$\bar{5}$	$\bar{5}$	$\bar{1}$	$\bar{11}$	$\bar{7}$
$\bar{7}$	$\bar{7}$	$\bar{11}$	$\bar{1}$	$\bar{5}$
$\bar{11}$	$\bar{11}$	$\bar{7}$	$\bar{5}$	$\bar{1}$

One observes that if we define a map from $\varphi : U_8 \rightarrow U_{12}$, by taking $\bar{1} \mapsto \bar{1}$, $\bar{3} \mapsto \bar{5}$, $\bar{5} \mapsto \bar{7}$, and $\bar{7} \mapsto \bar{11}$, that the Cayley tables match exactly. The function $\varphi : U_8 \rightarrow U_{12}$ defined above is another example of an isomorphism, and once again, because the (relabelled) tables match, $\varphi(ab) = \varphi(a)\varphi(b)$ for all $a, b \in U_8$.

Next consider the Cayley table for (the additive abelian group) \mathbb{Z}_4 .

\mathbb{Z}_4	$\bar{0}$	$\bar{1}$	$\bar{2}$	$\bar{3}$
$\bar{0}$	$\bar{0}$	$\bar{1}$	$\bar{2}$	$\bar{3}$
$\bar{1}$	$\bar{1}$	$\bar{2}$	$\bar{3}$	$\bar{0}$
$\bar{2}$	$\bar{2}$	$\bar{3}$	$\bar{0}$	$\bar{1}$
$\bar{3}$	$\bar{3}$	$\bar{0}$	$\bar{1}$	$\bar{2}$

It doesn't obviously match either of the previous examples, but if we interchange the third and fourth rows and columns, we see it matches with the first two examples. Consider one of them, U_5 .

U_5	$\bar{1}$	$\bar{2}$	$\bar{3}$	$\bar{4}$
$\bar{1}$	$\bar{1}$	$\bar{2}$	$\bar{3}$	$\bar{4}$
$\bar{2}$	$\bar{2}$	$\bar{4}$	$\bar{1}$	$\bar{3}$
$\bar{3}$	$\bar{3}$	$\bar{1}$	$\bar{4}$	$\bar{2}$
$\bar{4}$	$\bar{4}$	$\bar{3}$	$\bar{2}$	$\bar{1}$

That is if $\varphi : U_5 \rightarrow \mathbb{Z}_4$ is defined by $\varphi(\bar{1}) = \bar{0}$, $\varphi(\bar{2}) = \bar{1}$, $\varphi(\bar{3}) = \bar{3}$, and $\varphi(\bar{4}) = \bar{2}$, the tables will match, though this time the relation is that $\varphi(ab) = \varphi(a) + \varphi(b)$ for all $a, b \in U_5$. This is because Cayley tables correspond to the group operations, which in the case of U_5 is multiplication, and in the case of \mathbb{Z}_4 is addition.

6.3. A Couple of Non-abelian Groups

We investigate two general classes of non-abelian groups. The *symmetric group* (on n letters) is defined as the set of permutations of the set $\{1, 2, \dots, n\}$, that is, as a set, it is the set of functions

$$S_n = \{f : \{1, 2, \dots, n\} \rightarrow \{1, 2, \dots, n\} \mid f \text{ is a bijection}\}.$$

The group operation is function composition so that the group product $f * g$ is just the composite function $f \circ g$. The identity is the function e so that $e(k) = k$ for all k , $1 \leq k \leq n$, and every function in S_n has an inverse precisely because it is one-to-one and onto. We easily check that the order of S_n , $|S_n|$, is $n!$. It should be easy to check that S_1 is just the trivial group consisting of one element, and $S_2 \cong \mathbb{Z}_2$, so is abelian. Starting with $n = 3$, we can verify that S_n is non-abelian. For concreteness, let us consider the case of $n = 3$.

While a bit cumbersome, we will denote an element in S_3 by $f = \begin{bmatrix} 1 & 2 & 3 \\ a & b & c \end{bmatrix}$, the function f defined by $f(1) = a$, $f(2) = b$, $f(3) = c$.

The group operation is function composition, so $f * g$ (which we will write as the product fg) is the function whose action on k is $f(g(k))$. In permutation notation, this translates as follows:

$$f = \begin{bmatrix} 1 & 2 & 3 \\ 3 & 1 & 2 \end{bmatrix}, g = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 1 & 3 \end{bmatrix} \mapsto fg = \begin{bmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{bmatrix},$$

that is

$$f(g(1)) = f(2) = 1; \quad f(g(2)) = f(1) = 3; \quad f(g(3)) = f(3) = 2.$$

Exercise. Let

$$\sigma = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{bmatrix} \text{ and } \tau = \begin{bmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{bmatrix}.$$

Compute $\sigma, \sigma^2, \sigma^3, \tau, \tau^2, \tau^3, \sigma\tau, \sigma^2\tau$.



Exercise. Fill in the Cayley table for S_3 using the elements listed in the first row or column, and show that S_3 is non-abelian.

\circ	e	σ	σ^2	τ	$\sigma\tau$	$\sigma^2\tau$
e						
σ						
σ^2						
τ						
$\sigma\tau$						
$\sigma^2\tau$						


Another class of groups is the symmetries of a regular n -gon, denoted D_n , or in some texts D_{2n} since the order of the group is $2n$. The group D_n is called the *dihedral group* of order $2n$. Again, we consider only the special case $n = 3$.

The symmetries of an equilateral triangle form a finite group called D_3 . Consider two basic symmetries of such a triangle, the first a counter-clockwise rotation by 120 degrees (denoted R) and the second a flip (denoted F) about a vertical axis through the vertex

labeled 1.



The group operation is again function composition, so that RF means first act by F then R .



Exercise. Compute $R, R^2, R^3, F, F^2, F^3, RF, R^2F$.

Exercise. Fill in the Cayley table for D_3 using the elements listed along the first row or column.

\circ	e	R	R^2	F	RF	R^2F
e						
R						
R^2						
F						
RF						
R^2F						

Exercise. Notice that each symmetry can be thought of as a permutation of the three vertices. If we regard the numbers marking the vertices of the left-hand triangle as positions, then R can be described as the permutation $R = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{bmatrix}$, and $F = \begin{bmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{bmatrix}$. Describe R^2 , F , RF , R^2F in terms of the elements σ and τ used to define S_3 . Can you determine if $D_3 \cong S_3$?

6.4. Cyclic Groups and Direct Products

To go a bit further, we need to think about another aspect of the question of classification. Above we said that $U_5 \cong U_{10} \cong \mathbb{Z}_4$ and $U_8 \cong U_{12}$. Suppose G was another group of order 4. We want to say $G \cong X$, where X is only one particular group, but which one should we list? So part of the classification is that we start with the simplest objects, give them names, and then think about simple constructions which allow one to build more complex groups from simpler ones. The simplest objects are the cyclic groups, which we discuss now.

We need a small bit of terminology. Let G be a group, and let $g \in G$. We denote by $\langle g \rangle$ the set

$$\langle g \rangle = \{g^m \mid m \in \mathbb{Z}\} = \{\dots, g^{-3}, g^{-2}, g^{-1}, g^0 = e, g, g^2, g^3, \dots\}.$$

We easily check that $\langle g \rangle$ is a group called the *cyclic subgroup generated by g* . A group G is called *cyclic* if there exists a $g \in G$ with $G = \langle g \rangle$.

For example, the following are all cyclic:

$$U_{10} = \langle \bar{3} \rangle = \langle \bar{7} \rangle = \{\bar{3}, \bar{3}^2, \bar{3}^3, \bar{3}^4\} = \{\bar{7}, \bar{7}^2, \bar{7}^3, \bar{7}^4\},$$

$$U_5 = \langle \bar{2} \rangle = \langle \bar{3} \rangle,$$

$$\mathbb{Z}_4 = \langle \bar{1} \rangle = \{\bar{1}, \bar{1} + \bar{1} = \bar{2}, \bar{1} + \bar{1} + \bar{1} = \bar{3}, \bar{1} + \bar{1} + \bar{1} + \bar{1} = \bar{4} = \bar{0}\},$$

where we note the operation in \mathbb{Z}_4 is additive, so g^m means $g + g + \dots + g$ (m times). The other groups we considered (U_8 and U_{12}) are not cyclic, since all their elements have order 1 or 2, which means there is no element whose powers can fill out all of G .

We note that viewing \mathbb{Z}_n as an additive abelian group, it is always cyclic having (at least) the generator $\bar{1}$. So for each integer $n \geq 1$ there is a cyclic group of order n . In fact, as we show below, any two cyclic

groups of the same order are isomorphic, so we shall keep \mathbb{Z}_n as our distinguished cyclic group of order n .

To establish the result about there being (up to isomorphism) only one cyclic group of a given order, we state without proof an important result in group theory, a theorem of Lagrange, a special case of which we stated in Theorem 4.15.

Theorem 6.1 (Lagrange). *Let G be a finite group, and let $H \subseteq G$ be a subgroup (a subset that is also a group). Then $|H|$ divides $|G|$. That is, the order of a subgroup divides the order of the group.*

Using this we show the following.

Corollary 6.2. *Let G be a finite group, let $g \in G$, and put $H = \langle g \rangle$. Then $H = \{e = g^0, g, g^2, \dots, g^{d-1}\}$, where $d = |g|$. In particular, the order of an element divides the order of the group.*

Remark 6.3. Perhaps we should comment a bit about what this corollary says. We have already said that if you take any element g in a group G , you can construct the cyclic subgroup $H = \langle g \rangle$ which consists of all the powers of g . First it should be clear that since $H \subseteq G$ and G is finite, then so is H , implying that the element g has finite order (see Proposition 4.13; otherwise each of the powers g, g^2, \dots would be distinct making H an infinite group). So g has finite order d (the smallest positive integer so that $g^d = e$).

The first part of the corollary says that you get all the distinct elements of H by taking $H = \{e = g^0, g, g^2, \dots, g^{d-1}\}$, and since H is a subgroup of G having order d , Lagrange's theorem says that d (the order of H and of g) divides the order of G .

Proof. Let $d = |g|$, and let m be any integer. Using the division algorithm, write $m = dq + r$ where $0 \leq r < d$. Thus $g^m = g^{dq+r} = (g^d)^q g^r = g^r$, so every element of $\langle g \rangle$ is in $\{e = g^0, g, g^2, \dots, g^{d-1}\}$. It only remains to show that no two elements in the list $g^0, g, g^2, \dots, g^{d-1}$ are equal. If that were not true, then $g^i = g^j$ for $0 \leq i < j \leq d-1$, which would imply that $g^{j-i} = e$ so d would not be the smallest positive exponent so that $g^d = e$. This would be a contradiction. The rest of the corollary follows from Lagrange's theorem and Remark 6.3. \square

Corollary 6.4. *Let G be a finite group with prime order p . Then G is cyclic.*

Proof. Let $g \in G$ with $g \neq e$, and put $H = \langle g \rangle$. Then $|H| > 1$, and by Lagrange, $|H|$ divides $|G| = p$. This implies $|H| = p$, and since $H \subseteq G$, this means H is all of G , so $G = H = \langle g \rangle$. \square

Proposition 6.5. *Let $G_1 = \langle g_1 \rangle$ and $G_2 = \langle g_2 \rangle$ both be cyclic groups of order n . Then $G_1 \cong G_2$, that is they are isomorphic.*

Proof. $G_1 = \{e, g_1, g_1^2, \dots, g_1^{n-1}\}$ and $G_2 = \{e, g_2, g_2^2, \dots, g_2^{n-1}\}$. Sending $g_1^k \mapsto g_2^k$ shows that the Cayley tables match, establishing the isomorphism. \square

Corollary 6.6. *Let G be a cyclic group of order n . Then $G \cong \mathbb{Z}_n$.*

Now that we have some basic examples of groups, we can ask for simple constructions to build new groups out of old ones. One simple construction is the notion of a *direct product* of groups. Suppose we have groups H and K with the group operations in H, K denoted (for clarity) by $h_1 \circ h_2$ and $k_1 * k_2$, respectively. Using these two groups, we want to build a new group $G = H \times K$ as follows. As a set, G is the Cartesian product of H and K , that is the set of ordered pairs $G = \{(h, k) \mid h \in H, k \in K\}$. We make G into a group by giving it the binary operation

$$(h_1, k_1)(h_2, k_2) = (h_1 \circ h_2, k_1 * k_2).$$

The inverse of (h, k) is (h^{-1}, k^{-1}) and the identity is (e_H, e_K) . The axioms are easily verified.

Returning to our classification problem, consider the Cayley table for $\mathbb{Z}_2 \times \mathbb{Z}_2$ (without the bars over the numbers for notational ease)

$\mathbb{Z}_2 \times \mathbb{Z}_2$	(0, 0)	(0, 1)	(1, 0)	(1, 1)
(0, 0)	(0, 0)	(0, 1)	(1, 0)	(1, 1)
(0, 1)	(0, 1)	(0, 0)	(1, 1)	(1, 0)
(1, 0)	(1, 0)	(1, 1)	(0, 0)	(0, 1)
(1, 1)	(1, 1)	(1, 0)	(0, 1)	(0, 0)

We see this table matches those for U_8 and U_{12} . In terms of isomorphisms, we have

$$U_{10} \cong U_5 \cong \mathbb{Z}_4, \text{ and } U_8 \cong U_{12} \cong \mathbb{Z}_2 \times \mathbb{Z}_2.$$

Indeed “up to isomorphism” there are only two groups of order 4, \mathbb{Z}_4 and $\mathbb{Z}_2 \times \mathbb{Z}_2$, meaning any group of order 4 is isomorphic to exactly one of these.

A very large project (to which thousands of journal pages have been devoted) is to create a list of all the isomorphism classes of finite groups. We shall give an algorithm in the next section for how to do this for abelian groups, but the question of classifying all non-abelian groups is still open.

Example 6.7. Groups of small order. Let G be a group of order n .

- $n = 1$: $G = \{e\}$; this is called the trivial group.
- $n = 2, 3$: Both numbers are prime, so the groups are cyclic and isomorphic to \mathbb{Z}_2 and \mathbb{Z}_3 , respectively.
- $n = 4$: We have seen \mathbb{Z}_4 and $\mathbb{Z}_2 \times \mathbb{Z}_2$ are not isomorphic. It turns out that every group of order four is isomorphic to one of these.
- $n = 5$: Prime order, so cyclic, $G \cong \mathbb{Z}_5$.
- $n = 6$: The first non-abelian case. Clearly, \mathbb{Z}_6 is possible, and it turns out $\mathbb{Z}_2 \times \mathbb{Z}_3 \cong \mathbb{Z}_6$. On the other hand S_3 (the group of permutations of three objects) and D_3 (the set of symmetries of an equilateral triangle) are groups of order 6, but are not abelian, so they are certainly not isomorphic to \mathbb{Z}_6 . However, it turns out that $S_3 \cong D_3$, and every group of order 6 is isomorphic to exactly one of \mathbb{Z}_6 or S_3 .
- $n = 7$: Prime order, so $G \cong \mathbb{Z}_7$.
- $n = 8$: There are five isomorphism classes of groups of order 8: $\mathbb{Z}_8, \mathbb{Z}_2 \times \mathbb{Z}_4, \mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2$ are the abelian ones. The symmetries of the square D_4 is another. The last is called the quaternion group, usually denoted Q_8 .

6.5. Fundamental Theorem of Finite Abelian Groups

Though the proof is well beyond the scope of this text, the structure of finite abelian groups is particularly easy to describe. We take a somewhat iterative approach to the final result. As a preliminary step, we talk about what appears to be a digression: the partitions of a positive integer n .

We say that a set of positive integers $\{n_1, \dots, n_k\}$ is a partition of the positive integer n if $n = n_1 + \dots + n_k$ and $n_1 \geq n_2 \geq \dots \geq n_k \geq 1$. We denote by $\pi(n)$ the number of partitions of n . Thus, $\pi(5) = 7$, since there are seven partitions of the number 5:

$$\begin{aligned} 5 = 5; & \quad 4 + 1; \quad 3 + 2; \quad 3 + 1 + 1; \quad 2 + 2 + 1; \\ & \quad 2 + 1 + 1 + 1; \quad 1 + 1 + 1 + 1 + 1. \end{aligned}$$

Our first theorem is for abelian groups of prime-power order.

Theorem 6.8. *Let p be a prime, and let $n \geq 1$. Then up to isomorphism, there are $\pi(n)$ abelian groups of order p^n .*

Said another way, there are only $\pi(n)$ distinct Cayley tables you can write down for an abelian group of order p^n . This means that any group of order p^n is isomorphic to exactly one of the $\pi(n)$ candidates and shares the appropriate Cayley table. What is striking about this result is that it depends only on the exponent n and not on the prime p . Moreover, the proof (not given here) shows not only that the partition function $\pi(n)$ counts the number of abelian groups, but that the partitions of n identify the isomorphism classes. In particular, different partitions correspond to nonisomorphic groups. We demonstrate this in the example below.

Example 6.9. We characterize all the abelian groups of order p^5 in terms of partitions as follows.

$$\begin{aligned}
 5 &\longleftrightarrow \mathbb{Z}_{p^5} \\
 4 + 1 &\longleftrightarrow \mathbb{Z}_{p^4} \times \mathbb{Z}_p \\
 3 + 2 &\longleftrightarrow \mathbb{Z}_{p^3} \times \mathbb{Z}_{p^2} \\
 3 + 1 + 1 &\longleftrightarrow \mathbb{Z}_{p^3} \times \mathbb{Z}_p \times \mathbb{Z}_p \\
 2 + 2 + 1 &\longleftrightarrow \mathbb{Z}_{p^2} \times \mathbb{Z}_{p^2} \times \mathbb{Z}_p \\
 2 + 1 + 1 + 1 &\longleftrightarrow \mathbb{Z}_{p^2} \times \mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_p \\
 1 + 1 + 1 + 1 + 1 &\longleftrightarrow \mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_p
 \end{aligned}$$

Below is one version of the Fundamental Theorem of Finite Abelian Groups.

Theorem 6.10 (Fundamental Theorem). *Let $N \geq 2$ be an integer with prime factorization $N = p_1^{e_1} \cdots p_r^{e_r}$. Then every abelian group G of order N is isomorphic to a direct product $G \cong G(p_1) \times \cdots \times G(p_r)$ where $G(p_i)$ is an abelian group of order $p_i^{e_i}$, described in the previous theorem. Thus, up to isomorphism, there are $\pi(e_1)\pi(e_2) \cdots \pi(e_r)$ abelian groups of order N .*

It is useful in analyzing abelian groups to note that as a consequence of the Chinese Remainder Theorem, we have the following.

Theorem 6.11. $\mathbb{Z}_m \times \mathbb{Z}_n \cong \mathbb{Z}_{mn}$ if and only if $\gcd(m, n) = 1$.

Example 6.12. Classify all abelian groups of order p^2q^3 , where p and q are distinct primes.

Up to isomorphism, there are $\pi(2) = 2$ abelian groups of order p^2 : \mathbb{Z}_{p^2} and $\mathbb{Z}_p \times \mathbb{Z}_p$. Up to isomorphism, there are $\pi(3) = 3$ abelian groups of order q^3 : \mathbb{Z}_{q^3} , $\mathbb{Z}_{q^2} \times \mathbb{Z}_q$, and $\mathbb{Z}_q \times \mathbb{Z}_q \times \mathbb{Z}_q$. Thus there are $\pi(2)\pi(3) = 6$ distinct isomorphism classes of abelian groups of order

p^2q^3 , and they are

$$\begin{aligned} \mathbb{Z}_{p^2} \times \mathbb{Z}_{q^3} &\cong \mathbb{Z}_{p^2q^3} \\ \mathbb{Z}_{p^2} \times \mathbb{Z}_{q^2} \times \mathbb{Z}_q &\cong \mathbb{Z}_q \times \mathbb{Z}_{p^2q^2} \\ \mathbb{Z}_{p^2} \times \mathbb{Z}_q \times \mathbb{Z}_q \times \mathbb{Z}_q &\cong \mathbb{Z}_q \times \mathbb{Z}_q \times \mathbb{Z}_{p^2q} \\ \mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_{q^3} &\cong \mathbb{Z}_p \times \mathbb{Z}_{pq^3} \\ \mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_{q^2} \times \mathbb{Z}_q &\cong \mathbb{Z}_{pq} \times \mathbb{Z}_{pq^2} \\ \mathbb{Z}_p \times \mathbb{Z}_p \times \mathbb{Z}_q \times \mathbb{Z}_q \times \mathbb{Z}_q &\cong \mathbb{Z}_p \times \mathbb{Z}_{pq} \times \mathbb{Z}_{pq} \end{aligned}$$

(the left-hand column being the one from Theorem 6.10 and the right-hand column leveraging Theorem 6.11).

There is an alternate statement of the Fundamental Theorem for Finite Abelian Groups which is useful for showing that all the groups of the form U_p (p a prime) are cyclic. Recall that above we showed that every group of prime order is cyclic, but U_p has order $p - 1$ which is not prime except when $p = 3$. This alternate version of the Fundamental Theorem follows.

Theorem 6.13 (Fundamental Theorem). *Let G be a finite abelian group of order $N \geq 2$. Then there are uniquely determined integers: $t \geq 1$ and $n_1, \dots, n_t \geq 2$, so that $G \cong \mathbb{Z}_{n_1} \times \dots \times \mathbb{Z}_{n_t}$ and $n_1 \mid n_2 \mid \dots \mid n_t$. Necessarily, $N = n_1 n_2 \dots n_t$.*

In Example 6.12, the right-hand column gives the isomorphism class of the group using the characterization in Theorem 6.13.

Exercise. We know that U_n is a finite abelian group. For $5 \leq n \leq 15$, use your knowledge of these groups to characterize them as in the Fundamental Theorem. For example, U_3 is a group of order 2, a prime, so U_3 is a cyclic group of order 2, that is $U_3 \cong \mathbb{Z}_2$. The group U_8 is an abelian group of order 4, so by the Fundamental Theorem it is isomorphic to either $\mathbb{Z}_2 \times \mathbb{Z}_2$ or to \mathbb{Z}_4 . We easily check for all $a \in U_8$ that $a^2 = 1$, so U_8 is not cyclic, and so $U_8 \cong \mathbb{Z}_2 \times \mathbb{Z}_2$.

In the next chapter, we will encounter abelian groups coming from elliptic curves.

6.6. Primitive Roots

In this section we want to talk about the discrete logarithm problem as a precursor to the Diffie–Hellman key exchange and to a method of encryption called the ElGamal method.

We begin with an important result in algebra.

Theorem 6.14. *Let p be a prime, and let U_p be the set of reduced residues modulo p . Then U_p is a cyclic group.*

Proof. We know that $U_2 = \{\bar{1}\}$ is the trivial group and is cyclic as it is generated by the identity, so we assume that $p \geq 3$. Then U_p is a finite abelian group of order $p - 1 \geq 2$, so by the Fundamental Theorem for Finite Abelian Groups, we know that there are uniquely determined integers: $t \geq 1$ and $n_1, \dots, n_t \geq 2$ with $n_1 \mid n_2 \mid \dots \mid n_t$, so that $U_p \cong \mathbb{Z}_{n_1} \times \dots \times \mathbb{Z}_{n_t}$. Now the condition $n_1 \mid \dots \mid n_t$ means that every element in U_p has order dividing n_t . If $t > 1$, then since $p - 1 = n_1 \cdot n_2 \cdot \dots \cdot n_t$, we know $n_t < p - 1$, so there can be no element of order $p - 1$. On the other hand if $t = 1$, then $U_p \cong \mathbb{Z}_{n_1}$ is a cyclic group, which is our goal.

While, strictly speaking, we have not proven the result we are about to invoke, it is something you learned a long time ago: that a polynomial of degree n with coefficients in a field can have at most n roots in that field. In our case, the field in question is \mathbb{F}_p , and we see that since all elements in U_p have order dividing n_t , every element of U_p is a root of the polynomial $x^{n_t} - 1$. But if $t > 1$, we have said that $n_t < p - 1$, so that means the polynomial $x^{n_t} - 1$ has more roots than its degree, a contradiction. So we conclude that $t = 1$, which means that U_p is cyclic. \square

So for a prime p , the group U_p is cyclic, and so is generated by an element g , say $U_p = \langle g \rangle$, in our previous notation for cyclic groups. Any such generator for U_p is called a *primitive root* modulo p . So if g is a primitive root mod p , then

$$U_p = \{g, g^2, \dots, g^{p-1} = 1\}.$$

Said another way, for each integer a with $\gcd(a, p) = 1$, there is a unique k with $1 \leq k \leq p - 1$ so that

$$a \equiv g^k \pmod{p}.$$

The integer k is called the *index* or *discrete logarithm* of a to the base g modulo p , written $k = \text{ind}_g a$, analogous to a logarithm. Technically, $\text{ind}_g a$ is only determined modulo $(p - 1)$, but this poses no difficulty since $g^{p-1} = \bar{1}$.

As a small example, let's examine $U_{11} = \{1, 2, 3, 4, 5, 6, 7, 8, 9, 10\}$; note we have chosen to write the elements of U_{11} as integers k rather than \bar{k} to make the notation a bit lighter. It has primitive roots 2, 6, 7, 8, but let us fix $g = 2$ to generate U_{11} . Actually, there are always $\phi(p - 1)$ primitive roots, and when we have found one of them, say g , the others are all of the form g^k where $\gcd(k, p - 1) = 1$. This follows

Table 6.1

$2^1 \equiv 2 \pmod{11}$	$\text{ind}_2 2 = 1$	$ 2 = 10$
$2^2 \equiv 4 \pmod{11}$	$\text{ind}_2 4 = 2$	$ 4 = 2^2 = \frac{10}{\gcd(2,10)} = 5$
$2^3 \equiv 8 \pmod{11}$	$\text{ind}_2 8 = 3$	$ 8 = 2^3 = \frac{10}{\gcd(3,10)} = 10$
$2^4 \equiv 5 \pmod{11}$	$\text{ind}_2 5 = 4$	$ 5 = 2^4 = \frac{10}{\gcd(4,10)} = 5$
$2^5 \equiv 10 \pmod{11}$	$\text{ind}_2 10 = 5$	$ 10 = 2^5 = \frac{10}{\gcd(5,10)} = 2$
$2^6 \equiv 9 \pmod{11}$	$\text{ind}_2 9 = 6$	$ 9 = 2^6 = \frac{10}{\gcd(6,10)} = 5$
$2^7 \equiv 7 \pmod{11}$	$\text{ind}_2 7 = 7$	$ 7 = 2^7 = \frac{10}{\gcd(7,10)} = 10$
$2^8 \equiv 3 \pmod{11}$	$\text{ind}_2 3 = 8$	$ 3 = 2^8 = \frac{10}{\gcd(8,10)} = 5$
$2^9 \equiv 6 \pmod{11}$	$\text{ind}_2 6 = 9$	$ 6 = 2^9 = \frac{10}{\gcd(9,10)} = 10$
$2^{10} \equiv 1 \pmod{11}$	$\text{ind}_2 1 = 10$	$ 1 = 2^{10} = \frac{10}{\gcd(10,10)} = 1$

from a bit more algebra than we have at our disposal, but which also gives the following useful fact.

We know all the elements in U_p have order dividing $p - 1$, and we know that g has order $p - 1$. In general, algebra shows that the order of g^k , denoted $|g^k|$, is $\frac{p-1}{\gcd(k, p-1)}$. Our entries in Table 6.1 list the elements of U_{11} as $\{2^1, 2^2, \dots, 2^{10}\}$ (modulo 11), from which we extract the index, ind_2 , and the order of the element.

6.7. Diffie–Hellman Key Exchange

Recall that in an online transaction between you and a vendor, almost all encryption is done via a symmetric-key cryptosystem. We have already discussed the problem of how your computer and that of an online vendor's establish the protocol, but once done the shared key (which should be known only to your computer and the vendor's) must be generated. How do you proceed? Not surprisingly, there are a number of protocols, but we choose to give one of the simpler ones, known as static Diffie–Hellman, though we note that generally this static version has been replaced by ephemeral Diffie–Hellman to advance what is known as *forward secrecy*. In static Diffie–Hellman one finds embedded in the vendor's security certificate a large prime p and a primitive root g . So the game begins. The vendor has already picked a random integer b with $1 < b < p - 1$, and has $g^b \pmod{p}$ embedded as part of their certificate. Your computer generates a random integer a with $1 < a < p - 1$, and sends $g^a \pmod{p}$ to the vendor. With your integer a you compute $(g^b)^a \equiv g^{ab} \pmod{p}$. With its integer b , the vendor computes $(g^a)^b = g^{ab} \pmod{p}$. So you each share the common key $K \equiv g^{ab} \pmod{p}$. Someone who wants to compromise your secure exchange must have access to the key K , but has only been able to see (if at all) g^a and g^b modulo p . Determining $K \equiv g^{ab} \pmod{p}$ from the given data is the Diffie–Hellman problem, which of course could be solved by solving the discrete logarithm problem, that of determining a, b from $g^a, g^b \pmod{p}$, which is thought to be very hard.

Remark 6.15. In practice of course, one wants to pick a very large prime p , but how large? The National Institute of Standards (NIST)

makes recommendations for key sizes to protect data by various algorithms. For example, AES is one of the most common symmetric-key encryption methods in use today. It accepts key sizes of (at least) 80, 112, 128, 192, or 256 bits; longer keys correspond to greater security. In a table published in [NC09], the recommended bit sizes for primes to be used in either RSA or Diffie–Hellman are much longer; their recommended bit sizes range among 1024, 2048, 3072, 7680, and 15360 bits. Now a prime p approximately 15360 bits long has roughly 4624 digits! A prime with 3072 bits has roughly 925 digits. These are large primes, and while primality testing can be done in polynomial time, these times are significant.

Now a practical consideration. You have a prime. Just how are you going to go about finding a primitive root, that is a generator for U_p ? Start at 2 and keep testing until you have an element of the correct order? Probably not. To make Diffie–Hellman secure, you don't really need a generator for the whole group, you just need an element which generates a large subgroup, say half of it. How can you arrange that?

The following is a practical method of picking a prime and an element g whose order is either $p - 1$ (a primitive root) or $(p - 1)/2$, half the order of the group. Pick a prime q of approximately the right size. Actually pick lots of them, and consider $p = 2q + 1$. When you find a prime q so that $p = 2q + 1$ is also prime, stop; such a p is called a *safe prime*. Now, the order of U_p is $p - 1 = 2q$, and every element in the group has order dividing $2q$, so it is one of 1, 2, q , or $2q$. There is only one element of order 1 in a group (the identity), and it turns out that there is only one element of order 2 in U_p , namely $-1 \equiv p - 1 \pmod{p}$. Every other element g with $2 \leq g \leq p - 2$ has order q or $2q$, so choose $g = 2$, and even if it is not a primitive root, it is an element that will generate a (sub)group of size $q = |U_p|/2$.

6.8. ElGamal Encryption

Having introduced primitive roots and the Diffie–Hellman key exchange, we introduce a related public-key cryptosystem called ElGamal. As in the section on RSA, we will have two players, Alice and

Bob, and we shall describe in brief how Bob can send an encrypted message to Alice.

Suppose p is a large prime, and consider the group of units U_p , which we know is cyclic of order $p - 1$. For simplicity, assume that g is a generator for U_p , i.e., a primitive root modulo p . One can modify the process below if we simply take an element of very large order instead.

Alice chooses an integer a (her private key) satisfying $1 < a < p-1$ and generates $g^a \pmod{p}$. Her public key consists of the data U_p , g , and $g^a \pmod{p}$. As with RSA, we assume that a plaintext message that Bob wishes to send to Alice has been converted to a numerical equivalent integer M with $1 \leq M \leq |U_p| = p - 1$. Bob chooses a random integer k and computes $g^k \pmod{p}$, from Alice's public g . Then (as elements of U_p) he computes the product $M \cdot (g^a)^k$ using his plaintext M , random k , and Alice's public g^a . He now sends the ordered pair $(g^k, M \cdot g^{ak})$ to Alice. Using her private key a , Alice can compute $(g^k)^a = g^{ak}$, and since the inverse of g^{ak} in U_p is g^{p-1-ak} , she can multiply $M \cdot g^{ak}$ by the inverse of g^{ak} to retrieve M . As in the Diffie–Hellman exchange, someone who could solve the discrete logarithm problem could deduce Alice's secret key (a) from the public information g and g^a .

Chapter 7

Curves in Affine and Projective Space

7.1. Affine and Projective Space

In Chapter 2 we made some conjectures about the number of points of intersection of two plane curves, and we seemed to have settled on a conjecture that said that if there were only finitely many points of intersection, the product of the degrees of the curves was an upper bound for their number. But even with the simple example of a circle and a parabola, it seemed possible to have anywhere from zero to four points of intersection. By broadening our setting a bit, we can do much better.

The curves we considered in Chapter 2 were curves in \mathbb{R}^2 , and as such it is easy to see that the real numbers themselves create difficulties in obtaining a consistent answer to the number of points of intersection. For example, if we intersect a line (say the x -axis) with a quadratic (say $y = x^2 - a$), we expect at most two points of intersection, the roots. Indeed over \mathbb{R} , there can be two real points of intersection (e.g., if $a = 4$), one (e.g., if $a = 0$), or none (e.g., if $a = -4$). On the other hand, if instead of considering these curves in \mathbb{R}^2 , we considered them in \mathbb{C}^2 , then there would always be two points of intersection (at least if counted with multiplicity). The proof is simply that the points of intersection have x -coordinates which are the

roots of the quadratic, and the quadratic formula will always give two roots over \mathbb{C} (although it may be a root occurring with multiplicity 2, e.g., $a = 0$). So the field over which we view the curve plays a role. On the other hand, sometimes we are motivated by concerns other than maximizing the number of points of intersection, in which case we may focus on rational points or even points in a finite field, like \mathbb{F}_p .

But changing the field alone does not solve all of our problems. For example, when we consider two lines in the plane, do they always intersect in a single point? Well, no, but we have to be a little careful in what we are asking. The two lines could be equal, so let's dismiss that case. If they are distinct lines, they either intersect in a unique point or are parallel and do not intersect in the plane. And here changing the field from \mathbb{R} to \mathbb{C} does not help. On the other hand, anyone who has stood on railroad tracks has had the impression that parallel lines do intersect, just "at infinity". Expanding our notion of the plane by adding points at infinity is another of the ways we wish to broaden our view of plane curves. This expanded plane will be called the projective plane, but adding these extra points takes a bit of care.

Let's begin by establishing notation to help us distinguish some of our known settings from ones we wish to introduce. For concreteness we will define objects using the real numbers, but the notions easily generalize to other fields such as \mathbb{C} , \mathbb{Q} , or \mathbb{F}_p . We have always called \mathbb{R} the real line, \mathbb{R}^2 the plane, and \mathbb{R}^3 (real) 3-space. Typical elements of \mathbb{R} are denoted a, b, c, \dots . Elements of \mathbb{R}^2 are ordered pairs (a, b) of real numbers, and elements of \mathbb{R}^3 are ordered triples (a, b, c) . So in formalizing our definitions, we shall refer to \mathbb{R} as the *affine line*, \mathbb{R}^2 the *affine plane*, and \mathbb{R}^3 *affine 3-space*. Of course, we can talk about a generic affine n -space as well. We give *affine n -space* (over \mathbb{R}) the formal definition and notation

$$\mathbb{A}^n(\mathbb{R}) = \mathbb{R}^n = \{(a_1, \dots, a_n) \mid a_i \in \mathbb{R}\},$$

the set of ordered n -tuples of real numbers. So the real line is $\mathbb{A}^1(\mathbb{R})$, the plane is $\mathbb{A}^2(\mathbb{R})$, and so on. This notation is convenient for when we write something like $\mathbb{A}^2(\mathbb{Q})$ as we are simply looking at the set of rational points in the affine plane. So this definition makes sense if

we replace \mathbb{R} by any field F , e.g., \mathbb{Q} , \mathbb{R} , \mathbb{C} , or a finite field like \mathbb{F}_p . For example, $\mathbb{A}^2(\mathbb{F}_p)$ is the set of ordered pairs (a, b) where $a, b \in \mathbb{F}_p$.

In order to explain how to add points to affine space to make it bigger (in a useful manner), we shall first give a construction that will embed an affine space of a given dimension into a corresponding projective space and then return to suggest why this particular construction is very natural in the context of studying curves.

The notion of a point in projective space relies on the concept of an equivalence relation we introduced in Chapter 3, when we reviewed how we think of the rational numbers and how we work with congruence classes in the context of modular arithmetic. For now, simply recall that an equivalence class is a set of things that appear different but really represent the same object, just as the fractions $2/4$, $(-3)/(-6)$, $m/2m$ ($m \neq 0$) all represent the same object we usually denote by $1/2$. Why choose $1/2$ as our favorite representative of the equivalence class? Well, that representative is special in some sense that is meaningful to us. So as we define how to view affine space as a subset of projective space, we will have to find a representative of the equivalence class of a point in projective space that is also meaningful to us.

We start with a careful definition of the *projective line*. Consider the set of all nonzero vectors in the affine plane $\mathbb{A}^2(\mathbb{R})$. Geometrically, we can think of these vectors as having tails at the origin $(0, 0)$ and tips at arbitrary points $(a, b) \neq (0, 0)$ in the plane. So we have a set

$$S = \{(a, b) \in \mathbb{A}^2(\mathbb{R}) \mid (a, b) \neq (0, 0)\}.$$

We introduce an equivalence relation on S saying that $(a, b) \sim (c, d)$ if the line determined by the vector (a, b) (i.e., the line through $(0, 0)$ and (a, b) , which is why we needed $(a, b) \neq (0, 0)$) is the same as the line determined by the vector (c, d) . For example, the line determined by the vector $(2, 4)$ is the same as the one determined by $(1, 2)$ or by $(-3, -6)$. Those with some background in vector geometry will immediately recognize this as saying that $(a, b) \sim (c, d)$ if and only if there is a nonzero real number t so that $(c, d) = (at, bt) = t(a, b)$, that is, the vectors (a, b) and (c, d) are parallel. This means that the set of points in S which lie in the equivalence class of (a, b) is




Figure 7.1

precisely the set of points (excluding the origin) on the line through the origin and the point (a, b) . So the set of equivalence classes of the elements of S are these “lines” through the origin. Now, just as with rational numbers or congruence classes, we want to find a nice representative of each equivalence class. Consider the line $y = 1$ in the plane (Figure 7.1). Every line through the origin (except the x -axis) intersects the line $y = 1$ in a unique point. First note that the line determined by (a, b) is the x -axis if and only if $b = 0$, so if (a, b) determines a line through the origin other than the x -axis, we know $b \neq 0$, so $(a, b) \sim (a/b, 1)$. Said almost equivalently, every line of the form $y = mx$ with $m \neq 0$ (which excludes the x -axis as well as the y -axis which has undefined slope) intersects the line $y = 1$ at the point $(1/m, 1)$.

So every line through the origin (except the x -axis) determines a point $(r, 1)$ on the line $y = 1$, and every point $(r, 1)$ determines a line through the origin except for the x -axis. In fact this is a one-to-one correspondence between the points $r \in \mathbb{A}^1(\mathbb{R})$ and the lines through the origin (except for the x -axis). So let $[a, b]$ denote the equivalence

class of the point (vector) $(a, b) \in S$. Then the set of equivalence classes of S consists of the set

$$\{[a, b] \mid a, b \in \mathbb{R}, (a, b) \neq (0, 0)\} = \{[r, 1] \mid r \in \mathbb{R}\} \cup \{[1, 0]\},$$

where the last statement comes from the fact that the only line excluded in our discussion was the x -axis, which is determined by any vector of the form $(a, 0)$ with $a \neq 0$, and it is clear that $(a, 0) \sim (1, 0)$, so $[a, 0] = [1, 0]$.

We define the projective line $\mathbb{P}^1(\mathbb{R})$ to be this set of equivalence classes, that is

$$\begin{aligned} \mathbb{P}^1(\mathbb{R}) &= \{[a, b] \mid a, b \in \mathbb{R}, (a, b) \neq (0, 0)\} \\ &= \{[r, 1] \mid r \in \mathbb{A}^1(\mathbb{R})\} \cup \{[1, 0]\}. \end{aligned}$$

Now of course, it is trivial to identify the affine line $\mathbb{A}^1(\mathbb{R})$ with $\{[r, 1] \in \mathbb{P}^1(\mathbb{R})\}$ via $r \mapsto [r, 1]$, so we think of

$$\mathbb{P}^1(\mathbb{R}) = \{[r, 1] \mid r \in \mathbb{A}^1(\mathbb{R})\} \cup \{[1, 0]\} = \text{“}\mathbb{A}^1(\mathbb{R})\text{”} \cup \{[1, 0]\},$$

that is, the projective line is a copy of the affine line together with a single point (at infinity). The “at infinity” part seems natural since the x -axis can be thought of as the limiting line through $(0, 0)$ and $(r, 1)$ as $r \rightarrow \infty$.

Now that we have taken our time to describe the projective line, we will be a bit briefer in the description of higher-dimensional projective spaces. Still, the projective plane is crucial to our efforts, so we will still be quite careful, but the analogy should seem clear. To define the *projective plane* $\mathbb{P}^2(\mathbb{R})$, we start with the set of nonzero vectors in $\mathbb{R}^3 = \mathbb{A}^3(\mathbb{R})$, which again we think of as vectors based at $(0, 0, 0)$ with tip at $(a, b, c) \neq (0, 0, 0)$. We define an equivalence relation by saying that triples $(a, b, c) \sim (a', b', c')$ if there is a nonzero real number t so that $(a', b', c') = t(a, b, c) = (at, bt, ct)$. In analogy with the case of the projective line, the equivalence class $[a, b, c]$ will be the set of points on the line through the origin and the point (a, b, c) (excluding the origin). Similarly to the previous case, we divide the lines into two subsets, those that pass through the plane $z = 1$ and those that do not. The lines that do not pass through $z = 1$ are determined by vectors whose z -coordinate equals zero. If we denote by $[a, b, c]$ the

equivalence class of (a, b, c) , then

$$\begin{aligned}\mathbb{P}^2(\mathbb{R}) &= \{[a, b, c] \mid a, b, c \in \mathbb{R}, (a, b, c) \neq (0, 0, 0)\} \\ &= \{[a, b, 1] \mid (a, b) \in \mathbb{A}^2(\mathbb{R})\} \cup \{[a, b, 0] \in \mathbb{P}^2(\mathbb{R})\}.\end{aligned}$$

As in the case of the projective line, it is clear that we can identify $\mathbb{A}^2(\mathbb{R})$ with $\{[a, b, 1] \in \mathbb{P}^2(\mathbb{R})\}$ via $(a, b) \mapsto [a, b, 1]$, so we see that the affine plane is identified with a natural subset of our projective plane. But what about the piece that is left over? Is this another copy of the plane? Actually, no because

$$\{[a, b, 0] \in \mathbb{P}^2(\mathbb{R})\} = \{[a, b, 0] \mid (a, b, 0) \neq (0, 0, 0)\}.$$

So what is left over is actually a copy of the projective line $\mathbb{P}^1(\mathbb{R})$, which can be easily seen as

$$\{[a, b, 0] \mid (a, b, 0) \neq (0, 0, 0)\} = \{[r, 1, 0] \mid r \in \mathbb{A}^1(\mathbb{R})\} \cup \{[1, 0, 0]\},$$

an affine line (at infinity) plus a point at infinity. So the projective plane $\mathbb{P}^2(\mathbb{R})$ is a copy of the affine plane, together with a projective line at infinity.

We now give the general definition of projective space for arbitrary dimension n and over any field F , where we used $F = \mathbb{R}$ in the examples above. For $n \geq 1$, we define *projective n -space* over a field F , $\mathbb{P}^n(F)$, as follows: Let

$$S = \{(a_1, a_2, \dots, a_{n+1}) \in \mathbb{A}^{n+1}(F) \mid (a_1, a_2, \dots, a_{n+1}) \neq (0, \dots, 0)\}.$$

Define a relation on S by $(a_1, \dots, a_{n+1}) \sim (b_1, \dots, b_{n+1})$ if and only if there is a nonzero scalar $t \in F$ so that $(b_1, \dots, b_{n+1}) = t(a_1, \dots, a_{n+1})$. We easily check that \sim is an equivalence relation, and we denote by $[a_1, \dots, a_{n+1}]$ the equivalence class of (a_1, \dots, a_{n+1}) in S . Projective n -space is the set of equivalence classes

$$\mathbb{P}^n(F) = \{[a_1, \dots, a_{n+1}] \mid (a_1, \dots, a_{n+1}) \in S\}.$$

As in the special cases above, we see that

$$\begin{aligned}\mathbb{P}^n(F) &= \{[a_1, \dots, a_{n+1}] \mid (a_1, \dots, a_{n+1}) \in S\} \\ &= \{[a_1, \dots, a_n, 1] \mid (a_1, \dots, a_n) \in \mathbb{A}^n(F)\} \\ &\quad \cup \{[a_1, \dots, a_n, 0] \in \mathbb{P}^n(F)\}.\end{aligned}$$

The first set can be identified with affine n -space $\mathbb{A}^n(F)$, and the second set can be identified with a copy of $(n-1)$ -dimensional projective space $\mathbb{P}^{n-1}(F)$, a so-called *hyperplane at infinity*.

7.2. Curves in the Affine and Projective Plane

Starting with a given affine space, we have seen how to create a projective space that contains a copy of the affine space naturally as a subset. For example, we view $\mathbb{A}^2(\mathbb{R}) \hookrightarrow \mathbb{P}^2(\mathbb{R})$ by associating $(a, b) \in \mathbb{A}^2(\mathbb{R})$ with $[a, b, 1] \in \mathbb{P}^2(\mathbb{R})$. One advantage of projective space is that there is some extra room in the space to find perhaps missing points of intersections of curves. One somewhat awkward aspect is that the points in projective space are represented as equivalence classes of points in affine space, that is, our association $(a, b) \mapsto [a, b, 1]$ is a map that takes a single point in $\mathbb{A}^2(\mathbb{R})$ to an entire equivalence class of points in $\mathbb{A}^3(\mathbb{R})$ represented by $(a, b, 1)$. So what's the problem?

Naively, we start with a curve we understand in affine space, say the set of points in $\mathbb{A}^2(\mathbb{R})$ which satisfy $y - x^2 = 0$, a parabola. We want all the points of the parabola that we already know about in affine space to still be part of the set of points in $\mathbb{P}^2(\mathbb{R})$ which characterize the parabola. So if $(2, 4)$ is a point on the parabola in $\mathbb{A}^2(\mathbb{R})$, we need $[2, 4, 1]$ to be a point on the projective version of the parabola. But what can this mean? There is an extra coordinate in $[2, 4, 1]$, which means that to accommodate curves in the projective plane, we need to introduce an extra variable, so we should have an equation in x , y , and z which describes our parabola. And it would also be natural if, when $z = 1$, our new equation looked like the old one, $y - x^2 = 0$. This doesn't seem so hard. One solution would be $yz - x^2 = 0$, but so would $yz^2 - x^2 = 0$ or even $yz^2 - x^2z = 0$. There seem to be many possibilities; let's see if we can narrow the field.

We will start by saying that the first choice, $yz - x^2 = 0$, is the equation we really want, and while it is appealing in that it is the simplest of the three, we want to be sure it is really the correct one. Let's dismiss the choice $yz^2 - x^2 = 0$. One problem is that for a point to be on a projective curve means all the representatives of the point satisfy the equation. So for $[2, 4, 1]$ to be a point on the curve

means not simply that $(2, 4, 1)$ is a solution, but so is $(2t, 4t, t)$ for any nonzero real number t . Substituting $(2t, 4t, t)$ into $yz^2 - x^2 = 0$ gives $4t^3 - 4t^2 = 0 = 4t^2(t - 1)$. Certainly this fails to be true for most values of t . But what was wrong? The powers of t did not match. Do they in the other candidates?

$$\begin{aligned}yz - x^2 = 0 &\mapsto 4t^2 - 4t^2 = 0, \\ yz^2 - x^2z = 0 &\mapsto 4t^3 - 4t^3 = 0.\end{aligned}$$

Well these both pass that test, but is there something else to distinguish $yz - x^2 = 0$ from $yz^2 - x^2z = 0$? Well, the second equation factors as $z(yz - x^2) = 0$ which has our first candidate as part of it. Sure, we could go for simplicity, but there is actually more going on here. To understand it, let's go back to affine space where we have more insight.

The question boils down to whether there is a difference between the set of solutions to $y - x^2 = 0$ and $x(y - x^2) = 0$. The first one we agree is the parabola, but what about the second? Well, since the product must equal zero, either $x = 0$ or $y - x^2 = 0$, so we get all the points on the parabola, but we also get all the points where $x = 0$, namely the y -axis, so the set of solutions to $x(y - x^2) = 0$ is the union of the y -axis and the parabola. The same thing is happening in projective space. We get extra solutions coming from $z = 0$, that is the set of points $[a, b, 0]$ in the projective plane, which is the projective line at infinity. So $yz - x^2 = 0$ was indeed the correct choice.

Now the general paradigm is quite simple, and it is called *homogenization*, and what is being made homogeneous (uniform) is the degree of each summand. So for example if we have a complicated equation like

$$f(x, y) = 3x^5y^2 - 2xy^3 + 2x^2 + 7y^3 = 0,$$

which describes a curve in the affine plane, we homogenize the polynomial $f(x, y)$ by multiplying each summand by the smallest power of z so all terms have the same degree. Simply look at the original, and determine the highest degree term (in our case, that degree is 7) and multiply by powers of z to make all terms have degree 7, so

$$F(x, y, z) = 3x^5y^2 - 2xy^3z^3 + 2x^2z^5 + 7y^3z^4$$

is our appropriately homogenized polynomial. Every affine point (a, b) on the original curve corresponds to the point $[a, b, 1]$ on the associated projective curve

$$F(x, y, z) = 3x^5y^2 - 2xy^3z^3 + 2x^2z^5 + 7y^3z^4 = 0.$$

So now we have a larger space in which to work, and we know how to take a curve in the affine plane and write down the equation which describes the corresponding curve in the projective plane. Finally, it is time to see that we have actually gained something in this process. Let us show that any two distinct lines in the plane intersect in a unique point; in particular, let's show that two parallel lines intersect in the projective plane.

We shall prove some of the general results here, leaving the remaining arguments as an exercise. First we show by example that two lines in the plane which are not parallel intersect in a unique point in the projective plane, and it corresponds to the point of intersection they had in affine plane. Then we show that two lines in affine plane which are parallel, have a unique point of intersection in the projective plane, this time in the line at infinity.

First we start with two lines $y - 2x + 5 = 0$ and $y - 5x - 7 = 0$, the first having slope 2, the second slope 5, so we know that they will intersect. A little algebra shows they intersect at $(-4, -13)$. The homogenized equations are $y - 2x + 5z = 0$ and $y - 5x - 7z = 0$. To find all the points of intersection, we divide the search into those points in the copy of the affine plane (i.e., where $z = 1$) and those in the line at infinity (i.e., where $z = 0$). When $z = 1$, the algebra implied above shows that the only point is $[-4, -13, 1]$, so now we look for extra points in the line at infinity, so we set $z = 0$, and then must solve simultaneously $y - 5x = 0$ and $y - 2x = 0$. Perhaps you are going to point out that these are two nonparallel lines and so will intersect! But where? Clearly, the point is $(x, y) = (0, 0)$, but because this is in the line at infinity, the corresponding projective point would have to be $[0, 0, 0]$. If you go back to our original definition of $\mathbb{P}^2(\mathbb{R})$, we started with the set $\mathbb{A}^3(\mathbb{R}) \setminus (0, 0, 0)$, that is every point $[a, b, c]$ in the projective plane must have at least one nonzero coordinate. So

we gain no more points in the line at infinity, thus these two lines have only the affine point of intersection.

Now we turn to the more interesting case, that of two parallel lines. The general equation of a line in the affine plane, say $\mathbb{A}^2(\mathbb{R})$, is $ax + by + c = 0$ and a line parallel but not equal to it will have the equation $ax + by + d = 0$ with $c \neq d$. We also have that at least one of a, b is nonzero so that the equations actually define lines. The homogenized equations are $ax + by + cz = 0$ and $ax + by + dz = 0$. We know there is no solution in the copy of the affine plane (i.e., when $z = 1$), so we look in the line at infinity, by setting $z = 0$; thus any point of intersection will have the form $[x, y, 0]$. We see both equations reduce to $ax + by = 0$. Because this is a general equation, we have to handle two separate cases. If $a \neq 0$, then $x = -by/a$, so $[x, y, 0] = [-b/ay, y, 0] = [-b/a, 1, 0]$. If $b \neq 0$, then $y = -ax/b$, so $[x, y, 0] = [x, -ax/b, 0] = [1, -a/b, 0]$. So in either case, there is a single point of intersection (in the line at infinity), and we note that when a and b are both nonzero, that $[-b/a, 1, 0] = [1, -a/b, 0]$, so there is no ambiguity in representing that point in the projective plane.

It is important to note for our use in defining a group law on an elliptic curve that, in the notation above, every vertical line has the form $ax + by + c = 0$ where $b = 0$, which means it intersects the line at infinity at the point $[0, 1, 0]$. We shall also see that this point is the unique point at infinity on an elliptic curve $y^2 = x^3 + ax^2 + bx + c$, so in particular, every vertical line in the plane intersects the elliptic curve at the point $[0, 1, 0]$ (in the line at infinity).

7.3. Rational Points on Curves

Now let's consider another motivation for projective space and the notion of homogenization. In section 2.4, we exploited a correspondence between rational points on the unit circle $x^2 + y^2 = 1$ and Pythagorean triples. We extended that correspondence to the Fermat curves $x^n + y^n = z^n$ for $n > 2$. In view of Wiles's proof of the Fermat conjecture, we know that there are only a few rational points on $x^n + y^n = 1$ (when $n > 2$). But the correspondence between rational points on $x^n + y^n = 1$ and integral points on $x^n + y^n = z^n$ remains

valid and instructive, so let's push the correspondence a bit harder so as to provide a natural connection to projective space. While the discussion here is quite straightforward, it would be hard to improve on the outline in Appendix A of [ST92] which we amplify, but largely follow.

Let $(a/b, c/d)$ be a rational point on the curve $x^n + y^n = 1$ ($n \geq 2$) with $a, b, c, d \in \mathbb{Z}$, $\gcd(a, b) = \gcd(c, d) = 1$, and without loss of generality $b, d > 0$. Then

$$(a/b)^n + (c/d)^n = 1 \quad \text{implies that} \quad (ad)^n + (bc)^n = (bd)^n.$$

In the second equality, we see that b^n divides two of the summands, so it must divide the third $(ad)^n$. Also d^n divides two of the summands, so it must divide the third $(bc)^n$. Thus

$$b^n \mid (ad)^n \text{ and } \gcd(a, b) = 1 \text{ implies } b^n \mid d^n \text{ which implies } b \mid d, \text{ and} \\ d^n \mid (bc)^n \text{ and } \gcd(c, d) = 1 \text{ implies } d^n \mid b^n \text{ which implies } d \mid b.$$

Since $b \mid d$ and $d \mid b$, we must have $b = \pm d$, but since both are positive, we have $b = d$. So all rational points on $x^n + y^n = 1$ have the special form $(a/c, b/c)$ with $\gcd(a, c) = \gcd(b, c) = 1$ and $c > 0$ (yes; we changed the notation). Thus we have one-half of a correspondence between rational points on $x^n + y^n = 1$ and integral points on $x^n + y^n = z^n$:

$$\left(\frac{a}{c}, \frac{b}{c}\right) \text{ on } x^n + y^n = 1 \mapsto (a, b, c) \text{ with } a^n + b^n = c^n.$$

Conversely, an integer solution $a^n + b^n = c^n$ with $c \neq 0$ corresponds to a rational point $(a/c, b/c)$ on $x^n + y^n = 1$.

But this correspondence between rational points and integral points is far from one-to-one. In particular, if (a, b, c) (with $c \neq 0$) satisfies $a^n + b^n = c^n$, then so too does every point of the form (at, bt, ct) , $t \neq 0$. But all of these points correspond to exactly the same rational point $(a/c, b/c)$, so if we want to get a one-to-one correspondence, we would be forced to identify all the points (at, bt, ct) , with $t \neq 0$. In section 2.4, and in the case of $n = 2$, we resolved this ambiguity with the notion of a primitive Pythagorean triple. But it should be clear from all the work done in the previous section that this identification

corresponds more generally to viewing the triple (a, b, c) as a point $[a, b, c]$ in the projective plane.

This is also an excellent time to motivate the extra points in projective space. Recall that the projective plane contained a copy of the affine plane as well as a projective line at infinity. An issue not previously addressed is to consider the solutions to $a^n + b^n = c^n$ where $c = 0$. We can easily dismiss the case when $a = b = c = 0$ as uninteresting or irrelevant. However, when n is odd, there are nontrivial solutions, e.g., $a^n + (-a)^n = 0$ for any nonzero a . Under the conjectured correspondence these would seem to correspond to “rational” points (loosely) $(a/0, -a/0) = (\infty, \infty)$, certainly not what we would consider a typical point on the curve. But indeed, we shall see that all these solutions reduce to one extra point in the projective line at infinity.

To reiterate what we said above, if $(a, b, c) \in \mathbb{A}^3(\mathbb{Z})$ (or $\mathbb{A}^3(\mathbb{Q})$) is a nonzero solution to $x^n + y^n = z^n$, then so is (at, bt, ct) for every nonzero scalar t . That means that every point which lies in the equivalence class $[a, b, c] \in \mathbb{P}^2(\mathbb{Q})$ is also a solution, so it makes sense to talk about the projective point $[a, b, c]$ as a solution to $x^n + y^n = z^n$. Now let’s try to put all this together. We began by looking at rational points on $x^n + y^n = 1$. These had the form $(a/c, b/c)$ for $c > 0$. The homogenization of $x^n + y^n = 1$ is $x^n + y^n = z^n$, and those original solutions to the affine equation correspond to the projective points $[a/c, b/c, 1] = [a, b, c]$, that is, those points of $x^n + y^n = z^n$ which lie in the affine part of the projective plane. The solutions to $x^n + y^n = z^n$ with $c = 0$ (n odd) correspond to the single point $[a, -a, 0] = [1, -1, 0]$ which lies in the projective line at infinity in the projective plane. So projective space gives us a single container for all the solutions in which we really don’t have to distinguish the cases if we don’t want to.

Now let’s take a look at the type of curves that form the heart of this book, elliptic curves. We said that an elliptic curve for us will be the set of points (x, y) which satisfy an equation of the form $y^2 = x^3 + ax^2 + bx + c$, where the cubic is nonsingular, that is, it has distinct roots. By now, we know enough to homogenize the equation and look for solutions in projective space.

The homogenized equation is

$$y^2z = x^3 + ax^2z + bxz^2 + cz^3,$$

and the original (affine) points on the curve all have the form $[a, b, 1]$. What do we gain by viewing things in projective space? Are there additional solutions?

Any new solutions to this equation would have the form $[x, y, 0]$ which reduces the homogeneous equation to $x^3 = 0$. That means that there is only one extra point on the elliptic curve living in the line at infinity, namely $[0, y, 0] = [0, 1, 0]$.

Exercise. Consider the points of intersection of the affine curves $x = y^2$ and $y = -3$. As this is the intersection of a line and a conic, we expect at most two points, and indeed there is only one affine point $(9, -3)$. Find the points of intersection of the corresponding projective curves.

Exercise. Find the points of intersection of the parallel lines $y = 3x$ and $y = 3x + 1$ in $\mathbb{P}^2(\mathbb{R})$.

Exercise. Consider the intersection of the cubic $y = x^3$ and the line $y = x + 6$. We would like to see three points of intersection, but where are they?

One thing we learn from these exercises is that, given two curves defined over \mathbb{Q} or \mathbb{R} , or \mathbb{C} having degree m and n , respectively, to have any hope of finding an environment in which these curves consistently intersect in mn points, we must look in $\mathbb{P}^2(\mathbb{C})$, that is, the projective plane over the algebraically closed field \mathbb{C} .

7.4. The Group Law for Points on an Elliptic Curve

We have seen various sets endowed with an operation which gave the set the structure of a group: given two elements in the set, there is a rule by which to produce a third element of the set; and this operation has an identity and inverses, and it satisfies an associative

law. In this section we shall take the set of points on an elliptic curve and define an operation on them that will endow the set with the structure of a group. This is actually a remarkable achievement with broad ramifications, which are not at all obvious. Before digging in, we make a few general observations.

Suppose we start with an elliptic curve E given by the equation $y^2 = x^3 + ax^2 + bx + c$ where the coefficients are taken from some fixed field F . The choice of field may not be obvious from the defining equation. For example, suppose we start with an elliptic curve given by the equation $y^2 = x^3 - x$. This could be defined over \mathbb{Q} , \mathbb{R} , \mathbb{C} , or even \mathbb{F}_p . And so it makes sense to talk about the set of points on this elliptic curve whose coordinates both lie in a specific field F ; we shall denote this set of points by $E(F)$. So of course $E(\mathbb{Q}) \subset E(\mathbb{R}) \subset E(\mathbb{C})$, and $E(\mathbb{F}_p) \subset E(F)$ for any field F containing \mathbb{F}_p . It is interesting that considering elliptic curves over these different fields leads to distinctly different mathematical endeavors. We shall see that cryptographic applications are primarily interested in elliptic curves over finite fields. Elliptic curves over \mathbb{C} were perhaps the first to be studied and connected to the study of elliptic functions, like the Weierstrass \wp -function, which not only gives an interpretation of the group law on $E(\mathbb{C})$, but describes $E(\mathbb{C})$ as a complex torus. We talk a bit more about this in Appendix A, and we also consider the set of rational points $E(\mathbb{Q})$, which has no geometric structure, but an algebraic one. A theorem of Mordell tells us this set of points forms a finitely generated abelian group. This type of group is more general than the groups we have looked at so far; it simply means $E(\mathbb{Q}) \cong \mathbb{Z}^r \times E_0$ where E_0 is a finite abelian group. The group $E(\mathbb{Q})$ is intimately related to the Birch and Swinnerton-Dyer conjecture, which is one of the Millennium Problems (<http://www.claymath.org/millennium-problems>) whose correct solution will earn the solver one million dollars!

An important point we want to make is that the group law we will define respects the underlying field. That is, if you have an elliptic curve defined over a field F and start with two points $P, Q \in E(F)$, the group law will produce a third point $P \oplus Q$ both of whose coordinates are also in F . This may not seem so special now, but

perhaps it will, given that our insight for how to add points in $E(F)$ (for any field F) will come from geometric insight gleaned by looking at $E(\mathbb{R})$. However, once we have extracted the algebraic rules which define the group law based upon that geometric insight, we will see that the group law makes sense over any field, including a finite field where most cryptographic applications lie.

To gain this geometric insight, we first consider the set of real points on an elliptic curve. This set of points will either have one or two components, as shown in Figures 7.2 and 7.3 of $y^2 = x^3 + 2$ and $y^2 = x^3 - x$. This fact follows from the assumption that the cubic is nonsingular (has distinct roots in \mathbb{C}), and a cubic with real coefficients has at least one real root, so either all three roots are real or the other two roots are complex conjugates of each other. Figures 7.2 and 7.3 display these two possibilities.




Figure 7.2




Figure 7.3

Let's start with some basic observations. Suppose we are given two points P, Q on our elliptic curve. Euclidean geometry says that there is a unique line that passes through those two points. Our observations about Bézout's theorem suggest that a line and a cubic should intersect in three points (at least in $\mathbb{P}^2(\mathbb{C})$ and accounting for multiplicity). But a group law for the curve must take two points and produce a third on the curve. Could it be this simple? Actually, no, but it is not hard to see that the obvious guess is wrong, and investigating this first guess will give us a bit more insight.

Let $E(\mathbb{R})$ denote the set of real points on an elliptic curve. We have observed that under the right circumstances, a line and a cubic should intersect in at most three points, so given points P and Q in $E(\mathbb{R})$, consider the line through P and Q . If $P = Q$, we consider the tangent line to the curve at P . Let $P * Q$ denote the third point of intersection (which may, on occasion, turn out to be either P or Q). At least $(P, Q) \mapsto P * Q$ is a binary operation on $E(\mathbb{R})$. Figure 7.4 illustrates a typical situation.




Figure 7.4

First let's see that the operation $*$ is commutative, that is, $P*Q = Q*P$. Well, this is quite easy since the line through P and Q is the same as the line through Q and P , so where it intersects the cubic again is obviously the same. So far, so good.

The first wrinkle appears because there can be no element which acts as the identity for this operation. To see this, suppose there is a point which we will denote by $\mathbf{0}$ in $E(\mathbb{R})$ so that $\mathbf{0}*P = P$ for all points P on the curve. To say that $\mathbf{0}*P = P$ is to say that the line through $\mathbf{0}$ and P intersects the elliptic curve in a third point which happens to be P . This means the line is tangent to the curve at P . Since this happens for every P , it says that every tangent line to the curve at a point P intersects the elliptic curve at the same point $\mathbf{0}$, which is visibly not the case.

It is also easy to see that $*$ is not an associative operation. Consider points P and Q on the curve as in the figure above where P , Q and $P*Q$ are distinct points. What is $P*(P*Q)$? Well, it is the third point of intersection of the line through P and $P*Q$ with the

elliptic curve which is visibly Q , so $P * (P * Q) = Q$. If the operation $*$ were associative, then $Q = P * (P * Q) = (P * P) * Q$. That says take the tangent line to the curve at P , and find its point of intersection, $P * P$, with the curve. For the equality $Q = (P * P) * Q$ to hold, a line through $P * P$ and the point Q (as in Figure 7.4) would have to be tangent at Q , and while it certainly fails for this choice of Q , it would have to be true for all Q , which is clearly impossible.

So our first guess was not correct. But the good news is that it is not that far off, and actually there are many related ways in which to define a group operation. Fortunately, they all give rise to isomorphic groups, which allows us a degree of flexibility.

One deficiency we noted in our first attempt at defining a group law was that there was no identity element for the operation, and this is obviously an important matter. Perhaps there is a matter of even more concern. The applications we have in mind will have us look at $E(\mathbb{Q})$ and $E(\mathbb{F}_p)$, and we have no idea whether there are even any elements in this set, to say nothing of its elements forming a group! For example, some curves have no rational points:

Exercise. Show that the curve $x^2 + y^2 = 3$ (a conic) has no rational points, even when we extend the search to $\mathbb{P}^2(\mathbb{Q})$.

You see, we have been relying on our geometric intuition by looking at $E(\mathbb{R})$ where there are visibly lots of points, but we have not considered whether it is possible that $E(\mathbb{Q})$ or $E(\mathbb{F}_p)$ might be empty. For the moment, let's defer that issue. It will turn out that viewing our curve in projective space will allow us to show that $E(F)$ is nonempty no matter what field F we choose, and it will also provide a natural candidate for the element $\mathbf{0}$ that we wish to designate as the identity element of our group. So for now, we need only assume we have fixed an element $\mathbf{0}$ in our set $E(F)$. To move forward, we return to the case where $F = \mathbb{R}$, give our revised binary operation, and outline the ideas which show that the operation gives $E(\mathbb{R})$ the structure of a group. Details on the group law can be found in [ST92], but pictorially it is given in Figure 7.5 (which is a rendition of the cover of their text).




Figure 7.5

Now in words, the variation we need is to take two points P, Q and to find the third point of intersection with the elliptic curve, and to label it $P * Q$. Then using our distinguished point $\mathbf{0}$, we define

$$P \oplus Q := \mathbf{0} * (P * Q),$$

that is, $P \oplus Q$ is the third point of intersection of the line through $\mathbf{0}$ and $P * Q$ with the elliptic curve. As $P * Q = Q * P$, it follows that $P \oplus Q = Q \oplus P$, so we have a commutative operation.

Let's first see that our choice of $\mathbf{0}$ makes sense as the additive identity of the group. We want to show that $P \oplus \mathbf{0} = P$ for all points P , so we consider the line through P and $\mathbf{0}$ as pictured in Figure 7.6.

The point $P * \mathbf{0}$ is the third point of intersection, and we now take the line through $\mathbf{0}$ and $P * \mathbf{0}$. Its intersection with the curve is clearly P , which establishes the result.

To understand how additive inverses work, first construct the tangent line to the elliptic curve at the point $\mathbf{0}$ as in Figure 7.7; the




Figure 7.6




Figure 7.7

condition that our elliptic curve is nonsingular guarantees its existence. Let Q denote the third point of intersection of the tangent line and the elliptic curve.

Next choose any point P and draw line through P and Q . The claim is that $P * Q$ is the additive inverse $-P$ of P . To check we need to compute $P \oplus (P * Q)$. As above we first compute the point of intersection $P * (P * Q)$ of the line through P and $P * Q$ with the elliptic curve. With the help of the diagram, we see this is the point Q . Then by our above definition, $P \oplus (P * Q)$ is the point of intersection of the line through $\mathbf{0}$ and the point $Q = P * (P * Q)$. But this point of intersection is $\mathbf{0}$ (the point of tangency counts as two points of intersection in the context of Bézout's theorem), so indeed, $P * Q$ is the additive inverse of P .

The only thing remaining in order to confirm that $E(\mathbb{R})$ is an abelian group with $\mathbf{0}$ as the identity is to verify that this addition law is associative, and this is not at all a trivial matter. While we will not give a detailed proof, we will indicate several ways it can be verified and hopefully give some appreciation of why this is a complicated, yet very interesting result. Perhaps Figure 7.8 would help as a start.

Admittedly it is a somewhat daunting image, but let's see what is involved. We want to show that $(P \oplus Q) \oplus R = P \oplus (Q \oplus R)$ for any points P, Q, R on the elliptic curve. Backing up one step, we need to show that the line through $\mathbf{0}$ and $(P \oplus Q) * R$ intersects the elliptic curve at the same place as the line through $\mathbf{0}$ and $P * (Q \oplus R)$. This would be trivial if we show that $(P \oplus Q) * R = P * (Q \oplus R)$, which is evident at the bottom right of the figure.

To ease your way into the figure, start with the points P, Q, R and compute $P \oplus Q$ and $Q \oplus R$, each of which requires two lines. Computing $(P \oplus Q) * R$ and $P * (Q \oplus R)$ each requires one line, and now we are in the lower right corner. The line between these common points and $\mathbf{0}$ intersects the curve at $(P \oplus Q) \oplus R = P \oplus (Q \oplus R)$.

Now a picture does not a proof make, and while we can turn the pictures characterizing the other properties of a group into actual proofs, the associative law is much more complicated. One way to deal with this complication will come in a few pages when we write




Figure 7.8

down formulas for the coordinates of $P \oplus Q$ in terms of the coordinates of P and Q . Then tedious algebraic computations can verify that the associative law holds. Also in Appendix A (where we discuss elliptic curves over \mathbb{C} and their connection to elliptic functions), another proof is outlined. For now, we take time to give enough context to understand how the proof of associativity comes about and to give connections of this result to other results from classical geometry of curves.

It probably comes as no surprise, but to motivate this, we make a slight digression. We know that two points in the plane determine a unique line which contains them. What are some generalizations? In calculus, if you study Simpson's rule, you learn that three points in the plane (no two on a vertical line) lie on the graph of a uniquely determined quadratic, $y = ax^2 + bx + c$; we note that a can be zero if the points are colinear. Can we ask for more? Yes! And the answer is called *Lagrange interpolation*. The statement is that given $n + 1$ points in the plane $\mathbb{A}^2(\mathbb{R})$, no two on a vertical line, there is a unique polynomial $p(x) = c_n x^n + \cdots + c_1 x + c_0$ of degree at most n , so that

all the points lie on the graph of $y = p(x)$. There are two aspects to proving such a statement: the existence of a polynomial, and that it is unique. Let's make existence an exercise and talk about uniqueness, since that bears at least suggestively on the problem of the associative law.

Exercise. We want to define a polynomial p of degree (at most) n which passes through $(a_1, b_1), \dots, (a_{n+1}, b_{n+1})$ where we assume all the a_i 's are distinct. Suppose we could define polynomials q_i having degree at most n so that $q_i(a_i) = 1$ and $q_i(a_j) = 0$ for $i = 1, \dots, n+1$ and $j \neq i$. Show that we may take $p = b_1q_1 + \dots + b_{n+1}q_{n+1}$. To construct the q_i , let's provide some guidance. You may or may not recall that if a polynomial $q(x)$ with real coefficients has a root a , then $(x - a)$ is a factor, that is $q(x) = (x - a)q_0(x)$. Even if you don't recall that fact, it should be clear that the polynomial $q(x) = (x - a_2)(x - a_3) \cdots (x - a_{n+1})$ has the property that $q(a_i) = 0$ for $i = 2, \dots, n + 1$, and that $q(a_1) = (a_1 - a_2)(a_1 - a_3) \cdots (a_1 - a_{n+1}) \neq 0$ precisely because the a_i 's are all distinct. So we may take $q_1(x) = q(x)/q(a_1)$ as the first element of our set. The others are similarly constructed.

Now as we said, it is the uniqueness that we want to discuss. The fact upon which we shall rely is that a nonzero polynomial $p(x)$ with coefficients in a field having degree n has at most n roots in the field. One proof is a consequence of establishing a division algorithm for such polynomials, but of course it is also a special case of Bézout's theorem since the roots correspond to the points of intersection of the curve of degree n , $y - p(x) = 0$ and the line that is the x -axis. Suppose we have two solutions to the Lagrange interpolation problem, that is, we have two polynomials p_1 and p_2 of degree at most n , satisfying $p_1(a_i) = p_2(a_i)$ for $i = 1, \dots, n+1$. Consider the polynomial $p(x) = p_1(x) - p_2(x)$. It has the property that its degree is at most n , yet $p(a_i) = 0$ for all i , that is, it has $n + 1$ roots, more than its degree. The only way out is that $p(x)$ is the zero polynomial, meaning $p_1(x) = p_2(x)$, so the solution is unique.

Exercise. Show that in contrast to the uniqueness result which Lagrange interpolation provides, there are an infinite number of polynomials of degree $n + 1$ which pass through the $n + 1$ given points.

Now that we have made this digression, we come back to the associative law. Replacing Lagrange interpolation, we have Bézout's theorem which talks about the number of points of intersection of two plane curves. We have agreed that the upper bound for the number of points of intersection for *nice* curves is the product of the degrees of the curves, and we suggested the Bézout bound is achieved in $\mathbb{P}^2(\mathbb{C})$. First we comment about what “nice” means. In a nutshell, if the curves are given as zero sets $Z(f)$ and $Z(g)$, then the polynomials f and g cannot have any common factors. Consider an example where this condition fails. The curves given as the zero sets of $f(x, y) = y^2 - x^2$ and $g(x, y) = y^3 - x^3$ have degrees 2 and 3, respectively, so the Bézout bound would be 6. But the polynomials that define these two curves share a common factor, $y^2 - x^2 = (y - x)(y + x)$ and $y^3 - x^3 = (y - x)(y^2 + xy + y^2)$, so we see that every point on the line $y = x$ lies on both curves, giving them an infinite number of points of intersection. When we preclude this from happening (and we view things in $\mathbb{P}^2(\mathbb{C})$), the curves will intersect in precisely $\deg(f) \deg(g)$ points, counted with appropriate multiplicities. One way to ensure there are no common factors is to require one of the curves, say $Z(f)$, to be given by an irreducible polynomial f which is not a divisor of g .

The theorem is invoked to establish the validity of the associative law for the addition of points on elliptic curves is the following (see Chapter 5, §6 of [Ful69]).

Theorem. *Let C, C_1, C_2 be cubic curves in $\mathbb{P}^2(\mathbb{C})$, and assume that C is irreducible. Suppose that C and C_1 intersect in the nine points P_1, \dots, P_9 (and that C has tangent lines at each P_i). If C and C_2 intersect in nine points, eight of which are among the P_i , then the ninth point of intersection is also among the P_i .*

So in loose analogy with Lagrange interpolation where a polynomial of degree n is completely determined by any $n + 1$ points through

which it passes, the intersections of these pairs of cubic curves are completely constrained once there are eight common points of intersection.

In its application to the associative law for the group of points on an elliptic curve, the curve C is the elliptic curve which is irreducible and nonsingular, meeting the criteria of the theorem. The other cubics (remember a cubic is just a homogeneous polynomial of degree 3) arise as the product of three lines (degree 1 curves) coming from our diagram trying to compute $P \oplus Q \oplus R$ in two different ways.

Before leaving this thread, we comment that results closely related to the theorem above (and proven in [Ful69]) give two classical results in projective geometry: Pascal's and Pappus's theorems which involve the intersection of a cubic and a conic. We state them for interest and reference.

Theorem (Pascal). *Let C be a conic (e.g., an ellipse, hyperbola, parabola). Choose any six distinct points on C , and join them with line segments in any order to form a hexagon. Then each pair of opposite sides of the hexagon (extended if necessary) intersect in a point, and those three points all lie on a straight line, called the Pascal line of the hexagon.*

An example is shown in Figure 7.9.

Now the reader may wonder what happens if, for example, we inscribe a hexagon in an ellipse in such a way that the opposite sides are parallel. Where do the lines extending opposite sides intersect? And then we remember this is a theorem about projective geometry, so that answer should now be clear.

Pappus's theorem is, in a sense, a degenerate version of Pascal's theorem in which the conic is reducible (the product of two lines).

Theorem (Pappus). *Consider two lines in the plane and choose points A, E, C on one and D, B, F on the other. The opposite sides of the hexagon $ABCDEF$ intersect in three colinear points.*

An example is shown in Figure 7.10.




Figure 7.9

Now that we have given reasonable justification that our addition law makes the set of real points on an elliptic curve into an abelian group, we return to the question of what choice to make for our distinguished point $\mathbf{0}$, given that we have essentially no intuition about $E(F)$ for arbitrary fields F . In what follows, we point out a number of nice coincidences which suggest a very natural choice for $\mathbf{0}$. One thing that works in our favor is that is we have taken a definition of an elliptic curve as the set of solutions to $y^2 = x^3 + ax^2 + bx + c$, where the cubic is nonsingular. This is not the most general setting we could adopt, but for us there is no real loss.

We showed by homogenizing that equation, that the point $[0, 1, 0]$ is on every such elliptic curve; it is the only point on each curve that lies in the line at infinity, and it is clearly a point that is in $E(F)$ for any field F . So we have at least one distinguished point on every single one of our elliptic curves over any field, so it should not come as a surprise then that we will choose $[0, 1, 0]$ for our identity $\mathbf{0}$.




Figure 7.10

There are other interesting consequences of choosing the point $[0, 1, 0]$ as the identity. Consider Figure 7.11.

When we intersect our elliptic curve with vertical lines such as the ones through P_1, P_2 and Q_1, Q_2 , what is the third point of intersection? It does not appear to be anything in the affine plane, but when we talked about lines in the projective plane, we noted that they intersect the line at infinity in a unique point. We showed that the line $ax + by + c = 0$ intersects the line at infinity at $[0, 1, 0]$ if $b = 0$ (a vertical line), at $[1, 0, 0]$ if $a = 0$ (a horizontal line), or at




Figure 7.11

$[-b/a, 1, 0] = [1, -a/b, 0]$ if $a, b \neq 0$. In particular, every vertical line intersects the elliptic curve at the point $\mathbf{0} = [0, 1, 0]$ at infinity.

So if we take $\mathbf{0} = [0, 1, 0]$ as our identity, then the general addition rule we gave above says that for points P_1 and P_2 on the curve which also lie on a vertical line, $P_1 * P_2 = \mathbf{0}$, and so $P_1 \oplus P_2 = \mathbf{0} * (P_1 * P_2) = \mathbf{0} * \mathbf{0}$. Well, what is $\mathbf{0} * \mathbf{0}$? This is a bit subtle without formal definitions of intersection multiplicity, but we wish to claim that $\mathbf{0} * \mathbf{0} = \mathbf{0}$, that is, the third point of intersection of the line through $\mathbf{0}$ and $\mathbf{0}$ (the tangent line) with the elliptic curve is also $\mathbf{0}$, so $\mathbf{0}$ is a triple point. A plausible explanation is the following. All of the lines we explored previously have the property that they intersect the elliptic curve either in three points in the affine plane ($\mathbb{A}^2(\mathbb{C})$), or in a unique point in the line at infinity and at two points in the affine plane (counting points of tangency with multiplicity two). But $\mathbf{0} * \mathbf{0}$ is the point of intersection of the line through $\mathbf{0}$ and $\mathbf{0}$ with the elliptic curve. But since two of the points of intersection are already in the line at infinity, the line in question cannot be one of the lines we considered before.

So that line is not an affine line, but instead the line at infinity, which intersects the elliptic curve only in one point, $\mathbf{0}$. All of this is to say that $P_1 \oplus P_2 = \mathbf{0}$. So in general, points $P = (x, y)$ and $Q = (x, -y)$ (on the elliptic curve), which are symmetric across the x -axis, sum to $\mathbf{0}$.

Now we recast the general addition law we gave before in this new context, with the distinguished point $\mathbf{0}$ (the identity) located at the point at infinity on the curve. Pictorially, the addition law is shown in Figure 7.12.

As before, we take points P and Q and find their point of intersection with the curve, denoted $P * Q$. Then we take the line through $\mathbf{0}$ and $P * Q$ and designate the third point of intersection as $P \oplus Q$. In this special case, the line through $\mathbf{0}$ and $P * Q$ is a vertical line, so we know where it intersects the elliptic curve: if $P * Q = (x, y)$, then $P \oplus Q = (x, -y)$. Shortly, we will write down formulas for the coordinates for $P \oplus Q$ given the coordinates of P and Q , but first we




Figure 7.12

firm up our geometric intuition about special cases of the addition law and shift our view fully to projective space.

As usual, we will draw pictures in $\mathbb{A}^2(\mathbb{R})$, but the discussion will reflect what happens in $\mathbb{P}^2(\mathbb{C})$. We have set $\mathbf{0} = [0, 1, 0]$, and for a point $P = [x, y, 1]$ on the elliptic curve (and in the affine plane), its reflection across the x -axis $Q = [x, -y, 1]$ is also a point on the elliptic curve. Since these two points lie on a vertical line, they sum to $\mathbf{0}$, that is $Q = -P = [x, -y, 1]$ is the additive inverse of $P = [x, y, 1]$.

Sometimes it is convenient to think of the group law on the elliptic curve in terms of the three points of intersection of a line with the elliptic curve. For example, in the generic case, we have something like Figure 7.13.

What is $P \oplus Q \oplus (P * Q)$? Our definition of $P \oplus Q$ is $\mathbf{0} * (P * Q)$, that is the third point of intersection of the (vertical) line through $\mathbf{0}$ and $P * Q$ with the elliptic curve. We have just seen that this is the




Figure 7.13




Figure 7.14

point $-(P * Q)$, so

$$P \oplus Q \oplus (P * Q) = (P \oplus Q) \oplus (P * Q) = -(P * Q) \oplus (P * Q) = \mathbf{0},$$

and under our group law, the sum of any three colinear points which lie on the elliptic curve is $\mathbf{0}$.

This remains true in all special cases we have seen. For example, in the case of a vertical line as in Figure 7.14, we have seen that $P_1 \oplus P_2 = \mathbf{0} = Q_1 \oplus Q_2$, but if we reinterpret this as the sum of the three points of intersection of a vertical line with the elliptic curve, the equations would be

$$P_1 \oplus P_2 \oplus \mathbf{0} = P_1 \oplus P_2 = \mathbf{0} = Q_1 \oplus Q_2 \oplus \mathbf{0} = Q_1 \oplus Q_2.$$




Figure 7.15

The remaining special cases regard tangent lines which have an intersection point of multiplicity 2 with the elliptic curve. So in Figure 7.15 we have that $Q \oplus Q \oplus R = \mathbf{0}$ and $P \oplus P \oplus \mathbf{0} = \mathbf{0}$, with the second equation generally written as $2P = \mathbf{0}$.

Remark 7.1. This is a good point at which to reflect upon the Barchet duplication formula introduced in Chapter 2. Using our current notation, Barchet took a point P on the curve $y^2 = x^3 + k$ and found a formula for $P * P$. In looking at our example from that chapter, the intent was to iterate to process, generating a sequence of points, $P, P_2 := P * P, P_4 := P_2 * P_2 = (P * P) * (P * P)$, and so on.

Our interest with elliptic curves is analogous. We want to start with a single point P and compute $2P = P \oplus P$, then $3P = P \oplus 2P = P \oplus P \oplus P$, and in general $kP = \underbrace{P \oplus P \oplus \cdots \oplus P}_{k \text{ summands}}$. Sometimes for some $k > 1$, we find that $kP = \mathbf{0}$ as in our case of the vertical tangent above, where $2P = \mathbf{0}$. If this occurs, we say that P is a *torsion point* on the curve, and its order is the order of the point as an element of

the group, namely the smallest positive k for which $kP = \mathbf{0}$. If kP is never equal to zero for $k = 1, 2, \dots$, we say that P is a point of *infinite order*.

7.5. A Formula for the Group Law on an Elliptic Curve

While our investigation of elliptic curves has been mostly theoretical so far, our applications require that we do arithmetic with the points on an elliptic curve and, even more importantly, we need to be able to have a computer do the computations. So that means we need to write down the addition law for points on an elliptic curve in terms of formulas that can be coded.

For the most part, we have thus far considered rational, real, and complex points on an elliptic curve, but for applications we will need to consider elliptic curves defined over a finite field, such as \mathbb{F}_p , with p a prime. The outline we give here is for primes $p \neq 2, 3$. For finite fields that contain \mathbb{F}_2 or \mathbb{F}_3 , the equation describing an elliptic curve needs to be generalized slightly from what appears below, as does the notion of the discriminant (see below). A careful exposition is Chapter VI of [Kob87a].

So for a field F containing either \mathbb{Q} , or \mathbb{F}_p with $p \neq 2, 3$, we have said that an elliptic curve can be given by an equation of the form $y^2 = x^3 + Ax^2 + Bx + C$, and indeed we claim we can further assume the curve has the form $y^2 = x^3 + ax + b$ with $a, b \in F$. The latter observation follows from the easy exercise below.

Exercise. Consider the polynomial $z^3 + Az^2 + Bz + C$, and let $z = x - A/3$. Show that under this substitution $z^3 + Az^2 + Bz + C$ becomes $x^3 + ax + b$ with

$$a = \frac{1}{3}(3B - A^2), \quad b = \frac{1}{27}(2A^3 - 9AB + 27C).$$

Associated to any polynomial $p(x) = a_n x^n + \dots + a_1 x + a_0$ with coefficients in a field F is a quantity called its *discriminant*, denoted Δ or $\Delta(p)$ which is described in terms of the roots of p . Let's start with a polynomial with coefficients in a field $F \subseteq \mathbb{C}$, so we know that

p has n roots r_1, \dots, r_n in \mathbb{C} (which do not necessarily have to be distinct). The discriminant of $p(x) = a_n x^n + \dots + a_1 x + a_0$ is defined to be

$$\Delta = a_n^{2n-2} \prod_{1 \leq i < j \leq n} (r_i - r_j)^2.$$

We remark that an analogous statement for an arbitrary field F can be made if we understand the notion of an algebraically closed field containing F . The complex numbers \mathbb{C} play this role for fields like \mathbb{Q} and \mathbb{R} .

Exercise. Using the quadratic formula to find expressions for the roots, show that the discriminant of the quadratic $ax^2 + bx + c$ is $\Delta = b^2 - 4ac$.

Hopefully, the discriminant tells us something interesting about the polynomial in question. For example if our quadratic $ax^2 + bx + c$ has real coefficients, what does the discriminant $\Delta = b^2 - 4ac$ tell us? By examining the quadratic formula, we see that if $\Delta > 0$, the polynomial has two distinct real roots. If it is negative, the quadratic has two complex conjugate roots, and if it is zero, it has a double real root.

One thing that is evident from the formula above is that if the polynomial has a multiple root, then we will have $r_i = r_j$ for some $i \neq j$ which means the discriminant Δ will equal zero. It is also a fact that the discriminant, while defined in terms of roots which may not lie in the original field F , can actually be expressed in terms of the coefficients of the original polynomial just as $\Delta = b^2 - 4ac$ is expressed in terms of the coefficients of $ax^2 + bx + c$.

Exercise. In the exercise above, we have shown how to transform a polynomial of the form $z^3 + Az^2 + Bz + C$ into one of the form $x^3 + ax + b$. Show that these two related cubics have exactly the same discriminant. *Hint:* The roots of the two polynomials are related by a simple formula.

Given this preamble, we state without proof that the discriminant of the cubic $x^3 + ax + b$ is $\Delta = 4a^3 + 27b^2$, and the curve $y^2 = x^3 + ax + b$ defines an elliptic curve when the cubic is nonsingular

(has distinct roots). We see now that this condition is determined by the discriminant Δ being nonzero.

We have seen that the projective (homogenized) curve $y^2z = x^3 + axz^2 + bz^3$ has one point $\mathbf{0} = [0, 1, 0]$ at infinity (where all vertical lines intersect the curve), and all other points on the curve are affine. Since $\mathbf{0}$ is the identity of the group (so we know $P \oplus \mathbf{0} = P$ for any point P), we need only consider how to find the sum of two points P_1 and P_2 in the affine plane, say $P_i = [x_i, y_i, 1]$, $i = 1, 2$. From the geometric description of the group law, we let L be the line between the two points (the tangent line if $P_1 = P_2$). We have already seen that if P_1 and P_2 lie on vertical line (and are points on the curve), then $P_1 \oplus P_2 = \mathbf{0}$, so we assume the line L is not vertical. The slope of the line is

$$m = \begin{cases} \frac{y_2 - y_1}{x_2 - x_1} & \text{if } x_1 \neq x_2, \\ \frac{3x_1^2 + a}{2y_1} & \text{if } x_1 = x_2, \end{cases}$$

where the slope of the tangent line is obtained by implicit differentiation of the equation $y^2 = x^3 + ax + b$ and evaluation at (x_1, y_1) . So the equation of the line L through P_1 and P_2 is

$$y - y_1 = m(x - x_1).$$

Solving for y and substituting into $y^2 = x^3 + ax + b$, we obtain

$$(m(x - x_1) + y_1)^2 = x^3 + ax + b.$$

Expanding and collecting terms, we find

$$x^3 - m^2x^2 + Cx + D = 0,$$

where expressions for C and D are not needed for what we do below.

Now the roots of this cubic are precisely the x -coordinates of the three points of intersection, x_1 , x_2 (possibly equal), and x_3 the x -coordinate of the third point of intersection of the line and the elliptic curve. Thus,

$$\begin{aligned} x^3 - m^2x^2 + Cx + D &= (x - x_1)(x - x_2)(x - x_3) \\ &= x^3 - (x_1 + x_2 + x_3)x^2 + Cx + D, \end{aligned}$$

from which we conclude (comparing coefficients of x^2) that $x_3 = m^2 - x_1 - x_2$. Since $y_3 = m(x_3 - x_1) + y_1$, the point of intersection

of the line L with the curve is (x_3, y_3) , so $P_1 \oplus P_2 = [x_3, -y_3, 1]$ by our geometric rule.

We summarize all of our observations.

Summary of the Group Law for $E : y^2 = x^3 + ax + b$.

We have $\mathbf{0} = [0, 1, 0]$ is the identity for the group, so $\mathbf{0} \oplus P = P = P \oplus \mathbf{0}$ for any point P on the curve.

- (1) If $P = [x, y, 1]$ is an affine point on the elliptic curve, then its inverse is $-P = [x, -y, 1]$.
- (2) If $P_1 = [x_1, y_1, 1]$ and $P_2 = [x_2, y_2, 1]$ are any two affine points on the elliptic curve, then $P_2 = -P_1$ if and only if $P_1 \oplus P_2 = \mathbf{0}$, which is true if and only if P_1 and P_2 are the two affine points of intersection of some vertical line with the elliptic curve, so $x_1 = x_2$, $y_1 = -y_2$.
- (3) If $P_2 \neq -P_1$, then $P_1 \oplus P_2 = P_3 = [x_3, y_3, 1]$, where

$$\begin{aligned} x_3 &= m^2 - x_1 - x_2, \\ y_3 &= -m(x_3 - x_1) - y_1, \text{ and} \\ m &= \begin{cases} \frac{y_2 - y_1}{x_2 - x_1} & \text{if } x_1 \neq x_2, \\ \frac{3x_1^2 + a}{2y_1} & \text{if } x_1 = x_2. \end{cases} \end{aligned}$$

These rules make the set of points on the curve into an abelian group under \oplus . Note that if the coordinates of P_1 and P_2 all lie in a fixed field F (e.g., \mathbb{Q} or \mathbb{F}_p), then the coordinates of $P_1 \oplus P_2$ also lie in that same field since the computations for x_3 and y_3 only involve arithmetic in F . So we write $E(F)$ for the group of F -rational points on the curve, and what we have said above is that this set of points is an abelian group under the operation of \oplus , with the point at infinity, $\mathbf{0} = [0, 1, 0]$, as the identity element of the group.

We note that the formulas for the slope given above still make sense even when we are in a finite field containing \mathbb{F}_p . Both formulas require us to divide, so we should check that we are in situations where the denominator is nonzero. If it is, this means the denominator is a unit in F , so has a multiplicative inverse. So in the first expression

for m , we see the denominator $x_2 - x_1$ is not zero precisely because we are in the case that $x_1 \neq x_2$, but what about the expression $2y_1$ in the case $x_1 = x_2$? If $F \supseteq \mathbb{F}_2$, we would have $2y_1 = 0$, which would cause a problem and is one of the reasons these formulas apply only to fields containing \mathbb{Q} , or \mathbb{F}_p with $p \neq 2, 3$. So since $p \neq 2, 3$, we have that 2 is a unit in F , so it creates no problem, but couldn't $y_1 = 0$? Not in the case where this formula is valid: Consider our curve $y^2 = x^3 + ax + b$. For a given value of x , there are at most two solutions to the equation which differ by a sign. But if $y_1 = 0$ and $x_1 = x_2$, then y_2 is also zero, so $P_1 = [x_1, 0, 1] = P_2 = -P_2$, and this case is explicitly excluded in item (3) above.

Since our interest in applications is to consider elliptic curves over finite fields, let's work through some examples. Consider curves $y^2 = x^3 + ax + b$ where a, b run over \mathbb{F}_7 with $\Delta = 4a^3 + 27b^2 \neq 0$. Since there are seven potential values for each of a, b , there are at most 49 possible elliptic curves of this form over \mathbb{F}_7 , although there are definitely fewer since $\Delta = 4a^3 + 27b^2$ can be zero. For each fixed curve, $y^2 = x^3 + ax + b$, there are seven values for x , and for each x , 0, 1, or 2 values of y which satisfy the equation, plus the point at infinity, so $E(\mathbb{F}_7)$ has size no larger than 15 and generally less. We will discuss the expected size of this group in just a bit.

Example 7.2. Consider the example of the cubic curve $E : y^2 = x^3 + 5x + 2$ over the field \mathbb{F}_7 . The discriminant Δ of the curve is

$$\begin{aligned} \Delta &= 4a^3 + 27b^2 = 4 \cdot 5^3 + 27 \cdot 2^2 \\ &\equiv 4 \cdot (-2)^3 + 6 \cdot 2^2 \equiv -8 \equiv 6 \not\equiv 0 \pmod{7}, \end{aligned}$$

so this is an elliptic curve over \mathbb{F}_7 . We check directly that

$$\begin{aligned} E(\mathbb{F}_7) &= \{\mathbf{0}, [0, \pm 3, 1], [1, \pm 1, 1], [3, \pm 3, 1], [4, \pm 3, 1]\} \\ &= \{\mathbf{0}, [0, 3, 1], [0, 4, 1], [1, 1, 1], [1, 6, 1], [3, 3, 1], [3, 4, 1], [4, 3, 1], [4, 4, 1]\}. \end{aligned}$$

Thus $E(\mathbb{F}_7)$ is an abelian group of order 9 which, by the Fundamental Theorem on Finite Abelian Groups, is isomorphic either to $\mathbb{Z}_3 \times \mathbb{Z}_3$ or to \mathbb{Z}_9 .

Consider the point $P = [4, 4, 1]$ on the curve. To find the order of P , we need to find the smallest positive integer k so that $kP =$

$P \oplus P \oplus \cdots \oplus P$ (k times) = $\mathbf{0}$. We begin by computing $2P$ using our formulas, though we know that by Lagrange's theorem, the order of P must be 3 or 9.

To double the point P , we use the slope of the tangent line in our formulas above:

$$m = \frac{3x_1^2 + a}{2y_1} = \frac{3 \cdot 4^2 + 5}{2 \cdot 4} = \frac{4}{1} = 4 \text{ in } \mathbb{F}_7.$$

Thus, $2P = [x_3, y_3, 1] = [1, 1, 1]$ (by our formulas). Similarly, we use the formulas to compute:

$$\begin{aligned} P &= [4, 4, 1], \\ 2P &= [1, 1, 1], \\ 3P &= P \oplus 2P = [3, 4, 1] \neq \mathbf{0}, \\ 4P &= 2(2P) = [0, 3, 1], \\ 5P &= P \oplus 4P = [0, 4, 1] = [0, -3, 1] = -4P, \\ 6P &= 2P \oplus 4P = [3, 3, 1] = [0, -4, 1] = -3P, \\ 7P &= P \oplus 6P = [1, 6, 1] = [1 - 1, 1] = -2P, \\ 8P &= 2(4P) = [4, 3, 1] = [4, -4, 1] = -P, \\ 9P &= P \oplus 8P = P \oplus -P = [0, 1, 0] = \mathbf{0}. \end{aligned}$$

Since we know that the order of the point P is either 3 or 9, and $3P \neq \mathbf{0}$, we know that P has order 9, so the group, $E(\mathbb{F}_7)$, is cyclic of order 9, with $P = [4, 4, 1]$ one of the possible generators. There are five other generators.

Example 7.3. Consider another example over \mathbb{F}_7 , this time the curve given by $E : y^2 = x^3 + 3x$. The discriminant of the curve is

$$\Delta = 4a^3 + 27b^2 = 4 \cdot 3^3 + 0 \equiv 3 \pmod{7},$$

so this equation defines an elliptic curve over \mathbb{F}_7 . We check directly that

$$\begin{aligned} E(\mathbb{F}_7) &= \{\mathbf{0}, [0, 0, 1], [1, \pm 2, 1], [2, 0, 1], [3, \pm 1, 1], [5, 0, 1]\} \\ &= \{\mathbf{0}, [0, 0, 1], [1, 2, 1], [1, 5, 1], [2, 0, 1], [3, 1, 1], [3, 6, 1], [5, 0, 1]\}. \end{aligned}$$

This time $E(\mathbb{F}_7)$ is an abelian group of order 8, so isomorphic to one of \mathbb{Z}_8 , $\mathbb{Z}_4 \times \mathbb{Z}_2$, or $\mathbb{Z}_2 \times \mathbb{Z}_2 \times \mathbb{Z}_2$.

We check that there is no point having order 8, so the group is not cyclic. With $P = [1, 2, 1]$, we compute

$$\begin{aligned} P &= [1, 2, 1], \\ 2P &= [2, 0, 1], \\ 3P &= P \oplus 2P = [1, 5, 1] = [1, -2, 1] = -P, \\ 4P &= \mathbf{0}, \end{aligned}$$

so we know that $E(\mathbb{F}_7) \cong \mathbb{Z}_4 \times \mathbb{Z}_2$. If we put $Q = [5, 0, 1]$ (or any or any of the remaining points), we can generate the rest of the group. Note that since Q has y -coordinate equal to 0, when we double the point $2Q$, we are taking a vertical tangent, and we have seen above that $2Q = \mathbf{0}$. We finish the group with

$$\begin{aligned} Q &= [5, 0, 1], \\ P \oplus Q &= [3, 6, 1], \\ 2P \oplus Q &= [0, 0, 1], \\ 3P \oplus Q &= [3, 1, 1]. \end{aligned}$$

Exercises. Consider the following elliptic curves over \mathbb{F}_7 , and determine the set of points on the curve and its structure as an abelian group. Write out how to determine all the elements of the group in terms of the generators you choose.

- (1) $E : y^2 = x^3 + 3x + 6$.
- (2) $E : y^2 = x^3 + 2$.
- (3) $E : y^2 = x^3 + 4$.

7.6. The Number of Points on an Elliptic Curve

A question which is important as we look towards how to use elliptic curves in a cryptographic setting is how large $E(\mathbb{F}_p)$ can be. For simplicity, let's restrict our discussion to primes $p > 3$. Let $x_0 \in \mathbb{F}_p$. When we substitute x_0 into $y^2 = f(x)$, there are several possibilities. If $f(x_0)$ is not a square in \mathbb{F}_p , then we have no points of the form $[x_0, y, 1]$ on the curve. If $f(x_0) = 0$, we get one, namely $[x_0, 0, 1]$, and if $f(x_0) = y_0^2$ is a nonzero square, we get two, $[x_0, \pm y_0, 1]$.

Consider the multiplicative group U_p , which has order $p - 1$. We claim that exactly half of the elements in U_p are squares (modulo p), and the other half are nonsquares. Let's make the proof of this fact more of a leisurely stroll than a succinct proof.

If you want to know which elements in U_p are squares, the simplest way is to square all the elements and see which residues remain. While earlier we expressed the elements of U_p as the set of positive residues $U_p = \{\overline{1}, \overline{2}, \dots, \overline{p-1}\}$, it is often convenient to list them as $U_p = \{\overline{1}, \overline{2}, \dots, \overline{\frac{p-1}{2}}, -\overline{\frac{p-1}{2}}, \dots, -\overline{2}, -\overline{1}\}$. Of course, we see that $k^2 \equiv (-k)^2 \pmod{p}$ for $k = 1, 2, \dots, (p-1)/2$. These represent all possible squares in U_p . So there are at most $(p-1)/2$ squares in U_p .

To see that there are exactly $(p-1)/2$ squares, we show that there is no redundancy in the list. For suppose that $a, b \in \mathbb{Z}$ and $a^2 \equiv b^2 \pmod{p}$. This means that $a^2 - b^2 \equiv 0 \pmod{p}$, or that $p \mid (a^2 - b^2) = (a-b)(a+b)$. But p is a prime, and whenever a prime divides a product of two integers, it must divide one of them (see the exercise just before Theorem 3.15). So this means p divides $(a-b)$ or $(a+b)$. In terms of congruences, that says that $a \equiv \pm b \pmod{p}$. So for a fixed $k \in \{1, 2, \dots, \frac{p-1}{2}\}$, the only way for $a^2 \equiv k^2 \pmod{p}$ is if $a \equiv \pm k \pmod{p}$. This says that $\{\overline{k^2} \mid k = 1, 2, \dots, \frac{p-1}{2}\}$ are all the distinct squares in U_p . So among the elements of U_p , half of them are squares and the complementary half are nonsquares.

Of course another way to see this is to recall the U_p is cyclic, say $U_p = \langle g \rangle$. So g is a primitive root modulo p and $U_p = \{g, g^2, \dots, g^{p-1}\}$. Clearly, those elements g^k with k even are squares. Can we check that those with k odd are not? If $g^k = g^{2\ell}$ with k odd, then $g^{2\ell-k} = 1$ in U_p , which means (by Lagrange) that $|g| = p-1 \mid (2\ell-k)$ or $2\ell-k \equiv 0 \pmod{p-1}$. But since p is odd, $p-1$ is even, so that $2\ell-k \equiv 0 \pmod{p-1}$ implies $2\ell-k \equiv 0 \pmod{2}$, which is clearly false.

Indeed, the squares and nonsquares of U_p or of \mathbb{F}_p are so important in number theory, they have been given special names. If $a \in U_p$, we say a is a *quadratic residue* if a is a square in U_p , and it is a *quadratic nonresidue* otherwise.

Returning to our question of the size of $E(\mathbb{F}_p)$, we first offer a heuristic: Suppose that as x runs over all the values of \mathbb{F}_p , the values

of $f(x)$ are uniformly distributed modulo p . So we expect some x_0 for which $f(x_0) = 0$ (which will yield one point on the curve), $(p-1)/2$ of the points x in U_p yielding $f(x)$ a quadratic residue, contributing another $(p-1)$ points on the curve, plus the point at infinity for a total estimate of $p+1$ points.

With that as a guess, based on a uniform distribution, we write $\#E(\mathbb{F}_p) = p+1 + (\text{error term})$. That's all well and good, but is there any truth in this heuristic? That is answered by a theorem of Hasse which says the error is bounded in absolute value by $2\sqrt{p}$.

Theorem 7.4 (Hasse).

$$-2\sqrt{p} < \#E(\mathbb{F}_p) - (p+1) < 2\sqrt{p}.$$

Culling from some historical facts in [CP05], we note that there is a theorem of Deuring (1941) [Deu41] which says that if we let $E_{a,b}$ denote the elliptic curve $y^2 = x^3 + ax + b$, then for any integer m with $p+1 - 2\sqrt{p} < m < p+1 + 2\sqrt{p}$, there exists $a, b \in \mathbb{F}_p$ so that $\#E(\mathbb{F}_p) = m$. A 1987 theorem of Hendrik Lenstra [Len87] says there are actually many of them.

Hasse's theorem gives teeth to the above heuristic which we want to exploit in giving a description of elliptic curve cryptography. We set the stage here and continue the discussion in the next chapter. The idea is, given an elliptic curve E over a finite field \mathbb{F}_p , we want the probability that a randomly chosen x in \mathbb{F}_p is the x -coordinate of a point on E . If $N = \#E(\mathbb{F}_p)$, then with at most four exceptions, every point (x, y) is paired with a distinct point $(x, -y)$ on the curve, so $N/2$ is roughly the number of x -coordinates of points on the curve, meaning $N/2p$ is roughly the probability that a randomly chosen x will be the x -coordinate of a point on the curve. Now by Hasse's theorem we see that

$$p+1 - 2\sqrt{p} < N < p+1 + 2\sqrt{p},$$

so

$$\frac{1}{2} + \frac{1}{2p} - \frac{1}{\sqrt{p}} < \frac{N}{2p} < \frac{1}{2} + \frac{1}{2p} + \frac{1}{\sqrt{p}},$$

which (for large p) gives a probability of approximately $1/2$.

Chapter 8

Applications of Elliptic Curves

We have seen an implementation of a public-key cryptosystem given by RSA, whose security is based on the difficulty of factoring an integer n that is the product of two large primes. We have also discussed the Diffie–Hellman key-exchange protocol and the ElGamal public-key cryptosystem whose security rests, at least indirectly, on the difficulty of solving the discrete logarithm problem.

Diffie–Hellman and ElGamal rely on the cyclic group structure of U_p , where a generator (primitive root) plays a pivotal role. In Chapter 4, we discussed Pollard’s $p - 1$ method for factoring integers, which relies only on Fermat’s little theorem, or more generally the theorem of Lagrange, which says the order of an element divides the order of the group. In both cases the group we leveraged was U_p . Elliptic curves over finite fields offer us a broad new collection of abelian groups whose structure we can hope to exploit in a fashion analogous to what we have done with U_p .

On January 15, 2009, the NSA posted to their website a three-page document titled “The Case for Elliptic Curve Cryptography—NSA/CSS” [NC09]. It is no longer accessible on their site, but was captured by the Internet Archive (WayBack Machine); see [NC09] for a working URL.

To quote the conclusion of that document,

Elliptic Curve Cryptography provides greater security and more efficient performance than the first generation public-key techniques (RSA and Diffie–Hellman) now in use. As vendors look to upgrade their systems, they should seriously consider the elliptic curve alternative for the computational and bandwidth advantages they offer at comparable security.

In the sections that follow, we shall sample a few aspects of elliptic curve cryptography. In the final section we comment on developments that are perhaps even more interesting: notably that the document [NC09] was removed from the NSA/CSS website sometime after June 15, 2015. In its place are recommendations that supersede those concerning the use of elliptic curves. They have to do with NSA recommendations for a post-quantum computing world [Sch15a], [KM15]. We shall comment more on these latest recommendations at the end of this chapter. But for now we keep the thread moving along.

8.1. Elliptic Curves and Factoring

We begin with a brief discussion of the impact of the arithmetic of elliptic curves on cryptography, not only as a means to introduce new methods of encryption, but also on its use to break older methods of encryption, notably as a means to factor integers. We begin with the later.

Factorization is very much an art with integers having a certain composition being amenable to factorization by highly tailored methods. For example, we saw this was the case with Pollard’s $p - 1$ method which worked very well on integers that were divisible by a prime p for which the factors of $p - 1$ were all small. The elliptic curve method of factorization (ECM) was developed by Hendrik Lenstra, and is currently one of the tools at the cutting edge in factorization techniques. Crandall and Pomerance [CP05] put it very nicely:

A subexponential factorization method of great elegance and practical importance is the elliptic

curve method (ECM) of H. Lenstra. The elegance will be self-evident. The practical importance lies in the fact that unlike the quadratic sieve (QS) or the number field sieve (NFS), ECM complexity to factor a number n depends strongly on the size of the least prime factor to divide n , and only weakly on the size of n itself. For this reason, many factors of truly gigantic numbers have been uncovered in recent years; many of these numbers lying well beyond the range of QS or NFS.

In a real sense, ECM is a natural and broad generalization of Pollard's $p - 1$ method, so we begin our introduction with a comparison of the basic mechanics of Pollard's $p - 1$ and the ECM as algorithms to factor a composite integer n ; details of the ECM will follow. Recall that since factoring is a recursive process, our goal is, given a composite integer n , to find a divisor a of n with $1 < a < n$.

Pollard's $p - 1$ method uses Fermat's little theorem (Theorem 4.18) to attempt to factor a composite integer n as follows: Choose an integer a with $1 < a < n$, so $\gcd(a, n) < n$. If a and n are not relatively prime, then $\gcd(a, n)$ represents a nontrivial factor of n , and we have succeeded in our effort. We therefore assume that $\gcd(a, n) = 1$ and hence for any prime $p \mid n$, we know $a \in U_p$, so that $a^{p-1} = 1$ in U_p . This means that $p \mid \gcd(a^{p-1} - 1, n)$, so in particular $\gcd(a^{p-1} - 1, n) > 1$. We look for evidence of such a p as follows. We know that $a^{p-1} = 1$ in U_p implies $a^k = 1$ in U_p for any k divisible by $p - 1$, so $\gcd(a^{p-1} - 1, n) > 1$ implies $\gcd(a^k - 1, n) > 1$ as well. So we hope that n is divisible by a prime p for which $p - 1$ is the product of small primes to small powers. If so, then simply by looking at small values of k , we compute $\gcd(a^k - 1, n)$ hoping for a nontrivial gcd. If n is indeed divisible by a prime p for which $p - 1$ is the product of small primes to small factors, then we should be able to find such a k .

In describing the ECM, we replace U_p and the condition $a^{p-1} = 1$ with another finite abelian group and related condition. For the analogue of U_p we want to consider an elliptic curve E and look at its set of points $E(\mathbb{F}_p)$ over the finite field \mathbb{F}_p with p elements. But

as with the Pollard method, where we actually start with an element $a \in U_n$ (and hence in U_p for any prime $p \mid n$), for ECM we start with a curve E which would be an elliptic curve over any finite field \mathbb{F}_p with $p \mid n$. In particular, we consider the elliptic “pseudo-curve” E and its set of points $E(\mathbb{Z}_n)$. Analogous to choosing $a \in U_n$, we choose a point $P \in E(\mathbb{Z}_n)$. We also have to remember that U_p is a multiplicative group, while $E(\mathbb{F}_p)$ (for $p \mid n$) is an additive one, so the condition $a^k = 1$ in U_p translates to $k \cdot P = \underbrace{P \oplus P \oplus \cdots \oplus P}_{k \text{ times}} = \mathbf{0}$ in

$E(\mathbb{F}_p)$, where $\mathbf{0}$ is the identity of the group, our distinguished point at infinity. As with Pollard’s method, we try values of k which are the product of small primes to small powers, and we win if there is a prime $p \mid n$ with $\#E(\mathbb{F}_p)$ the product of small primes to small powers. A more succinct comparison is given by:

Pollard’s $p - 1$	versus	Lenstra’s ECM
$a \in U_p$	\longleftrightarrow	$P \in E(\mathbb{F}_p)$
$a^k = 1$ in U_p (if $\#U_p = p - 1 \mid k$)	\longleftrightarrow	$kP = \mathbf{0}$ in $E(\mathbb{F}_p)$ (if $\#E(\mathbb{F}_p) \mid k$)
We win if there is a prime $p \mid n$ with $\#U_p$ being the product of small primes to small powers.	\longleftrightarrow	We win if there is a prime $p \mid n$ with $\#E(\mathbb{F}_p)$ being the product of small primes to small powers.

One big difference between these methods is that in the case of Pollard’s $p - 1$, for each prime p , there is only one group to exploit associated to p , namely U_p . In the case of the ECM, for each p , we have an enormous supply of elliptic curves E whose group of points $E(\mathbb{F}_p)$ are finite abelian groups with potentially vastly different structures. In particular, for each $a, b \in \mathbb{F}_p$ with $\Delta = 4a^3 + 27b^2 \neq 0$, $y^2 = x^3 + ax + b$ defines an elliptic curve $E_{a,b}$ over \mathbb{F}_p .

Next, we discuss some practical aspects of computing with the group law on $E(\mathbb{F}_p)$, in particular, how to compute kP , that is, the point P on the elliptic curve added to itself k times. Just as with

modular exponentiation, this can be done efficiently by using the binary expansion of k : $k = k_0 + k_1 2^1 + k_2 2^2 + \cdots + k_r 2^r$, with each k_i equal to 0 or 1. We know how to add points on an elliptic curve; doubling a point $P \mapsto 2P$ is done via the tangent line, while if $P \neq Q$, the sum $P \oplus Q$ is determined using the line through P and Q (see formulas below). So we precompute:

$$\begin{aligned} P_0 &= P, \\ P_1 &= 2P_0 = 2P, \\ P_2 &= 2P_1 = 2^2 P, \\ P_3 &= 2P_2 = 2^3 P, \\ &\vdots \\ P_r &= 2P_{r-1} = 2^r P. \end{aligned}$$

Then $kP = k_0 P_0 \oplus k_1 P_1 \oplus \cdots \oplus k_r P_r$. Note that the k_i are either 0 or 1, so we are just adding the points P_i where k_i is nonzero.

To implement the procedure above, we review the group law for points on the curve $E_{a,b} : y^2 = x^3 + ax + b$.

Summary of Group Law. Let $\mathbf{0} = [0, 1, 0]$ (the point at infinity), and let $P_i = [x_i, y_i, 1]$ be any affine point on the curve.

- (1) $P \oplus \mathbf{0} = \mathbf{0} \oplus P = P$ for all points P on the elliptic curve.
- (2) $-P_i = [x_i, -y_i, 1]$; $-\mathbf{0} = \mathbf{0}$.
- (3) $P_1 \oplus P_2 = \mathbf{0}$ if and only if $P_2 = -P_1$ (meaning P_1 and P_2 lie on a vertical line).
- (4) If $P_2 \neq -P_1$, then $P_1 \oplus P_2 = P_3 = [x_3, y_3, 1]$, where

$$\begin{aligned} x_3 &= m^2 - x_1 - x_2, \\ y_3 &= -[m(x_3 - x_1) + y_1], \text{ and where} \\ m &= \begin{cases} \frac{y_2 - y_1}{x_2 - x_1} & \text{if } x_1 \neq x_2, \\ \frac{3x_1^2 + a}{2y_1} & \text{if } x_1 = x_2. \end{cases} \end{aligned}$$

Let's take a second look at Example 7.2.

Example 8.1. Consider the cubic curve $E : y^2 = x^3 + 5x + 2$ over the field \mathbb{F}_7 . The discriminant of the curve is

$$\begin{aligned}\Delta &= 4a^3 + 27b^2 = 4 \cdot 5^3 + 27 \cdot 2^2 \\ &\equiv 4 \cdot (-2)^3 + 6 \cdot 2^2 \equiv -8 \equiv 6 \not\equiv 0 \pmod{7},\end{aligned}$$

so this is an elliptic curve over \mathbb{F}_7 . We checked that

$$\begin{aligned}E(\mathbb{F}_7) &= \{\mathbf{0}, [0, \pm 3, 1], [1, \pm 1, 1], [3, \pm 3, 1], [4, \pm 3, 1]\} \\ &= \{\mathbf{0}, [0, 3, 1], [0, 4, 1], [1, 1, 1], [1, 6, 1], [3, 3, 1], [3, 4, 1], [4, 3, 1], [4, 4, 1]\},\end{aligned}$$

and that $E(\mathbb{F}_7) \cong \mathbb{Z}_9$. We took the point $P = [4, 4, 1]$ as a generator of the group. To do our precomputation P_0, P_1, \dots, P_r , we begin by doubling our point P .

To do so, we use the slope of the tangent line in our formulas above:

$$m = \frac{3x_1^2 + a}{2y_1} = \frac{3 \cdot 4^2 + 5}{2 \cdot 4} = \frac{4}{1} = 4 \text{ in } \mathbb{F}_7.$$

Thus, $2P = [x_3, y_3, 1] = [1, 1, 1]$ (by our formulas). Similarly, we use the formulas to compute

$$\begin{aligned}4P &= 2(2P) = [0, 3, 1] \quad \text{and} \\ 8P &= 2(4P) = [4, 3, 1] = [4, -4, 1] = -P.\end{aligned}$$

Filling things out, we have

$$\begin{aligned}P &= [4, 4, 1], \\ 2P &= [1, 1, 1], \\ 4P &= 2(2P) = [0, 3, 1], \\ 8P &= 2(4P) = [4, 3, 1] = [4, -4, 1] = -P, \\ 16P &= 2(8P) = 2(-P) = -2P = 7P = [1, -1, 1] = [1, 6, 1] \\ 32P &= 2(16P) = -4P = 5P = [0, -3, 1] = [0, 4, 1] \\ 64P &= 2(32P) = 10P = P \oplus 9P = P \oplus \mathbf{0} = P = [4, 4, 1],\end{aligned}$$

and now of course the list of points will cycle through the same sequence as we consider $2^k P$ for $k \geq 6$.

Remark 8.2. It is obvious that the set of points $\{P, 2P, 4P, \dots\}$ is a finite set since $E(\mathbb{F}_p)$ is a finite group, say $\{P, 2P, \dots, 2^h P\}$ is the set of distinct points. It is also clear that the list will cycle through $\{2^j P, 2^{j+1} P, \dots, 2^h P\}$, where j is the integer satisfying $2^{h+1} P = 2^j P$, $0 \leq j \leq h$. In the example above, we see that $j = 0$. It is also possible to have $j = h$ when $2^h P = \mathbf{0}$.

In analogy with Pollard's method, we choose $P \in E(\mathbb{Z}_n)$ (so P would be a point on the elliptic curve $E(\mathbb{F}_p)$ for any prime $p \mid n$). If $\#E(\mathbb{Z}_p) \mid k$, then $kP = \mathbf{0}$, so we hope that finding an integer k with $kP = \mathbf{0}$ suggests the existence of such a prime p . But the goal is not really determining whether for a small value of k that $kP = \mathbf{0}$ or not; Lenstra's method is much more clever. One computes $kP = k_0 P_0 \oplus k_1 P_1 \oplus \dots \oplus k_r P_r$ as before with the binary expansion of k , but for each sum (or doubling for that matter), one must compute a slope, $m = (y_2 - y_1)/(x_2 - x_1)$ or $m = (3x_1^2 + a)/2y_1$. Now in reality, we are working with an elliptic pseudocurve, $E_{a,b}(\mathbb{Z}_n) : y^2 = x^3 + ax + b$, which we are treating as an elliptic curve over \mathbb{Z}_n . When n is not prime, not all nonzero elements of \mathbb{Z}_n have multiplicative inverses (indeed only those relatively prime to n), and so in computing the inverses associated to the slopes m , we can detect a factor of n by noting the failure of a denominator in m have an inverse modulo n , so at each multiplication we check $\gcd(x_2 - x_1, n)$ or $\gcd(2y_1, n)$. Producing a nontrivial gcd gives us a factor of n ; otherwise, we continue the computation.

Lenstra's ECM algorithm. Given a composite integer n to factor, we perform the following steps.

- (1) Check $\gcd(n, 6) = 1$ (2's and 3's make life with elliptic curves more difficult, but then it is not so difficult to pull out the factors of 2 and/or 3 from n as an initial step).
- (2) Check that n is not a perfect power, i.e., $n \neq m^k$ for some m and k . This is easy and quick to do. Just check (using real-valued functions) that none of $\sqrt[n]{n}$, $\sqrt[3]{n}$, \dots , $\sqrt[\ell]{n}$ are integers for $\ell = \lceil \ln n / \ln 2 \rceil$ (this guarantees that $\sqrt[\ell]{n} < 2$).
- (3) Choose a bound B (say $B \approx 10000$).
- (4) Choose a curve $E_{a,b}(\mathbb{Z}_n) : y^2 = x^3 + ax + b$ and a point $P = [x, y, 1]$ on the curve as follows:

- (a) Choose random integers $x, y, a \in [0, n - 1]$.
 (b) Compute $b \equiv (y^2 - x^3 - ax) \pmod{n}$.
 (c) Compute $d = \gcd(4a^3 + 27b^2, n)$. If $d = n$, start over choosing a new x, y, a . If $1 < d < n$, then d is a proper factor of n , and we have succeeded. Otherwise $d = 1$ which means we have an elliptic pseudocurve over \mathbb{Z}_n (in particular an honest elliptic curve over \mathbb{F}_p for any prime $p \mid n$), and a point $P = [x, y, 1]$ on that curve.
- (5) Compute highest prime powers less than or equal to the bound B :

$$2^{a_2}, 3^{a_3}, \dots, p_r^{a_r} \leq B.$$

- (6) Technically, we are hoping that if $k = 2^{a_2} 3^{a_3} \dots p_r^{a_r}$, then $kP = \mathbf{0}$, but we will actually compute kP in stages hoping for a failure anywhere along the way. For example:

$$\begin{aligned} P &\mapsto 2P \mapsto 4P \mapsto \dots \mapsto 2^{a_2} P \mapsto 3 \cdot 2^{a_2} P \mapsto 3^2 \cdot 2^{a_2} P \mapsto \dots \\ &\dots \mapsto 3^{a_3} \cdot 2^{a_2} P \mapsto \dots \mapsto p_r \cdot (p_{r-1}^{a_{r-1}} \dots 3^{a_3} \cdot 2^{a_2} P) \mapsto \dots \\ &\dots \mapsto p_r^{a_r} \cdot (p_{r-1}^{a_{r-1}} \dots 3^{a_3} \cdot 2^{a_2} P) = kP. \end{aligned}$$

At each addition or doubling, we are looking to find a slope that cannot be computed by failure of the gcd of the denominator and n to equal one. If the gcd is one, we continue the arithmetic; if the gcd is a proper divisor of n , we return the factor; if the gcd is n , we can increase the bound B or try another curve.

It turns out that the computational complexity of the ECM is related to the size of the smallest prime factor which divides n and not very much to n itself. This means the ECM can be effective in finding a divisor of enormous composites (those with at least one not-so-large factor), but it is worse at factoring RSA composites which are the product of two primes of roughly the same size.

8.2. Elliptic Curves and Cryptography

While the mathematical prerequisites begin to rise a bit more steeply to describe fully how we can generalize Diffie–Hellman and ElGamal

to the setting of elliptic curves, we can give the highlights. More details can be found in [Kob84] and [Kob87b].

8.2.1. Embedding Plaintext into an Elliptic Curve. As in all previous cryptosystems we have considered, we assume our plaintext has been converted to an integer equivalent m with $0 \leq m < M$. Choose κ to be an integer so that the probability that one fails to embed the plaintext m into a chosen elliptic curve is less than $2^{-\kappa}$. Choose a prime p so that $p > M\kappa$. Then it is possible to write every integer ℓ with $1 \leq \ell \leq M\kappa$ uniquely as $\ell = m\kappa + j$ with $0 \leq m < M$ and $1 \leq j \leq \kappa$. And since $M\kappa < p$, we can think of all the integers $m\kappa + j$ (or more properly the residues $\overline{m\kappa + j}$ modulo p) with $0 \leq m < M$ and $1 \leq j \leq \kappa$ as distinct elements of the finite field \mathbb{F}_p .

We can embed our plaintext message m as a point P_m on an elliptic curve $E_{a,b} : y^2 = f(x) = x^3 + ax + b$ over \mathbb{F}_p as follows. For each j with $1 \leq j \leq \kappa$, we are going to test $x = x(j) = m\kappa + j$ to see if $\overline{f(x)}$ is the x -coordinate of a point on the elliptic curve $E_{a,b}(\mathbb{F}_p)$. We have observed that all such $\overline{f(x)}$ are distinct elements of \mathbb{F}_p . We compute $\overline{f(x)}$ and ask if it is a square in \mathbb{F}_p ; we showed in the previous chapter that there is roughly a 50% chance that this is so. If it is a square, we find a $\overline{y} \in \mathbb{F}_p$ so that $\overline{y}^2 = \overline{f(x)}$ (see Chapter II.2 of [Kob84]), and we have our point $P_m = (\overline{x}, \overline{y})$. If $\overline{f(x)}$ is not a square, we increment the value of j by one and test again. Since we have κ integers $x = m\kappa + j$, the probability that we will fail to produce a point P_m on the curve is approximately $2^{-\kappa}$.

Now given a point $P_m = (\overline{x}, \overline{y})$, we consider the integer $x = m\kappa + j$ as above with $\overline{x} = \overline{m\kappa + j}$. Notice that $x - 1$ satisfies

$$m\kappa \leq x - 1 \leq m\kappa + (\kappa - 1),$$

so

$$m \leq \frac{x-1}{\kappa} \leq m + \frac{\kappa-1}{\kappa} < m+1,$$

which means m , the plaintext message, is recoverable from P_m as $m = \lfloor \frac{x-1}{\kappa} \rfloor$.

So now we have a means of taking a plaintext message and embedding it into an elliptic curve and, conversely, given a point on

the curve, we can extract the plaintext message. Next we talk about encryption.

8.2.2. Analogues of Diffie–Hellman and ElGamal. With both the original versions of Diffie–Hellman and ElGamal, we chose a prime p , and used the group U_p and an element g which was either a primitive root or simply generated a very large subgroup of U_p .

In the elliptic curve setting, we have far more freedom. Given a prime p , we choose $a, b \in \mathbb{F}_p$ so that $E : y^2 = x^3 + ax + b$ is an elliptic curve. So now $E(\mathbb{F}_p)$ has replaced U_p . What do we do about an analogue of g ? Fortunately, with these encryption schemes based upon elliptic curves, a great deal of information about the elliptic curve and the finite field is intended to be public. So documents like [Bro10], give lists of recommended primes (having a prescribed (bit)size), and parameters $a, b \in \mathbb{F}_p$ which will provide an elliptic curve $E = E_{a,b}(\mathbb{F}_p)$ and a point G (a basepoint) on E together with the order of G and the cofactor h , ($h = \#E(\mathbb{F}_p)/|G|$), so $1/h$ is the proportion of $E(\mathbb{F}_p)$ which G generates. So we proceed supplied with all these data.

As in the classical case, Alice and Bob choose two integers a, b with $1 < a, b < |G|$. Alice sends aG to Bob; Bob sends bG to Alice. Each computes $baG = abG$ as their shared key for any secret-key cryptosystem based on elliptic curves.

Alice publishes the finite field \mathbb{F}_p , the chosen elliptic curve E defined over \mathbb{F}_p , the basepoint G on $E(\mathbb{F}_p)$, and her multiple aG as her public key which can be used in an ElGamal scheme as follows: Bob wishes to send a message m to Alice that he has embedded as the point P_m in $E(\mathbb{F}_p)$. As in the classical case, he chooses a random integer k and sends the ordered pair $(kG, P_m \oplus k(aG))$ to Alice. Upon receipt, Alice computes $a(kG)$ and computes $[P_m \oplus k(aG)] \oplus -akG = P_m$. Given the point P_m , she recovers the plaintext message m .

8.3. Remarks on a Post-Quantum Cryptographic World

As we said in the opening of this chapter, up until June 2015, the NSA was actively promoting ECC over first-generation methods such as

RSA and Diffie–Hellman; see [NC09]. In August 2015 that changed when the NSA released a policy statement which came as a surprise to many [IN15], including the statement:

IAD (Information Assurance Directorate at the NSA) will initiate a transition to quantum resistant algorithms in the not too distant future. Based on experience in deploying Suite B, we have determined to start planning and communicating early about the upcoming transition to quantum resistant algorithms. Our ultimate goal is to provide cost effective security against a potential quantum computer. We are working with partners across the USG, vendors, and standards bodies to ensure there is a clear plan for getting a new suite of algorithms that are developed in an open and transparent manner that will form the foundation of our next suite of cryptographic algorithms.

Until this new suite is developed and products are available implementing the quantum resistant suite, we will rely on current algorithms. For those partners and vendors that have not yet made the transition to Suite B elliptic curve algorithms, we recommend not making a significant expenditure to do so at this point but instead to prepare for the upcoming quantum resistant algorithm transition.

What explains this rather abrupt change in policy by the NSA? What exactly is a quantum-resistant algorithm? And just how concerned should everyone be? To offer possible explanations for the change in NSA policy, Koblitz and Menezes in [KM15] gave a brief history of the development of ECC, and they evaluate many of the possible reasons for this shift in policy. Independent of the reasons, we want to understand what it is that the NSA is suggesting, so we should say a few words about quantum computers and the meaning of a quantum-resistant algorithm. A good starting place is the quantum computing Wikipedia page [Wik16] from which we borrow.

Computers, as we know them today, store data in bit strings, strings of zeros and ones. We have referred to key lengths for various cryptographic systems in terms of the number of bits in their key. Imagine that you have an incredibly primitive computer with a memory capacity of 3 bits. What possible states could the memory be in? With a triple (b_0, b_1, b_2) and $b_i = 0$ or 1, there are $2^3 = 8$ possible strings or states for the memory to be in, and at any given time the memory is in exactly one of these eight states. Quantum computing is based on the physics of quantum mechanics, and the analogue of our 3-bit computer is a three quantum bit (qubit) computer whose memory can be in a superposition of those eight states. In fact, the state of our 3-qubit computer is described as a point on the 8-dimensional complex unit sphere, that is a state has the form (b_0, \dots, b_7) where each $b_i \in \mathbb{C}$ and $\sum_{i=0}^7 |b_i|^2 = 1$. Now each b_i corresponds to one of the classical 8-bit states 000 to 111, and the magnitude $|b_i|^2$ corresponds to the probability of being in the i th state.

From even this brief characterization, one can sense the enormous difference in capacity between a 3-bit machine and a 3-qubit one. If the quantum computer could be scaled to the size of modern computers, its computing power would be enormous. In a vague sense, you can think of quantum computing as a vast parallelization of standard computing. So we have some sense of its capabilities, but what explains the NSA's new perspective?

We recall that the security of RSA is based upon the supposition that factoring integers is a computationally hard problem and, analogously, that solving the discrete log problem for elliptic curves, on which the security of ECC is based, is also a computationally hard problem. In 1994 Peter Shor [Sho94] showed that both of these computational problems could be solved in polynomial time (that is to say very quickly) if one had access to a quantum computer of appropriate size. So the takeaway is that large quantum computers compromise RSA and ECC.

On the other hand, if one surveys the web for the current state of quantum computers, it would be hard to find a site that did not refer to it as still being in its infancy. Security expert Bruce Schneier [Sch15a], [Sch15b] quips that “the largest number to date that has

been factored by a quantum computer is 143.” He continues, “So while a practical quantum computer is still science fiction, it’s not *stupid* science fiction.” Thus RSA and ECC are not in immediate danger, so why the abrupt change in policy? One consideration is that the NSA tries to be very forward looking. Their needs for encryption are not short-term as in an electronic banking transaction; they often have the need to be able to encrypt documents which will remain secure for decades to come, so they need to anticipate what computing power could exist in 10, 20, even 30 or more years.

We have known for decades that a practical quantum computer can compromise many traditional cryptographic schemes. What else can a quantum computer do? Perhaps there are undiscovered algorithms for breaking other encryption schemes with a quantum computer. For simplicity, let’s consider the analysis of a brute-force attack on a keyspace which would apply to any encryption scheme. If we have a key for a symmetric-key encryption scheme consisting of k bits, the keyspace has size 2^k . As computing power has increased, the value of k has increased gradually so that an exhaustive search of 2^k keys remains infeasible. Algorithms for quantum computers exist which can reduce the search to the square-root of that number of keys, effectively reducing the keyspace to size $2^{k/2}$. For the time being, the NSA has addressed this concern by publishing an updated table of key lengths that should be used with first- and second-generation cryptographic systems, taking into account known attacks. Even though no practical quantum computer currently exists, to counteract the potential exhaustive search approach above, recommended key lengths have been doubled.

It remains somewhat of a puzzle why the NSA has changed its tune about the (economic) value of changing to elliptic curve-based cryptographic methods over first-generation methods, especially given that even the most optimistic estimates for the development of a practical quantum computer are more than 15 years away, and given that most users of encryption need short-term security, not security that will endure for decades. Whatever the reason, the NSA is clearly concerned, and it is leading a march to develop quantum-resistant algorithms. The appendix to [KM15] includes a number of candidates.

Deeper Results and Concluding Thoughts

A major goal of this text was to develop sufficient mathematics to understand the basics surrounding the arithmetic of elliptic curves and their applications to cryptography. As a byproduct, we have revealed some of the vista that represents modern mathematics, and at the same time hopefully left ample hooks for tangential explorations.

Still, the story of elliptic curves has barely been touched, and so as a means to bring closure to an early motivating topic in the text (congruent numbers) and to point to many other related topics on the horizon, we explore a bit more about elliptic curves, but now defined over the rational and complex numbers. The intent here is to be significantly more telegraphic in our exposition, leaving only a few bread crumbs to follow if inclined.

A.1. The Congruent Number Problem and Tunnell's Solution

In the first two chapters we defined congruent numbers and established a relationship between Pythagorean triples and congruent numbers. Moreover, by parametrizing the rational points on the unit circle, we were able to list all Pythagorean triples, and if we let the process continue forever, eventually all congruent numbers would be listed. But this is rather unsatisfactory since even if one knew a

given number was a congruent number, there is no way of telling how long one would wait before it was listed as corresponding to some Pythagorean triple.

All that changed in 1983 with Tunnell's elegant answer to the congruent number problem. What is most elegant about it is that the answer is just as easy to understand as the statement of the problem itself. The mathematics that underlies his answer, however, is quite deep, yet we shall at least broach these topics.

Tunnell's answer is given by the following theorem (see [Kob84]) which determines whether n is a congruent number by comparing the representation numbers of two ternary quadratic forms.

Theorem A.1 (Tunnell). *Let n be a square-free positive integer.*

- (1) *Suppose that n is a congruent number.*

If n is odd, then

$$\begin{aligned} \#\{(x, y, z) \in \mathbb{Z}^3 : n = 2x^2 + y^2 + 32z^2\} \\ = \frac{1}{2} \#\{(x, y, z) \in \mathbb{Z}^3 : n = 2x^2 + y^2 + 8z^2\}. \end{aligned}$$

If n is even, then

$$\begin{aligned} \#\{(x, y, z) \in \mathbb{Z}^3 : \frac{n}{2} = 2x^2 + y^2 + 32z^2\} \\ = \frac{1}{2} \#\{(x, y, z) \in \mathbb{Z}^3 : \frac{n}{2} = 2x^2 + y^2 + 8z^2\}. \end{aligned}$$

- (2) *Conversely, if the weak Birch–Swinnerton–Dyer conjecture is true for elliptic curves of the form $E_n : y^2 = x^3 - n^2x$, then these equalities of cardinalities imply that n is a congruent number.*

For example, on page 5 of [Kob84] the author gives an example by Zagier which shows that 157 is a congruent number. The right triangle which demonstrates that 157 is congruent has a hypotenuse whose length is a rational number which in reduced form has a denominator with 45 digits. Clearly, this would not be so easy to find using our enumeration method. On the other hand, we consider Tunnell's theorem.

The number 157 is odd and square-free, so we need only verify the the number of integer solutions to $2x^2 + y^2 + 32z^2 = 157$ is one-half

the number of integer solutions to $2x^2 + y^2 + 8z^2 = 157$. We claim that there are no solutions to either equation. Note that if $2x^2 + y^2 + 8z^2 = 157$ has no solutions, then neither does $2x^2 + y^2 + 32z^2 = 157$, since $32z^2 = 8(2z)^2$. Consider the equality $2x^2 + y^2 + 8z^2 = 157$ as a congruence modulo 8. We obtain $2x^2 + y^2 \equiv 5 \pmod{8}$ which implies that y is odd, hence $y^2 \equiv 1 \pmod{8}$. It is also true that $2x^2 \equiv 0, 2 \pmod{8}$, so that $2x^2 + y^2 + 8z^2 \equiv 1, 3 \pmod{8}$, so there can be no solution. Thus both sets have the same (zero) cardinality, so by the theorem, 157 is a congruent number.

Oh yes—what about this Birch–Swinnerton-Dyer conjecture, and where did elliptic curves come in? That will take us a bit longer to explain. We follow Koblitz [Kob84] here. He begins with a series of propositions to connect congruent numbers to points on elliptic curves.

Proposition A.2. *Let $n \geq 1$ be a square-free integer. Let $X, Y, Z \in \mathbb{Q}$ with $X < Y < Z$. There is a one-to-one correspondence between right triangles with sides X, Y and hypotenuse Z having area n , and rational numbers x so that $x, x + n, x - n$ are all squares in \mathbb{Q} . The correspondence is given by*

$$X, Y, Z \mapsto x = (Z/2)^2$$

$$x \mapsto X = \sqrt{x+n} - \sqrt{x-n}, \quad Y = \sqrt{x+n} + \sqrt{x-n}, \quad Z = 2\sqrt{x}.$$

Proof. Let's first see that the numbers do what is claimed. Given X, Y, Z with $X^2 + Y^2 = Z^2$ and $XY/2 = n$, we see that

$$X^2 + Y^2 \pm 4 \frac{1}{2} XY = Z^2 \pm 4n, \text{ or}$$

$$(X \pm Y)^2 = Z^2 \pm 4n, \text{ or}$$

$$\left(\frac{X \pm Y}{2}\right)^2 = \left(\frac{Z}{2}\right)^2 \pm n.$$

So $x = (Z/2)^2$ is obviously a square, and hence so is $x \pm n = (Z/2)^2 \pm n = ((X \pm Y)/2)^2$.

Conversely, given $x \in \mathbb{Q}$ with $x, x \pm n$ all squares, we put $X = \sqrt{x+n} - \sqrt{x-n}$, $Y = \sqrt{x+n} + \sqrt{x-n}$, and $Z = 2\sqrt{x}$, all of which

are now rational numbers by the assumption. We see that

$$\begin{aligned}\frac{1}{2}XY &= \frac{1}{2}(x+n - (x-n)) = n \text{ and} \\ X^2 + Y^2 &= 2(x+n + x-n) = 4x = Z^2.\end{aligned}$$

So X, Y, Z are the sides of a rational right triangle with area n . It is trivial to check that $X < Y < Z$.

Now we see that there is a one-to-one correspondence. Let $x \longleftrightarrow X, Y, Z$ and $x' \longleftrightarrow X', Y', Z'$.

Suppose that $(X, Y, Z) = (X', Y', Z')$. The fact that $Z = 2\sqrt{x} = Z' = 2\sqrt{x'}$ implies $x = x'$ (since both are positive). Conversely, suppose that $x = x'$. Then $Z = 2\sqrt{x} = Z' = 2\sqrt{x'}$ and hence

$$\begin{aligned}Z'^2 &= X'^2 + Y'^2 = X^2 + Y^2 = Z^2 \text{ and} \\ \frac{1}{2}XY &= n = \frac{1}{2}X'Y'.\end{aligned}$$

Geometrically, we are looking at the intersection of the circle $X^2 + Y^2 = Z^2$ (with $Z = Z'$ fixed), and the hyperbola $\frac{1}{2}XY = n$. The typical situation is pictured in Figure A.1.

So there are only four possible points (X, Y) that work, and the constraints $0 < X < Y$ mean there is only one point. The proof is complete. \square

We saw above that $(\frac{X+Y}{2})^2 = (\frac{Z}{2})^2 \pm n$. Multiplying the two expressions together yields $(\frac{X^2-Y^2}{4})^2 = (\frac{Z}{2})^4 - n^2$, which says the curve $u^4 - n^2 = v^2$ has a rational solution $u = (X^2 - Y^2)/4, v = (Z/2)$. Multiplying the equation $u^4 - n^2 = v^2$ by u^2 yields $u^6 - n^2u^2 = (uv)^2$. So putting

$$x = (Z/2)^2 = u^2, \quad y = (uv) = (X^2 - Y^2)Z/8,$$

we have a rational point on the elliptic curve $y^2 = x^3 - n^2x$. Conversely, we have the following.

Proposition A.3. *Let (x, y) be a rational point on the elliptic curve $y^2 = x^3 - n^2x$. Suppose that x is a square and has an even denominator. Then putting*

$$X = \sqrt{x+n} - \sqrt{x-n}, \quad Y = \sqrt{x+n} + \sqrt{x-n}, \quad Z = 2\sqrt{x}$$

produces a rational right triangle with area n .




Figure A.1

Denote by E_n the elliptic curve $y^2 = x^3 - n^2x$. Critical to moving forward is Mordell's important theorem describing the structure of the group of rational points on E_n , denoted $E_n(\mathbb{Q})$.

Theorem A.4 (Mordell). $E_n(\mathbb{Q}) \cong \mathbb{Z}^r \oplus E_n(\mathbb{Q})_{tor}$.

Here $E_n(\mathbb{Q})_{tor}$ is the *torsion* subgroup, that is the set of points in $E_n(\mathbb{Q})$ having finite order. The integer $r \geq 0$ is called the rank of the elliptic curve and $r > 0$ if and only if there are infinitely many rational points on E_n . The rank for this finitely generated abelian group is the analogue of dimension for a vector space. More precisely, the theorem says there are r points $P_1, \dots, P_r \in E_n(\mathbb{Q})$ so that for any point $P \in E_n(\mathbb{Q})$, P can be written uniquely as

$P = k_1P_1 \oplus k_2P_2 \oplus \cdots \oplus k_rP_r \oplus Q$ for (unique) integers k_1, \dots, k_r and a unique torsion point Q .

It is clear from the graph of the elliptic curve that the points $[-n, 0, 1]$, $[0, 0, 1]$, and $[n, 0, 1]$ all have order 2 (they're on the x -axis with a vertical tangent). This means that these three points (together with the identity $\mathbf{0}$) are all elements of $E_n(\mathbb{Q})_{\text{tor}}$. On the other hand, we have the following theorem.

Theorem A.5. $\#E_n(\mathbb{Q})_{\text{tor}} = 4$ for all square-free positive n .

Knowing the Fundamental Theorem of Finite Abelian Groups, we then deduce the following.

Corollary A.6. $E_n(\mathbb{Q})_{\text{tor}} = \{\mathbf{0}, [-n, 0, 1], [0, 0, 1], [n, 0, 1]\} \cong \mathbb{Z}_2 \times \mathbb{Z}_2$.

Finally, we come to the theorem which ties these ideas together.

Theorem A.7. *A positive, square-free integer n is a congruent number if and only if $E_n(\mathbb{Q})$ has positive rank, which is to say if and only if it has infinitely many rational points.*

Proof. We sketch the proof. If n is a congruent number, then we have seen there exists a point $(x, y) \in E_n(\mathbb{Q})$ with x a positive square. By inspection, such a point is not in the torsion subgroup, so is a point of infinite order. Conversely, if P is a point of infinite order in $E_n(\mathbb{Q})$, then using our doubling formula, we easily check that $2P$ has as its x -coordinate a square with even denominator which by previous work shows n is a congruent number. \square

This ends the so-called easy part of Tunnell's proof and occupies only the first chapter of Koblitz's book [Kob84]. We go a bit further to describe the Birch–Swinnerton-Dyer conjecture, but we do so more to advertise the role and independent interest of complex analysis than to just define the analytic objects involved with the Birch and Swinnerton-Dyer conjecture.

A.2. A Digression on Functions of a Complex Variable

At first blush, one might presume there to be little difference in studying differentiable functions $f : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ and $f : \mathbb{C} \rightarrow \mathbb{C}$, but there is, and the differences are dramatic. We point out three important distinctions.

First, if a function $f : \mathbb{C} \rightarrow \mathbb{C}$ has a continuous first derivative, it is infinitely differentiable, and indeed it can be expressed as a power series and is said to be analytic. The same is certainly not true even for functions $f : \mathbb{R} \rightarrow \mathbb{R}$, as $f(x) = x^{5/3}$ shows. Indeed, this remarkable property is related to another that says that if one knows the values of the analytic function $f : \mathbb{C} \rightarrow \mathbb{C}$ on the boundary of a nice region, then the values of f on the interior of the region are determined. Geometrically, this implies a certain rigidity. We can't see the graph of a function $f : \mathbb{C} \rightarrow \mathbb{C}$ (it lives in $\mathbb{C}^2 \cong \mathbb{R}^4$), but if we could and if the same were true of functions $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ (whose graph is a surface), we would see the following. Imagine the graph of a function defined on the closed unit disk and which is zero on the boundary. Surely you could draw lots of surfaces like that. If the function was an analytic function, there would only be one graph since the values on the boundary determine the values inside. Strange? Yes.

Second, it is often the case that we have two definitions of an analytic function or, more precisely, two analytic functions whose definitions agree on a nice set. Then they agree everywhere they are both defined; this is known as the identity theorem. First we give an example of how this does not happen for differentiable real-valued functions. Consider two functions $f, g : \mathbb{R} \rightarrow \mathbb{R}$ defined by

$$f(x) = \begin{cases} (x-1)^4 & \text{if } x \geq 1, \\ 0 & \text{otherwise,} \end{cases} \quad g(x) = \begin{cases} (x+1)^4 & \text{if } x \leq -1, \\ 0 & \text{otherwise.} \end{cases}$$

These are both continuously differentiable functions, defined on all of \mathbb{R} whose values agree for all $x \in [-1, 1]$, but which are clearly not the same wherever both are defined.

This can't happen in the complex case, and this turns out to be very handy. Consider the Riemann zeta function

$$\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s},$$

probably one of the most famous functions in all of mathematics. You studied this function in calculus when you worked with infinite series. You studied so-called p -series which have the form $\sum_{n=1}^{\infty} \frac{1}{n^p}$, and showed that these series converge when $p > 1$. Recall that $p = 1$ corresponds to the harmonic series which diverges, and for $p < 1$ the series diverges by comparison.

What the calculus result really says is that $\zeta(s)$ is a function whose domain in \mathbb{R} is $(1, \infty)$. It doesn't take much more effort to look at series in \mathbb{C} , and we find that the actual domain of $\zeta(s)$ is the right half-plane $\Re(s) > 1$ and defines an analytic function there. Now here comes the rub. There is a famous conjecture (better than Fermat's) which says that (except for some trivial cases) the function $\zeta(s)$ is zero only when $\Re(s) = 1/2$. This is the famous Riemann hypothesis. There is just one problem. The function $\zeta(s) = \sum_{n=1}^{\infty} \frac{1}{n^s}$ isn't defined where the hypothesis is telling us to look. That's where the identity theorem comes in. Suppose with a bit more math, you could define an analytic function $Z(s)$ whose domain was all of \mathbb{C} except for the point $s = 1$, and for which $Z(x) = \zeta(x)$ for all real $x > 1$. Then $Z(s) = \zeta(s)$ for all complex points $\Re(s) > 1$ and $Z(s)$ defines what is called an analytic continuation of the zeta function $\zeta(s)$. The new function $Z(s)$ is the one to which the Riemann hypothesis refers.

The third distinction between real and complex analytic functions is Liouville's theorem, which says that any function $f : \mathbb{C} \rightarrow \mathbb{C}$ which is analytic in all of \mathbb{C} (called an entire function) and bounded must be a constant. This certainly doesn't happen for real-valued functions, e.g., $\sin x$, or even more complicated ones, such as $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined by

$$f(x, y) = \begin{cases} \frac{\sin(x^2+y^2)}{x^2+y^2} & (x, y) \neq (0, 0), \\ 1 & (x, y) = (0, 0). \end{cases}$$

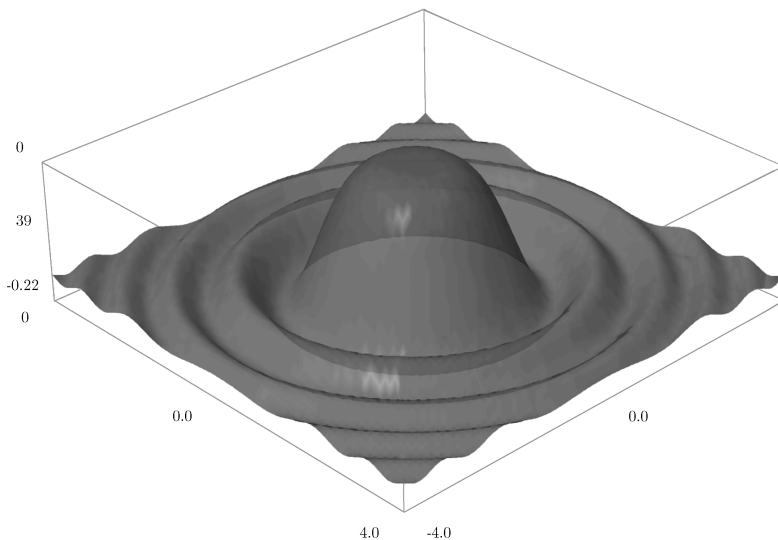


Figure A.2

The function is infinitely differentiable and bounded between -1 and 1 , but it's obviously not constant. Its graph is pictured in Figure A.2.

A.3. Return to the Birch and Swinnerton-Dyer Conjecture

We start with the elliptic curve $E_n : y^2 = x^3 - n^2x$. Notice that while we have been interested in the set of rational points $E_n(\mathbb{Q})$, it also makes sense to think about $E_n(\mathbb{F}_p)$ at least for primes $p \nmid n$. Since, over the finite field, the number of points is finite, we can count them and record the information as follows. For $p \nmid 2n$, let $\#E_n(\mathbb{F}_p) = p + 1 - a_p$, the shape here influenced by Hasse's theorem. One forms an “ L -function” associated to the curve E_n ,

$$L(E_n, s) = \prod_{p \nmid 2n} \left(\frac{1}{1 - a_p p^{-s} + p^{1-2s}} \right).$$

As mysterious as this looks, it is just a complex-valued function, very much like the Riemann zeta function, which is defined on the half-plane $\Re(s) > 3/2$. As with the Riemann zeta function, $L(E_n, s)$ has an analytic continuation (which we also call $L(E_n, s)$) to the whole complex plane, and in particular is defined at $s = 1$. Since the function is analytic at $s = 1$, it makes sense to talk about its order of vanishing. For example, for real-valued functions, $f(x) = x^2$ is nonzero at $x = 1$ so has zero order of vanishing. The function $f(x) = (x - 1)^k(x^2 + 3)$ has k th order vanishing at $x = 1$. So to proceed, we write $L(E_n, s) = (s - 1)^k g(s)$ where $g(1) \neq 0$. Then k is the order of vanishing.

One version of the Birch and Swinnerton-Dyer conjecture is that if $L(E_n, s) = (s - 1)^k g(s)$ with $g(1) \neq 0$, and $E_n(\mathbb{Q}) = \mathbb{Z}^r \oplus E_n(\mathbb{Q})_{tor}$, then $r = k$, that is the (algebraic) rank of the elliptic curve is the order of vanishing of its L -function at $s = 1$, its so-called analytic rank.

We bring this all the way back to the congruent number problem. We knew that n being a congruent number depended upon $E_n(\mathbb{Q})$ having infinitely many rational points. This is the same as saying the rank r is positive, which (given the Birch and Swinnerton-Dyer conjecture) is simply to say $L(E_n, 1) = 0$. Actually half of this connection is known. The Coates–Wiles theorem says that if $r \geq 1$, then $L(E_n, 1) = 0$. The converse is the weak version of the Birch and Swinnerton-Dyer conjecture.

Finally connecting the vanishing of $L(E_n, s)$ at $s = 1$ to the formulas in Tunnell’s theorem is where things *really* get exciting, but to talk about that we need modular forms and the Shimura lift, something well beyond the scope of this book.

A.4. Elliptic Curves over \mathbb{C}

Here we give a sketch that an elliptic curve over \mathbb{C} is a torus. In a sense this topic will also revisit many of the ideas we have developed in the text: equivalence relations, groups, projective curves, complex variables, and more. Again, this presentation will be highly telegraphic.




Figure A.3

Consider a piece of the complex plane shown in Figure A.3 with distinguished vectors ω_1 and ω_2 . These two vectors form two sides of what we will call the fundamental parallelogram Π . The vertex diagonally opposite the origin is $\omega_1 + \omega_2$.

There are several things to observe about this image. First is that Π tiles the plane by translation; you can see the translated parallelograms outlined by dashed lines. Second the vertices of all the parallelograms are precisely the set of points $\Lambda = \{a\omega_1 + b\omega_2 \mid a, b \in \mathbb{Z}\}$. More importantly, Λ is an abelian group under addition.

Now consider the set of points in the fundamental parallelogram $\{a\omega_1 + b\omega_2 \mid 0 \leq a, b \leq 1\}$, but let's also call this Π , blurring the distinction between the parallelogram and the region it bounds.

Let's define a relation on the points in \mathbb{C} . We shall say $z_1 \sim z_2$ if and only if $z_2 - z_1 \in \Lambda$. It is trivial to check that this is reflexive, symmetric, and transitive, hence it is an equivalence relation. One checks that every point in \mathbb{C} is equivalent (can be translated by integer multiples of ω_1 and ω_2) to a point in Π . For example, if $\omega_1 = 1$ and $\omega_2 = i$, then $-4.5 + 26.3i \sim .5 + .3i \in \Pi$. Moreover, two points in Π are equivalent if and only if they are on the boundary of Π : related by $z \sim z + \omega_i$. This has an important geometric (or really topological) interpretation. What we have said is that opposite sides of the parallelogram bounding Π should be identified. If you think of the parallelogram as a sheet of paper, when we identify two opposite edges, we roll the paper up into a cylinder gluing the two edges together. Now imagine the cylinder long and flexible. We could then fold the two ends up and glue them together forming a donut, mathematically known as a torus. It turns out that the torus is a group, and as a group it is isomorphic to the group of points on a complex elliptic curve.

First, let's identify the group. We shall define \mathbb{C}/Λ (read " \mathbb{C} mod Λ ") to be the set of equivalence classes under our equivalence relation defined above:

$$\mathbb{C}/\Lambda = \{[z] \mid z \in \mathbb{C}\}.$$

This is just like defining \mathbb{Z}_n from the equivalence relation on \mathbb{Z} with $a \sim b$ if and only if $a \equiv b \pmod{n}$. In particular, we can define a group law on \mathbb{C}/Λ by defining $[z] + [w] = [z + w]$. The identity is $[0]$ and the inverse of $[z]$ is $[-z]$. Moreover, there is a one-to-one correspondence between the elements of this group and the points on the torus. Most people identify the two.

Now we need an elliptic curve and a map from the torus to the elliptic curve. We start in what appears a roundabout manner. We all know that periodic functions on the real line have interesting properties, starting with sine and cosine and progressing to the theory of Fourier series, which is a good deal more versatile than the Taylor series you studied in calculus. But for now, we shall settle for the statement that periodic functions are important.

What about periodic functions in the complex plane? What kinds of functions exist that satisfy $f(z + \omega_1) = f(z)$ and $f(z + \omega_2) = f(z)$?

These are called doubly periodic functions. Naively, we might look for analytic, doubly periodic functions, but this turns out to be uninteresting for the following reason. The first important observation is that if f is doubly periodic, then all the values $f(z)$ are determined by $z \in \Pi$. Now even a continuous function on a closed bounded region like Π (the fancy word is *compact*) achieves an absolute maximum and a minimum, so the function is bounded on all of \mathbb{C} . So if in addition f is analytic, then it is a bounded, entire function, which by Liouville, must be constant. This is what we meant by uninteresting.

If f were analytic, it would have a power series $\sum_{n=0}^{\infty} a_n z^n$, but if it is not, it can still have a series expansion. It's just that there may be some negative exponents: $\sum_{n=\mu}^{\infty} a_n z^n$ for $\mu < 0$. Functions like this are called meromorphic.

Now we define yet another object without foreshadowing, but not unexpectedly.

Let $\mathcal{E}_\Lambda = \{f : \mathbb{C} \rightarrow \mathbb{C} \mid f \text{ is meromorphic and } f(z + \lambda) = f(z) \text{ for all } \lambda \in \Lambda\}$. It is clear that constant functions are in \mathcal{E}_Λ , so the set is nonempty. If $f, g \in \mathcal{E}_\Lambda$, then so is $f \pm g$, $f \cdot g$ and f/g as long as $g \neq 0$. But this makes \mathcal{E}_Λ a field, called the field of elliptic functions with period lattice Λ . Something else is true:

Proposition A.8. *If $f \in \mathcal{E}_\Lambda$, then so is its derivative f' .*

Proof. We really need only show that f' is also periodic. From calculus, we observe that

$$\begin{aligned} f'(z + \lambda) &= \lim_{h \rightarrow 0} \frac{f(z + \lambda + h) - f(z + \lambda)}{h} \\ &= \lim_{h \rightarrow 0} \frac{f(z + h) - f(z)}{h} = f'(z), \end{aligned}$$

where at the second equality we used that f was periodic (i.e., in \mathcal{E}_Λ). \square

There is a great deal more one could say, but we're almost at the end, so we'll push on. An extremely important example of an elliptic function is the Weierstrass function $\wp(z)$. We skip its formal

definition and simply say that it has a series expansion of the form

$$\wp(z) = \frac{1}{z^2} + a_2 z^2 + a_4 z^4 + a_6 z^6 + \cdots .$$

Equally important is its derivative

$$\wp'(z) = \frac{-2}{z^3} + 2a_2 z + 4a_4 z^3 + 6a_6 z^5 + \cdots .$$

We consider $\wp'(z)^2$ and $\wp(z)^3$ and compare

$$\begin{aligned} \wp'(z)^2 &= \frac{4}{z^6} - \frac{8a_2}{z^2} - 16a_4 + z^2(\cdots) \in \mathcal{E}_\Lambda, \\ \wp(z)^3 &= \frac{1}{z^6} + \frac{3a_2}{z^2} + 3a_4 + z^2(\cdots) \in \mathcal{E}_\Lambda. \end{aligned}$$

We compute

$$\wp'(z)^2 - 4\wp(z)^3 = \frac{-20a_2}{z^2} - 28a_4 + z^2(\cdots), \text{ so}$$

$$\wp'(z)^2 - 4\wp(z)^3 + 20a_2\wp(z) = -28a_4 + z^2(\cdots) \in \mathcal{E}_\Lambda.$$

What is the point of this? Well both sides of the last equation are elliptic functions in \mathcal{E}_Λ , but from the right-hand side we see that the function is actually analytic, and by Liouville, analytic elliptic functions are constant, so the right-hand side is just $-28a_4$. Putting this all together, we see that

$$\wp'(z)^2 = 4\wp(z)^3 - 20a_2\wp(z) - 28a_4.$$

If we let $x = \wp(z)$ and $y = \wp'(z)$, then $y^2 = 4x^3 - 20a_2x - 28a_4$. The point (x, y) is a point on an elliptic curve!

We define a function $\Phi : \mathbb{C}/\Lambda \rightarrow \mathbb{P}^2(\mathbb{C})$ whose image is the elliptic curve $E : y^2 = 4x^3 - 20a_2x - 28a_4$ by

$$\Phi(z) = \begin{cases} [\wp(z), \wp'(z), 1] & z \neq 0, \\ [0, 1, 0] & z = 0. \end{cases}$$

Then $\Phi : \mathbb{C}/\Lambda \rightarrow E(\mathbb{C})$ is an isomorphism of groups. If $z_1 \mapsto P_1 = [\wp(z_1), \wp'(z_1), 1]$ and $z_2 \mapsto P_2 = [\wp(z_2), \wp'(z_2), 1]$, then $z_1 + z_2 \mapsto P_1 \oplus P_2$ (as we would define the sum of points on the elliptic curve), which equals $[\wp(z_1 + z_2), \wp'(z_1 + z_2), 1]$. Using that \mathbb{C}/Λ is a group and Φ is a group homomorphism provides an independent proof that

addition of points on the elliptic curve is an associative operation: if $P_i = \Phi(z_i)$ for $i = 1, 2, 3$, then

$$\begin{aligned} P_1 \oplus (P_2 \oplus P_3) &= \Phi(z_1) \oplus (\Phi(z_2) \oplus \Phi(z_3)) = \Phi(z_1) \oplus \Phi(z_2 + z_3) \\ &= \Phi(z_1 + (z_2 + z_3)) = \Phi((z_1 + z_2) + z_3) \\ &= \Phi(z_1 + z_2) \oplus \Phi(z_3) \\ &= (\Phi(z_1) \oplus \Phi(z_2)) \oplus \Phi(z_3) = (P_1 \oplus P_2) \oplus P_3. \end{aligned}$$

Answers to Selected Exercises

B.1. Chapter 2

Exercise (page 16). Analogous to what we did above, find a parametrization for the points on the circle $x^2 + y^2 = 2$, and extract a characterization of the rational points. To start, project from the rational point $(1, 1)$. *Note:* Projecting onto the x - or y -axis does not work as expected, as not all lines from $(1, 1)$ to points on the circle intersect those axes. Instead, try to project onto the line $y = -x$.

Solution. We project from $(1, 1)$ onto line $y = -x$. Consider the line between $(1, 1)$ and the point $(t, -t)$. If $t \neq 1$, the line has slope $m = (1+t)/(1-t)$, and the equation of the line is $y - 1 = m(x - 1)$. We'll worry about $t = 1$ later. Let (a, b) be the other point of intersection of the line and the circle, so that $b = m(a - 1) + 1$ and

$$b^2 = 2 - a^2 = (m(a - 1) + 1)^2 = m^2(a - 1)^2 + 2m(a - 1) + 1.$$

Subtracting 1 from both sides yields

$$1 - a^2 = m^2(a - 1)^2 + 2m(a - 1) = m^2(1 - a)^2 - 2m(1 - a).$$

Since we are assuming $a \neq 1$, we can divide by $1 - a$ on both sides, yielding

$$1 + a = m^2(1 - a) - 2m.$$

We solve for a and b :

$$a = \frac{m^2 - 2m - 1}{1 + m^2} = \frac{-1 + 2t + t^2}{1 + t^2},$$

$$b = m(a - 1) + 1 = \frac{t^2 - 2t - 1}{1 + t^2},$$

so

$$(a, b) = \left(\frac{-1 + 2t + t^2}{1 + t^2}, \frac{t^2 - 2t - 1}{1 + t^2} \right).$$

Having produced the formula, we see that it is actually valid when $t = 1$, yielding the point $(1, -1)$.

We note that rational values of t give rise to rational points on the circle via the formula above. Given a rational point (a, b) on the circle, we note that the slope m of the line through $(1, 1)$ and (a, b) will be rational. But $m = (1+t)/(1-t)$ which means that $t = (m-1)/(1+m)$, so m rational implies t is rational. So as in the case of the unit circle there is a one-to-one correspondence (except for $(1,1)$) between rational points (a, b) on the circle and rational values of t . \square

Exercise (page 16). Now consider the issue of rational points on $x^2 + y^2 = 3$. In contrast to the examples above, prove that there are no rational points on this curve, and describe the crucial difference between this example and the one before.

Solution. If there were a rational point $(a/b, c/d)$ on the curve, substituting and clearing denominators would produce an integer equation of the form $A^2 + B^2 = 3C^2$. There is no loss to assume A, B, C have no common factor. Any integer n has the form $n = 3k + r$ with $r = 0, 1, 2$ by the division algorithm, so $n^2 = 9k^2 + 6kr + r^2 \equiv 0, 1 \pmod{3}$. Since the right-hand side of the equation is congruent to zero modulo 3, our only possibility is that $A \equiv B \equiv 0 \pmod{3}$. But this means $9 \mid 3C^2$ or that $3 \mid C^2$. This forces $3 \mid C$, a contradiction since we were assuming A, B, C have no common factor. \square

Exercise (page 26). Find a square-free congruent number not in the list above, showing all work to obtain it.

Solution. Running the code for $n = 9$ yields the following table, which reveals 14 and 390 as new congruent numbers.

m	n	A	B	C	CN
1	2	3	4	5	6
2	3	5	12	13	30
1	4	15	8	17	15
3	4	7	24	25	21
2	5	21	20	29	210
4	5	9	40	41	5
1	6	35	12	37	210
5	6	11	60	61	330
2	7	45	28	53	70
4	7	33	56	65	231
6	7	13	84	85	546
1	8	63	16	65	14
3	8	55	48	73	330
5	8	39	80	89	390
7	8	15	112	113	210

□

Exercise (page 28). Find all the rational points on the curve $x^n + y^n = 1$ where n is an integer, $n > 2$.

Solution. Let $(\frac{a}{b}, \frac{c}{d})$ be a rational point on $x^n + y^n = 1$, i.e., $a, b, c, d \in \mathbb{Z}$ and $b, d \neq 0$. Then $(\frac{a}{b})^n + (\frac{c}{d})^n = 1$ implies (clearing denominators) that $(ad)^n + (bc)^n = (bd)^n$. But that means that (ad, bc, bd) is an integral solution to Fermat's equation $x^n + y^n = z^n$. By Fermat's Last Theorem, we know there can only be a solution to $x^n + y^n = z^n$ ($n > 2$) if one of x, y , or z is zero.

The answer depends on the parity of n . For n even, the only rational points are $(\pm 1, 0)$ and $(0, \pm 1)$. For n odd, the only rational points are $(1, 0)$ and $(0, 1)$. □

Exercise (page 34). Let V be the set of functions $f : \mathbb{R} \rightarrow \mathbb{R}$ which satisfy the differential equation $f'' - f = 0$. Show that V is a vector space over \mathbb{R} and, assuming its dimension is 2, find a basis for V .

Solution. The proof is exactly as in the previous examples, and a basis is $\{e^x, e^{-x}\}$. \square

Exercise (page 38). Using the ideas above, prove the Bachet duplication formula for $y^2 = x^3 + k$, $k \neq 0$.

$$(x, y) \mapsto \left(\frac{x^4 - 8kx}{4y^2}, \frac{-x^6 - 20kx^3 + 8k^2}{8y^3} \right).$$

We outline some useful steps.

- (1) Use implicit differentiation to derive a formula for the slope of the tangent line to the curve $y^2 = x^3 + k$, which is valid at any point (x, y) where $y \neq 0$.

Solution to (1). If we differentiate $y^2 = x^3 + k$ with respect to x , we obtain

$$2yy' = 3x^2, \text{ so } y' = \frac{3x^2}{2y}. \quad \square$$

- (2) Now write down the equation of the tangent line to the curve at the point (a, b) where we assume $b \neq 0$. It will be convenient if you use m for the slope for the time being until you have need to use its actual value.

Solution to (2). Of course this is just

$$y = m(x - a) + b, \text{ where } m = \frac{3a^2}{2b}. \quad \square$$

- (3) Now we want to find the point(s) of intersection of the tangent line with the cubic, and this requires a little work. Substituting the expression for y given by the line into the equation defining the cubic results in an equation of the form $f(x) = 0$ where f is a polynomial of degree three. Your job is to factor the polynomial since its roots are the x coordinates corresponding to the points of intersection. Here you catch a bit of a break. For sure one of the roots is a , which means $(x - a)$ is a factor. But it should not be too much

of a surprise that a is (at least) a double root since the line is tangent to the curve at $x = a$ (much like $y = (x - a)^r$ is tangent to the x -axis at $x = a$ and the root a has *multiplicity* r). After you factor out the first of the $(x - a)$'s, it would be a good time to put in the real value of m to see what simplifies.

Solution to (3). We begin with

$$\begin{aligned} [m(x - a) + b]^2 &= y^2 = x^3 + k, \text{ which becomes} \\ m^2(x - a)^2 + 2mb(x - a) + b^2 &= x^3 + k, \text{ or} \\ x^3 - m^2(x - a)^2 - 2mb(x - a) - b^2 + k &= 0. \end{aligned}$$

Noting that (a, b) is a point on the cubic means that $b^2 = a^3 + k$, meaning the term $-b^2 + k = -a^3$. Thus the expression above becomes

$$\begin{aligned} 0 &= x^3 - a^3 - m^2(x - a)^2 - 2mb(x - a) \\ &= (x - a) [x^2 + ax + a^2 - m^2(x - a) - 2mb] \\ &= (x - a) \left[x^2 + ax + a^2 - \left(\frac{3a^2}{2b} \right)^2 (x - a) - 2b \frac{3a^2}{2b} \right] \\ &= (x - a) \left[x^2 + ax - 2a^2 - (x - a) \left(\frac{3a^2}{2b} \right)^2 \right] \\ &= (x - a) \left[(x - a)(x + 2a) - (x - a) \left(\frac{3a^2}{2b} \right)^2 \right] \\ &= (x - a)^2 \left[x + 2a - \frac{9a^4}{4b^2} \right] \\ &= (x - a)^2 \left(x - \frac{a(9a^3 - 8b^2)}{4b^2} \right) \\ &= (x - a)^2 \left(x - \frac{a(9a^3 - 8(a^3 + k))}{4b^2} \right) \\ &= (x - a)^2 \left(x - \frac{a^4 - 8ak}{4b^2} \right). \end{aligned}$$

So the other point of intersection occurs when $x = \frac{a^4 - 8ak}{4b^2}$, and substituting into the equation of the line produces $y = \frac{-a^6 - 20ka^3 + 8k^2}{8b^3}$, which provides the Bachet formula. \square

Exercise (page 40). Properties of rational lines in the plane.

- (1) Is every point on a rational line a rational point?

Solution to (1). The answer is no. For example, $y = x$ (i.e., $Z(y - x)$) is certainly a rational line, and yet $(\sqrt{2}, \sqrt{2})$ is a point on the line.

- (2) If a line passes through at least two rational points, is it a rational line? What about lines if we only know one rational point through which they pass?

Solution to (2). If a nonvertical line passes through two rational points, then the point-slope formula will prove the line is a rational line. But of course the line $y = \sqrt{2}x$ is not a rational line, but passes through the rational point $(0, 0)$. Similarly, if a vertical line passes through two rational points, it says only that the line has the form $x = r$ where r is a rational number, so is rational.

- (3) Consider two distinct rational lines that intersect. Do they intersect in a rational point?

Solution to (3). Given two distinct lines that intersect, think about how you solve for the point of intersection. Perhaps you know something fancy like Gaussian elimination. If not, you will learn it in a linear algebra class, but with just two variables. In essence you take one equation, say $ax + by + c = 0$, and solve for either x or y . For example, if $b \neq 0$, one could write $y = -ax/b - c/b$. Note that the coefficients $-a/b$ and $-c/b$ are still rational numbers. Then you would substitute this expression into the second equation and solve for the remaining variable—in this case x , always doing arithmetic which only involved the rational numbers. That would produce a rational x -coordinate which, when put into one of the rational lines, would produce a rational y -coordinate, hence a rational point of intersection. The case where $a \neq 0$ is analogous. \square

Exercise (page 40). Characterizing the intersection of lines and conics.

- (1) In how many points can two arbitrary lines (in the plane) intersect?

Solution to (1). If we draw upon our experience in \mathbb{R}^2 , then the answer should be 0, 1, or an infinite number. Two distinct lines could be parallel or intersect, or the two given lines may actually be the same line. \square

- (2) In how many points can a line and a conic intersect?

Solution to (2). If we think simply of a line and a parabola, we have no trouble drawing pictures (or generating equations) where the answer is 0, 1, or 2. How about an infinite number? Well, how did we define a conic? In just the same way as the set of points in the plane which satisfy an equation of the form $ax^2 + bxy + cy^2 + dx + ey + f = 0$. So what about the line $y = x$ and the conic $x^2 - y^2 = 0$? Hmm, something seems not quite fair. The problem is that the conic $x^2 - y^2 = 0$ factors as $(x - y)(x + y) = 0$. We can take care of this potential degeneracy. \square

Exercise (page 42). In the questions below, we assume all the plane curves are irreducible, meaning they are the zero sets of polynomials $f(x, y)$ where $f(x, y)$ is irreducible. It follows (from abstract algebra) that two distinct irreducible curves can only intersect in a finite number of points. The questions below try to get at discovering what that number might be.

For all the problems below, consider your curves in \mathbb{R}^2 . Can you come up with examples which suggest answers to the following questions? Can you prove any of your assertions?

- (1) In how many points can two (distinct) conics intersect?
- (2) In how many points can a conic and a cubic intersect?
- (3) In how many points can two (distinct) cubics intersect?
- (4) What would be your guess for a generalization?
- (5) Consider the intersection of a rational line with a rational conic.

- (a) Are the point(s) of intersection necessarily rational? Give a proof or provide a counterexample.
- (b) Now let's suppose that the line intersects the conic in two points, one of which is rational. Is the second necessarily rational? Give a proof or a counterexample.

To assist with your intuition, a few curves to consider are illustrated in Figure B.1. The first set fixes a parabola and slides the circle up the y -axis. The second set is a cubic and quartic, and a cubic and conic.

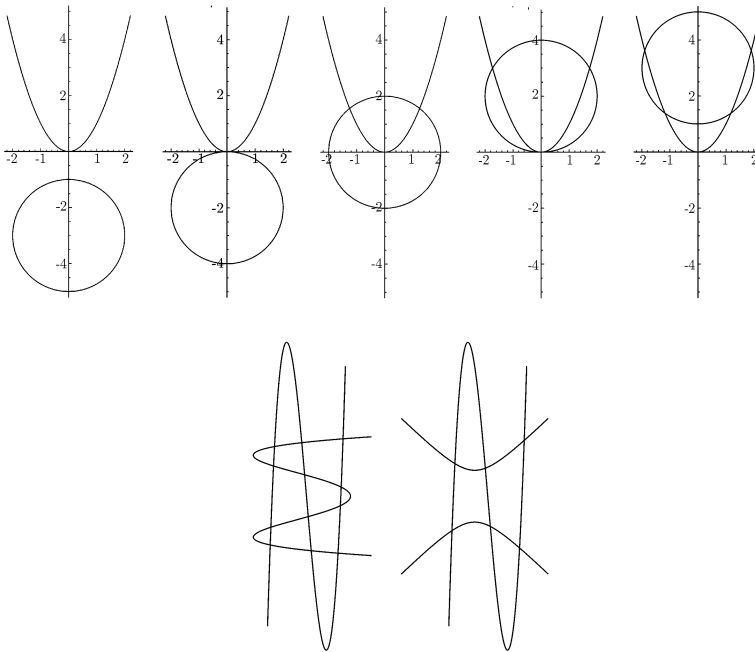


Figure B.1

Solution. For (5), it is clear the answer to the first question is no, since the line $y = x$ intersects the unit circle $x^2 + y^2 = 1$ at the points $\pm(\sqrt{2}/2, \sqrt{2}/2)$ which are not rational points.

The answer to the second question is yes, but it takes some thought. First consider the more general situation of an arbitrary

conic given by $h(x, y) = ax^2 + bxy + cy^2 + dx + ey + f = 0$ and a line which intersects the conic at a point (x_0, y_0) . Based upon the exercises above, we would conjecture that the line can intersect the conic in at most two points, but perhaps we should prove that. The line that intersects the conic at (x_0, y_0) either has the form $x = x_0$ or $y = px + q$. Substituting $y = px + q$ into the equation for the conic will produce a quadratic equation in the variable x , namely

$$Q(x) = h(x, px + q) = \alpha x^2 + \beta x + \gamma.$$

If the line was $x = x_0$, then substitution will produce a quadratic in the variable y . We discuss the first case; the second is completely analogous. Now the key observation is that the roots of this quadratic are the only possible x -coordinates for points of intersection of the line with the conic. We see this as follows. If (x_1, y_1) is any point that lies in the intersection of the line and the conic, then $y_1 = px_1 + q$ and $h(x_1, y_1) = 0$. Putting them together, we see that $0 = h(x_1, y_1) = h(x_1, px_1 + q) = Q(x_1) = \alpha x_1^2 + \beta x_1 + \gamma$, so there are at most two x -coordinates for points of intersection, and each x -coordinate gives only one y -coordinate using $y = px + q$, so there are at most two points of intersection.

It is interesting to think about the case in which there might be only one point of intersection. A tangent line might come to mind, but that should mean that x_0 is a double root of the quadratic. But is it possible that there is only one point of intersection? (For example if the coefficient $\alpha = 0$.) Consider

$$x^2 - 4y^2 + 3y - x - 1 = 0,$$

which has a rational point $(2, 1)$. The line $y = x/2$ intersects the conic at the point $(2, 1)$, and it is clearly not tangent to the conic. Where else does the line intersect the conic? This is another example of where the introduction of projective space becomes relevant.

Returning to our example, we are assuming that the conic and line are rational, there are two points of intersection (x_0, y_0) and (x_1, y_1) , and that (x_0, y_0) is rational. We may as well assume that $x_0 \neq x_1$, precluding the easy case of a tangent line to the conic.

Now in the analysis above, the line $y = px + q$ is rational, so substitution into $h(x, y)$ (which has rational coefficients) produces

the quadratic $Q(x) = \alpha x^2 + \beta x + \gamma$, with rational coefficients, and because there are two distinct roots of this quadratic, we know that $\alpha \neq 0$.

There are a number of ways in which to see that this second root x_1 is also rational, one involving the so-called division algorithm for polynomials with rational coefficients (which is part of an abstract algebra class), but we pursue the end result directly. The quadratic does factor—if you are nervous, it certainly factors over the complex numbers—so let's write

$$\begin{aligned} Q(x) &= \alpha x^2 + \beta x + \gamma = \alpha(x - x_0)(x - x_1) \\ &= \alpha(x^2 - (x_0 + x_1)x + x_0x_1). \end{aligned}$$

So comparing the coefficients of x in the expressions, we see that

$$-\beta = \alpha(x_0 + x_1).$$

Since $\alpha \neq 0$, we have that $x_1 = -\beta/\alpha - x_0$ is rational, and substituting that value into the equation for the rational line shows that y_1 is also rational. \square

Exercise (page 44). As a simple example, show that the curve $y = (x - a)^k$ intersects the x -axis with multiplicity k at $x = a$ and with multiplicity 0 at all other points $x = b$.

Solution. If we follow the paradigm, we put $f(x, y) = y - (x - a)^k$. The line in question is $y = 0$, so $h(x) = f(x, 0) = -(x - a)^k$. We have $h(x) = (x - a)^k \cdot (-1) = (x - a)^k q(x)$ with $q(x) = -1$, so $q(a) \neq 0$, which establishes the result at $x = a$. At $x = b$ (with $b \neq a$), we have $h(x) = -(x - a)^k = (x - b)^0 q(x)$ with $q(x) = -(x - a)^k$, and we note $q(b) = -(b - a)^k \neq 0$ which establishes the second result. \square

Exercise (page 44). Next, let's gain a little more insight by examining the case of zeroes of orders 1 and 2. Let $h(x)$ be a polynomial of degree $n \geq 2$ with coefficients in a field F , and let $a \in F$. Then the following hold:

- (1) $h(x) = (x - a)q(x) + h(a)$ for some polynomial q having coefficients in F .
- (2) $h(a) = 0$ if and only if $h(x) = (x - a)q(x)$.

(3) h has a double root at a if and only if $h(a) = h'(a) = 0$.

Solution. The first statement is completely general and is a consequence of the division algorithm in polynomial rings, a fact you learn in an abstract algebra course, but which is probably known to you from high school. Loosely speaking, when you divide one polynomial by another, you get a quotient and remainder with the remainder having degree less than the degree of the polynomial by which you divided. Thus if we divide $h(x)$ by $x - a$, we get $h(x) = (x - a)q(x) + r(x)$ for some polynomials q, r . Since the degree of r is less than 1, r is a constant which we evaluate by plugging in $x = a$: $h(a) = (a - a)q(a) + r = r$.

The second statement is now immediate from the first, but in our earlier terminology it says that h has a zero of order at least one if and only if $h(a) = 0$.

The third item is where more interest lies. The polynomial h has a double root at $x = a$, which means $h(x) = (x - a)^2q(x)$ for some polynomial q (and $q(a) \neq 0$). Obviously $h(a) = 0$ and the product rule for derivatives shows that $h'(a) = 0$ as well: $[h'(x) = (x - a)^2q'(x) + 2(x - a)q(x)]$.

Conversely, suppose that $h(a) = h'(a) = 0$. Since $h(a) = 0$, we know that $h(x) = (x - a)q_1(x)$. Now $h'(x) = (x - a)q_1'(x) + q_1(x)$, so $h'(a) = 0$ means that $q_1(a) = 0$, which means that $q_1(x) = (x - a)q(x)$ for some polynomial q . Putting things together, we see that $h(x) = (x - a)q_1(x) = (x - a)^2q(x)$, which means h has a double root at $x = a$. \square

Exercise (page 44). Establish the following generalization of the work we have started above. Show that h has a zero of order k at $x = a$ if and only if $h(a) = h'(a) = \dots = h^{(k-1)}(a) = 0$ and $h^{(k)}(a) \neq 0$, where $h^{(i)}$ is the i th derivative of h . *Hint:* Taylor polynomials are your friend.

Solution. Since h is a polynomial of degree n , its Taylor series about $x = a$ is a polynomial of degree n of the form

$$h(x) = h(a) + h'(a)(x-a) + \frac{h''(a)}{2!}(x-a)^2 + \dots \\ \dots + \frac{h^{(n)}(a)}{n!}(x-a)^n.$$

Certainly, if $h(a) = h'(a) = \dots = h^{(k-1)}(a) = 0$, we have

$$h(x) = \frac{h^{(k)}(a)}{k!}(x-a)^k + \dots + \frac{h^{(n)}(a)}{n!}(x-a)^n \\ = (x-a)^k \left[\frac{h^{(k)}(a)}{k!} + \dots + \frac{h^{(n)}(a)}{n!}(x-a)^{n-k} \right] \\ = (x-a)^k q(x),$$

where

$$q(x) = \frac{h^{(k)}(a)}{k!} + \dots + \frac{h^{(n)}(a)}{n!}(x-a)^{n-k}$$

and

$$q(a) = \frac{h^{(k)}(a)}{k!} \neq 0$$

by assumption.

Conversely, consider $h(x) = (x-a)^k q(x)$ for some polynomial q satisfying $q(a) \neq 0$. There are many ways to proceed, but all boil down to the fact that if we write any polynomial $h(x) = b_0 + b_1(x-a) + \dots + b_n(x-a)^n$, then by induction we prove that $b_k = \frac{h^{(k)}(a)}{k!}$, that is any polynomial expansion is the Taylor expansion. Write $q(x) = q(a) + \dots + \frac{q^{(n-k)}(a)}{(n-k)!}(x-a)^{n-k}$ as a Taylor polynomial. Then

$$h(x) = q(a)(x-a)^k + \dots + \frac{q^{(n-k)}(a)}{(n-k)!}(x-a)^n \\ = h(a) + h'(a)(x-a) + \frac{h''(a)}{2!}(x-a)^2 + \dots \\ + \frac{h^{(n)}(a)}{n!}(x-a)^n.$$

From the uniqueness, it is clear that $h(a) = h'(a) = \dots = h^{(k-1)}(a) = 0$ and $h^{(k)}(a) \neq 0$. \square

Exercise (page 44). Now consider $y^2 = g(x)$ where g is a cubic, that is the zero set of $f(x, y) = y^2 - g(x)$. We want to see that a nonvertical tangent has multiplicity at least two at the point of tangency.

Solution. Let (a, b) be a point on the curve $y^2 = g(x)$, that is, if $f(x, y) = y^2 - g(x)$, then $f(a, b) = 0$. To compute the slope of the tangent line, we differentiate $y^2 = g(x)$ implicitly. We obtain $2yy' = g'(x)$, so $dy/dx = g'(x)/2y$, and $m = g'(a)/2b$ is the slope of the tangent line as long as we stay away from the points on the x -axis where the tangent line is vertical. The equation of the tangent line is $y = m(x - a) + b$. To find the multiplicity of the point of intersection of $y^2 = g(x)$ and the tangent line consider $h(x) = f(x, m(x - a) + b) = [m(x - a) + b]^2 - g(x)$ and claim $h(x) = (x - a)^2q(x)$, that is h had at least a double root at $x = a$.

To verify this we need only check that $h(a) = h'(a) = 0$. We see that $h(a) = b^2 - g(a) = 0$ since (a, b) is a point on the curve. We compute $f'(x) = g'(x) - 2m[m(x - a) + b]$, so that $h'(a) = 2mb - g'(a) = 0$ from our computation of m . \square

B.2. Chapter 3

Exercise (page 55). Let $p > 1$ be a prime. Show the following.

- For any integer n , $\gcd(p, n) = 1$ or p .
- For integers m, n , if $p \mid mn$, then either $p \mid m$ or $p \mid n$.

Solution. The gcd is a positive integer which is a common divisor, so the gcd of a prime p and an integer n must be a divisor of p . But since p is prime, its only positive divisors are 1 and p which already gives the statement. We can amplify it by saying that $\gcd(p, n) = p$ if and only if $p \mid n$. For the second assertion, assuming that $p \mid mn$, either $p \mid m$ in which case we are done, or $p \nmid m$ which from above implies that $\gcd(p, m) = 1$. The assertion now follows from Corollary 3.14 that $p \mid n$. \square

Exercise (page 62). Can you find integers x, y, z so that $987654319 = x^2 + y^2 + z^2$? *Hint:* Determine the possible values of $x^2 + y^2 + z^2 \pmod{8}$.

Solution. One easily checks that for an integer m , $m^2 \equiv 0, 1, 4 \pmod{8}$ are the only possibilities, so that $x^2 + y^2 + z^2$ can never be congruent to $987654319 \equiv 7 \pmod{8}$. \square

Exercise (page 62) (A precursor to the Chinese Remainder Theorem). Find the smallest number of marbles in a jar so that one remains if the marbles are taken out 2, 3, 5 at a time, but none remain if taken out 11 at a time.

Solution. Really, you have been asked to solve the following system of congruences:

$$\begin{aligned}x &\equiv 1 \pmod{2}, \\x &\equiv 1 \pmod{3}, \\x &\equiv 1 \pmod{5}, \\x &\equiv 0 \pmod{11}.\end{aligned}$$

Now the first three say that $(x - 1)$ is divisible by 2, 3, and 5, and since these numbers are coprime in pairs, it follows that $(x - 1)$ is divisible by 30, so we are reduced to solving $x \equiv 1 \pmod{30}$ and $x \equiv 0 \pmod{11}$. You can consider the sequence satisfying the first congruence and look for the first instance where a number is divisible by 11: 1, 31, 61, 91, 121, \dots , and we have a winner at 121. \square

Exercise (page 66). Explore an encryption scheme known as ROT13; it is a shift cipher. What can you say about the encryption and decryption functions E and D ?

Solution. In ROT13, the shift is by 13, that is $C \equiv P + 13 \pmod{26}$, but since $-13 \equiv +13 \pmod{26}$, the encryption and decryption algorithms are identical. \square

Exercise (page 71). Let $m, n > 1$ be coprime integers, and let a, b be arbitrary integers. Then the system of congruences

$$\begin{aligned}x &\equiv a \pmod{m}, \\x &\equiv b \pmod{n}\end{aligned}$$

has a unique solution modulo mn .

As a generous hint, note that since $\gcd(m, n) = 1$, Bézout's identity says there exists $u, v \in \mathbb{Z}$ so that $mu + nv = 1$. Show that the

number $bmu + anv$ is a solution to the system, and then prove it is unique modulo mn .

Solution. As the hint suggests, since $\gcd(m, n) = 1$, Bézout says there exists $u, v \in \mathbb{Z}$ so that $mu + nv = 1$. This means that $mu \equiv 1 \pmod{n}$ and $nv \equiv 1 \pmod{m}$. We consider $x_0 = bmu + anv$ as a candidate solution for the system of congruences. Using the congruences we deduced from Bézout, we see

$$\begin{aligned}x_0 &\equiv bmu + anv \equiv 0 + a \cdot 1 \equiv a \pmod{m}, \\x_0 &\equiv bmu + anv \equiv b \cdot 1 + 0 \equiv b \pmod{n}.\end{aligned}$$

So x_0 is a solution to the system. If x_0 and x_1 are two solutions, then $x_0 \equiv a \equiv x_1 \pmod{m}$ and $x_0 \equiv b \equiv x_1 \pmod{n}$ means that m and n both divide $x_0 - x_1$. Since $\gcd(m, n) = 1$, we have that the product $mn \mid (x_0 - x_1)$ which is to say $x_0 \equiv x_1 \pmod{mn}$, so there is a unique solution modulo mn . \square

Exercise (page 71). Explain how to use the above version of the CRT to solve a system

$$\begin{aligned}x &\equiv a \pmod{\ell}, \\x &\equiv b \pmod{m}, \\x &\equiv c \pmod{n},\end{aligned}$$

where $\ell, m, n > 1$ are integers that are coprime in pairs.

Solution. Use the CRT on the first two congruences as above to produce the single congruence $x \equiv x_0 \pmod{\ell m}$. Now the system has been reduced to two congruences, and since $\gcd(\ell m, n) = 1$, we may apply the CRT once again. Thus any system of congruences where the moduli are coprime in pairs can be solved simultaneously. \square

B.3. Chapter 4

Exercise (page 83). It is actually not difficult to show that there is a one-to-one correspondence between partitions of a set and equivalence relations on the set. While we have seen the example that the equivalence relation of congruence modulo n gives rise to the partition of the integers into congruence classes, consider the equivalence

relation associated to the partitions of M&M's we gave above. Then see if you can prove the general statement.

Solution. Let S be a nonempty set. Suppose we have a partition of S given by nonempty, pairwise-disjoint subsets $\{A_i\}_{i \in I}$. For $x, y \in S$, we define a relation by saying that $x \sim y$ if and only if $x, y \in A_i$ for some i . Since the $\{A_i\}$ form a partition of S , every element of S is in one and only one subset A_i . So the reflexive part is easy since x (and x) is in some A_i . If $x \sim y$ then x, y are in some A_i , hence so are y and x , which gives the symmetric part. And the transitive part says suppose that $x, y \in A_i$ and $y, z \in A_j$. Since $y \in A_i \cap A_j$ and we are dealing with a partition, we must have $A_i = A_j$, so $x, y, z \in A_i$; in particular $x \sim z$. So a partition gives rise to an equivalence relation.

Conversely, suppose that we have an equivalence relation on the set S . The partition will consist of all of the equivalence classes. Since $x \in [x]$, every element of S is in some class, and we know that $[x] \cap [y] \neq \emptyset$ if and only if $[x] = [y]$, so the equivalence classes are pairwise disjoint making the set of them a partition. \square

Exercise (page 93). Let p be a prime. Determine the value of $\phi(p^r)$ for any positive integer r . *Hint:* It may be easier to count the number of elements of $a \in \mathbb{Z}_{p^r}$ which are not relatively prime to p^r and use that to determine the value of the function. Of course be sure to check your answer against a few examples you can compute by hand.

Solution. For a prime p , $\phi(p^r)$ is the number of integers k , with $1 \leq k \leq p^r$ and $\gcd(k, p^r) = 1$. Now $\gcd(k, p^r) = 1$ if and only if $\gcd(k, p) = 1$. Perhaps it is easy to count the complementary set, those k with $\gcd(k, p^r) > 1$, but this is the same as those k with $\gcd(k, p) > 1$, and since p is a prime, this is just the number of k with $p \mid k$. In the collection of integers $1 \leq k \leq p^r$, every p th integer is divisible by p , so there are $p^r/p = p^{r-1}$ integers divisible by p . That means that $\phi(p^r) = p^r - p^{r-1} = p^{r-1}(p - 1)$. \square

Exercise (page 94). It is easy to show that in general $\phi(mn) \neq \phi(m)\phi(n)$ for general m, n , but what is remarkable is the when $\gcd(m, n) = 1$, $\phi(mn) = \phi(m)\phi(n)$. The function ϕ is an example of a *multiplicative* function in number theory. Perhaps more surprising is that this is a direct consequence of the Chinese Remainder

Theorem. Give a proof that ϕ is multiplicative using the following idea: Suppose that $m, n \geq 2$, and $\gcd(m, n) = 1$. Show that there is a bijection between the sets U_{mn} and $U_m \times U_n$ (ordered pairs (a, b) with $a \in U_m, b \in U_n$). Note that U_{mn} has cardinality $\phi(mn)$ and $U_m \times U_n$ has cardinality $\phi(m) \cdot \phi(n)$. To establish the bijection, define a map $F : U_{mn} \rightarrow U_m \times U_n$ by $F([a]_{mn}) = ([a]_m, [a]_n)$. You need to show this map is well-defined, one-to-one, and onto. Then deduce the result.

Some of these words may be new to you, so here are some definitions.

- We have encountered the term well-defined before. In this context it means that if $[a]_{mn} = [b]_{mn}$, then $F([a]) = F([b])$.
- The map F is one-to-one (injective) if $F([a]) = F([b])$ implies $[a]_{mn} = [b]_{mn}$.
- The map F is onto (surjective) if given $([b]_m, [c]_n) \in U_m \times U_n$, there exists $[a]_{mn} \in U_{mn}$ so that $F([a]) = ([b], [c])$.
- A map is bijective if it is one-to-one and onto.
- If $f : S \rightarrow T$ is a bijection, then S and T are said to have the same cardinality (size), and the result you are to prove is simply that when $\gcd(m, n) = 1$, the size of U_{mn} and $U_m \times U_n$ is the same.

Solution. To show the map is well-defined, we show that if $[a]_{mn} = [b]_{mn}$, then $F([a]) = F([b])$. This means that if $[a]_{mn} = [b]_{mn}$, then $[a]_m = [b]_m$ and $[a]_n = [b]_n$. But this is obvious, since the initial condition means that $mn \mid (a - b)$ so of course $m, n \mid (a - b)$.

To show the map is one-to-one is the reverse, that if $F([a]) = F([b])$ implies $[a]_{mn} = [b]_{mn}$. The condition $F([a]) = F([b])$ translates to $m \mid (a - b)$ and $n \mid (a - b)$. Since $\gcd(m, n) = 1$, we deduce $mn \mid (a - b)$, hence $[a]_{mn} = [b]_{mn}$, so F is one-to-one.

To see that F is onto, we take $([b]_m, [c]_n) \in U_m \times U_n$. We need to show there is an a so that $a \equiv b \pmod{m}$ and $a \equiv c \pmod{n}$, but this is exactly the Chinese Remainder Theorem. So such an a exists, and the map F is onto.

Thus F is one-to-one and onto, or a bijection, which means the sets U_{mn} and $U_m \times U_n$ have the same cardinality. The cardinality of U_{mn} is by definition $\phi(mn)$. The sets U_m and U_n have cardinalities $\phi(m)$ and $\phi(n)$, respectively, so the set of ordered pairs $U_m \times U_n$ has cardinality $\phi(m)\phi(n)$. \square

B.4. Chapter 5

Exercise (page 111). Suppose we choose primes p and q , so that $n = pq = 59753237$. With the knowledge of those primes, we compute $\phi(n) = (p-1)(q-1) = 59737740$ and choose the common encryption exponent $e = 2^{16} + 1 = 65537$ (the last known Fermat prime).

- (1) Find the primes p and q ; this is not necessary to break the code, but it reinforces that knowing $\phi(n)$ is equivalent to factoring n .

Solution to (1). We know that $n = pq$, $\phi(n) = (p-1)(q-1) = pq - (p+q) + 1 = n - (p+q) + 1$, so $n - \phi(n) + 1 = p+q = 15498$. On the other hand (assuming $p > q$), $(p-q) = \sqrt{(p-q)^2} = \sqrt{(p+q)^2 - 4n} = \sqrt{15498^2 - 4 \cdot 59753237} = \sqrt{1175056} = 1084$, which yields $p = 8291$ and $q = 7207$. \square

- (2) Find the decryption exponent.

Solution to (2). Euclid's algorithm for showing

$$\gcd(65537, \phi(n)) = 1$$

yields $65537(-9103267) + 59737740(9987) = 1$, so

$$d \equiv -9103267 \equiv 50634473 \pmod{\phi(n)}. \quad \square$$

- (3) Using the base 27 encoding scheme as above, decrypt the message consisting of two blocks of numerical ciphertext, i.e., given as $C = P^e \pmod{n}$

$$10881312 \quad 41465338.$$

Solution to (3). $10881312^d \equiv 6391358 \pmod{n}$ and $41465338^d \equiv 9436302 \pmod{n}$. Expanding in base 27, we

have

$$\begin{aligned}6391358 &= 12(27^4) + 0(27^3) + 19(27^2) \\ &\quad + 7(27^1) + 26(27^0) \mapsto \text{'MATH}\sqcup\text{'}, \\ 9436302 &= 17(27^4) + 20(27^3) + 11(27^2) \\ &\quad + 4(27^1) + 18(27^0) \mapsto \text{'RULES'}. \quad \square\end{aligned}$$

Exercise (page 112). Compute the probability that a plaintext message M is not prime to $n = pq$. If we wanted to ensure that our messages were always relatively prime to n , what could be done?

Solution. The probability that a plaintext message M not prime to $n = pq$ is $1/p + 1/q - 1/pq$, less than 1% of the time for p and q with at least 200 digits.

One could add filler to the message to make it relatively prime, and indeed padding is necessary for security reasons in real implementations of RSA. \square

Project (page 126). Suppose you want a 1024-bit RSA modulus, so you want two primes 512-bits long. About how many primes are there of that approximate size? What are the chances that randomly chosen odd integers of that length will be prime?

Solution. There are 2^{512} integers with 512 bits, 2^{511} of them odd, so those are the only candidates for primes, but how many primes are there really in that range? If their number were very small, we might be able to exhaustively try all the primes in trying to factor a 1024-bit modulus. There is an answer to this question given by a famous theorem in number theory called the *prime number theorem* which says there are approximately 2^{503} primes in that range. So I guess we are not going to brute-force our way to finding a factorization of n .

What are the odds that a randomly chosen 512-bit odd integer is prime? About $2^{503}/2^{511} = 1/256$, so picking primes in this range is quite easy. \square

Exercise (page 126). The observations above also provide an answer to another important security question. We know that people often use $e = 65537$ as an encryption exponent. What if they also chose the same value of n ? What would be the security implications?

Solution. Now that we have a bit of data, what is the probability of that happening? Of course to understand that, we have to know how people find large primes, the subject of primality testing. In general candidates for large primes are simply randomly chosen integers which are quickly tested for primality. The prime number theorem tells us that approximately .39% of all the odd 512-bit integers are primes, so on average, how many odd numbers in this range do we need to test to find one prime? Two primes? Since there are so many primes in this range, the chance that ones chosen randomly by one user would be the same as those chosen randomly would be astronomically small, except for one matter. How does one choose a random integer?

The problem is if random number generators fail to be random!! Past problems noted in [Sch12] are an interesting starting point. \square

B.5. Chapter 6

Exercise (page 132). Let

$$\sigma = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{bmatrix} \text{ and } \tau = \begin{bmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{bmatrix}.$$

Compute $\sigma, \sigma^2, \sigma^3, \tau, \tau^2, \tau^3, \sigma\tau, \sigma^2\tau$.

Solution.

$$\begin{aligned} \sigma &= \begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{bmatrix}, & \sigma^2 &= \begin{bmatrix} 1 & 2 & 3 \\ 3 & 1 & 2 \end{bmatrix}, & \sigma^3 &= \begin{bmatrix} 1 & 2 & 3 \\ 1 & 2 & 3 \end{bmatrix}, \\ \tau &= \begin{bmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{bmatrix}, & \tau^2 &= \begin{bmatrix} 1 & 2 & 3 \\ 1 & 2 & 3 \end{bmatrix}, & \tau^3 &= \tau = \begin{bmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{bmatrix}, \\ \sigma\tau &= \begin{bmatrix} 1 & 2 & 3 \\ 2 & 1 & 3 \end{bmatrix}, & \sigma^2\tau &= \begin{bmatrix} 1 & 2 & 3 \\ 3 & 2 & 1 \end{bmatrix}. \end{aligned} \quad \square$$

Exercise (page 132). Fill in the Cayley table for S_3 using the elements listed in the first row or column, and show that S_3 is non-abelian.


Solution.


\circ	e	σ	σ^2	τ	$\sigma\tau$	$\sigma^2\tau$
e	e	σ	σ^2	τ	$\sigma\tau$	$\sigma^2\tau$
σ	σ	σ^2	e	$\sigma\tau$	$\sigma^2\tau$	τ
σ^2	σ^2	e	σ	$\sigma^2\tau$	τ	$\sigma\tau$
τ	τ	$\sigma^2\tau$	$\sigma\tau$	e	σ^2	σ
$\sigma\tau$	$\sigma\tau$	τ	$\sigma^2\tau$	σ	e	σ^2
$\sigma^2\tau$	$\sigma^2\tau$	$\sigma\tau$	τ	σ^2	σ	e

We check that $\sigma\tau \neq \tau\sigma$. □

Exercise (page 133). Compute $R, R^2, R^3, F, F^2, F^3, RF, R^2F$.

Solution.





□

Exercise (page 133). Fill in the Cayley table for D_3 using the elements listed along the first row or column.

Solution.

\circ	e	R	R^2	F	RF	R^2F
e	e	R	R^2	F	RF	R^2F
R	R	R^2	e	RF	R^2F	F
R^2	R^2	e	R	R^2F	F	RF
F	F	R^2F	RF	e	R^2	R
RF	RF	F	R^2F	R	e	R^2
R^2F	R^2F	RF	F	R^2	R	e

□

Exercise (page 134). Notice that each symmetry can be thought of as a permutation of the three vertices. If we regard the numbers marking the vertices of the left-hand triangle as positions, then R can be described as the permutation $R = \begin{bmatrix} 1 & 2 & 3 \\ 2 & 3 & 1 \end{bmatrix}$, and $F = \begin{bmatrix} 1 & 2 & 3 \\ 1 & 3 & 2 \end{bmatrix}$.

Describe R^2 , F , RF , R^2F in terms of the elements σ and τ used to define S_3 . Can you determine if $D_3 \cong S_3$?

Solution. Yes, they are isomorphic via $R^a F^b \mapsto \sigma^a \tau^b$. \square

Exercise (page 140). We know that U_n is a finite abelian group. For $5 \leq n \leq 15$, use your knowledge of these groups to characterize them as in the fundamental theorem. For example, U_3 is a group of order 2, a prime, so U_3 is a cyclic group of order 2, that is $U_3 \cong \mathbb{Z}_2$. The group U_8 is an abelian group of order 4, so by the fundamental theorem, it is isomorphic to either $\mathbb{Z}_2 \times \mathbb{Z}_2$ or to \mathbb{Z}_4 . We easily check for all $a \in U_8$ that $a^2 = 1$, so U_8 is not cyclic, and so $U_8 \cong \mathbb{Z}_2 \times \mathbb{Z}_2$.

Solution.

$$U_5 = \langle 2 \rangle = \langle 3 \rangle \cong \mathbb{Z}_4,$$

$$U_6 = \langle 5 \rangle \cong \mathbb{Z}_2,$$

$$U_7 = \langle 3 \rangle = \langle 5 \rangle \cong \mathbb{Z}_6 \cong \mathbb{Z}_2 \times \mathbb{Z}_3,$$

$$U_8 = \langle 3 \rangle \times \langle 5 \rangle \cong \mathbb{Z}_2 \times \mathbb{Z}_2,$$

$$U_9 = \langle 2 \rangle = \langle 5 \rangle \cong \mathbb{Z}_6,$$

$$U_{10} = \langle 3 \rangle = \langle 7 \rangle \cong \mathbb{Z}_4,$$

$$U_{11} = \langle 2 \rangle = \langle 5 \rangle = \langle 6 \rangle = \langle 7 \rangle \cong \mathbb{Z}_{10},$$

$$U_{12} = \langle 5 \rangle \times \langle 7 \rangle \cong \mathbb{Z}_2 \times \mathbb{Z}_2,$$

$$U_{13} = \langle 2 \rangle \cong \mathbb{Z}_{12} \cong \mathbb{Z}_4 \times \mathbb{Z}_3,$$

$$U_{14} = \langle 3 \rangle \cong \mathbb{Z}_6,$$

$$U_{15} = \langle 2 \rangle \times \langle 11 \rangle \cong \mathbb{Z}_4 \times \mathbb{Z}_2. \quad \square$$

B.6. Chapter 7

Exercise (page 159). Consider the points of intersection of the affine curves $x = y^2$ and $y = -3$. As this is the intersection of a line and a conic, we expect at most two points, and indeed there is only one affine point $(9, -3)$. Find the points of intersection of the corresponding projective curves.

Solution. The projective curves are $xz = y^2$ and $y = -3z$. Equating the two gives $xz = 9z^2$ or $z(9z - x) = 0$. So either $z = 0$ or $x = 9z$

or both. First we conclude no solutions result from both conditions being true. $z = 0$ implies $x = 9z = 0$ and $y = -3z = 0$, and $[0, 0, 0]$ is not a point in projective space. So we have two cases $z \neq 0$ and $x = 9z$, or $z = 0$. So $x = 9z, y = -3z$ and $z \neq 0$ gives us the single projective point $[9z, -3z, z] = [9, -3, 1]$, which corresponds to our affine solution $(9, -3)$. That leaves the case where $z = 0$. In that case we have $y^2 = xz = 0 = -3z$, so we have $y = z = 0$ and x is arbitrary (but not zero). So we gain one more point $[x, 0, 0] = [1, 0, 0]$ of intersection which lies on the line at infinity. \square

Exercise (page 159), Find the points of intersection of the parallel lines $y = 3x$ and $y = 3x + 1$ in $\mathbb{P}^2(\mathbb{R})$.

Solution. Of course there are no points of intersection in the affine plane, so we look projectively. The corresponding projective lines are $y = 3x$ and $y = 3x + z$. Equating, we see $3x = 3x + z$, so $z = 0$ (which is good since it says the only possible solutions are on the line at infinity since we know there are no affine solutions). So we have $z = 0$ and $y = 3x$, which gives the single point $[x, 3x, 0] = [1, 3, 0]$ as the point of intersection of these projective lines. \square

Exercise (page 159). Consider the intersection of the cubic $y = x^3$ and the line $y = x + 6$. We would like to see three points of intersection, but where are they?

Solution. We see immediately that $x^3 = x + 6$ is equivalent to the equation $x^3 - x - 6 = (x - 2)(x^2 + 2x + 3) = 0$, and the quadratic has no real roots, though it has the complex roots $-1 \pm i\sqrt{2}$. At any rate the point $(2, 8)$ is a point on the affine curve, so $[2, 8, 1]$ should be a point on the projective curves $yz^2 = x^3$ and $y = x + 6z$. Looking at the line at infinity $z = 0$, we see $x = 0$ and hence $y = 0$, so there are no projective points on two curves which were not affine. Indeed we see that the solutions are $(2, 8)$ if we look in $\mathbb{A}^2(\mathbb{R})$, though we get three points $(2, 8), (1 \pm i\sqrt{2}, 7 \pm i\sqrt{2})$ if we look in $\mathbb{A}^2(\mathbb{C})$. \square

Exercise (page 185). Consider the following elliptic curves over \mathbb{F}_7 , and determine the set of points on the curve and its structure as an abelian group. Write out how to determine all the elements of the group in terms of the generators you choose.

$$(1) E : y^2 = x^3 + 3x + 6.$$

Solution to (1). $\Delta \equiv 2 \pmod{7}$; $E(\mathbb{F}_7) \cong \mathbb{Z}_4$, cyclic of order 4;

$$\begin{aligned} E(\mathbb{F}_7) &= \langle [6, 3, 1] \rangle = \langle P \rangle = \{P, 2P, 3P, 4P = \mathbf{0}\} \\ &= \{[6, 3, 1], [3, 0, 1], [6, 4, 1], [0, 1, 0]\}. \quad \square \end{aligned}$$

$$(2) E : y^2 = x^3 + 2.$$

Solution to (2). $\Delta \equiv 3 \pmod{7}$; $E(\mathbb{F}_7) \cong \mathbb{Z}_3 \times \mathbb{Z}_3$;

$E(\mathbb{F}_7)$

$$\begin{aligned} &= \langle [6, 1, 1] \rangle \times \langle [0, 4, 1] \rangle = \langle P \rangle \times \langle Q \rangle = \{P, 2P, 3P\} \times \{Q, 2Q, 3Q\} \\ &= \{[6, 1, 1], [6, 6, 1], [0, 1, 0]\} \times \{[0, 4, 1], [0, 3, 1], [0, 1, 0]\} \\ &= \{\mathbf{0}, P, 2P, Q, 2Q, P \oplus Q, P \oplus 2Q, Q \oplus 2P, 2P \oplus 2Q\} \\ &= \{[0, 1, 0], [6, 1, 1], [6, 6, 1][0, 4, 1], \\ &\quad [0, 3, 1], [3, 1, 1], [5, 1, 1], [5, 6, 1], [3, 6, 1]\} \quad \square \end{aligned}$$

$$(3) E : y^2 = x^3 + 4.$$

Solution to (3). $\Delta \equiv 5 \pmod{7}$; $E(\mathbb{F}_7) \cong \mathbb{Z}_3$;

$$\begin{aligned} E(\mathbb{F}_7) &= \langle [0, 2, 1] \rangle = \langle P \rangle = \{P, 2P, \mathbf{0}\} \\ &= \{[0, 2, 1], [0, 5, 1], [0, 1, 0]\}. \quad \square \end{aligned}$$

Bibliography

- [Bro10] Daniel R. L. Brown, *Standards for efficient cryptography 2 (sec 2); Recommended elliptic curve domain parameters*, <http://www.secg.org/sec2-v2.pdf>, 2010, URL Date: 2010-01-27; Accessed: 2016-06-27.
- [CP05] Richard Crandall and Carl Pomerance, *Prime numbers. A computational perspective*, second ed., Springer, New York, 2005, MR2156291 (2006a:11005)
- [Deu41] Max Deuring, *Die Typen der Multiplikatorenringe elliptischer Funktionenkörper*, Abh. Math. Sem. Univ. Hamburg **14** (1941), no. 1, 197–272. MR3069722
- [DH76] Whitfield Diffie and Martin E. Hellman, *New directions in cryptography*, IEEE Trans. Information Theory **IT-22** (1976), no. 6, 644–654. MR0437208 (55 #10141)
- [Ful69] William Fulton, *Algebraic curves: An introduction to algebraic geometry*, first ed., W. A. Benjamin, Reading, MA, 1969.
- [Gir15] Damien Giry, *Cryptographic key length recommendation*, <http://www.keylength.com/>, 2015, URL Date: 2015-02-26; Accessed: 2015-07-06.
- [IN15] IAD-NSA, *Commercial national security algorithm suite*, <https://www.iad.gov/iad/programs/iad-initiatives/cnsa-suite.cfm>, 2015, URL Date: 2015-08-19; Accessed: 2016-06-27.
- [JJ98] Gareth A. Jones and J. Mary Jones, *Elementary number theory*, Springer Undergraduate Mathematics Series, Springer-Verlag London Ltd., London, 1998. MR1610533 (2000b:11002)

- [JNT07] Antoine Joux, David Naccache, and Emmanuel Thomé, *When e -th roots become easier than factoring*, Advances in Cryptology—ASIACRYPT 2007, Lecture Notes in Comput. Sci., vol. 4833, Springer, Berlin, 2007, pp. 13–28. MR2565721 (2011b:11169)
- [KKM11] Ann Hibner Koblitz, Neal Koblitz, and Alfred Menezes, *Elliptic curve cryptography: The serpentine course of a paradigm shift*, J. Number Theory **131** (2011), no. 5, 781–814. MR2772472 (2012b:14052)
- [KM15] Neal Koblitz and Alfred Menezes, *A riddle wrapped in an enigma*, Cryptology ePrint Archive, Report 2015/1018, 2015, <http://eprint.iacr.org/2015/1018.pdf>.
- [Kob84] Neal Koblitz, *Introduction to elliptic curves and modular forms*, Graduate Texts in Mathematics, vol. 97, Springer-Verlag, New York, 1984. MR766911 (86c:11040)
- [Kob87a] ———, *A course in number theory and cryptography*, Graduate Texts in Mathematics, vol. 114, Springer-Verlag, New York, 1987. MR910297 (88i:94001)
- [Kob87b] ———, *Elliptic curve cryptosystems*, Math. Comp. **48** (1987), no. 177, 203–209. MR866109
- [Lab12] RSA Laboratories, *PKCS#1 v2.2: RSA cryptography standard*, <https://www.emc.com/collateral/white-papers/h11300-pkcs-1v2-2-rsa-cryptography-standard-wp.pdf>, 2012, URL Date: 2012-10-27; Accessed: 2015-07-01.
- [Len87] H. W. Lenstra, Jr., *Factoring integers with elliptic curves*, Ann. of Math. (2) **126** (1987), no. 3, 649–673. MR916721 (89g:11125)
- [Mer82] Ralph C. Merkle, “Secure communications over insecure channels”, *Secure Communications and Asymmetric Cryptosystems*, AAAS Sel. Sympos. Ser., vol. 69, Westview, Boulder, CO, 1982, pp. 181–196. MR668724
- [NC09] NSA-CSS, *The case for elliptic curve cryptography - NSA/CSS*, https://www.nsa.gov/business/programs/elliptic_curve.shtml, 2009, URL Date: 2009-01-12; Accessed: 2015-06-25; Archived at http://web.archive.org/web/20150627183730/https://www.nsa.gov/business/programs/elliptic_curve.shtml.
- [NIS12a] NIST, *Recommendation for applications using approved hash algorithms*, <http://csrc.nist.gov/publications/nistpubs/800-107-rev1/sp800-107-rev1.pdf>, 2012, URL Date: 2012, August; Accessed: 2015-07-02.

- [NIS12b] ———, *Recommendation for key management—part 1: General (revision 3)*, http://csrc.nist.gov/publications/nistpubs/800-57/sp800-57_part1_rev3_general.pdf, 2012, URL Date: 2012-07; Accessed: 2015-07-06.
- [NIS12c] ———, *Secure hash standard SHS*, <http://csrc.nist.gov/publications/fips/fips180-4/fips-180-4.pdf>, 2012, URL Date: 2012, March; Accessed: 2015-07-02.
- [Ros05] Kenneth H. Rosen, *Elementary number theory and its applications*, fifth ed., Addison-Wesley, Reading, MA, 2005.
- [RSA78] R. L. Rivest, A. Shamir, and L. Adleman, *A method for obtaining digital signatures and public-key cryptosystems*, *Comm. ACM* **21** (1978), no. 2, 120–126. MR700103 (83m:94003)
- [S⁺15] W. A. Stein et al., *Sage Mathematics Software (Version 6.5)*, The Sage Development Team, 2015, <http://www.sagemath.org>.
- [Sch12] Seth D. Schoen, *Understanding common factor attacks: An RSA-cracking puzzle*, <http://www.loyalty.org/~schoen/rsa/>, 2012, URL Date: 2012; Accessed: 2015-07-06.
- [Sch15a] Bruce Schneier, *NSA plans for a post-quantum world*, <https://www.lawfareblog.com/nsa-plans-post-quantum-world>, 2015, URL Date: 2015-08-21; Accessed: 2016-06-28.
- [Sch15b] ———, *NSA plans for a post-quantum world*, https://www.schneier.com/blog/archives/2015/08/nsa_plans_for_a.html/, 2015, URL Date: 2015-08-21; Accessed: 2016-06-28.
- [Sho94] P. W. Shor, *Algorithms for quantum computation: Discrete logarithms and factoring*, Proceedings of the 35th Annual Symposium on Foundations of Computer Science (Washington, DC, USA), SFCS '94, IEEE Computer Society, 1994, pp. 124–134.
- [ST92] Joseph H. Silverman and John Tate, *Rational points on elliptic curves*, Undergraduate Texts in Mathematics, Springer-Verlag, New York, 1992. MR1171452 (93g:11003)
- [Tun83] J. B. Tunnell, *A classical Diophantine problem and modular forms of weight 3/2*, *Invent. Math.* **72** (1983), no. 2, 323–334. MR700775 (85d:11046)
- [Wik16] Wikipedia, *Quantum computing—Wikipedia, the free encyclopedia*, 2016 [online; accessed 28-November-2016].

Index

- abelian, 76
- affine space, 148
- associative, 74

- basis, 32
- binary operation, 73
- birthday paradox, 119

- Cayley table, 128
- Chinese Remainder Theorem, 70
- ciphertext, 63
- closed, 31
- collision, 118
- commutative, 74
- composite, 55
- congruence modulo n , 80
- conic, 40
- coprime, 51
- cubic, 40
- cyclic, 134

- degree, 40
- dihedral group, 132
- dimension, 32
- direct product, 136
- discriminant, 179
- divides, 47
- division algorithm, 47

- elliptic curve, 3
- equivalence class, 79

- equivalence relation, 79
- Euler's totient function, 86

- forward secrecy, 143
- Fundamental Theorem of Finite Abelian Groups, 139
- Fundamental Theorem of Arithmetic, 23, 41, 55

- genus, 28
- greatest common divisor, gcd, 51
- group, 75

- hash function, 114
- homogenization, 154
- homomorphism, 130

- identity, 74
- inverse, 75
- irreducible, 41
- isomorphic, 128
- isomorphism, 130

- keyspace, 64

- Lagrange interpolation, 168
- line, 40
- linear combination, 29, 32

- modulus, 108
- multiplicity, 38, 223

natural numbers, 46

one-way function, 118

order of a group, 88

order of an element, 88

partition, 80, 81

plaintext, 63

preimage resistance, 118

prime, 55

Prime Number Theorem, 237

primitive Pythagorean triple, 16

primitive root, 141

projective plane, 151

projective n -space, 152

projective line, 149

pseudoprime, 91

quadratic nonresidue, 186

quadratic residue, 186

rational curve, 40

reflexive, 78

relatively prime, 16, 51

ring, 76

safe prime, 144

second preimage resistance, 118

shift cipher, 64

symmetric, 78

symmetric group, 131

textbook RSA, 108, 112

torsion point, 178

transitive, 78

vector spaces, 30

well-defined, 83

well-ordered, 48

zero set, 39