Number Theory I

Number theory is the study of the integers. Number theory is right at the core of mathematics; even Ug the Caveman surely had some grasp of the integers— at least the positive ones. In fact, the integers are so elementary that one might ask, "What's to study?" There's 0, there's 1, 2, 3 and so on, and there's the negatives. Which one don't you understand? Doesn't math become easy when we don't have to worry about nasty numbers like $\sqrt{7}$, $1/\pi$, and i? We can even forget about fractions!

All the variables in these notes represent integers.

1   Divisibility

The true nature of number theory emerges from the first definition. We say that a *divides* b if there is an integer k such that $ak = b$. This is denoted $a \mid b$. For example:

$$7 \mid 63 \quad \text{because} \quad 7 \cdot 9 = 63$$

A consequence of this definition is that every number divides zero since $a \cdot 0 = 0$ for every integer a. If a divides b, then b is a *multiple* of a. For example, 63 is a multiple of 7.

This seems simple enough, but let's play with this definition. The Pythagoreans, an ancient sect of mathemtical mystics, said that a number is *perfect* if it equals the sum of its positive integeral divisors, excluding itself. For example, $6 = 1 + 2 + 3$ and $28 = 1 + 2 + 4 + 7 + 14$ are perfect numbers. On the other hand, 10 is not perfect because $1 + 2 + 5 = 8$, and 12 is not perfect because $1 + 2 + 3 + 4 + 6 = 16$. Euclid characterized all the *even* perfect numbers around 300 BC. But is there an *odd* perfect number? More than two thousand years later, we still don't know! All numbers up to about 10^{300} have been ruled out, but no one has proved that there isn't an odd perfect number waiting just over the horizon.

So a half-page into number theory, we've strayed past the outer limits of human knowledge. This is pretty typical; number theory is full of questions that are easy to pose, but incredibly difficult to answer. Interestingly, computer scientists have found ways to turn these difficulties to their advantage. Every time you buy a book from Amazon, check your grades on WebSIS, or use a PayPal account, you are relying on number theoretic algorithms.

DON'T PANIC— we're going to stick to some relatively benign parts of number theory. We won't put any of these super-hard unsolved problems on exams!

1.1 Facts About Divisibility

The lemma below states some basic facts about divisibility that are *not* difficult to prove:

Lemma 1. *The following statements about divisibility hold.*

1. *If $a \mid b$, then $a \mid bc$ for all c.*

2. *If $a \mid b$ and $b \mid c$, then $a \mid c$.*

3. *If $a \mid b$ and $a \mid c$, then $a \mid sb + tc$ for all s and t.*

4. *For all $c \neq 0$, $a \mid b$ if and only if $ca \mid cb$.*

Proof. We'll only prove parts (2) and (4); the other proofs are similar.

Proof of (2): Since $a \mid b$, there exists an integer k_1 such that $ak_1 = b$. Since $b \mid c$, there exists an integer k_2 such that $bk_2 = c$. Substituting ak_1 for b in the second equation gives $ak_1k_2 = c$, which implies that $a \mid c$.

Proof of (4): We must show that $a \mid b$ implies $ca \mid cb$ and vice-versa.

- First, suppose $a \mid b$. This means $ak = b$ for some k. Multiplying both sides by c gives $cak = cb$ for some k. This implies $ca \mid cb$.

- Now, suppose $ca \mid cb$. Then $cak = cb$ for some k. We can divide both sides by c since c is nonzero, so $ak = b$ for some k. This means $a \mid b$.

$\square$

A number $p > 1$ with no positive divisors other than 1 and itself is called a ***prime***. Every other number greater than 1 is called ***composite***. For example, 2, 3, 5, 7, 11, and 13 are all prime, but 4, 6, 8, and 9 are composite. The number 1 is considered neither prime nor composite. This is just a matter of definition, but reflects the fact that 1 does not behave like a prime in many contexts, such as the Fundamental Theorem of Arithmetic, which we'll come to shortly.

1.2 When Divisibility Goes Bad

As you learned in elementary school, if one number does *not* evenly divide another, then there is a "remainder" left over. More precisely, if you divide n by d, then you get a quotient q and a remainder r. This basic fact is the subject of a useful theorem:

Theorem 2 (Division Theorem). *Let n and d be integers such that $d > 0$. Then there exists a unique pair of integers q and r such that $n = qd + r$ and $0 \leq r < d$.*

Famous Problems in Number Theory

Fermat's Last Theorem Do there exist positive integers x, y, and z such that

$$x^n + y^n = z^n$$

for some integer $n > 2$? In a book he was reading around 1630, Fermat claimed to have a proof, but not enough space in the margin to write it down. Wiles finally solved the problem in 1994, after seven years of working in secrecy and isolation in his attic.

Goldbach Conjecture Is every even integer greater than or equal to 4 the sum of two primes? For example, $4 = 2 + 2$, $6 = 3 + 3$, $8 = 3 + 5$, etc. The conjecture holds for all numbers up to 10^{16}. In 1939 Schnirelman proved that every even number can be written as the sum of not more than 300,000 primes, which was a start. Today, we know that every even number is the sum of at most 6 primes.

Twin Prime Conjecture Are there infinitely many primes p such that $p + 2$ is also a prime? In 1966 Chen showed that there are infinitely many primes p such that $p + 2$ is the product of at most two primes. So the conjecture is known to be *almost* true!

Primality Testing Is there an efficient way to determine whether n is prime? An amazingly simple, yet efficient method was finally discovered in 2002 by Agrawal, Kayal, and Saxena. Their paper began with a quote from Gauss emphasizing the importance and antiquity of the problem even in his time— two centuries ago.

Factoring Given the product of two large primes $n = pq$, is there an efficient way to recover the primes p and q? The best known algorithm is the "number field seive", which runs in time proportional to:

$$e^{1.9(\ln n)^{1/3}(\ln \ln n)^{2/3}}$$

This is infeasible when n has a couple hundred digits or more.

As an example, suppose that $a = 10$ and $b = 2716$. Then the quotient is $q = 271$ and the remainder is $r = 6$, since $2716 = 271 \cdot 10 + 6$.

The remainder r in the Division Theorem is denoted n **rem** d. In other words, n rem d is the remainder when n is divided by d. For example, 32 rem 5 is the remainder when 32 is divided by 5, which is 2. Similarly, -11 rem $7 = 3$, since $-11 = (-2) \cdot 7 + 3$. There is a rem operator built into many programming languages. For example, the expression "32 % 5" evaluates to 2 in Java, C, and C++. However, all these languages treat negative numbers strangely.

There are a couple naming problems related to the Division Theorem. First, the theorem is often called the "Division Algorithm", even though it is not an algorithm in the modern sense. Second, *some* people use the notation "mod" (which is short for "modulo") instead of "rem". This is unfortunate, because "mod" has been used by mathematicians for centuries in a confusingly similar context, which we'll come to shortly. So we'll stick to rem here.

We're not going to prove the Division Theorem, but there is an important feature that you should notice. The theorem asserts that the quotient q and remainder r *exist* and also that these values are *unique*. Thus, the Division Theorem is one example of an "existence and uniqueness" theorem; there are many others. Not surprisingly, the proof of such a theorem always has two parts:

- A proof that something exists, such as the quotient q and remainder r.

- A proof that nothing else fits the bill; that is, there is no other quotient q' and remainder r'.

We'll prove a famous "existence and uniqueness" theorem in this way shortly.

2   Die Hard

> **Simon:** On the fountain, there should be 2 jugs, do you see them? A 5-gallon and a 3-gallon. Fill one of the jugs with exactly 4 gallons of water and place it on the scale and the timer will stop. You must be precise; one ounce more or less will result in detonation. If you're still alive in 5 minutes, we'll speak.
>
> **Bruce:** Wait, wait a second. I don't get it. Do you get it?
>
> **Samuel:** No.
>
> **Bruce:** Get the jugs. Obviously, we can't fill the 3-gallon jug with 4 gallons of water.
>
> **Samuel:** Obviously.
>
> **Bruce:** All right. I know, here we go. We fill the 3-gallon jug exactly to the top, right?
>
> **Samuel:** Uh-huh.
>
> **Bruce:** Okay, now we pour this 3 gallons into the 5-gallon jug, giving us exactly 3 gallons in the 5-gallon jug, right?
>
> **Samuel:** Right, then what?
>
> **Bruce:** All right. We take the 3-gallon jug and fill it a third of the way...
>
> **Samuel:** No! He said, "Be precise." Exactly 4 gallons.
>
> **Bruce:** Shit. Every cop within 50 miles is running his ass off and I'm out here playing kids games in the park.
>
> **Samuel:** Hey, you want to focus on the problem at hand?

This is from the movie *Die Hard 3: With a Vengeance*. Samuel L. Jackson and Bruce Willis have to disarm a bomb planted by the diabolical Simon Gruber. Fortunately, they find a solution in the nick of time. (No doubt reading the script helped.) On the surface, *Die Hard 3* is just a B-grade action movie; however, I think the inner message of the film is that everyone should learn at least a little number theory.

3   Die Once and For All

Unfortunately, Hollywood never lets go of a gimmick. They're planning to keep the *Die Hard* series going with:

Die Hard 4: Die Hardest Bruce goes on vacation and— shockingly— happens into a terrorist plot. To save the day, he must make 3 gallons using 21 and 26 gallon jugs.

Die Hard 5: Die of Old Age Bruce must save his assisted living facility from a criminal
   mastermind by forming 2 gallons with 899 and 1147 gallon jugs.

Die Hard 6: Die Once and For All Bruce has to make 4 gallons using 3 and 6-gallon jugs.

It would be nice if we could solve all these silly water jug questions at once. In particular,
how can one form g gallons using jugs with capacities a and b?

   That's where number theory comes in handy.

3.1 Finding an Invariant Property

Suppose that we have water jugs with capacities a and b. Let's carry out a few arbitrary
operations and see what happens. The state of the system at each step is described below
with a pair of numbers (x, y), where x is the amount of water in the jug with capacity a
and y is the amount in the jug with capacity b.

$$
\begin{aligned}
(0,0) &\rightarrow (a, 0) && \text{fill first jug} \\
&\rightarrow (0, a) && \text{pour first into second} \\
&\rightarrow (a, a) && \text{fill first jug} \\
&\rightarrow (2a - b, b) && \text{pour first into second} \\
&\rightarrow (2a - b, 0) && \text{empty second jug} \\
&\rightarrow (0, 2a - b) && \text{pour first into second} \\
&\rightarrow (a, 2a - b) && \text{fill first} \\
&\rightarrow (3a - 2b, b) && \text{pour first into second}
\end{aligned}
$$

Of course, we're making some assumptions about the relative capacities of the two jugs
here. But another point leaps out: at every step, the amount of water in each jug is of the
form

$$s \cdot a + t \cdot b$$

for some integers s and t. An expression of this form is called a ***linear combination*** of a
and b. This sounds like an assertion that we might be able to prove by induction!

Lemma 3. *Suppose that we have water jugs with capacities a and b. Then the amount of water in
each jug is always a linear combination of a and b.*

Proof. We use induction. Let $P(n)$ be the proposition that after n steps, the amount of
water in each jug is a linear combination of a and b.

Base case. $P(0)$ is true, because both jugs are initially empty, and $0 \cdot a + 0 \cdot b = 0$.

Inductive step. Now we must show that $P(n)$ implies $P(n + 1)$ for $n \geq 0$. So assume that
after n steps the amount of water in each jug is a linear combination of a and b. There are
two cases:

- If we fill a jug from the fountain or empty a jug into the fountain, then that jug is empty or full. The amount in the other jug remains a linear combination of a and b. So $P(n + 1)$ holds.

- Otherwise, we pour water from one jug to another until one is empty or the other is full. By our assumption, the amount in each jug is a linear combination of a and b before we begin pouring:

$$j_1 = s_1 \cdot a + t_1 \cdot b$$
$$j_2 = s_2 \cdot a + t_2 \cdot b$$

  After pouring, one jug is either empty (contains 0 gallons) or full (contains a or b gallons). Thus, the other jug contains either $j_1 + j_2$ gallons, $j_1 + j_2 - a$, or $j_1 + j_2 - b$ gallons, all of which are linear combinations of a and b.

The claim follows by the principle of induction. $\qquad\square$

This theorem has an important corollary.

Corollary 4. *Bruce dies.*

Proof. In Die Hard 6, Bruce has water jugs with capacities 3 and 6 and must form 4 gallons of water. However, the amount in each jug is always of the form $3s + 6t$ by Lemma 3. This is always a multiple of 3 by part (3) of Lemma 1, so he can not measure out 4 gallons. $\quad\square$

Lemma 3 isn't very satisfying. We've just managed to recast a pretty understandable question about water jugs into a complicated question about linear combinations. This might not seem like progress. Fortunately, linear combinations are closely related to something more familiar and that will help us solve the water jug problem.

4   The Greatest Common Divisor

The *greatest common divisor* of a and b is exactly what you'd guess: the largest number that is a divisor of both a and b. It is denoted $\gcd(a, b)$. For example, $\gcd(18, 24) = 6$.

Probably some junior high math teacher made you compute greatest common divisors for no apparent reason until you were blue in the face. But, amazingly, the greatest common divisor actually turns out to be quite useful for reasoning about the integers. Specifically, the quantity $\gcd(a, b)$ is a valuable piece of information about the relationship between the numbers a and b. So we'll make arguments about greatest common divisors all the time.

4.1   Linear Combinations and the GCD

The theorem below relates the greatest common divisor to linear combinations. This theorem is *very* useful; take the time to understand it and then remember it!

Theorem 5. *The greatest common divisor of a and b is equal to the smallest positive linear combination of a and b.*

For example, the greatest common divisor of 52 and 44 is 4. And, sure enough, 4 is a linear combination of 52 and 44:

$$6 \cdot 52 + (-7) \cdot 44 = 4$$

Furthermore, no linear combination of 52 and 44 is equal to a smaller positive integer.

Proof. Let m be the smallest positive linear combination of a and b. We'll prove that $m = \gcd(a, b)$ by showing both $\gcd(a, b) \leq m$ and $m \leq \gcd(a, b)$.

First, we show that $\gcd(a, b) \leq m$. By the definition of common divisor, $\gcd(a, b) \mid a$ and $\gcd(a, b) \mid b$. Therefore, for every pair of integers s and t:

$$\gcd(a, b) \mid sa + tb$$

Thus, in particular, $\gcd(a, b)$ divides m, and so $\gcd(a, b) \leq m$.

Now, we show that $m \leq \gcd(a, b)$. We do this by showing that $m \mid a$. A symmetric argument shows that $m \mid b$, which means that m is a common divisor of a and b. Thus, m must be less than or equal to the *greatest* common divisor of a and b.

All that remains is to show that $m \mid a$. By the Division Algorithm, there exists a quotient q and remainder r such that:

$$a = q \cdot m + r \qquad\qquad \text{(where } 0 \leq r < m)$$

Recall that $m = sa + tb$ for some integers s and t. Substituting in for m and rearranging terms gives:

$$a = q \cdot (sa + tb) + r$$
$$r = (1 - qs)a + (-qt)b$$

We've just expressed r as a linear combination of a and b. However, m is the *smallest* positive linear combination and $0 \leq r < m$. The only possibility is that the remainder r is not positive; that is, $r = 0$. This implies $m \mid a$. □

The proof notes that every linear combination of a and b is a multiple of $\gcd(a, b)$. Conversely, since $\gcd(a, b)$ is a linear combination of a and b, every multiple of $\gcd(a, b)$ is as well. This establishes a corollary:

Corollary 6. *Every linear combination of a and b is a multiple of $\gcd(a, b)$ and vice versa.*

Now we can restate the water jugs lemma in terms of the greatest common divisor:

Corollary 7. *Suppose that we have water jugs with capacities a and b. Then the amount of water in each jug is always a multiple of $\gcd(a, b)$.*

For example, there is no way to form 4 gallons using 3 and 6 gallon jugs, because 4 is not a multiple of $\gcd(3, 6) = 3$.

4.2   Properties of the Greatest Common Divisor

We claimed that greatest common divisors are powerful tools for reasoning about the integers. So we'll often make use of some basic gcd facts:

Lemma 8. *The following statements about the greatest common divisor hold:*

1. *Every common divisor of a and b divides $\gcd(a, b)$.*

2. *$\gcd(ka, kb) = k \cdot \gcd(a, b)$ for all $k > 0$.*

3. *If $\gcd(a, b) = 1$ and $\gcd(a, c) = 1$, then $\gcd(a, bc) = 1$.*

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

5. *$\gcd(a, b) = \gcd(b, a \text{ rem } b)$.*

Here's the trick to proving these statements: translate the gcd world to the linear combination world using Theorem 5, argue about linear combinations, and then translate back using Theorem 5 again.

Proof. We prove only parts (3) and (4).

Proof of (3): The assumptions together with Theorem 5 imply that there exist integers s, t, u, and v such that:

$$sa + tb = 1$$
$$ua + vc = 1$$

Multiplying these two equations gives:

$$(sa + tb)(ua + vc) = 1$$

The left side can be rewritten as $a \cdot (asu + btu + csv) + b \cdot c(tv)$. This is a linear combination of a and bc that is equal to 1, so $\gcd(a, bc) = 1$ by Theorem 5.

Proof of (4): Theorem 5 says that $\gcd(ac, bc)$ is equal to a linear combination of ac and bc. Now $a \mid ac$ trivially and $a \mid bc$ by assumption. Therefore, a divides *every* linear combination of ac and bc. In particular, a divides $\gcd(ac, bc) = c \cdot \gcd(a, b) = c$. The first equality uses part (2) of this lemma, and the second uses the assumption that $\gcd(a, b) = 1$. $\qquad\square$

Part (5) of the lemma is useful for quickly computing the greatest common divisor of two numbers. For example, we could compute the greatest common divisor of 1147 and 899 by repeatedly applying part(5):

$$
\begin{aligned}
\gcd(1247, 899) &= \gcd(899, \underbrace{1247 \text{ rem } 899}_{=248}) \\
&= \gcd(248, \underbrace{899 \text{ rem } 248}_{=155}) \\
&= \gcd(155, \underbrace{248 \text{ rem } 155}_{=93}) \\
&= \gcd(93, \underbrace{155 \text{ rem } 93}_{=62}) \\
&= \gcd(62, \underbrace{93 \text{ rem } 62}_{=31}) \\
&= \gcd(31, \underbrace{62 \text{ rem } 31}_{=0}) \\
&= \gcd(31, 0) \\
&= 31
\end{aligned}
$$

This is called **Euclid's algorithm**. The last equation might look wrong, but 31 is a divisor of both 31 and 0 since every integer divides 0.

This calculation, together with Corollary 7, implies that there is no way to measure out 2 gallons of water using jugs with capacities 1247 and 899; we can only obtain multiples of 31 gallons. This is good news– Bruce won't even survive Die Hard 5!

Let's see if Bruce can possibly make 3 gallons using 21 and 26-gallon jugs. First, we compute the greatest common divisorof 21 and 26 using Euclid's algorithm:

$$\gcd(26, 21) = \gcd(21, 5) = \gcd(5, 1) = 1$$

Now 3 is a multiple of 1, so we can't *rule out* the possibility that Bruce can form 3 gallons. On the other hand, we don't know he *can* do it either.

4.3 One Solution for All Water Jug Problems

Can Bruce form 3 gallons using 21 and 26-gallon jugs? This question is not so easy to answer without some number theory.

Corollary 6 says that 3 can be written as a linear combination of 21 and 26, since 3 is a multiple of $\gcd(21, 26) = 1$. In other words, there exist integers s and t such that:

$$3 = s \cdot 21 + t \cdot 26$$

We don't know what the coefficients s and t are, but we do know that they exist.

Now the coefficient s could be either positive or negative. However, we can readily transform this linear combination into an equivalent linear combination

$$3 = s' \cdot 21 + t' \cdot 26$$

where the coefficient s' is positive. The trick is to notice that if we increase s by 26 in the original equation and decrease t by 21, then the value of the expression $s \cdot 21 + t \cdot 26$ is unchanged overall. Thus, by repeatedly increasing the value of s (by 26 at a time) and decreasing the value of t (by 21 at a time), we get a linear combination $s' \cdot 21 + t' \cdot 26 = 3$ where the coefficient s' is positive. Notice that t' must be negative; otherwise, this expression would be much greater than 3.

Now here's how to form 3 gallons using jugs with capacities 21 and 26:

- Repeat s' times:

  - Fill the 21-gallon jug.

  - Pour all the water in the 21-gallon jug into the 26-gallon jug. Whenever the 26-gallon jug becomes full, empty it out.

At the end of this process, there must be exactly 3 gallons in the 26-gallon jug! Here's why: we've taken $s' \cdot 21$ gallons of water from the fountain, we've poured out some multiple of 26 gallons, and in the end the 26-gallon jug holds somewhere between 0 and 26 gallons. Furthermore, we know:

$$s' \cdot 21 + t' \cdot 26 = 3$$

Thus, we must have emptied the 26-gallon jug exactly $-t'$ times; if we had emptied it fewer times, then there woule be more than 26 gallons left. And we did not withdraw enough water from the fountain to empty the 26-gallon jug more than $-t'$ times. Thus, by the equation above, there must be exactly 3 gallons left.

Remarkably, we don't even need to know the coefficients s' and t' in order to use this strategy! Instead of repeating the outer loop s' times, we could just repeat *until we obtain 3 gallons*, since that must happen eventually. Of course, we have to keep track of the amounts in the two jugs so we know when we're done. Here's the solution that approach

gives:

$(0,0) \xrightarrow{\text{fill 21}} (21,0) \xrightarrow{\text{pour 21 into 26}} (0,21)$

$\xrightarrow{\text{fill 21}} (21,21) \xrightarrow{\text{pour 21 into 26}} (16,26) \xrightarrow{\text{empty 26}} (16,0) \xrightarrow{\text{pour 21 into 26}} (0,16)$

$\xrightarrow{\text{fill 21}} (21,16) \xrightarrow{\text{pour 21 into 26}} (11,26) \xrightarrow{\text{empty 26}} (11,0) \xrightarrow{\text{pour 21 into 26}} (0,11)$

$\xrightarrow{\text{fill 21}} (21,11) \xrightarrow{\text{pour 21 into 26}} (6,26) \xrightarrow{\text{empty 26}} (6,0) \xrightarrow{\text{pour 21 into 26}} (0,6)$

$\xrightarrow{\text{fill 21}} (21,6) \xrightarrow{\text{pour 21 into 26}} (1,26) \xrightarrow{\text{empty 26}} (1,0) \xrightarrow{\text{pour 21 into 26}} (0,1)$

$\xrightarrow{\text{fill 21}} (21,1) \xrightarrow{\text{pour 21 into 26}} (0,22)$

$\xrightarrow{\text{fill 21}} (21,22) \xrightarrow{\text{pour 21 into 26}} (17,26) \xrightarrow{\text{empty 26}} (17,0) \xrightarrow{\text{pour 21 into 26}} (0,17)$

$\xrightarrow{\text{fill 21}} (21,17) \xrightarrow{\text{pour 21 into 26}} (12,26) \xrightarrow{\text{empty 26}} (12,0) \xrightarrow{\text{pour 21 into 26}} (0,12)$

$\xrightarrow{\text{fill 21}} (21,12) \xrightarrow{\text{pour 21 into 26}} (7,26) \xrightarrow{\text{empty 26}} (7,0) \xrightarrow{\text{pour 21 into 26}} (0,7)$

$\xrightarrow{\text{fill 21}} (21,7) \xrightarrow{\text{pour 21 into 26}} (2,26) \xrightarrow{\text{empty 26}} (2,0) \xrightarrow{\text{pour 21 into 26}} (0,2)$

$\xrightarrow{\text{fill 21}} (21,2) \xrightarrow{\text{pour 21 into 26}} (0,23)$

$\xrightarrow{\text{fill 21}} (21,23) \xrightarrow{\text{pour 21 into 26}} (18,26) \xrightarrow{\text{empty 26}} (18,0) \xrightarrow{\text{pour 21 into 26}} (0,18)$

$\xrightarrow{\text{fill 21}} (21,18) \xrightarrow{\text{pour 21 into 26}} (13,26) \xrightarrow{\text{empty 26}} (13,0) \xrightarrow{\text{pour 21 into 26}} (0,13)$

$\xrightarrow{\text{fill 21}} (21,13) \xrightarrow{\text{pour 21 into 26}} (8,26) \xrightarrow{\text{empty 26}} (8,0) \xrightarrow{\text{pour 21 into 26}} (0,8)$

$\xrightarrow{\text{fill 21}} (21,8) \xrightarrow{\text{pour 21 into 26}} (3,26) \xrightarrow{\text{empty 26}} (3,0) \xrightarrow{\text{pour 21 into 26}} (0,3)$

The same approach works regardless of the jug capacities and even regardless the amount we're trying to produce! Simply proceed as follows:

- Repeat until the desired amount of water is obtained:

  - Fill the smaller jug.
  - Pour all the water in the smaller jug into the larger jug. Whenever the larger jug becomes full, empty it out.

By the same reasoning as before, this method eventually generates every multiple of the greatest common divisor of the jug capacities— all the quantities we can possibly produce. No ingenuity is needed at all!

5   The Fundamental Theorem of Arithemtic

We now have almost enough tools to prove something that you probably already know.

Theorem (Fundamental Theorem of Arithmetic). *Every positive integer n can be written in a unique way as a product of primes:*

$$n \;=\; p_1 \cdot p_2 \cdots p_j \qquad\qquad (p_1 \le p_2 \le \ldots \le p_j)$$

Notice that the theorem would be false if 1 were considered a prime; for example, 15 could be written as $3 \cdot 5$ or $1 \cdot 3 \cdot 5$ or $1^2 \cdot 3 \cdot 5$. Also, we're relying on a standard convention: the product of an empty set of numbers is defined to be 1, much as the sum of an empty set of numbers is defined to be 0. Without this convention, the theorem would be false for $n = 1$.

There is a certain wonder in the Fundamental Theorem, even if you've known it since the crib. Primes show up erratically in the sequence of integers. In fact, their distribution seems almost random:

$$2, 3, 5, 7, 11, 13, 17, 19, 23, 29, 31, 37, 41, 43, \ldots$$

Basic questions about this sequence have stumped humanity for centuries. And yet we know that every natural number can be built up from primes in *exactly one way*. These quirky numbers are the building blocks for the integers. The Fundamental Theorem is not hard to prove, but we'll need a couple preliminary facts.

Lemma 9. *If p is a prime and $p \mid ab$, then $p \mid a$ or $p \mid b$.*

Proof. The greatest common divisor of a and p must be either 1 or p, since these are the only divisors of p. If $\gcd(a, p) = p$, then the claim holds, because a is a multiple of p. Otherwise, $\gcd(a, p) = 1$ and so $p \mid b$ by part (4) of Lemma 8. $\square$

A routine induction argument extends this statement to the fact we assumed last time:

Lemma 10. *Let p be a prime. If $p \mid a_1 a_2 \ldots a_n$, then p divides some a_i.*

Now we're ready to prove the Fundamental Theorem of Arithemtic.

Theorem 11 (Fundamental Theorem of Arithmetic). *Every positive integer n can be written in a unique way as a product of primes:*

$$n \;=\; p_1 \cdot p_2 \cdots p_j \qquad\qquad (p_1 \leq p_2 \leq \ldots \leq p_j)$$

Proof. We must prove two things: (1) every positive integer can be expressed as a product of primes, and (2) this expression is unique.

First, we use strong induction to prove that every positive integer n is a product of primes. As a base case, $n = 1$ is the product of the empty set of primes. For the inductive step, suppose that every $k < n$ is a product of primes. We must show that n is also a product of primes. If n is itself prime, then this is true trivially. Otherwise, $n = ab$ for some $a, b < n$. By the induction assumption, a and b are both products of primes. Therefore, $a \cdot b = n$ is also a product of primes. Thus, the claim is proved by induction.

Second, we use the well-ordering principle to prove that every positive integer can be written as a product of primes in a unique way. The proof is by contradiction: assume, contrary to the claim, that there exist positive integers that can be written as products of

The Prime Number Theorem

Let $\pi(x)$ denote the number of primes less than or equal to x. For example, $\pi(10) = 4$ because 2, 3, 5, and 7 are the primes less than or equal to 10. Primes are very irregularly distributed, so the growth of π is similarly erratic. However, the Prime Number Theorem give an approximate answer:

$$\lim_{x \to \infty} \frac{\pi(x)}{x / \ln x} = 1$$

Thus, primes gradually taper off. As a rule of thumb, about 1 integer out of every $\ln x$ in the vicinity of x is a prime.

The Prime Number Theorem was conjectured by Legendre in 1798 and proved a century later by de la Vallee Poussin and Hadamard in 1896. However, after his death, a notebook of Gauss was found to contain the same conjecture, which he apparently made in 1791 at age 15. (You sort of have to feel sorry for all the otherwise "great" mathematicans who had the misfortune of being contemporaries of Gauss.)

In late 2004 a billboard appeared in various locations around the country:

$$\left\{ \begin{array}{c} \text{first 10-digit prime found} \\ \text{in consecutive digits of } e \end{array} \right\} \textbf{. com}$$

Substituting the correct number for the expression in curly-braces produced the URL for a Google employment page. The idea was that Google was interested in hiring the sort of people that could and would solve such a problem.

How hard is this problem? Would you have to look through thousands or millions or billions of digits of e to find a 10-digit prime? The rule of thumb dervied from the Prime Number Theorem says that among 10-digit numbers, about 1 in

$$\ln 10^{10} \approx 23$$

is prime. This suggests that the problem isn't really so hard! Sure enough, the first 10-digit prime in consecutive digits of e appears quite early:

$$e = 2.718281828459045235360287471352662497757247093699959574966$$
$$9676277240766303535475945713821785251664\underline{2742746639}19320030$$
$$5992181741359662904357290033429526059563073813232862794434\ldots$$

primes in more than one way. By the well-ordering principle, there is a smallest integer with this property. Call this integer n, and let

$$n = p_1 \cdot p_2 \cdots p_j$$
$$= q_1 \cdot q_2 \cdots q_k$$

be two of the (possibly many) ways to write n as a product of primes. Then $p_1 \mid n$ and so $p_1 \mid q_1 q_2 \cdots q_k$. Lemma 10 implies that p_1 divides one of the primes q_i. But since q_i is a prime, it must be that $p_1 = q_i$. Deleting p_1 from the first product and q_i from the second, we find that n/p_1 is a positive integer *smaller* than n that can also be written as a product of primes in two distinct ways. But this contradicts the definition of n as the smallest such positive integer. $\square$