Arithmetic in the Ring of Integers.

The set of integers is usually denoted $\mathbb{Z}$. The properties of the integers I'm going to discuss involve **divisibility**. If a and b are elements of $\mathbb{Z}$— the shorthand for this is $a, b \in \mathbb{Z}$ — then a **divides** b if there is an integer c such that $ac = b$. For example, 4 divides 12 because $4 \cdot 3 = 12$.

You probably learned how to perform "long division" in grade school. You know that if a and b are integers and $b > 0$, you can divide a by b to get a quotient q and a remainder r. Long division is just a recipe for producing q and r.

The first thing I'm going to prove is that there really is such a q and r, and that they're *unique* — a division problem has only one answer.

You may say: "Why do you have to prove this? I know long division works." Well, what you *really* know is that it's worked in every problem you've tried. However, *in mathematics you can't prove a statement about infinitely many cases by simply trying a few examples* (unless, of course, you can prove that your examples really exhaust the possibilities). Let me give you one example which shows why you have to be careful.

Plug positive integers into $1141n^2 + 1$ — $n = 1$, $n = 2$, and so on. Do you ever get a perfect square?

If you try a few cases by hand, you may be led to conclude that you never get a perfect square. In fact, you can use a program like *Mathematica* to try all the n's from 1 to (say) 1000000. Still no perfect square.

Now one million cases is more than you've probably ever tried for *anything* in the real world. And in the real world, you'd probably conclude based on far fewer trials that the thing was impossible.

However, one million numbers is essentially nothing, because the number of integers is infinite. There is no reason why a solution to a problem has to be small by human standards. And indeed, there is an n for which $1141n^2 + 1$ is a perfect square. The smallest such n is a little bigger than 10^{26}!

Back to our discussion of division. To prove that division works, I will use the following result:

Well-Ordering Principle. The positive integers $\mathbb{Z}^+$ are **well-ordered** — that is, every nonempty subset of the positive integers has a smallest element.

This is a deep property of $\mathbb{Z}^+$; to prove it would require a long excursion into set theory which would not be appropriate here.

Example. There is no positive integer between 0 and 1.

This is hardly a surprise — I will use the proof to illustrate the use of well-ordering.

This will be a **proof by contradiction**: I'll assume that the result I want to prove is *false*, then derive a contradiction. The contradiction means that my original assumption (that the result was false) must be false, so the result must be true.

Suppose, then, that there is a positive integer between 0 and 1. Then the set of positive integers between 0 and 1 contains at least one element — it's nonempty. I can apply well-ordering to this set, and therefore conclude that there is a *smallest* positive integer between 0 and 1. Call it x.

Now $0 < x < 1$, so multiplying through by x I get $0 < x^2 < x < 1$. (I'm assuming some elementary properties of inequalities here.) Now x^2 is an integer, it lies between 0 and 1, and it's smaller than x. This is a contradiction, because x was supposed to be the smallest positive integer between 0 and 1. $\times$

This contradiction means that my original supposition is false — and hence that there is no positive integer between 0 and 1. $\square$

Here is a more useful result: our algorithm for division.

Theorem. (The Division Algorithm) Let $a, b \in \mathbb{Z}$, $b > 0$. There are unique integers q and r such that

$$a = bq + r, \qquad \text{where } 0 \le r < b.$$

Proof. There are two things to do. First, I will show that there is *at least* one pair q, r which works. Then I will show that there is *only* one pair q, r which works.

To see the picture, take $a > 0$. Then the idea is to repeatedly subtract b from a until you get a negative number. At that point, you know you subtracted off one b too many, and one less than the number you subtracted is q.

This suggests the following trick. Look at the set

$$S = \{a - bq \mid q \in \mathbb{Z} \text{ and } a - bq \geq 0\}.$$

First, S is nonempty. To show this, I'll consider the three possibilities for a. If $a > 0$, then $a - b \cdot 0 = a > 0$ is in the set.

If $a < 0$, take $q = 2a$. Then

$$a - bq = a - 2ab = a(1 - 2b).$$

$a < 0$ and $b > 0$ implies $b \geq 1$. Then $2b > b \geq 1$, so $1 - 2b < 0$. Therefore, $a - bq = a(1 - 2b) > 0$, since the right side is the product of negative numbers.

Finally, if $a = 0$, take $q = -1$. Then $a - bq = b > 0$.

Thus, S is nonempty. If $0 \in S$, then $a - bq = 0$ for some q. I can use that q and take $r = 0$, and I have a q and r that work.

Suppose then that $0 \notin S$. Then S consists of positive integers. By well-ordering, it has a smallest element r. Thus, $a - bq = r$ for some q. $r > 0$ by assumption; is $r < b$? Suppose not. If $r \geq b$, then $r - b \geq 0$, so

$$a - b(q + 1) = r - b \geq 0.$$

Thus, $a - b(q + 1) \in S$. But $a - b(q + 1) = r - b < r$, which contradicts the fact that r is the smallest element of S. I've proved by contradiction that $r < b$. Thus, $0 < r < b$, and I have a q and r that work.

Finally, I need to show that q and r are unique. Suppose I had two pairs $\{q, r\}$, $\{q', r'\}$ which satisfied the conclusion. Then

$$bq + r = a = bq' + r', \quad \text{so} \quad b(q - q') = r - r'.$$

Assume without loss of generality that $r \geq r'$. (I can do this because if it's the other way around, I can just switch the names of the two pairs and start over again.) Now $0 \leq r < b$ and $0 \leq r' < b$, so the difference of r and r' must surely be less than b. That is $0 \leq r - r' < b$.

But the left side of $b(q - q') = r - r'$ is a nonnegative multiple of b — if it's nonzero, it's a positive multiple of b, i.e. greater than or equal to b. This is impossible because of the restriction on $r - r'$.

The only way out is if $b(q - q') = 0$. Since $b > 0$, this means $q = q'$. In addition, $r - r' = 0$, so $r = r'$. That is, the two pairs were the same to begin with. □

I want to discuss greatest common divisors next, but before that, I want to record an easy property of divisibility. First, some more shorthand. I'll write $a|b$ to mean that a divides b. (Don't confuse this with the "|" that's used in defining sets.)

Lemma. If $a|b$ and $a|c$, then $a|sb + tc$ for all $s, t \in \mathbb{Z}$.

Proof. $a|b$ means $b = ae$ for some e; $a|c$ means $c = af$ for some f. So

$$sb + tc = s(ae) + t(af) = a(se + tf).$$

Therefore, $a|sb + tc$. □

Definition. If m and n are integers, the **greatest common divisor** (m, n) of m and n is the largest integer which divides m and n. $(0, 0)$ is undefined. I'll often get lazy and abbreviate "greatest common divisor" to "gcd".

Having defined the gcd, it would be helpful to know that it exists! First, $1|m$ and $1|n$, so the set of positive common divisors of m and n is nonempty. Since any common divisor of m and n must be less than or equal to $\max(|m|, |n|)$ (use the Division Algorithm!), the set of positive common divisors is bounded above. Hence, there is a large positive common divisor, and this is the greatest common divisor of m and n.

Let $m > 0$. Then $m|a$ if and only if $m|(-a)$. It follows that $(m, n) = (|m|, |n|)$. So if I want to compute (m, n), it does not hurt to assume m and n are positive.

I'm going to give an algorithm which computes the gcd. It isn't efficient, but it's easy to see that it does what I want.

Algorithm: Let m and n be nonnegative integers, not both 0. $\{m, n\}$ will denote the *set* consisting of m and n (so $\{m, n\} = \{n, m\}$). At each step, replace the pair with the pair obtained by subtracting the smaller number from the larger. (If the two numbers are equal, then do $\{m, m\} \to \{m, 0\}$.) When one of the elements is 0, stop.

Symbolically,

$$\{m, n\} \to \begin{cases} \{m - n, n\} & \text{if } m \geq n > 0 \\ \{m, n\} & \text{if } m = 0 \text{ or } n = 0 \end{cases}$$

At each step, the pair decreases in "size". If I wanted to be precise, I could take as a measure of "size" the sum of the two numbers. Since the numbers in a pair are nonnegative, the procedure must terminate in a pair $\{m, 0\}$, where $m > 0$.

Claim: Each step of the algorithm "preserves the greatest common divisor". That is, the output pair always has the same greatest common divisor as the input pair.

Proof. Suppose $m \geq n > 0$. The input is $\{m, n\}$ and the output is $\{m - n, n\}$ (so I'm taking $n \neq m$).

I will show that these two pairs have the same set of *common divisors*. It follows from this that the two pairs must have the same *greatest* common divisor.

First, I'll show that a common divisor of m and n is a common divisor of $m - n$ and n. If $k|m$ and $k|n$, then $k|m - n$ and $k|n$.

Conversely, if $k|m - n$ and $k|n$, then $k|(m - n) + n = m$. So k divides m and n.

Hence, the pairs $\{m, n\}$ and $\{m - n, n\}$ have the same set of common divisors. Hence, they must have the same greatest common divisor. $\square$

The algorithm terminates in a pair $\{m, 0\}$ with $m > 0$. Clearly, $(m, 0) = m$. Since the greatest common divisor never changes, m must also be the greatest common divisor of the original pair of numbers. Therefore, this algorithm finds the greatest common divisor of two numbers.

Next, observe that repeated subtraction is the same as division. This means that I may condense repeated subtractions into a single division via the Division Algorithm, and the resulting procedure will be equivalent. Thus, if $m > n$, the Division Algorithm gives

$$m = nq + r, \quad \text{where } 0 \leq r < n.$$

This represents a replacement $\{m, n\} \to \{n, r\}$. In terms of division, the procedure terminates when the remainder is 0, and the last nonzero remainder is the greatest common divisor of the original pair. When you use division to find the greatest common divisor, the algorithm is usually referred to as the **Euclidean algorithm**.

Example. To find $(51, 36)$, write:

$$\begin{array}{ccccc} 51 & = & 1 \cdot 36 & + & 15 \\ & \swarrow & & \swarrow & \\ 36 & = & 2 \cdot 15 & + & 6 \\ & \swarrow & & \swarrow & \\ 15 & = & 2 \cdot 6 & + & 3 \\ & \swarrow & & \swarrow & \\ 6 & = 2 \cdot 3 & + & 0 & \end{array}$$

Hence, $(51, 36) = 3$.

The subtraction algorithm would look like this:

$$\{51, 36\} \to \{15, 36\} \to \{21, 15\} \to \{6, 15\} \to \{9, 6\} \to \{3, 6\} \to \{3, 3\} \to \{3, 0\} \ \square$$

Terminology. If x and y are things, a **linear combination** of x and y is something of the form $ax + by$, where a and b are numbers. (The kind of "number" depends on the context.)

Proposition. (m,n) is a linear combination of m and n: $(m,n) = am + bn$ for some integers a, b.

Warning: a and b are not unique.

Proof. Return to our subtraction algorithm. I'll show that at each step, the current pair is a linear combination of the original pair. This will be our first example of **Mathematical Induction** (though I will do it somewhat informally).

To do this, I'll show that the *first step* of the procedure satisfies the claim. Next, *assume* that the procedure is known to satisfy the claim up to a certain point. Then I'll show that it still satisfies the claim at the next step.

First, observe that the result is true for the first (new) pair. For $\{m,n\}$ produces either $\{m-n,n\}$ or $\{m,n\}$, and in either case the members of the output pair are linear combinations of m and n. The first step works.

Now suppose the current pair is $\{m', n'\}$ and I know each element is a linear combination of m and n. Say

$$m' = am + bn, \qquad n' = cm + dn.$$

Is the same thing true when I take the next step?

If one of m', n' is 0, there is no next step — I'm finished. In this case, each element is a linear combination of m and n by assumption.

Otherwise, the next pair is $\{m'-n', n'\}$. The elements of this pair are

$$m' - n' = (a-c)m + (b-d)n, \qquad n' = cm + dn.$$

Here also each element is a linear combination of m and n.

This proves that each pair we obtain in our subtraction algorithm will be a linear combination of the elements in the original pair.

In particular, this is true at the point where the algorithm terminates — that is, when I get a pair $\{\tilde{m}, 0\}$. However, the greatest common divisor of $\{\tilde{m}, 0\}$ is $\tilde{m}$. Hence, the greatest common divisor of the original numbers is a linear combination of the original numbers. □

Example. Let m and n be nonnegative integers, not both 0. You can use the Euclidean algorithm to find number a and b such that $(m,n) = am + bn$. Here's the computation for $(51, 36)$:

$$\begin{array}{ccccc}
51 & = & 1 \cdot 36 & + & 15 \\
 & \swarrow & & \swarrow & \\
36 & = & 2 \cdot 15 & + & 6 \\
 & \swarrow & & \swarrow & \\
15 & = & 2 \cdot 6 & + & 3 \\
 & \swarrow & & \swarrow & \\
6 & = 2 \cdot 3 & + & 0 &
\end{array}$$

The third equation says $3 = 15 - 2 \cdot 6$.

By the second equation, $6 = 36 - 2 \cdot 15$, so

$$3 = 15 - 2 \cdot (36 - 2 \cdot 15) = 5 \cdot 15 - 2 \cdot 36.$$

The first equation says $15 = 51 - 36$, so

$$3 = 5 \cdot (51 - 36) - 2 \cdot 36 = 5 \cdot 51 - 7 \cdot 36.$$

I've expressed the greatest common divisor 3 as a linear combination of the original numbers 51 and 36.

□