

7.1 Program Correctness

Showing that a program is correct means that it does what it is supposed to do. More formally, our goal is to prove that a program satisfies its *specification*, that is, it correctly realizes the prescribed relationship between inputs and outputs. In other words, for each input, the specification tells us what the program should output as a response.

There are two parts to correctness of a program.

1. *Partial correctness*: If the program ever returns a result, it is the correct result.
2. *Termination*: The program returns.

Today we prove the correctness of the grade school multiplication algorithm.

7.2 Grade School Multiplication Algorithm

Let's start with the specification. As input, the program receives two positive integers, a and b . As output, it should return their product, i.e., $a \cdot b$.

7.2.1 Binary Representation of Integers

To make the analysis easier, we will work with binary representations of numbers instead of decimal representations. As we will see, this makes the grade school multiplication algorithm easier to describe.

In the usual decimal representation of a number, we represent a $(k + 1)$ -digit integer n as a sequence of digits between 0 and 9 and write it as $d_k d_{k-1} \dots d_1 d_0$ with $d_i \in \{0, 1, \dots, 9\}$ for $i \in \{0, 1, \dots, k\}$. Another way to think about n is as a sum of powers of 10, that is

$$n = \sum_{i=0}^k d_i \cdot 10^i, \quad d_i \in \{0, 1, \dots, 9\}. \quad (7.1)$$

For example, the sum of the form (7.1) corresponding to the integer 14376 is

$$1 \cdot 10^4 + 4 \cdot 10^3 + 3 \cdot 10^2 + 7 \cdot 10 + 6 \cdot 10^0$$

(so $d_4 = 1$, $d_3 = 4$, $d_2 = 3$, $d_1 = 7$, and $d_0 = 6$).

To obtain a description of the form (7.1) from some integer n , start with $n_0 = n$. Take the last digit of n_0 , that is, take the remainder after dividing n_0 by 10, set d_0 to that remainder, and then subtract d_0 from n_0 . Notice that $n_0 - d_0$ is divisible by 10. Dividing $n_0 - d_0$ by 10 gives us an integer n_1 . Next, repeat the process with n_1 . Take the last digit of n_1 by finding the remainder after dividing n_1 by 10, make the remainder d_1 , subtract d_1 from n_1 , divide the difference by 10, and get n_2 . Keep repeating this process until you end with $n_{k+1} = 0$.

i	n_i	d_i	$n_i - d_i$
0	14376	6	14370
1	1437	7	1430
2	143	3	140
3	14	4	10
4	1	1	0
5	0		

Table 7.1: Obtaining the decimal representation of $n = 14376$.

For example, if we do this with $n = 14376$, we get the values in Table 7.1.

Of course, for decimal numbers, this is a silly procedure since we can just read off the numbers d_0 through d_k from the decimal representation of n right away; however, it gives us a key insight into how to find the binary representation of n . To get the binary representation of n , we cannot just read the bits b_i off of the decimal representation of n , but we can apply the algorithm we described, except with 2 in place of 10, and with the remainders being either 0 or 1 instead of being one of 0 through 9. In the end, we obtain a representation of n of the form

$$\sum_{i=0}^l b_i \cdot 2^i, \quad b_i \in \{0, 1\}, \quad (7.2)$$

and can write the binary representation of n as $b_l b_{l-1} \dots b_1 b_0$.

Table 7.2 shows how we obtain the binary representation of $n = 75$, 1001011.

i	n_i	b_i	$n_i - b_i$
0	75	1	74
1	37	1	36
2	18	0	18
3	9	1	8
4	4	0	2
5	2	0	1
6	1	1	0
7	0		

Table 7.2: Obtaining the binary representation of $n = 75$. Reading the b_i column from bottom to top gives us the binary representation of n , 1001011.

To give some reasoning behind why the algorithm for obtaining a binary representation works, rewrite (7.2) as

$$b_0 + \sum_{i=1}^l b_i 2^i = b_0 + 2 \left(\sum_{i=1}^l b_i 2^{i-1} \right).$$

We see that b_0 is the remainder after dividing n_0 by 2, and we get n_1 by subtracting the remainder and dividing the difference by 2. Then we continue the process with

$$n_1 = \sum_{i=1}^l b_i 2^{i-1} = \sum_{i=0}^{l-1} b_{i+1} 2^i.$$

7.2.2 Description of the Algorithm

Now let's review the grade school multiplication algorithm. We write the two numbers we multiply, a and b , above each other. We multiply a by the last digit of b , and write down the result below b . Then we multiply a by the next to last digit of b , and write down the result on the next line, shifted one digit to the left. In general, when we write down the result of multiplying a by some digit of b , we write down the result so that its last digit is in the same column as the digit of b we used to produce the result. If some digit of b is zero, we simply skip it and don't write anything down. Finally, we add up all the intermediate results we wrote down to get the product $a \cdot b$. We show this for $a = 14376$ and $b = 108$ in Figure 7.1a.

$\begin{array}{r} 14376 \\ \cdot 2108 \\ \hline 115008 \\ 14376 \\ 28752 \\ \hline 30304608 \end{array}$	$\begin{array}{r} 10011 \\ \cdot 1101 \\ \hline 10011 \\ 10011 \\ 10011 \\ \hline 11110111 \end{array}$	$\begin{array}{r} 10011 \\ \cdot 1101 \\ \hline 10011 \\ 10011 \\ \hline 1011111 \\ 10011 \\ \hline 11110111 \end{array}$
(a) Using decimal representation	(b) Using binary representation	(c) Adding immediately

Figure 7.1: Grade school multiplication algorithm.

Using binary representations instead of decimal representations greatly simplifies our rules for multiplication. Our multiplication table goes down from being 10×10 to being just 2×2 . This makes the description of the “college version” of the grade school algorithm very easy. In each step, we look at one bit of b . If the bit is 1, we copy down a so that its last bit lines up with the bit of b we're currently considering. Then we perform binary addition of the intermediate results to get the result. In the example in Figure 7.1b, we multiply $19 \cdot 13$ in binary and get 247.

In order to analyze the algorithm, we should define it more formally. In its current version, we need to keep track of possibly many intermediate results, one for each bit of b . That would make the analysis more complicated. Notice that it doesn't matter whether we first perform all the multiplications and then add together all the intermediate results, or whether we perform an addition step after a multiplication step. An example is shown in Figure 7.1c. This version looks like more work because we have to perform multiple additions, but has the advantage that we only need to keep track of one variable that holds the sum of the intermediate results obtained so far.

Finally, we formalize our algorithm as Algorithm 1. The variable y represents the part of the number b we still need to multiply a with. We multiply x by two in each iteration of the loop so as to simulate lining up the last bit of the intermediate result with the bit of b used in the current multiplication step. The variable p holds our running total from Figure 7.1c. Finally, the notation $\lfloor m \rfloor$ on line 7 means that we round m down to the nearest integer, and we read $\lfloor m \rfloor$ as the “floor of m ”. Thus, $\lfloor y/2 \rfloor$ means we divide y by two and round down, which corresponds to cutting off the last bit from y . We can dispose of that bit because we've already multiplied a with it and won't need it for anything else.

7.2.3 Correctness of the Algorithm

Recall that there are two conditions a correct algorithm must satisfy: the partial correctness condition and the termination condition. Let's see what they are in the case of Algorithm 1.

Algorithm 1: Binary Multiplication Algorithm

Input: a, b - positive integers we want to multiply**Output:** ab - product of a and b

```

(1)  $x \leftarrow a$ 
(2)  $y \leftarrow b$ 
(3)  $p \leftarrow 0$ 
(4) while  $y > 0$  do
(5)   if  $y$  is odd then  $p \leftarrow p + x$ 
(6)    $y \leftarrow \lfloor y/2 \rfloor$ 
(7)    $x \leftarrow 2x$ 
(8) end
(9) return  $p$ 

```

1. *Partial correctness:* When we reach line 10, $p = ab$.
2. *Termination:* We eventually reach line 10. In other words, the loop on line 4 ends after a finite number of iterations.

It is fairly easy to see what happens on the first three lines. We just initialize our variables. Hence, the crux of all our arguments about Algorithm 1's behavior will be in proving facts about the behavior of the loop on line 4.

We can view our algorithm as a system whose state is represented by the values of the variables x , y , and p . We are trying to prove some facts about the state of the algorithm after each repetition of the loop, that is, after each "time step". This sounds like a setting for the use of invariants, and indeed, we will prove certain *loop invariants* on our way towards a proof that Algorithm 1 works correctly. More formally, a loop invariant is a property that holds at the beginning and after any number of iterations of a loop. A loop invariant usually describes relationships among variables.

7.2.3.1 Partial Correctness

To prove partial correctness, we prove the following loop invariant.

Invariant 7.1. After n iterations of the loop on line 4, $ab = xy + p$.

Proof. Let x_n , y_n , and p_n be the values of x , y , and p after n iterations of the loop from line 4, respectively. We show that

$$ab = x_n y_n + p_n \tag{7.3}$$

for every natural number n .

We prove by induction that $(\forall n)P(n)$ where $P(n)$ is "equation (7.3) holds".

For the base case, $P(0)$, there haven't been any iterations of the loop yet, so x_0 , y_0 , and p_0 have the values from the first three lines of Algorithm 1. Thus, $x_0 = a$, $y_0 = b$, and $p_0 = 0$. We see that $x_0 y_0 + p_0 = ab + 0 = ab$, so the base case is proved.

Now we prove the inductive step $(\forall n)P(n) \Rightarrow P(n+1)$. Since there is an if statement in the loop body, we use a proof by cases.

Case 1: y_n is odd. In this case we execute the body of the if statement, and add x_n to p_n . Since nothing else happens to the value of p in the loop body, we get $p_{n+1} = p_n + x_n$. Next, since y_n is odd, we can write $y_n = 2k + 1$ for some integer k , namely $k = (y_n - 1)/2$. After line 7 executes, we

have $y_{n+1} = \lfloor y_n/2 \rfloor = \lfloor (2k+1)/2 \rfloor = \lfloor k + \frac{1}{2} \rfloor = k = (y_n - 1)/2$. Finally, the last line of the loop body doubles the value of x , so $x_{n+1} = 2x_n$.

Now we have found the values of x , y and p at the end of the $(n+1)$ st iteration of the loop, so we can verify that (7.3) holds after the loop is complete. We have

$$x_{n+1}y_{n+1} + p_{n+1} = 2x_n \frac{y_n - 1}{2} + p_n + x_n = x_n(y_n - 1) + x_n + p_n = x_n y_n + p_n,$$

and the right-hand side is ab by the induction hypothesis. Hence, if y_n is odd, the invariant is maintained.

Case 2: y_n is even. In this case the if statement body is skipped, so the value of p doesn't change, and we get $p_{n+1} = p_n$. Next, since y_n is even, we can write $y_n = 2k$ for some integer k , namely $k = y_n/2$. After line 7 executes, we have $y_{n+1} = \lfloor \frac{y_n}{2} \rfloor = \lfloor \frac{2k}{2} \rfloor = \lfloor k \rfloor = k = y_n/2$. It follows that $y_{n+1} = y_n/2$. Finally, the last line of the loop body doubles the value of x , so $x_{n+1} = 2x_n$.

Now we have found the values of x , y and p at the end of the $(n+1)$ st iteration of the loop, so we can verify that (7.3) holds after the loop is complete. We have

$$x_{n+1}y_{n+1} + p_{n+1} = 2x_n \frac{y_n}{2} + p_n = x_n y_n + p_n,$$

and the right-hand side is ab by the induction hypothesis. It follows that if y_n is even, the invariant is maintained.

This completes the proof of the induction step, and also of our proposition. \square

We get out of the loop if $y \leq 0$ at the end of some iteration. If we show that $y = 0$ after the last iteration of the loop, Invariant 7.1 implies that $ab = xy + p = 0 + p = p$, which proves that the value of p after the loop terminates is ab . It suffices to show as another loop invariant that y never becomes negative.

Invariant 7.2. *After n iterations of the loop, $y \geq 0$.*

Proof. The proof goes by induction. Like an earlier proof, let y_n be the value of y after n iterations of the loop.

We see that $y_0 \geq 0$ because $y_0 = b > 0$, which proves the base case.

Now assume $y_n \geq 0$, and consider the $(n+1)$ st iteration of the loop. Inside of the loop body, y only changes on line 7. There, we divide y by 2 and round down the result to get y_{n+1} . Since $y_n \geq 0$, $y_n/2 \geq 0$ as well, and rounding down a non-negative number cannot round down below zero, which proves that $y_{n+1} \geq 0$. Hence, the inductive step is proved, and so is the invariant. \square

The loop condition implies that when the loop is over, $y \leq 0$. We just showed that $y \geq 0$ throughout the algorithm. It follows that when the loop terminates, $y = 0$, which means that the algorithm returns $p = ab$. Hence, partial correctness of Algorithm 1 is proved.

Let us offer some intuition behind Invariant 7.1. Again, it takes some ingenuity to come up with this invariant. One way of thinking is that in each iteration of the loop, we “shift” x to the left and y to the right. This corresponds to multiplying x by 2 and dividing y by 2. We cut off the last bit of y in the process, and if the last bit of y was 1, we compensate for this loss by adding x to our partial result p .

7.2.3.2 Termination

To complete our correctness proof for Algorithm 1, we show that the loop on line 4 eventually terminates. For this, we show that the value of y decreases in each iteration of the loop by at least one. Intuitively, this is true because y is a positive integer and we divide it by two in every step and round down, which should decrease its value. Now let's argue formally.

Proposition 7.3. *If $y > 0$, $\lfloor y/2 \rfloor \leq y - 1$.*

Proof. Write $y = 2k + r$ where $r \in \{0, 1\}$ and $k \in \mathbb{N}$. Then $\lfloor y/2 \rfloor = \lfloor (2k + r)/2 \rfloor = \lfloor k + r/2 \rfloor = k$.

To show $\lfloor y/2 \rfloor = k \leq y - 1 = 2k + r - 1$, just observe that the only way this would not hold for some $k \in \mathbb{N}$ and $r \in \{0, 1\}$ is if $k = r = 0$, but in that case we would have $y = 0$, which is a case our proposition doesn't consider. \square

Since y is initialized to b at the beginning of the algorithm and decreases by at least one in every iteration of the loop by Proposition 7.3, it follows that after at most b iterations of the loop, y becomes zero, at which point the loop terminates and the algorithm returns.

Now the proof of Algorithm 1's correctness is complete.

We now look at another correctness proof. We prove correctness of an algorithm for finding the greatest common divisor of two integers.

7.2.4 Greatest Common Divisor

We start by defining what we mean by greatest common divisor.

Definition 7.4. *The greatest common divisor of integers a and b , denoted $\gcd(a, b)$, is the largest integer d such that d divides both a and b . Furthermore, if c divides both a and b , then c also divides $\gcd(a, b)$.*

Because zero is divisible by every integer, $\gcd(0, 0)$ is undefined. In more generality, $\gcd(a, 0)$ is undefined if $a = 0$, and is a when $a > 0$. To see the latter, note that the largest divisor of a is a itself, and a also divides zero (because every integer does). Similarly, $\gcd(0, b)$ is undefined if $b = 0$, and is b otherwise.

We state and prove three properties of the greatest common divisor. We will use these properties to construct an algorithm for finding greatest common divisors.

Lemma 7.5. *Let $a, b \in \mathbb{N}$ with at least one of a, b nonzero. Then the following three properties hold.*

- (i) *If $a = b$, $\gcd(a, b) = a = b$.*
- (ii) *If $a < b$, $\gcd(a, b) = \gcd(a, b - a)$.*
- (iii) *If $a > b$, $\gcd(a, b) = \gcd(a - b, a)$.*

Proof. When $a = b$, a divides both a and b . Furthermore, no integer greater than a divides a , so $\gcd(a, b) = a = b$. This proves (i).

Now let's argue (ii). We show that d is a divisor of both a and b if and only if d is a divisor of both a and $b - a$. We do so by proving two implications.

First assume that d is a divisor of both a and b . Then there exist integers c_1 and c_2 such that $a = c_1d$ and $b = c_2d$. Therefore, we can write $b - a = (c_2 - c_1)d$, and we see that $b - a$ is a multiple

of d , which implies that d divides $b - a$. Since d also divides a by assumption, we have shown that d divides both a and $b - a$.

For the other direction, assume d is a divisor of both a and $b - a$. Then there exist integers c_1 and c_2 such that $a = c_1d$ and $b - a = c_2d$. Therefore, we can write $b = (b - a) + a = (c_2 + c_1)d$, and we see that b is a multiple of d , which implies that d divides b . Since d also divides a by assumption, we have shown that d divides both a and b .

Therefore, the sets $\{d \mid d \text{ divides } a \text{ and } d \text{ divides } b\}$ and $\{d \mid d \text{ divides } a \text{ and } d \text{ divides } b - a\}$ are the same, so they both have the same largest element. This largest element is the greatest common divisor of both a and b , and of a and $b - a$. This completes the proof of (ii).

We would argue (iii) the same way as (ii), except with the roles of a and b switched. This completes the proof of the lemma. \square

7.2.5 The GCD Algorithm

The observations proved as Lemma 7.5 are a good starting point for an algorithm that computes the greatest common divisor of two numbers. The basic idea is to use (ii) and (iii) of Lemma 7.5 to decrease the values of x and y used for the computation of the greatest common divisor with the hope that they eventually achieve values for which the greatest common divisor is easy to compute (such as case (i) of Lemma 7.5).

We describe the algorithm formally now, and show the specification together with the algorithm.

Algorithm 2: GCD Algorithm

Input: a, b positive integers

Output: $\gcd(a, b)$

```

(1)  $(x, y) \leftarrow (a, b)$ 
(2) while  $x \neq y$  do
(3)   | if  $x < y$  then  $y \leftarrow y - x$ 
(4)   |
(5)   |           else  $x \leftarrow x - y$ 
(6)   |
(7) end
(8) return  $x$ 

```

To argue correctness of Algorithm 2, we must show that the algorithm returns the greatest common divisor of its two inputs, a and b , and that it terminates. To that end, we define and prove some loop invariants.

Observe that the loop on line 2 terminates when $x = y$, and in that case $\gcd(x, y) = x$. At the beginning, we have $x = a$ and $y = b$, so $\gcd(x, y) = \gcd(a, b)$. If we show that we maintain $\gcd(a, b) = \gcd(x, y)$ throughout the algorithm, that will give us partial correctness. In fact, it is possible to show this, and we do so now.

Invariant 7.6. *After n iterations of the loop on line 2, $\gcd(a, b) = \gcd(x, y)$.*

Proof. We prove this invariant by induction on the number of iterations of the loop.

Let x_n and y_n be the values of x and y after n iterations of the loop, respectively.

For the base case, we have $x_0 = a$ and $y_0 = b$ from line 1, so $\gcd(x_0, y_0) = \gcd(a, b)$.

To prove the inductive step, assume that $\gcd(x_n, y_n) = \gcd(a, b)$.

If $x_n = y_n$, there isn't going to be another iteration of the loop. In this case we don't need to argue about the values of x and y after the $(n + 1)$ st iteration of the loop at all.

Now suppose that $x_n \neq y_n$. Then there will be another iteration of the loop, and we have two cases to consider.

Case 1: $x_n < y_n$. Then $x_{n+1} = x_n$ and $y_{n+1} = y_n - x_n$. We see that

$$\gcd(x_{n+1}, y_{n+1}) = \gcd(x_n, y_n - x_n) = \gcd(x_n, y_n) = \gcd(a, b). \quad (7.4)$$

The second equality in (7.4) follows by part (ii) of Lemma 7.5 and the last equality follows by the induction hypothesis. Therefore, $\gcd(x_{n+1}, y_{n+1}) = \gcd(a, b)$ in this case.

Case 2: $x_n > y_n$. Then $x_{n+1} = x_n - y_n$ and $y_{n+1} = y_n$. We see that

$$\gcd(x_{n+1}, y_{n+1}) = \gcd(x_n - y_n, y_n) = \gcd(x_n, y_n) = \gcd(a, b). \quad (7.5)$$

The second equality in (7.5) follows by part (iii) of Lemma 7.5 and the last equality follows by the induction hypothesis. Therefore, $\gcd(x_{n+1}, y_{n+1}) = \gcd(a, b)$ in this case as well.

This completes the proof of the induction step, and of the invariant. \square

To argue partial correctness, we observe that when the loop terminates, $\gcd(x, y) = \gcd(a, b)$ by Invariant 7.6. Next, since the loop terminated, $x = y$, so $\gcd(x, y) = x$. Therefore, $x = \gcd(a, b)$, and we return x , so we return $\gcd(a, b)$ as desired. This completes the proof of partial correctness.

Here we remark that in practice, we would not spell out the argument for Case 2 in the proof of Invariant 7.6 because it is almost the same as the argument for Case 1. In mathematical writing, we would just say that if $x_n > y_n$, an argument similar to the one for Case 1 shows that the invariant is maintained after $n + 1$ iterations of the loop in Case 2 as well.

We actually took this shortcut in the proof of Lemma 7.5 where we omitted the proof of part (iii). Make sure that when you take such a shortcut in mathematical writing, you at least check that the omitted proof goes through, for example by writing it down somewhere else.

We now argue that Algorithm 2 terminates by showing that the loop on line 2 terminates after some number of iterations. We prove an additional loop invariant in order to do so.

Invariant 7.7. *After n iterations of the loop, $x > 0$ and $y > 0$.*

Proof. We prove the invariant by induction. As in the proof of Invariant 7.6, let x_n and y_n be the values of x and y after n iterations of the loop, respectively.

Before the loop starts, $x_0 = a$, $y_0 = b$, and the specification tells us that $a, b > 0$. This proves the base case.

Now assume that $x_n > 0$ and $y_n > 0$ for some n . If $x_n = y_n$, there is not going to be an $(n + 1)$ st iteration of the loop. If $x_n \neq y_n$, there are two cases. If $x_n < y_n$, $x_{n+1} = x_n > 0$, and we subtract x from y to get $y_{n+1} = y_n - x_n$, which is greater than zero because $y_n > x_n$. Thus, $x_{n+1}, y_{n+1} > 0$ in this case. We can argue similarly in the case $x_n > y_n$, and we see that the invariant is maintained after the $(n + 1)$ st iteration of the loop. \square

To prove that an algorithm terminates, we often find a quantity, sometimes called a *potential*, that decreases in discrete steps over time, and show that it cannot go below a certain threshold. In our case, the potential is $x + y$, and decreases by either x or y (that's the "discrete step") after each iteration of the loop (i.e., "over time"). We use this quantity to prove termination.

When $x = y$, the algorithm terminates right away, so there is nothing to prove in this case. Now suppose $x \neq y$ after n iterations of the loop. That is, in our earlier notation, $x_n \neq y_n$. Then,

depending on which of x_n, y_n is greater, either $x_{n+1} = x_n - y_n$ and $y_{n+1} = y_n$, or $y_{n+1} = y_n - x_n$ and $x_{n+1} = x_n$. Invariant 7.7 tells us that $x_n, y_n > 0$, so either $x_{n+1} \leq x_n - 1$ or $y_{n+1} \leq y_n - 1$. In either case, $x_{n+1} + y_{n+1} \leq x_n + y_n - 1$, so $x + y$ decreases by at least 1 in each iteration of the loop.

Since $x_0 = a$ and $y_0 = b$, it follows that if the algorithm has gone through $m = a + b - 2$ iterations of the loop without terminating, $x_m + y_m \leq x_0 + y_0 - m = 2$. But Invariant 7.7 tells us that $x_m + y_m \geq 2$ after every iteration of the algorithm, so $x_m + y_m \geq 2$ (this states that our potential $x + y$ never goes below a certain threshold). Therefore, $x_m = y_m = 1$ since both x_m and y_m are positive integers, and the algorithm terminates after the m -th iteration.

This completes the proof of correctness of Algorithm 2.

7.2.6 A Remark about Writing Proofs

Observe that we cannot prove $x_n + y_n \geq 2$ only with Invariant 7.6. If we did not prove Invariant 7.7 beforehand, we would not be able to show that our potential $x + y$ is bounded below by the threshold value of 2.

So, what would have happened if we had never stated and proved Invariant 7.7 and if we had never stated that lemma? We would have still known that $x + y$ was decreasing. That would have led us, after some thinking, to the following thought process: The quantity $x + y$ decreases after every iteration, so it will become negative at some point. But if it gets negative, one of x or y would be negative. This can't happen, though, because we only subtract smaller values from larger ones in the loop. Actually, we can't even get x or y to be zero because that would mean x and y were equal before the iteration that supposedly sets one of them to zero. So we see that $x, y > 0$ at all times, and the smallest $x + y$ can be is 2. At this point, we would have written down the statement of Invariant 7.7 and continued with our termination proof.

The lesson to take away from this is that you should not give up if you get stuck on a proof. When you get stuck, think about why you are stuck. Maybe other facts you know could help you continue with the proof. There may be some facts you have not used in your argument yet. Also try to make more observations about the problem. Many proofs consist of multiple ingredients that need to be mixed together in the right way, and chances are you are not going to find all the ingredients and techniques right away.

As we said earlier in the course, this takes some ingenuity, so make sure you give yourself enough time, too. Also, if you are stuck on a proof too long, open your mind to the possibility that the fact you are trying to prove is wrong (maybe not if we tell you to prove a fact on a homework assignment, but it's a good thing to keep in mind).

7.2.7 On the Importance of a Precise Specification

Before we move on to recursion, we point out the importance of a precise specification. Suppose we allowed a or b to be any natural numbers. Consider the input $a = 0, b = 1$. Then $x = 0$ and $y = 1$ when we enter the loop for the first time. The condition of the if statement is satisfied, so we subtract x from y on line 3, which doesn't do anything since $x = 0$. Thus, x and y have not changed at all in this iteration of the loop, and the same thing will happen in subsequent iterations. Thus, our algorithm gets stuck in an infinite loop on this input.

Note that in the situation above, Invariant 7.7 doesn't hold when the program starts, so we cannot use it to prove termination on the bad input. If we allowed inputs to be zero, we would not be able to prove the base case in the inductive proof of Invariant 7.7, and would fail to prove correctness of Algorithm 2. But we did require $a, b > 0$, so we don't need to reason about inputs that don't satisfy that requirement, and our proof of Invariant 7.7 goes through.

The lesson to take away from this is that, in addition to specifying the input-output relationship, the specification also indicates which inputs the program was designed to work with. We only need to prove correctness on those inputs. The program may work on other inputs too, but isn't guaranteed to work—it could also wipe your hard drive on bad input. It is the responsibility of the user who runs the program to ensure that the inputs are valid.