2 Number Systems and Infinity

It is strange that we know so little about the properties of numbers. They are our handiwork, yet they baffle us; we can fathom only a few of their intricacies. Having defined their attributes and prescribed their behavior, we are hard pressed to perceive the implications of our formulas.
— James R. Newman

Our minds are finite, and yet even in these circumstances of finitude we are surrounded by possibilities that are infinite, and the purpose of human life is to grasp as much as we can out of the infinitude.
— Alfred North Whitehead

THERE ARE TWO complementary images that we should consider before starting this chapter. The first is how a painter or sculptor modifies a medium to create original structure from what was without form. The second is how sound waves propagate through a medium to travel from one point to another. Both images serve as metaphors for the motivation behind this and the next chapter. As for the first image, just as a painter adds pigment to canvas and a sculptor bends and molds clay, so a programmer twiddles bits within silicon. The second image relates to the way information within a computer is subject to the constraints of the environment in which it exists, namely, the computer itself.

The key word in both metaphors is "medium," yet there is a subtle difference in each use of the word. When a human programs a computer, quite often the underlying design of the program represents a mathematical process that is often creative and beautiful in its own right. The fact that good programs are logical by necessity does not diminish the beauty at all. In fact, the acts of blending colors, composing a fugue, and chiseling stone are all subject to their own logical rules, but since the result of these actions seems far removed from logic, it is easy to forget that the rules are really in place. Nevertheless, I would like you to consider the computer as a medium of expression just as you would canvas or clay.

As for the second metaphor, everything that is dynamic exists and changes in accordance with the environment in which it exists. The interactions among objects and environments are also governed by a well-defined set of rules. Similarly, programs executing inside of a computer are by definition following a logical path; thus one could think of a computer as a medium in which programs flow just as sound travels through matter.

There is also a sort of yin-yang duality in this idea that I find pleasing. One of the first items covered in an introductory physics course is the difference between potential and kinetic energy. You can think of potential energy as the energy stored in, say, a battery or a rock placed on top of a hill. Kinetic energy is energy that is in the process of being converted, as when the stored electricity in a battery drives a motor and when a rock rolls down a hill. Similarly, when a human designs a program, there exists a potential computation that is unleashed when the program executes within a computer. Thus, one can think of the computation as being kinetic and in motion. Moreover, just as a child with a firecracker can be surprised by the difference between potential and kinetic energy, so computer programmers are often surprised (even pleasantly) by the difference between potential and kinetic computation.

Now that we've agreed to look at the computer as a medium, and since this book is really about looking at the universe in terms of processes familiar to computer scientists, the next two chapters are devoted exclusively to the properties of numbers and computers.

2.1 Introduction to Number Properties

Sometime around the fifth century B.C., the Greek philosopher Zeno posed a paradox that now bears his name. Suppose that Achilles and a tortoise are to run a footrace. Let's assume that Achilles is exactly twice as fast as the tortoise. (Our tortoise is obviously a veritable Hercules among his kind.) To make things fair, the tortoise will get a head start of 1000 meters. After the start of the race, by the time Achilles runs 1000 meters, the tortoise is still ahead by 500 meters. However, Achilles is a far superior athlete, so he easily covers the next 500 meters. During this time, the tortoise has managed to go another 250 meters. We can repeat the process for an infinite number of time slices while always finding the tortoise just a bit ahead of Achilles. Will Achilles ever catch up to the tortoise?

Clearly, we know that something is amiss with the story, as common sense tells us that the world doesn't work this way and that there exists some distance in which Achilles should be able to overcome the tortoise and pass it. But what is that distance and how long does it take Achilles to finally reach the tortoise? There is an algebraic solution to the problem, but this doesn't directly address the paradox of Achilles always being somewhat behind the tortoise when we break the race up into small time slices.

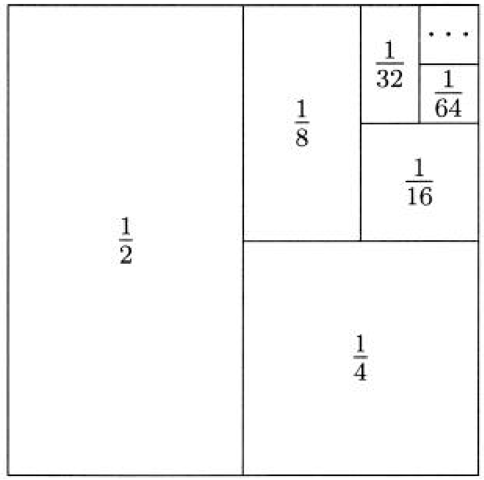

Figure 2.1 An infinite summation captured in a square

Let's add a little more information to the story and concentrate on the question of how long it will take Achilles to catch up to the tortoise. First, let's assume that Achilles can run 1000 meters in exactly one minute. After one minute, Achilles has traveled 1000 meters while the tortoise has covered 500. When Achilles travels the next 500 meters (and the tortoise another 250), one-half of a minute has passed. Similarly, each "time slice" that we are looking at will be exactly half the previous time.

Recall that it was earlier stated that we could look at an infinite number of time slices and always come to the conclusion that the tortoise was always slightly ahead of Achilles. However, just because there is an infinite number of time slices, it does not necessarily mean that the sum of all of the time slices (the total elapsed time) is also infinite. More specifically, what is the sum total of $1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \cdots$?

Forget for the moment that the 1 appears in the sum and just concentrate on the fractions. Another way of writing this is:

$$\sum_{i=1}^{\infty} \frac{1}{2^i} = \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \cdots.$$

At any step in the infinite sum we can represent the current running total by the area of a divided box whose total area is 1. At each step, we divide the empty portion of the box in half, mark one side as used and leave the other half for the next step. As Figure 2.1 illustrates, if we continue the process for an infinite number of steps, we will eventually fill the box. Therefore, the sum total of all of the infinite time slices is really equal to two minutes (1 + 1 = one minute for the infinite sum and the other minute that we originally ignored). Moreover, since we know that Achilles can run exactly 1000 meters a minute, we can conclude that Achilles and the tortoise will be tied if the track is 2000 meters in length. If the track is any length greater than 2000 meters, Achilles will win the race; any less, and Achilles will lose.

Zeno's paradox illustrates just one of the interesting aspects of numbers and infinity that will be highlighted in this chapter. To solve the paradox, we were required to examine the properties of an infinite summation of fractions (or rational numbers). In the remainder of this chapter, we will look at counting numbers, the rational numbers in more detail, and irrational numbers.

2.2 Counting Numbers

Consider the set of natural numbers: 1, 2, 3, ... We know that there is an infinite number of natural numbers. We can say the same thing about all of the even natural numbers. But are there more numbers than even numbers? Surprisingly, the size of the two sets is identical. The reason is that for every member in the set of natural numbers, there is a corresponding member in the set of even numbers. For example, we could construct what is known as a one-to-one mapping:

$$\begin{array}{cccccccc} 1 & 2 & 3 & 4 & 5 & 6 & & \\ \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \cdots . \\ 2 & 4 & 6 & 8 & 10 & 12 & & \end{array}$$

What about more complex sets, such as the set of all perfect cubes? Before answering this question, let's examine the first five perfect cubes in the context of the other natural numbers:

$\boxed{1}$ 2 3 4 5 6 7 $\boxed{8}$ 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26
$\boxed{27}$ 28 29 30 31 32 33 34 35 36 37 38 39 40 41 42 43 44 45 46 47 48 49
50 51 52 53 54 55 56 57 58 59 60 61 62 63 $\boxed{64}$ 65 66 67 68 69 70 71 72
73 74 75 76 77 78 79 80 81 82 83 84 85 86 87 88 89 90 91 92 93 94 95
96 97 98 99 100 101 102 103 104 105 106 107 108 109 110 111 112 113
114 115 116 117 118 119 120 121 122 123 124 $\boxed{125}$ $\cdots$.

Since the space between successive perfect cubes grows dramatically and perfect cubes become less common as we move down the list, you may think that there are far more natural numbers than perfect cubes. This is wrong. There are two reasons why the number of perfect cubes is equal to the number of natural numbers. First, the function to produce perfect cubes is invertible. If I tell you that I am looking at perfect cubes and you give me an example, say 2197, with some effort I can respond by saying that your number is the thirteenth perfect cube. Also, this function yields a one-to-one mapping between its argument and its result, just like the mapping from natural numbers to even numbers. A more general picture of a one-to-one mapping looks like:

$$\begin{array}{cccccccc} 1 & 2 & 3 & 4 & 5 & 6 & \\ \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \cdots, \\ \mathrm{f}(1) & \mathrm{f}(2) & \mathrm{f}(3) & \mathrm{f}(4) & \mathrm{f}(5) & \mathrm{f}(6) & \end{array}$$

where $f(x)$ is our mapping function. Depending on the circumstances, instead of talking about the natural numbers $\{1, 2, 3, \ldots\}$ it may be more appropriate for us to talk about integers $\{\ldots, -1, 0, 1, \ldots\}$ or the positive integers $\{0, 1, 2, \ldots\}$. It really doesn't matter which of these sets we are using, because all of them have same number of elements; that is, they all contain a *countably infinite* number of elements.

2.3 Rational Numbers

A rational number (or fraction) is a number that can be represented as the ratio of two natural numbers, such as a/b, with the understanding that the denominator, b, is never zero. One limiting aspect of the natural numbers is that for any two natural numbers, there is only a finite number of natural numbers between them. This is not so for the rational numbers. To convince yourself of this, you only need to take the average of any two different rational numbers. For example, given a_1/b_1 and a_2/b_2, we can compute the arithmetical mean or average as $(a_1b_2 + a_2b_1)/(2b_1b_2)$. Call this average a_3/b_3. We can repeat the process as long as we like by taking the average of a_1/b_1 and a_3/b_3, then a_1/b_1 and a_4/b_4, and so on.

Notice that there is no such thing as the smallest nonzero rational number, which implies that we simply cannot enumerate all of them by size. However, we can construct a simple procedure to enumerate all of the rationals based on another method. To do this, we will consider only rational numbers between 0 and 1 at first (excluding 0 and including 1), which implies that $a \leq b$. We can construct a triangular matrix that contains all of the rationals between 0 and 1 by having one row per denominator. In row b, there are exactly b columns, one for each value of a with $a \leq b$. The first few entries of the table look like this:

$$\begin{array}{cccccc} \frac{1}{1} & & & & & \\ \frac{1}{2} & \boxed{\frac{2}{2}} & & & & \\ \frac{1}{3} & \frac{2}{3} & \boxed{\frac{3}{3}} & & & \\ \frac{1}{4} & \boxed{\frac{2}{4}} & \frac{3}{4} & \boxed{\frac{4}{4}} & & \\ \frac{1}{5} & \frac{2}{5} & \frac{3}{5} & \frac{4}{5} & \boxed{\frac{5}{5}} & \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{array}$$

The boxed fractions are repeats and can be removed from the table so that all entries represent unique rational numbers. Now, if we read the table left-to-right and top-down (as one would read a book), all of the fractions between 0 and 1 will eventually be encountered. Thus, we could map each fraction between 0 and 1 to an odd natural number:

$$\begin{array}{cccccccccccc} 1 & 3 & 5 & 7 & 9 & 11 & 13 & 15 & 17 & 19 & \\ \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \updownarrow & \cdots . \\ \frac{1}{1} & \frac{1}{2} & \frac{1}{3} & \frac{2}{3} & \frac{1}{4} & \frac{3}{4} & \frac{1}{5} & \frac{2}{5} & \frac{3}{5} & \frac{4}{5} & \end{array}$$

For rational numbers greater than 1, we know that $a > b$. By taking the reciprocal of such a fraction, b/a, we are left with a number that is strictly greater than 0 and less than 1. Therefore, by the same process that allowed us to map the small fractions to the odd numbers, we can map the large fractions to the even numbers.

This leads us to a startling conclusion: There are as many natural numbers as fractions! The most important point about our construction is that it is one-to-one and invertible. Specifically, if you wanted to play devil's advocate and claim that the mapping failed, you would have to produce two rational numbers that mapped to the same natural number or one rational number that mapped to no natural number. Based on our method of construction, we are guaranteed that this will never happen.

2.4 Irrational Numbers

Fractions are known as rational numbers because they can be expressed as the ratio of two natural numbers. Irrational numbers, such as π and $\sqrt{2}$, are numbers that cannot be represented as the ratio of two natural numbers. If we represent a number by its decimal expansion, we find that rational numbers have a finite or a periodic decimal expansion, while irrational numbers have an infinite decimal expansion that has no pattern. For example, the rational number $\frac{1}{3}$ has the decimal expansion $0.\bar{3}$, where the bar over the last digit signifies that the expansion repeats forever. Moreover, there are numbers such as $0.123456\overline{789}$ that are also rational because the last four digits repeat. Whenever a number's decimal expansion falls into a pattern, it is always possible to convert the decimal expansion into a fraction.

Taking the analysis one step further, rational numbers can also be represented as a summation of fractions, such as $0.123 = \frac{1}{10} + \frac{2}{100} + \frac{3}{1000}$. What about the repeating fractions? It turns out that the repeating fractions require an infinite summation, but this is not a problem for us because the infinite series converges to a rational number. We saw this when we solved Zeno's paradox and computed that the footrace between Achilles and the tortoise would be tied two minutes into the race.

Infinite series of this type reveal a quirky aspect of rational numbers. Specifically, for any rational number we can construct multiple decimal expansions that

Is the Square Root of 2 Really Irrational?	**Digression 2.1**

Here is a great proof that $\sqrt{2}$ is irrational. It was first discovered by Pythagoras around the fifth century B.C. The technique is called a proof by contradiction and starts off with the assumption that $\sqrt{2}$ is actually rational. By making this assumption, we will be faced with an impossibility, which implies that $\sqrt{2}$ is in fact irrational.

Now if $\sqrt{2}$ is rational, then it is equal to some fraction, a/b. Let's take the square of the fraction that we know is equal to 2. We now have the equality $a^2/b^2 = 2$. Multiply each side by b^2 to get $a^2 = 2b^2$. Here comes the tricky part: We are going to take advantage of the fact that every natural number has a unique prime factorization. Taking the prime factorization of a and b, we know that the prime factorization of a^2 must have twice as many 2s as the factorization for a. The same thing applies to b^2 and b. Therefore, the prime factorizations of a^2 and b^2 must have an even number of 2s.

Now, looking at the equation $a^2 = 2b^2$, we know that the left side has an even number of 2s while the right side has an odd number of 2s. One side will have more than the other. We don't know which side, but we don't care. If we take the product of the smaller number of 2s and divide each side of $a^2 = 2b^2$ by that number, then one side will have at least one 2 in it, while the other will have none. Since 2 is the only even prime number and an odd number multiplied by an odd number always yields an odd result, we know that the side with the 2s must be an even number while the side with no 2 is an odd number. A contradiction! Therefore, it is impossible for $\sqrt{2}$ to be expressed as a fraction.

are clearly different but are numerically equivalent. As an example consider the equivalence of $1 = 0.\bar{9}$. It may seem counterintuitive to state that 1 and $0.\bar{9}$ are equal, but in fact they are because $\frac{9}{10} + \frac{9}{100} + \frac{9}{1000} + \cdots = 1$. This is not some subtle flaw in the properties of numbers, but an artifact of the different ways we can represent them.

Another way to represent a number is as a point on a number line. We are all familiar with the process of labeling a number line and placing a point on it to represent a particular value. This is easy enough for natural numbers and rationals, but where would you put a point on a number line for an irrational number? For example, suppose we want to put a point on a number line for the value of $\sqrt{2} = 1.41421356\cdots$. We could approximate $\sqrt{2}$ with three digits and place a point at 1.41 on the number line. However, we know that $\sqrt{2}$ is really a little bit to the right of 1.41, so we go one step further and put another point at 1.414. Once again our estimate is a bit short of the true location. It would seem that we could continue the process indefinitely, always failing to put a point on the correct location.

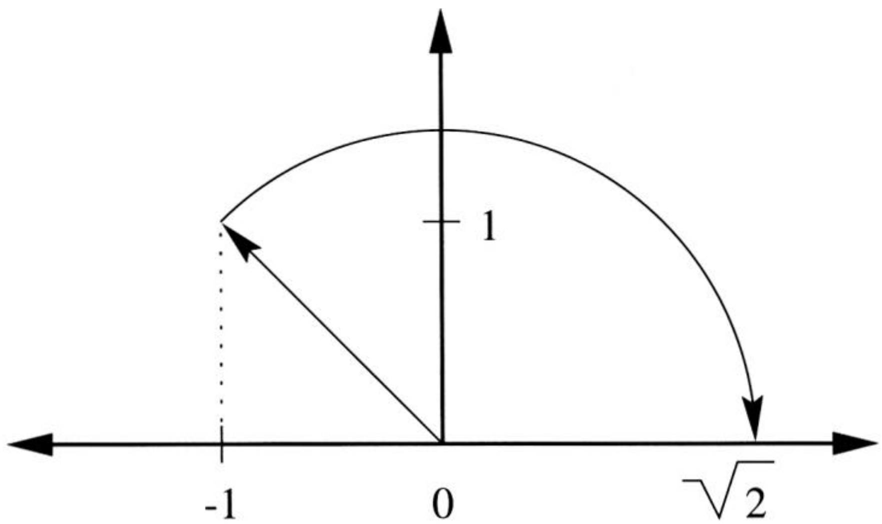

Figure 2.2 Isolating the square root of two on the number line

Does $\sqrt{2}$ really have a true location on the number line and can we find it? Our problem seems to be that $\sqrt{2}$ is always to the right of our best estimate, but we could approach the problem from another side, literally. Instead of using 1.41 as a first guess, we can use 1.42, which is just larger than $\sqrt{2}$. At the next step we use 1.415, and so on. Now we have $\sqrt{2}$ trapped. In fact, the infinite sequence of the $\sqrt{2}$ converges to a real location on the number line, just like the infinite series in Zeno's paradox. By approaching that point from each side, we can see that it can be isolated. Another method for isolating $\sqrt{2}$ is best illustrated with the diagram in Figure 2.2.

In Figure 2.2, I have constructed a triangle with two sides equal to 1 in length. We know that the third side must have a length equal to $\sqrt{1+1}$. Now take the arrow of length $\sqrt{2}$ and swing it around to the x-axis. Voila! We have found $\sqrt{2}$.

Up to this point, we found that natural numbers, integers, and fractions all contain the same infinite number of elements. Is there the same number of irrational numbers as well? Let's assume for the moment that there are. We will attempt to construct a one-to-one mapping, just as before, between the natural numbers and the irrational numbers to see what will happen. Below is a table with a natural number on the left side and some corresponding irrational number on the right side. Unfortunately, we cannot write out the full decimal expansion of the irrational numbers in the table, but this does not matter for our purposes. Moreover, instead of writing digits for the irrational numbers, we will use the notation x_{ij} to signify the jth digit of the irrational number that maps to i, and to keep things simple, we will worry only about irrational numbers between 0 and 1 so that each digit is to the right of the decimal point, as in $0.x_{i1}x_{i2}x_{i3}\cdots$.

$$\begin{array}{ccccccccc}
1 & \leftrightarrow & \boxed{x_{11}} & x_{12} & x_{13} & x_{14} & x_{15} & \cdots \\
2 & \leftrightarrow & x_{21} & \boxed{x_{22}} & x_{23} & x_{24} & x_{25} & \cdots \\
3 & \leftrightarrow & x_{31} & x_{32} & \boxed{x_{33}} & x_{34} & x_{35} & \cdots \\
4 & \leftrightarrow & x_{41} & x_{42} & x_{43} & \boxed{x_{44}} & x_{45} & \cdots \\
5 & \leftrightarrow & x_{51} & x_{52} & x_{53} & x_{54} & \boxed{x_{55}} & \cdots \\
\vdots & & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots
\end{array}$$

Boxes have been placed around the diagonal elements to highlight them. Now the important question is: Does this mapping work? Remarkably the mapping fails, but this is not a failure for us because we have found a deeper truth concerning the nature of irrational numbers. To see that the mapping fails, we must first agree that for it to work, there must be a place in the table for all of the irrational numbers. This seems reasonable, but consider an irrational number that is constructed as a sequence of digits that differ from the diagonal entries. If we represent this new number in the same way as the irrational numbers in our table, it will look like:

$$\text{NOT } x_{11} \quad \text{NOT } x_{22} \quad \text{NOT } x_{33} \quad \text{NOT } x_{44} \quad \text{NOT } x_{55} \quad \cdots.$$

If we try to find a place in our table for this irrational number, we are faced with the inevitable conclusion that our new number cannot exist in the table because it will always differ from each entry in the table by at least one digit. Specifically, if we assume it belongs in the table at, say, line 835, then by virtue of the way we constructed the number, the 835th digit must differ from the real entry at line 835.

We have just proved that the infinity of the irrational numbers is greater than the infinity of the natural numbers. In doing so, we followed in the footsteps of the great German mathematician Georg Cantor, who invented set theory and was the first to prove that not all infinities are equal. Since the size of the natural numbers is countably infinite, the size of all of the real numbers is properly referred to as *uncountably infinite*.

The next property of numbers that we will examine in this chapter concerns the density of the real number line. Let's start by considering only a mapping between the positive real numbers into a segment between 0 and 1. Including the negative real numbers or using a different segment is just as easy, but the examples will be clearer if we use this restriction. It turns out that any segment of finite length on the real number line is infinitely dense, in that you can squeeze all of the real numbers into it. As before, this is accomplished with a mapping function with a few special properties. The first property of our function is that it must be monotonically increasing. A monotone function is one that strictly increases or decreases but never both. This is a formal way of saying that the mapping function has no "humps." Readers familiar with calculus will recognize this property as a way of saying that the first derivative of a function is never 0. Further specifying that the function is increasing means that for any two arguments, a and b, if $a > b$,

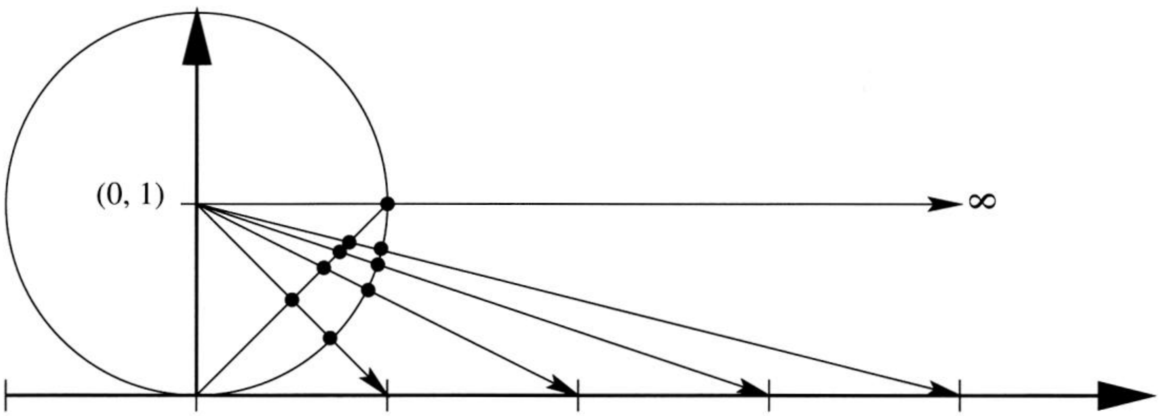

Figure 2.3 Mapping real numbers into a circle or a segment

then for the mapping function, $f(x)$, $f(a) > f(b)$. The second property is that the mapping function must asymptotically approach a constant value as the function argument approaches infinity. This means that if we increase the function argument, the function result will get larger, but each further increase will increase the function result by only a little bit more each time. The function result will get infinitely close to a constant value of our choice but never quite reach it. This is very similar in spirit to how an infinite series converges to a constant value.

Figure 2.3 demonstrates two mapping functions that meet our requirements. To see how they work, consider a vector starting at the coordinate (0, 1) and ending at some value on the number line. The number at the end of the vector is the one that we want to map into another value between 0 and 1. As the vector travels to its destination, it always passes through a fixed location on the unit circle centered at (0, 1). We can put a mark on the circle where the vector passes through. Notice that no matter how large the vector is, it will always hit somewhere in the bottom right quarter of the circle. Therefore, the x-coordinate of the marked points will always be between 0 and 1. In this way we have mapped all of the numbers on the number line to a corresponding point on the surface of the circle. Doing this type of mapping in higher dimensions is known as constructing a Riemann sphere.

Instead of mapping the vectors to a point on the surface of a circle, we can just as easily map them to a line segment that extends from (0, 0) to (1, 1). Figure 2.3 also shows this type of mapping, which is a bit easier to understand analytically. The line equation of the vector is $y = 1 - \frac{x}{c}$, where c is the number on the number line that we are mapping from. The line equation of the segment from (0, 0) to (1, 1) is simply $y = x$. By combining these two equations and solving for x, we end up with the equation $y = \frac{x}{x+1}$. Thus, if we plug any positive real value into the last equation, we will always get another number between 0 and 1. It is often tempting to try to imagine a size for infinity—but doing so is actually misleading. Using Figure 2.3 as a reference allows us to think of infinity as the direction, pointing due east.

Let's pause for a moment to consider what we have found so far. The natural numbers and integers occupy fixed locations on the number line at regular intervals. We know that there is the same number of rational numbers, but strangely enough there is a countably infinite number of rationals between any two points on the number line. We also isolated an irrational number on the number line, so we know that irrational numbers have a fixed place on the line as well. What is truly bizarre is that even though there is an infinite number of rational numbers between any two rational numbers on the number line, there are infinitely more irrational numbers than rational numbers in the same segment. Going back to the example of isolating $\sqrt{2}$ on the number line, there was a countably infinite number of rational number pairs that we could use to get increasingly closer to the real location of $\sqrt{2}$. Yet between any of those pairs of rational numbers there is an uncountably infinite number of irrational numbers. It is as if the points on the number line where the irrational numbers fall are like holes in which rational numbers are not. However, there are many, many more of these irrational holes than non-holes.

To fully appreciate how truly amazing rational and irrational numbers are, we will now use an irrational number as a sort of infinite memory. Because you are reading this book, I am going to assume that you have some general idea of how computers work. Specifically, you probably already know that computers store everything as one long sequence of zeros and ones. You also know that the number of natural numbers (or fractions) that you can store in your computer is determined solely by how much memory you have; that is, if you want to store twice as many numbers as you can now, you need to double your computer's memory. Therefore, no matter how many numbers you have, no matter how large the numbers are, as long as they are natural (or rational) and finite in size, there is some amount of memory that will do the job of storing the numbers.

Now, take all of your computer's memory and arrange it as one long line of zeros and ones: 0, 1, 1, 1, 0, 0, 0, 1, 1, 0, 1, ... Take this very long number and put a zero and a decimal point in front of it. We've just translated one huge number into a rational number between 0 and 1. By placing this single point at exactly the right spot on the number line, we can store an unlimited amount of information. Ah, if only it were so simple. In the real world, we simply don't have the precision required to put this method of storing memory into practice. We never will, either, but it's an interesting mental exercise to see that it can be done in theory in an idealized world. The point of this whole mental exercise is that in many ways an irrational number has as much "information" as an infinite number of natural numbers.

Figure 2.4 illustrates some of the properties that we have found so far about numbers. The set of all of the numbers on the number line is called the real numbers. The reals can be divided into rational numbers and irrational numbers. A proper

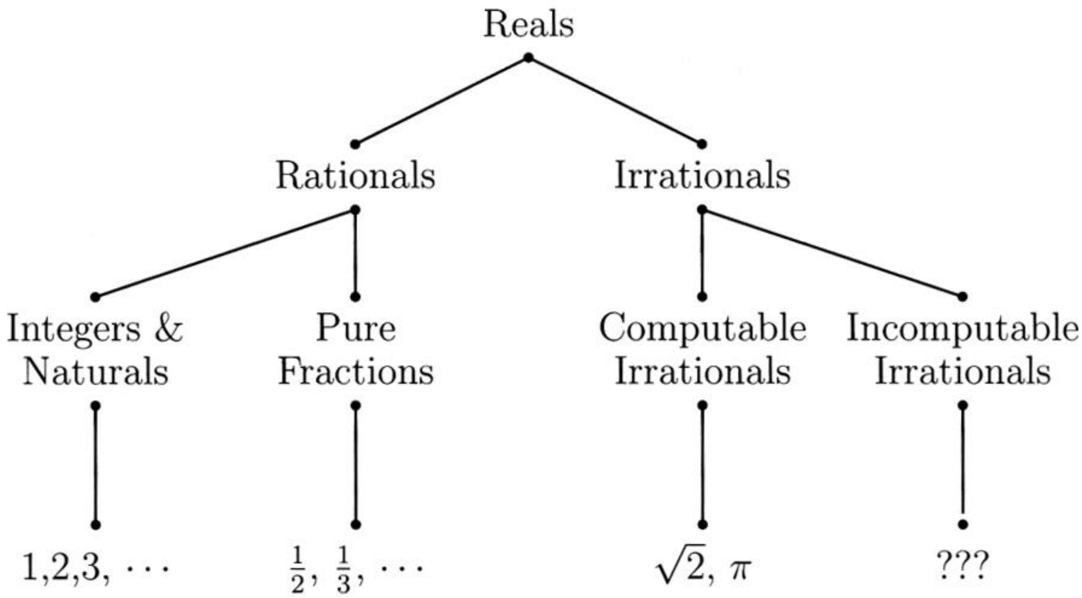

Figure 2.4 Subdivision of the real numbers

subset, A, of another subset, B, is a set such that all of the members of A are also members of B, but there exists some member of B that is not a member of A (for example, cars are a proper subset of vehicles). We can now say with certainty that the natural numbers are a proper subset of the rationals. Similarly, the irrationals have a proper subset known as the computable irrationals. You are already familiar with some computable irrational numbers, such as π and $\sqrt{2}$. In the next chapter we will define what it means to compute, which will give us some insight into what uncomputable irrational numbers look like.