Something from Nothing

Tom Davis
tomrdavis@earthlink.net
http://www.geometer.org/mathcircles
October 24, 2000

There are many different areas of mathematics, but the great majority can be constructed from “almost nothing”. We will show here exactly how that is done.

The foundations of almost all of mathematics can be solidly constructed from nothing more than set theory. One has to be a little careful with what, exactly, is meant by set theory, but it is nice to know that in a sense, everything can be made to depend on a theory that is so “simple”[1].

The topics covered in this paper can easily fill a full semester course at a university, so obviously we are omitting large amounts of material, and are presenting only those topics that lead in a sort of “bee line” to a fairly simple construction—that of the natural numbers: $\mathbb{N} = \{0, 1, 2, 3, 4, 5, \ldots\}$.

1 Naïve Set Theory

A great deal of set theory can be understood solely based on an intuitive understanding of the subject. We’ll begin with that and see what we can learn, but we will be careful to notice when and why there are problems.

Introductory (naïve) set theory basically considers a “set” to be a “collection” of “objects”. For example, we can think of the set of all elephants, or the set of even numbers, the set of English language sentences, or the set of all sets. Now of course it is difficult to give a precise mathematical definition of an elephant, but “the set of all elephants” will cause fewer problems than one of the others.

But let’s begin without worrying too much about the problems. If we simply consider a set to be a collection of objects, there are some obvious operations that we might want to perform on sets or combinations of sets.

1.1 Notation

Although most readers will be familiar with the following, what follows is a quick review of the basic notions and definitions of (naïve) set theory:

- **Set notation.** We will indicate a set with a pair of braces: “$\{$” and “$\}$”. $\{A, B, C\}$ represents the set containing the three elements A, B, and C (whatever A, B, and C may be). In cases like this (relatively small finite sets), it is simple just to list all the elements in a set.

  The order in which the elements are listed is unimportant: $\{A, B, C\} = \{B, A, C\} = \{C, B, A\}$, et cetera. Also, an object is either a member of a set or it is not. A set cannot, for example, contain two copies of the same object, so $\{A, A\}$ is just an inefficient way of writing $\{A\}$.

[1] Unfortunately, set theory turns out to be not at all “simple”, but at least it is just a single theory

When the meaning is obvious, we can even indicate large (or even infinite) sets using a similar notation. For example, to indicate the set consisting of all the integers from 1 and 1000000, we might write: $\{1, 2, 3, \ldots, 1000000\}$. Or to indicate the set consisting of all the positive multiples of 3 (an infinite set) we might write: $\{3, 6, 9, 12, \ldots\}$.

Sometimes it is best to describe the elements of a set in terms of some property that they satisfy. The following is a description of all the prime numbers:

$$\{x : x \text{ is prime}\}. \tag{1}$$

You can read the ":" character as "such that", so that the set in (1) can be read as "the set of all x such that x is prime". Of course, this assumes that you know what is meant by "x is prime".

- **Membership.** If you want to indicate that an object is a member (is an element of) a set, use the symbol $\in$. Thus we can write $x \in S$, or $A \in \{A, B, C\}$ to mean "x is a member of S", or in the second example, "A is a member of the set $\{A, B, C\}$".

  If we want to indicate the opposite (that a particular object is not a member of a particular set), use the symbol $\notin$, so we have $D \notin \{A, B, C\}$, and $15 \notin \{x : x \text{ is prime}\}$.

- **Union and intersection.** We can indicate the union and intersection of two sets with the symbols $\cup$ and $\cap$, respectively. If A and B are two sets, then $A \cup B$ indicates the set "A union B"—the set that consists of all the elements that are either in set A or in set B. Similarly, $A \cap B$ is read "A intersection B", which is the set of all elements that are in both A and B.

  For example, if $A = \{1, 3, 5, 7, 9\}$ and $B = \{5, 7, 9, 11, 13\}$, then $A \cup B = \{1, 3, 5, 7, 9, 11, 13\}$, and $A \cap B = \{5, 7, 9\}$.

- **Set difference.** We can "subtract" one set from another using the symbol "$-$". If A and B are sets, then $A - B$ represents the set consisting of all the elements that are members of A and are not members of B. This is a little different from the subtraction you may be used to in arithmetic, since B may contain elements that are not members of A. For example, if $A = \{1, 2, 3, 4\}$ and $B = \{3, 4, 5, 6, \ldots\}$, then even though B is an infinite set, $A - B$ still makes sense. $A - B = \{1, 2\}$—the set of all the items that are in A and are not in B.

- **Subset, equivalence.** If set B contains every element of set A, then we say that $A \subset B$ (read, "A is a subset of B". $A \subset B$ is equivalent to: "for every x, if $x \in A$, then $x \in B$. We say sets A and B are equivalent ($A = B$), if they contain exactly the same elements. The "axiom of extensionality" that we'll see later says that if $A \subset B$ and $B \subset A$, then $A = B$. Sometimes we write $B \supset A$ to mean exactly the same thing as $A \subset B$.

1.2 Examples

One very important set is the set with nothing in it. This set is called the "empty set", and is often denoted by the Greek letter ϕ: $\phi = \{\}$. It is important to realize that the empty set is not "nothing"—it is the set that contains nothing. If you think of a set as a box, and the elements of a set as the items you find in the box, then the empty set is like an empty box.

Another important idea is that the objects in a set can be almost anything, including other sets. The set

$$\{A, 1, 7, \{8, 9\}\}$$

is a set that contains four objects: it contains A, 1, 7, and a set—the set that contains 8 and 9.

The set $\{\{\}\}$ is also interesting. It is a set with one thing in it, and that thing is the empty set. It is *not* an empty box; it is a box that contains one thing—an empty box. Perhaps this is clearer if we write it in another way: $\{\phi\}$.

How many objects does the following set contain?

$$\{A, B, C, \{A, B\}, B, \{B, A\}\}.$$

The answer is 4. It contains A, B, C, and the set $\{A, B\}$. Note that B is listed twice in the set, and that since the order of listing elements is not important, $\{A, B\} = \{B, A\}$, so in effect that member is listed twice as well.

1.3 Problems with Naïve Set Theory: Russell's Paradox

For a long time people used the concepts of naïve set theory as listed above in fairly sloppy ways, and were able to make quite good use of them, but some very nasty problems began to emerge, and getting rid of those nasty problems turned out to be far more difficult than anyone expected.

Perhaps the most famous problem is called "Russell's paradox", named after Bertrand Russell, the man who finally showed a way around the problem.

Here's the basic idea. We've seen above that sets may contain other sets, and there is no reason that those sets they contain cannot contain still other sets, and so on. In fact, it might be useful to consider the set that contains all sets—it would be very large, but it would be useful, too. In fact, since it is a set, it would contain itself—it is the most obvious example of a set that contains itself, but there are lots of other examples of sets that contain themselves (in naïve set theory, that is).

So some sets (like the set consisting of all sets) contain themselves, and other sets (like almost all the other examples in this paper so far) do not contain themselves. A set either contains itself or it does not, so let's look at the following interesting set S:

$$S = \{x : x \text{ is a set and } x \notin x\}.$$

In other words, S is the set of all sets that do not contain themselves.

So the obvious question is, does S contain itself? Suppose it does. Well, then it must not be in S, since S consists of only those sets that do not contain themselves. So S must not contain itself. But then it must contain itself, since S is defined to include all the sets that do contain themselves. So if S contains itself, then it doesn't, and if it doesn't contain itself, then it does. It seems to be a hopeless contradiction.

There are a couple of other paradoxes that are related, but are not stated in such a strict mathematical form as the paradox above:

- In Seville, there's a barber who shaves all those people who do not shave themselves. Does the barber shave himself or not? This is known as the "Barber of Seville problem".

- Imagine a card. On one side is written, "The statement on the other side of this card is true." and on the other side is written, "The statement on the other side of this card is false."

Bertrand Russell, one of the most famous logicians ever, struggled with this problem for a long time. In his autobiography, he describes just how hard he found the problem. Every morning, he said, he would sit down at his desk with a blank piece of paper in front of him. At the end of the day, he would still be staring at the same blank sheet of paper.

Russell's final resolution to the problem is described in his "Principia Mathematica", written with Alfred North Whitehead, in which he introduced a "Theory of Types" to get around his paradox. The basic idea was this: sets cannot contain themselves. In fact, one has to be very careful about exactly what is a set and what is not a set. Sets can be built up from more primitive objects, but only in a very careful and controlled way. You cannot just say things like, "Consider the set of all sets." This is not a set at all, at least according to Russell.

2 Axiomatic Set Theory

Russell's "Principia Mathematica" is a difficult read, to say the least. Today we are lucky to have a much easier to understand method for dealing with set theory in a way that does not seem to lead to any contradictions[2].

What we'll do in this section is look at set theory based on a set of axioms, just like the rest of mathematics. There are various ways to do this, but perhaps the most intuitive is from the so-called Zermelo-Fraenkel axioms for set theory (sometimes called the "Z-F axioms" for short).

2.1 Making Language Precise

Since we got into trouble in the first place because of some sloppiness with language, one of the best ways to begin to study sets is with a precise description of the language that will be used to discuss them.

A natural language—any natural language: English, German, Italian, Russian, Chinese, Arabic, or anything else—has the same basic problem; none was designed specifically for precision, and hence it is easy to introduce ambiguity using any natural language as a starting point. Of course it's a lot of trouble to learn a new language, so the strategy we'll take in this paper is to show how such a new artificial, precise language can be put together. Then, rather than force you, the reader, to struggle through the rest of this paper using it, we will show a couple of examples of its use, but then return to English with the understanding that without too much trouble the statements in English can be converted to the more formal language.

What follows is one possible way to define a language suitable for talking about mathematics.

2.2 A Formal Language

First, we'll begin with a description of the sorts of symbols that can be used in such a language. Such symbols fall into the following categories:

- **Punctuation.** We will use parentheses and commas for various grouping operations. The more obvious uses will be to indicate the order of operation: $(3/6)/7$ indicates that 3 is to be divided by 6 first, and the result divided by 7. $3/(6/7)$ means that 6 is to be divided by 7, and then 3 is divided by the resulting number.

[2] We can only say "does not *seem* to lead to contradictions" because nobody knows for certain that it does not. There is a lot of evidence that it does not, but (according to another famous theorem of Kurt Gödel) it is impossible to prove that it does not.

Similarly, $F(x)$ will mean that the function (or predicate) F is applied to the variable x, $G(x, y, z)$ indicates that the function (or predicate) G is applied to three variables, x, y, and z.

- **Variables.** We need some way to talk about the objects of interest. In set theory, these objects will be sets; in number theory, they will be integers; in functional analysis, they will be functions.

  For these objects, we will use lower-case letters: a, b, c, $\ldots$. If we need more than 26 of them, we'll use subscripts: a_1, b_{17}, or z_{14641}, for example.

- **Constants.** In many fields, it is convenient to give specific names to objects that are extremely important. They will be used in exactly the same ways as the variables described above, since they stand for the same things—objects that the language is talking about.

  In set theory, for example, we use ϕ to indicate the empty set. In number theory, we might use $\mathbb{N}$ to indicate the set of natural numbers, and $\mathbb{Z}$ for the set of integers.

- **Predicates.** We will use the term "predicate" to stand for a property that variables may or may not satisfy. For example, in number theory, we may want to look at numbers to see if they are prime, or even, or that two numbers are relatively prime. We could use $P(x)$ to mean that x is prime; $O(x)$ to indicate that x is an odd number, and $R(x, y)$ to indicate that x and y are relatively prime.

  These predicates will have a truth value depending on the values of the variables. Continuing the example above, $P(7)$ is true, $P(9)$ is false, and $R(17, 43)$ is true.

  We will use upper-case letters for predicates like $A, B, C, \ldots$. As was the case with variables, if we run out of upper-case letters, we'll use subscripts: A_{99}, or X_{451}.

  There are a couple of predicates that are used so often that they have special symbols: $=$ (equality), and in set theory: $\in$ (is a member of). Note that we could use $E(x, y)$ to mean "x is equal to y", but we are so used to writing $x = y$ that it seems foolish to add another level of complexity. Similarly, there could be a predicate M, standing for "member", and we could write $M(x, s)$ to mean "x is a member of the set s". But $x \in s$ is much easier.

- **Logical operators.** Since the formulas and sentences of our language will generally have truth values (true or false—T or F), we need some method to combine truth values. This is done with logical operators. A variety of these could be used, but here we'll stick with the more common and useful ones: $\neg$ (not), $\wedge$ (and), $\vee$ (or), $\Rightarrow$ (implies), and $\Leftrightarrow$ (is equivalent to).

  All but $\neg$ (the "not" operator), combine two formulas. Not simply reverses the truth value, while the others have the following "truth tables", where the values of A are the column headers and the values of B label the rows:

$A \vee B$	F	T
F	F	T
T	T	T

$A \wedge B$	F	T
F	F	F
T	F	T

$A \Rightarrow B$	F	T
F	T	F
T	T	T

$A \Leftrightarrow B$	F	T
F	T	F
T	F	T

- **Quantifiers.** Finally, there are two quantifiers, $\exists$ (there exists) and $\forall$ (for all) that are used together with a variable to quantify the rest of the sentence.

  For example, $\forall s((\neg(s = \phi)) \Rightarrow (\exists x(x \in s)))$ is a statement in set theory that can be read, "For every s, if s is not equal to the empty set (ϕ), then there exists an object x such that x is a member of s. This statement is true of the theory of sets, but doesn't even make sense in other fields since the predicate $\in$ only makes sense in set theory.

2.3 Formal Language Grammar

There is, of course, a grammar associated with the language that describes the rules for forming sentences that are syntactically correct (and whose truth values can be evaluated). We will not specify a complete set of rules here; we will merely illustrate the general flavor of those rules by means of a small number of examples:

- $\forall x \forall y (x = y)$.

  For all x and for all y, x is equal to y. This sentence will only make sense in a system that contains zero or one objects, but the sentence is correctly formed, and evaluates to a truth value (which is usually false).

- $\forall x (O(x) \vee E(x))$.

  For all x either $O(x)$ is true or $E(x)$ is true. This is a reasonable and true statement about number theory, if $O(x)$ happens to mean that "x is odd", and $E(x)$ means that "x is even".

- $\forall P \forall x P(x)$.

  This is not a legal grammatical construct—the quantifiers can only quantify over variables, not over predicates.

- $\exists s \forall x (\neg(x \in s))$.

  "There exists an s such that for all x, it is not true that x is a member of s." This is an actual axiom of set theory, and states the existence of the empty set, ϕ.

- $\forall s \forall t (((\forall x ((x \in s) \Rightarrow (x \in t))) \wedge (\forall x ((x \in t) \Rightarrow (x \in s)))) \Rightarrow (s = t))$.

  This is another axiom from set theory, the "axiom of extensionality". It states that if s is a subset of t and if t is a subset of s, then $s = t$". But the "subset" operator isn't part of the language, so we use a phrase like $\forall x ((x \in s) \Rightarrow (x \in t))$ to mean that $s \subset t$.

Don't worry about the formal rules of this grammar—the point of this section is simply to convince you that such a formal grammar could be described. If you really want to dig into the details, look in any book on axiomatic set theory, or on formal logic.

In the sections that follow, we'll usually include both the formal statement and the (roughly equivalent) English statement to describe the axioms. But remember that all the axioms for set theory (and for almost any other field of mathematics) can be written in the same general way as shown in the valid examples above.

3 The Zermelo-Fraenkel Axioms of Set Theory

All of set theory can be based on a single predicate, $\in$.

- **The existence of the empty set.**

  $$\exists s \forall x (\neg(x \in s)).$$

  There is a set s so that for every item x, x is not in s. In other words, s contains nothing, but the first "$\exists s$" tells us that the set s itself exists.

Without too much difference this axiom can be replaced by the simpler axiom: $\exists x(x = x)$ which simply implies the existence of *some* set.

- **The axiom of extensionality.**

$$\forall z((z \in x \Leftrightarrow z \in y) \Rightarrow x = y).$$

What this axiom states is that if you are given two sets x and y, if every item z is either in both of them or in neither or them, then x and y are the same set. This is sometimes stated as, "if x is a subset of y and if y is a subset of x, then $x = y$.

- **The pairing axiom.**

$$\forall u \forall v \exists x \forall z((z \in x) \Leftrightarrow (z = u \vee z = v)).$$

This axiom states that given any two objects there is a set that contains exactly those two objects and no others.

- **The subset axioms.**

$$\exists y \forall x((x \in y) \Leftrightarrow ((x \in z) \wedge \mathcal{A}(x))).$$

There are two strange things about this. First, what is meant by $\mathcal{A}(x)$ (after all, the symbol $\mathcal{A}$ wasn't ever described in our discussion of the grammar), and why do we say "axioms" instead of "axiom"?

The two questions are related. The $\mathcal{A}(x)$ stands for *any* valid formula in the language that involves the variable x, and for every possible such formula, there is another axiom. Thus, there are an infinite number of subset axioms. All the axioms have the same form, but different formulas replacing the $\mathcal{A}(x)$. Such a collection of related axioms is often called an "axiom schema".

The idea behind it simply this: If you already have a set (called z in the axiom schema above), then for any description available in the language of a property that some variable may satisfy, the set of all objects that satisfy that condition make up a set.

For example, if you know that the natural numbers $\mathbb{N} = \{0, 1, 2, 3, \ldots\}$ form a valid set (we don't know this yet in our formal scheme so far), then for any property that natural numbers might satisfy, there exists a set of numbers satisfying that property.

To make the example concrete we can show that the set consisting of all even numbers exists. Simply let $\mathcal{A}(x)$ be the following: $\exists y(y + y = x)$ (where we assume you know what is meant by "$+$") in arithmetic. Then $\mathcal{A}(x)$ is simply a way to say "x is even".

- **The sum axiom.**

$$\forall x \exists y \forall z \forall w((w \in z \wedge z \in x) \Rightarrow w \in y).$$

This axiom allows us to form the union of sets. Note that this is an arbitrary union—a union of any number of sets. In the statement above, x is a collection of sets, and y is the union of all the sets in x. You can, of course, take the union of two particular sets (say x and y), since the pairing axiom shows the existence of a set u containing those two sets ($u = \{x, y\}$), and this sum axiom will then let you take the union of all the sets in u, which will amount to $x \cup y$.

- **The power set axiom.**

$$\forall x \exists y \forall w (w \subset x \Rightarrow w \in y).$$

  In other words, for any set x there is a set y that consists of all the subsets of x. Notice that the symbol "$\subset$" was used above, although it's not officially part of the language. This is done more and more frequently, and is valid, since we can replace "$w \subset x$" by "$\forall t(t \in w \Rightarrow t \in x)$".

  This may seem like an innocent axiom, but it is wildly powerful, and (as we will see later) posits the existence of huge sets. Even if we begin with a relatively small set consisting of only n elements, that set has 2^n subsets, so this power set axiom guarantees a set of size 2^n. If we simply begin with the empty set ϕ which has zero elements, the power set of ϕ has one element. The power set of the power set of ϕ has 2 elements; another application of the power set gives one with $2^2 = 4$ elements, and if we continue, we have sets of size $2^4 = 16$, then $2^{16} = 65536$, and $2^{65536} =$ a number with 19728 digits.

- **The axioms of replacement.**

$$(\forall x \forall y_1 \forall y_2 (\mathcal{A}(x, y_1) \wedge \mathcal{A}(x, y_2)) \Rightarrow (y_1 = y_2)) \Rightarrow \forall s \exists t (\forall x \forall y ((x \in s) \wedge \mathcal{A}(x, y)) \Rightarrow y \in t).$$

  The formal statement above is probably a bit tricky to understand, but what is says is this. If you have a set s and a function whose domain is elements of s, then there is a set that contains the range of the function. First, notice that this is an axiom schema, since $\mathcal{A}$ can be any valid formula in the language.

  In the formula above, think of the $\mathcal{A}(x, y)$ to mean that the function maps x to y. Then the part that says: "$\forall y_1 \forall y_2 (\mathcal{A}(x, y_1) \wedge \mathcal{A}(x, y_2))) \Rightarrow (y_1 = y_2))$" means that if y_1 and y_2 are both images of x, then they must be the same. This is how a function can be defined in this language[3].

  The second part of the formula says that if s is a set, then there is a set t that contains every y such that $\mathcal{A}(x, y)$ is true for some x in s.

- **The axiom of infinity.**

$$\exists z (\phi \in z \wedge \forall u (u \in z \Rightarrow \{u\} \in z).$$

  Again, we've been a little sloppy, and have included both the symbol ϕ and the set-defining braces "$\{$" and "$\}$". Both can be eliminated with a more complicated expression. For example, the "$\phi \in z$" can be replaced by "$\exists w((w \in z) \wedge \forall q \neg (q \in w))$". Replacing "$\{u\}$" is left as an exercise.

  What the axiom means is that there exists a set with an infinite number of members. We can see that ϕ is a member, so the second part tells us that $\{\phi\}$ is also a member. But since $\{\phi\}$ is a member, so also is $\{\{\phi\}\}$. And by similar reasoning, so is $\{\{\{\phi\}\}\}$, $\{\{\{\{\phi\}\}\}\}$, $\{\{\{\{\{\phi\}\}\}\}\}$, and so on.

There are usually two other axioms in Zermelo-Fraenkel set theory, but they are a bit technical, so we will not discuss them here. They are quite interesting, however, and you may want to look them up. They are usually called "the axiom of regularity", and "the axiom of choice".

[3]Functions are described in more detail in Section 5.

4 The Natural Numbers

Finally we have enough set-theoretic machinery to define the natural numbers solely in terms of sets. What we mean by the natural numbers is the set $\mathbb{N} = \{0, 1, 2, 3, 4, \ldots\}$. Of course at this point, we have no idea what the symbols "0", "1", "2", and so on mean. That's the purpose of this section.

There are lots of ways to do this, but most people have settled on a way that is fairly nice. Before we present the construction, let's look at two properties that must be satisfied, and also a few properties that would be very nice:

- **Required.** The natural numbers must form a set. This is obviously required or we won't be able to do much with them.
- **Required.** Each number must be a set. If not, we'll have to suppose the existence of something that's not a set, and the Z-F axioms do not guarantee the existence of anything other than sets.
- It should be easy to look at the set-theoretic version of a number and figure out what number it is.
- If would be nice if it's easy to compare the size of numbers by looking at a simple set-theoretic operation.
- It would be nice if the set corresponding to the natural number n contains n elements. This is perhaps the most important property, since, as we will see, it is easy to compare the sizes of two different sets by finding a function that matches them up, so for any other (finite) set we may come across, we can count the number of elements in it by finding a natural number set that can be matched with it, element by element.

The infinite set guaranteed by the axiom of infinity would serve as a model for the natural numbers. It is a set, and all its members are clearly sets. It is also easy to figure out what number it corresponds to—simply look at the number of nested parentheses. If a and b are two such sets, then we know that $a < b$, $a = b$, or $a > b$ if (as sets) $a \subset b$, $a = b$, or $a \supset b$, respectively.

The only thing wrong with this model is that every set in this version of the natural numbers except zero will have exactly one element.

But here is a method that does exactly what we want:

1. Let 0 (zero) be the empty set.
2. If we know the set-theoretic representation for n, the representation for $n + 1$ is $n \cup \{n\}$.

That's it! But let's see what this amounts to:

- From the first rule, we know that:
$$0 = \{\} = \phi.$$
- Since we now know what 0 is, apply rule 2 to obtain
$$1 = 0 \cup \{0\} = \phi \cup \{0\} = \{0\} = \{\{\}\} = \{0\}.$$
- Since we now know what 1 is, use rule 2 to get 2:
$$2 = 1 \cup \{1\} = \{\{\}\} \cup \{\{\{\}\}\} = \{\{\}, \{\{\}\}\} = \{0, 1\}.$$

- To get 3:

$$3 = 2 \cup \{2\} = \{\{\}, \{\{\}\}\} \cup \{\{\{\}, \{\{\}\}\}\} = \{\{\}, \{\{\}\}, \{\{\}, \{\{\}\}\}\} = \{0, 1, 2\}.$$

- And so on: $4 = \{0, 1, 2, 3\}, 5 = \{0, 1, 2, 3, 4\}, \ldots, n = \{0, 1, 2, \ldots, (n-1)\}$.

These are, of course, just the numbers. We haven't given any indication of how addition, multiplication, et cetera, can be defined in terms of them. For now, take it on faith, and we'll look at some the machinery that is used to make such definitions.

5 Functions

Perhaps more important even than the natural numbers (at least from the point of view of mathematics in general) is the concept of a function.

Recall that a function is a rule that assigns outputs uniquely to inputs. One useful way to think of a function is as a machine that takes values as inputs and generates an output that depends only on the input. If the same object is put in twice, the same output will be generated. (The same output can come from two inputs—the only requirement is that the same input always generates the same output.)

From your algebra class, you probably recall functions like $f(x) = x^2$, or $g(x) = |x|$ (the absolute value of x). In the first case, the function f takes a number as input and multiplies it by itself to make the output. In the second case, g returns the input value as it output if it is positive, and otherwise multiplies it by -1 before returning it.

But what we would like to do is to define a function as a set. In fact, this is just the tip of the iceberg—we are basically going to define every object and operation in mathematics as a set (That's of course where the title of this paper comes from—all of mathematics can be based solely on complicated sets that are built of nothing but the empty set.). What's more, we are going to use only "pure" sets that are built up from nothing but the empty set ϕ in the same sort of way that we built up the set $\mathbb{N}$ of natural numbers.

5.1 Definition of a function

Since we're going to think in terms of sets, we'd better stop thinking about a function as a machine. The following is a very set-theoretic way to define a function.

Suppose we know how to make an ordered pair of sets (we'll see how to do this later). We will write (x, y) to indicate the ordered pair whose first element is x and whose second element is y. Because it is an ordered pair, the only time $(x, y) = (y, x)$ is if x happens to be the same as y.

A function is just a set of ordered pairs where the first element of each ordered pair is the input to the function, and the second is the output. If we restrict the example function $f(x) = x^2$ above to the natural numbers, then we can write f as:

$$f = \{(0, 0), (1, 1), (2, 4), (3, 9), (4, 16), (5, 25), \ldots\}.$$

Of course not every collection of ordered pairs make a function. For example, $\{(0, 1), (0, 2)\}$ is not a function, since such a function would have two different outputs for the same input, 0.

We can write formally the fact that the set f is a function as follows:

$$\forall v \forall w (v \in f \wedge w \in f \wedge v = (x_1, y_1) \wedge w = (x_2, y_2) \wedge x_1 = x_2) \Rightarrow y_1 = y_2.$$

This simply says that if two ordered pairs are in the same function and they share the first element (the input to the function), then they have the same output.

5.2 Definition of an ordered pair

Given any two sets x and y, what do we mean by the "ordered pair" (x, y)? We can't just put them in a set like this: $\{x, y\}$, since $\{x, y\} = \{y, x\}$—there is no ordering implied in a set.

Here is the usual way it is done. Define $(x, y) = \{\{x\}, \{x, y\}\}$. On the surface, it looks like an ordered pair is always a set consisting of two sets, but one of those sets contains one object (the first entry in the ordered pair), and the other contains two objects, and the "other one" is the second element of the ordered pair.

In the previous paragraph, we said, "on the surface". This definition is a little deeper than it appears. First, what represents the ordered pair (x, x)? Well, it is $\{\{x\}, \{x, x\}\} = \{\{x\}, \{x\}\} = \{\{x\}\}$. Is this going to cause problems? Why not?

Would something like this work? Define $(x, y) = \{x, \{y\}\}$. What's wrong with this?

Try to think of other possible ways to define an ordered pair, and you'll find that it is trickier than you think.

But the bottom line is this: once we do have a good definition for an ordered pair in terms of sets, we can give a good definition for a function in terms of sets.

5.3 Trivial example of a function

Although it looks simple, and is easy to understand, since everything is built upon the empty set, the actual representation of more interesting structures can be quite complicated. Here, using what we have defined above, is the complete listing for a very simple function f that maps $0 \rightarrow 1$, $1 \rightarrow 2$, and $2 \rightarrow 3$.

First, in terms of ordered pairs, here it is:

$$f = \{(0, 1), (1, 2), (2, 3)\}.$$

Now replace the three ordered pairs by what their set-theoretic equivalent (and we'll put one on each line to make it clear what is happening):

$$\begin{aligned} f \quad = \quad & \{ \\ & \{\{0\}, \{0, 1\}\}, \\ & \{\{1\}, \{1, 2\}\}, \\ & \{\{2\}, \{2, 3\}\}, \\ & \}. \end{aligned}$$

Finally, remember that $0, 1, 2$, and 3 are also sets:

$$0 \quad = \quad \{\},$$

$$
\begin{aligned}
1 &= \{\{\}\}, \\
2 &= \{\{\},\{\{\}\}\}, \text{ and} \\
3 &= \{\{\},\{\{\}\},\{\{\},\{\{\}\}\}\},
\end{aligned}
$$

so we have:

$$
\begin{aligned}
f = \{ & \\
& \{\{\{\}\},\{\{\},\{\{\}\}\}\}, \\
& \{\{\{\{\}\}\},\{\{\{\}\},\{\{\},\{\{\}\}\}\}\}, \\
& \{\{\{\{\},\{\{\}\}\}\},\{\{\{\},\{\{\}\}\}\},\{\{\},\{\{\}\},\{\{\},\{\{\}\}\}\}\}\} \\
& \}.
\end{aligned}
$$

Or, to write it in a more confusing manner,

$$
\begin{aligned}
f = \; & \{\{\{\{\}\},\{\{\},\{\{\}\}\}\},\{\{\{\{\}\}\},\{\{\{\}\},\{\{\},\{\{\}\}\}\}\}, \\
& \{\{\{\{\},\{\{\}\}\}\},\{\{\{\},\{\{\}\}\}\},\{\{\},\{\{\}\},\{\{\},\{\{\}\}\}\}\}\}\}.
\end{aligned}
$$

This is truly something from nothing!

6 Functions of two variables

OK, we now know how to define the natural numbers and how to define functions, but what we need is a method to define operations like addition, multiplication, and so on, so that we can actually do something interesting with the natural numbers we've created.

What does it mean when we write things like "$1 + 2 = 3$, or "$3 \times 4 = 12$"?

Well, the easiest approach is simply to think of "$+$" and "$\times$" as functions that map pairs of natural numbers into other natural numbers. What's more, we already have a method to describe pairs of numbers—ordered pairs[4]. So just define addition and multiplication to be functions that map pairs of numbers into other numbers. Then "$+$" and "$\times$" will just be sets, and everything will be copacetic.

Here's what the beginnings of $+$ and $\times$ would look like:

$$
\begin{aligned}
+ &= \{((0,0),0),((0,1),1),((1,0),1),((0,2),2),((1,1),2),((2,0),2),((0,3),3),\ldots\} \\
\times &= \{((0,0),0),((0,1),0),((1,0),0),((0,2),0),((1,1),1),((2,0),0),((0,3),0),\ldots\}.
\end{aligned}
$$

7 More complicated structures

We have only defined the natural numbers (and we haven't even done a very complete job of that—lots of ideas have been omitted or skimmed). But the same sorts of operations can be used to define more and more complicated mathematical structures, but all based solely on the empty set and the axioms of Zermelo-Fraenkel set theory.

[4]Notice that in the case of addition and multiplication of natural numbers there is no real need to order the pairs since $n + m = m + n$ and $n \times m = m \times n$, but for operations where the order does make a difference, like subtraction, division, or exponentiation, the ordering is important

In this final section, we'll indicate roughly how such constructions are carried out, but the explanations will be even skimpier than what was shown above.

7.1 The integers

The natural numbers begin at zero and go up. One problem is that you cannot define the operation of subtraction on them. After all, what is $3 - 4$? Without negative numbers, there is no answer, so we need to somehow add the equivalent of negative numbers to what we already have.

This could be accomplished in a number of crude ways, but here's the usual way (which is not crude at all):

First, we'll consider all possible ordered pairs of natural numbers, and we will secretly think to ourselves that the ordered pair (m, n) represents the integer $m - n$. If we make n bigger than m, we've got our negative whole numbers, right?

Well, almost. The problem is that now there are an infinite number of ways to represent each negative number. -1, for example, can be represented in this way by $(0, 1)$, $(1, 2)$, $(2, 3)$, or even $(10000, 10001)$. We want to have our set $\mathbb{Z}$ of integers to contain only one copy of each, and since $(0, 1) \neq (1, 2)$ we can't just pile all the ordered pairs above into a set and call that set $\mathbb{Z}$.

Here's what's usually done. All the equivalent ordered pairs are lumped into a so-called "equivalence class", which is just a set of ordered pairs that are in some sense, equivalent. So then we can express the integers as follows:

$$\begin{aligned}
2 &= \{(2,0),(3,1),(4,2),(5,3),\ldots\} \\
1 &= \{(1,0),(2,1),(3,2),(4,3),\ldots\} \\
0 &= \{(0,0),(1,1),(2,2),(3,3),\ldots\} \\
-1 &= \{(0,1),(1,2),(2,3),(3,4),\ldots\} \\
-2 &= \{(0,2),(1,3),(2,4),(3,5),\ldots\}
\end{aligned}$$

and so on in both directions.

Then we can define addition as follows. Let z_1 and z_2 be two integers, so each of z_1 and z_2 is an infinite set of ordered pairs of natural numbers as in the examples above. In addition to simply listing the members of $\mathbb{Z}$, we need to define exactly what is meant by addition, subtraction and multiplication in this new structure. We assume that we do know how to add and multiply natural numbers.

Obviously, in normal math, we use the same symbols, $+$ and $\times$ to mean multiplication in the natural numbers and in the integers, but since we have defined those two systems to be very different, the set-theoretic versions will have to be different, and some theorems need to be proven to show that they "behave the same way".

To avoid confusion in what follows, let's (temporarily) use the symbols $+$ and $\times$ for the natural numbers, and we will use $\oplus$ and $\otimes$ for the functions we are trying to define over our newly created integers. We can also define $\ominus$, which is one of the main reasons we created the integers in the first place.

Suppose that a typical member of the set z_1 is (m_1, n_1), and that a member of z_2 is (m_2, n_2). To calculate $z_1 \oplus z_2$, calculate $(m_1 + m_2, n_1 + n_2)$, and then see which equivalence class it lies in among the integers. Of course we have to do some work to show that it will be in some equivalence class, and that moreover, if different representatives had been selected from z_1 and z_2, that the resulting integer will always lie in the same equivalence class.

In a similar way, we can define subtraction. Using the same representatives for z_1 and z_2, define $z_1 \ominus z_2$ to be the integer that contains $(m_1 + n_2, m_2 + n_1)$. As with the definition of $\oplus$, we again have to demonstrate that this definition makes sense.

Finally, multiplication can be done similarly, with similar work required to show that it makes sense. A typical representative of $z_1 \otimes z_2$ will be $(m_1 \times m_2 + n_1 \times n_2, m_1 \times n_2 + m_2 \times n_1)$. It's a good exercise to try to figure out why this works.

7.2 The rational numbers

The set of rational numbers (usually called $\mathbb{Q}$) can be built from the integers in the same general way. Ordered pairs of integers stand for fractions: the fraction $2/3$ might be represented by $(2, 3)$ or $(4, 6)$, or $(-2, -3)$, for example. In the same way that we dealt with alternative representations of integers, we can define rational numbers to be equivalence classes of ordered pairs of integers.

If q_1 is represented by a typical element (n_1, d_1) and z_2 by (n_2, d_2), then z_1 and z_2 will be in the same equivalence class if and only if $n_1 \otimes d_2 = n_2 \otimes d_1$. Since division by zero always leads to heartache, we will avoid any ordered pairs where the second element is zero.

Addition, subtraction, multiplication, and now division can be defined for the rationals in the obvious way, in terms of the operations on the integers. Since we're running out of characters to represent the operations, we'll go back to $+$, $-$, $\times$, and $/$ for rationals, and we'll use $\oplus$, $\otimes$, and $\ominus$ for the integers.

As before, let q_1 and q_2 have representatives (n_1, d_1) and (n_2, d_2). Then here are some typical representatives of rationals that are the results of the various arithmetic operations:

$$\begin{aligned}
q_1 + q_2 &\quad : \quad (n_1 \otimes d_2 \oplus n_2 \otimes d_1, d_1 \otimes d_2) \\
q_1 - q_2 &\quad : \quad (n_1 \otimes d_2 \ominus n_2 \otimes d_1, d_1 \otimes d_2) \\
q_1 \times q_2 &\quad : \quad (n_1 \otimes n_2, d_1 \otimes d_2) \\
q_1 / q_2 &\quad : \quad (n_1 \otimes d_2, n_2 \otimes d_1)
\end{aligned}$$

Note, of course, that for q_1/q_2 that n_2 had better not be zero.

It's again a good exercise to see why these definitions work.

7.3 The real numbers

One nice way to make the real numbers $\mathbb{R}$ from the rationals is via so-called "Dedekind cuts". If the entire set $\mathbb{Q}$ is divided into two subsets so that one contains only rationals smaller than those in the other, this does a perfectly good job of defining the reals. Thus to get $\sqrt{2}$, just split the rationals into all the numbers less than $\sqrt{2}$ and all those greater. There are some minor complications, but not many.

One complication is that there are two representatives of some reals—if we split at a rational number, that number can be put into the lower or the upper set with no difference in the result.

7.4 The complex numbers

Finally, the complex numbers can be defined as ordered pairs of real numbers, and there are basically zero difficulties with this.