Kevin Buzzard, k.buzzard@imperial.ac.uk, version of 07/12/15

The p-adic numbers.

Here are some further questions on the p-adic numbers and related matters. Some of them are precise questions with precise answers. Other questions are more open-ended. Some are easy, some are difficult. The first five are pretty much independent of each other, so you can do them in any order you like and skip questions if you like, but questions 6 to 9 are meant to be done in order and build up to a study of p-adic geometry.

If you get stuck, you could maybe ask on `math.stackexchange.com`.

1. Let's look at some other number systems, generalising the light bulb one.

In the "switching a light bulb on and off" number system, every whole number was equal to either 0 or 1, because the entire system only had two states – "on" and "off". I briefly mentioned another such system in the talk – let's look at such systems in more depth.

a) In "walk round an octogon mathematics", you walk round an octogon, and you count how many sides you've walked around. So the numbers 0 (where you start), 1, 2, 3, 4, 5, 6, 7 are all different, and then $8 = 0$ (because after you've walked round eight sides you're back where you started), and then $9 = 1$ and so on.

Try figuring out whether $1 + 2 + 4 + 8 + 16 + \cdots = -1$ in this system of mathematics (remember this is what algebra predicted the answer to be).

b) In "walk round a triangle mathematics", convince yourself that $1+2+4+8+16+\cdots$ doesn't add up to anything. Do this by working out where on the triangle you end up after 1 edge, $1+2$ edges, $1+2+4$ edges, $1+2+4+8$ edges, and convincing yourself that you're not going to stabilise. But that's OK – it doesn't add up to anything in the usual real numbers either.

c) Try parts (a) and (b) in "walk around a polygon with n sides" mathematics. For which values of n does $1 + 2 + 4 + 8 + \cdots$ add up? When it does add up, does it always add up to the answer predicted by algebra?

d) In "walk round an octogon mathematics", we can easily check that $3^2 = 1$, because $3^2 = 9$ but if you walk round nine sides, you end up in the same place as if you just walked round one side. Convince yourself that in this number system the quadratic equation $x^2 = 1$ actually has *four* roots (by finding them)! Here's a "proof" that it can only have two roots: if $x^2 = 1$ then $(x + 1)(x - 1) = x^2 - 1 = 0$, so either $x + 1 = 0$ or $x - 1 = 0$, so either $x = -1 = 7$ or $x = 1$. This proof works fine in our usual number system. Where does it go wrong in the octogon system?

e) For which "walking round a polygon with n sides" number systems will the proof in part (d) work? A related question – for which n is it always true that if $a \times b$ is a multiple of n, then either a or b must be a multiple of n? If you did the earlier parts of this question you will realise that this does not hold when $n = 8$. If you're not sure where to start, then try some small values of n, list every number in the number system and square them all and count how many solutions to $x^2 = 1$ there are. Can you spot any patterns in your answer? How many square roots can a general number have in these systems? Can you find a value of n for which 1 has 8 square roots? What about 9? Make some tables. Make some conjectures. Try to prove them. For which values of n does -1 have a square root? When does 2 have a square root? What about cube roots? Can a number have three cube roots? Four? Nine? Invent your own questions and then answer them – Gauss spent many years working on exactly these problems in the early 1800s and ended up inventing a whole bunch of number theory. You can learn it at university, but you'll understand it much better if you just invent it yourself.

2. Let's sum some more geometric progressions.

a) In the 10-adic numbers, convince yourself that $1+100+100^2+100^3+\cdots = 1+100+10000+1000000+\cdots$ adds up to something. Write down the 10-adic number it adds up to. You might have to invent some "recurring" notation analogous to recurring decimals in the usual real number system.

b) Now use the power of algebra to find a rational number that it adds up to.

c) Convince yourself that this rational number you got in part (b) equals the 10-adic number you got in part (a).

d) If you know the formula for the sum of an infinite geometric progression, convince yourself that this formula also agrees with your answers.

e) When does the formula for the sum of an infinite geometric progression apply, *in the usual real numbers*? When do you think the formula applies in the 10-adic numbers? Can you add up $1+20+400+\cdots$, all the powers of 20, in the 10-adic numbers, this way? What does the formula predict? Do you think it's right? How might you check? Can you check it to 5 significant figures? Can you check it completely?

3. It's possible to view fractions as 10-adic numbers.

a) Let's try and find a 10-adic number which is equal to $2/3$. Well, 3s don't go in to 2, but they do go into 12, which is quite close to 2 because 10 is small. So we could start by working out $12/3$, and this will probably be right to one significant figure.

We want more significant figures though, so perhaps next we should try $102/3$, and then $1002/3$. What about $100000002/3$? What about $10000000000000000000000000000002/3$? What is the 10-adic number which is equal to $2/3$? Can you prove your answer is right, by multiplying your 10-adic number by 3 and getting 2?

b) Can you do the same with other fractions? Here are some you could try. How about $1/3$? How about $1/7$? What about $-1/19$? Did you run into this number in question 2? What about $1/10$? Maybe one needs to allow 10-adic numbers to have a few digits after the decimal point? That's OK of course – as long as you don't add infinitely many. If you are allowed a few numbers after the decimal point, can you do $1/300$? How come some rational numbers need digits after the decimal point but others don't? Which ones do and which ones don't? Why is the usual real decimal expansion of $1/2$ equal to the 10-adic expansion, but for $1/3$ they're completely different? Was every 10-adic number you saw in this question ultimately periodic, i.e., did it eventually start repeating as you went further to the left? Do all fractions always give 10-adic numbers which ultimately recur? Is the same true in our usual number system? Why? Is it true that a 10-adic number ultimately starts repeating if and only if it is a fraction (a rational number)? Can you prove this? Is it true for the real numbers? Can you prove that?

4. Let's do trial and improvement to see if we can find a cube root of 7 in the 10-adic numbers.

a) A good choice for the first digit to the left of the decimal point would be 3, because $3^3 = 27$ and 27 is close to 7 in the 10-adic numbers in the sense that their difference is a multiple of 10, and 10 is small. Try some other one-digit numbers. Can you find another one whose cube ends in 7? Check that no other one-digit number works! So perhaps our cube root of 7, if it exists, looks like $\cdots$????3.

b) What about that second digit? We should try 13^3, then 23^3, and so on; cube the first few numbers until you find one ending in 07. You should stumble upon 43^3, which according to my calculator is 79507. This is great – it agrees with 7 to two significant figures (remember that digits to the *right* are more significant, so e.g. the 5 in 79507 is quite insignificant, and the 9 and the first 7 are even more insignificant). So perhaps $\sqrt[3]{7}$ =????43. If you can be bothered, try all the 2-digit numbers that end in 3: do any of the others cube to give something that ends in 07?

c) Now we're really going to need a calculator to find that third digit. Try cubing 143, 243, 343, all the way up to 943, and don't forget $43 = 043$: you'll find that all of them end in $\cdots 07$ but only one of them ends in $\cdots 007$. Which one is it? That's got the answer right to three significant figures.

d) How much further can you go? Can you find the first ten digits? Can you write a computer program to find the first 100 digits? What problems do you run into when doing this question?

Comments on question 4: There are two issues here – a theoretical one and a practical one. Let me talk about them both.

The practical question is this: your stupid calculator, when you cube a large whole number, decides (when the answer is too big to display) to just dump the numbers which are *classically* the least significant, and for this question this design implementation means that we can't see the numbers we care about. Assuming you don't have a 10-adic calculator, you're going to need another tool after a while. You could use a computer. The "correct" thing to do would be to use a mathematics package; for example pari-gp (`http://pari.math.u-bordeaux.fr/`) and sage (`http://www.sagemath.org/`) are free and run on Unix machines and, to a lesser extent, on Windows. Pari-gp even runs on Android nowadays – it's on the google play store (search for paridroid), which means that you can multiply very long numbers together exactly on an Android device, although the interface is currently horrible. If you use Windows or a Mac then you could either emulate a Unix environment using VirtualBox, or pay for the privilege of having a maths package – e.g. Maple, Mathematica or Matlab should I guess run on these machines, but you will need to buy a license. Another option would be to install python – this can multiply large integers together, and runs on any operating system (don't forget that to raise something to a power you use $**$). If you don't want to install software and can't find anyone who has a maths package or python installed, you can go for online solutions. I would recommend the sage math cloud at `cloud.sagemath.com` – you need to create a login, create a project, create a worksheet, and then type in what you want to know and click "run". This should work and will give you an interface where you can write loops in python and do lots of calculations. If that is too much hassle then I will finally confess that you can go to the website `wolframalpha.com` and type in your long integer multiplication there. The problem with wolframalpha is that you can't program it to do lots of these calculations at once so if you really want to get into this then you should find out more about computer algebra packages.

The theoretical question is: how do you know that this trial and improvement trick is going to keep working, i.e., that there will always be a next digit that works? That's an interesting question! Can you prove that it always works? The proof is slightly tricky but go ahead and have a try. But wait a second – did you once in your life unquestioningly *assume* that trial and improvement works for the usual real numbers – and now you are worried about it for the 10-adic numbers? That means that you're being illogical. The conclusion from this is definitely *not* that trial and improvement always works with the 10-adic numbers! The conclusion is that we *definitely need to check* that trial and improvement always works with the usual real numbers! You will learn the proof of this in your first year at university if you study mathematics – you will check that the real numbers "have no holes in" and in particular that every positive real number has a unique positive real nth root.

5. Let's find a non-obvious solution to $x^2 = x$ in the 10-adic numbers.

Did you ever choose a random button on your calculator and just press it again and again to see what happened? When I was a kid I noticed that "cos" was a good one to try – this was before I even learnt what the button did. Even before I knew the definition, I knew that there was a solution to $\cos(x) = x$ and that I could work it out by mashing the cos button.

Let's try button-mashing to work out some solutions to $x^2 = x$ in the 10-adic numbers.

a) Start with 5. Square it. Square it again.

b) Square it until your calculator starts being annoying and only telling you the first digits, not the last digits. Then have a look at the numbers you got. Do you think the numbers you're producing are getting closer and closer to a 10-adic number? What are the first few significant digits of this 10-adic number?

c) If you have access to a computer, try computing some more digits of this 10-adic number (see the comments to Q4 for extensive hints about how you can multiply large integers together). Can you figure out the first 20 digits this way? This might be harder than you think, unless you have a good idea. Can you spot how many extra digits of accuracy you get each time you square? Can you prove it?

d) Do you think that if you go on forever you will discover a 10-adic number that equals its own square? Can you prove it?

e) Does it disturb you that the quadratic equation $x^2 - x = 0$ seems to have more than two solutions in the 10-adic numbers? If you did Q1 then you might be happy with this idea already. If we factor $x^2 - x$ as $x(x - 1)$ then we seem to have found a 10-adic number x such that x isn't zero and $x - 1$ isn't zero, but $x(x - 1)$ is zero.

Comments on Q5: Here's an explanation of what the x in this question is trying to be in the 10-adic numbers. The reason this weirdness is happening is because 10 is not prime; if you tried the last two parts of Q1 then you might have realised that primes had something to do with things. I said that we should think of 10 as small, but the problem is that $10 = 2 \times 5$, so one of 2 and 5 should be small, but I didn't say which. Take the number x that this question is about. Write down a "normal" whole number n which is very close to x (i.e., just take as many digits of x as you managed to work out and regard this as a normal whole number). Check that this number n is divisible by a big power of 5 (so if 5 is small then n is small, i.e., n is close to zero). Now check that $n-1$ is divisible by a big power of 2 (so if 2 is small then $n-1$ is small, so n is close to 1). The following rather cool thing is happening: x is trying to be 0 in the 5-adic world, and it's trying to be 1 in the 2-adic world, and it's trying to be both at the same time. In fact the 10-adic numbers are in a precise sense the "product" of the 2-adic numbers and the 5-adic numbers. There are two solutions to $x^2 = x$ in the 2-adic numbers, and two in the 5-adic numbers (2 and 5 are prime so these systems behave more sensibly), so there are four solutions in the 10-adic numbers. We have 0, 1 and the value of x from this question – Can you spot the fourth one? Hint: there is a very easy trick to find it, if you know x. Here is a non-easy way. Take your favourite method of being able to multiply huge numbers together exactly, and then start with 6 and keep raising to the power 5 (this works because if $x^2 = x$ then $x^5 = x$ too). If you figure out the easy way and the non-easy way, check they agree!

6. Let's think more carefully about the way we can actually make the "size" of a number small.

Even though a number is not *equal to* a length in these number systems, numbers still *have a size*. All this p-adic business works best when p is prime, so let's forget the 10-adic numbers and set $p = 2$, and let's really try and build a system of sizes so that 2 has a small size but is not zero. Note: a size is like a length – it's a non-negative real number. So we will have to be careful to distinguish a number from its size.

Let's start by defining the size of every whole number – let's really do this properly. Here's the idea. Let's define the size of 2 to be very small – say the size of 2 is the real number 0.1. And let's say that every single other prime number has got size equal to 1.

Here's some notation: let's write $|x|$ for the size of the number x. So another way of saying what I just said is that $|2| = 0.1$ and $|3| = |5| = |7| = |11| = \cdots = 1$.

To work out the size of a general positive number, write it as the product of prime numbers, and multiply the sizes of the prime numbers together! For example, $|12|$, the size of 12, is 0.01, because $12 = 2 \times 2 \times 3$, so $|12| = |2| \times |2| \times |3| = 0.1 \times 0.1 \times 1 = 0.01$.

a) Work out the sizes of the numbers from 2 to 10. Which number is the smallest? Which ones are the biggest?

b) Check that if x, y are whole numbers at least 2, then $|x \times y| = |x| \times |y|$ – "the size of the product is the product of the sizes". Try it for the numbers in part (a), or convince yourself it's true in general.

c) I will make the usual mathematician's abuse of notation and drop the $\times$ signs from now on, and so we can rephrase part (b) as

$$|xy| = |x||y|. \tag{$*$}$$

Given that we want this formula ($*$) to hold in general, what do you think $|1|$ should be? (hint: set $x = 1$ and $y = 2$ because we know the size of 2). What do you think $|0|$ should be? (hint: set $x = 0$ and $y = 1$).

What do you think $|-1|$ should be? Hint: set $x = y = -1$ in ($*$) and remember that the size of something should be a *non-negative real number*.

If you've got these wrong then you're in real trouble, so here's the answers just to check you're on the right track: we should have $|1| = 1$, $|0| = 0$ and $|-1| = 1$.

d) Convince yourself that $|-x|$ should be the same as $|x|$ (just like in the real or complex numbers, if you know about them). Convince yourself that 0 is the only whole number with size 0, and everything else has positive size.

e) We now have a formula for the size of every whole number. What about fractions? What should the size of $1/3$ be? Hint: set $x = 1/3$ and $y = 3$ in ($*$). What about $1/2$? If you know the size of m and of n, and $n \neq 0$, how can we work out the size of m/n?

f) Prove that $|5/12| = 100$, that $|-4/3| = 0.01$. Write down a rational number with a size greater than a billion. Now write down a non-zero whole number with size less than a billionth.

g) Which positive whole numbers have got size 1, and which have size less than 1? Take a positive whole number and write it in binary. How can you tell just from looking at the binary expansion if it has size 1 or less than 1? Which positive whole numbers have size less than 0.1? Can you tell from looking at the binary expansion of a number whether or not it has size less than 0.1? Can you tell what the size of a positive whole number is *immediately* if you are given its binary expansion?

h) Although we won't be seeing p-adic numbers for a while (in fact we're currently actually in the process of re-inventing them more carefully), consider the 2-adic number $\cdots 11011010100101001010110101 00$ (written in base 2). Assume you did the previous part of this question. What do you think the size of this number should be? Do you think it's still true in the 2-adic numbers that $|xy| = |x||y|$? Can you prove it?

i) Convince yourself that all this would work with p-adic numbers for any prime number p. We can define $|p| = 0.1$ and $|q| = 1$ for all other prime numbers q, and then use the rule $|0| = 0$, everything else has positive size, and $|xy| = |x||y|$ to work out the size of every rational number. If $p = 7$ then what is $|98|$? What is $|100|$? What is $|1/1000|$? What is $|2/7|$? And so on.

Comment on Q6. Real numbers have sizes too: the size of the real number 5 is 5, but it's not quite that simple – the size of -3 is $+3$. Maybe you have met the notation $|x|$ for the size of the real number x – question 6 borrows this notation and uses it in a different context.

7. Let's define a distance function using the strange definition of size that we developed in the previous question.

We're going to work with the rational numbers in this question, like $-5/6$, 0 or 7. In general a rational number is a number expressible as a ratio a/b with a and b whole numbers and $b \neq 0$.

In the previous question we defined "the 2-adic norm" – that's a fancy way of describing our size function. We have $|2| = 0.1$, and $|q| = 1$ for all other prime numbers q, so $|12/13| = 0.01$ and $|-7/2| = 10$ and so on. This means in particular that we currently know the size of every rational number. If x and y are rational numbers, let's define the *distance* from x to y to just be $|y - x|$. Let's abbreviate this as $\mathrm{dist}(x, y)$.

[If you temporarily stop thinking about this strange size function and think about real numbers instead, with their usual notion of size, then you will see that the distance from a real number x to a real number y on the number line is also $|y - x|$, where now $|y - x|$ means the usual absolute value – so this notion of distance is probably a sensible one, even if we're using it in a silly context.]

Example: to work out $\mathrm{dist}(1, 3)$, the distance from 1 to 3, we work out $|3 - 1|$ which is $|2|$ which is 0.1. So 3 is quite close to 1. What is the distance between 15 and 23? It's $\mathrm{dist}(15, 23) = |23 - 15| = |8| = 0.001$. So 15 and 23 are *really* close!

a) What is $\mathrm{dist}(1, 9)$? What is $\mathrm{dist}(1, 2)$? What is $\mathrm{dist}(1, 1\frac{1}{32})$? What is $\mathrm{dist}(1, 1/3)$?

b) Which of the following numbers is nearest to 7: (i) 8, (ii) 9, (iii) 10?

c) Try and draw a picture of all the whole numbers from 0 to 10 on a piece of paper, making the distances between them be approximately the strange definition of distance we're using in this question. Add some more if you're brave. Where does 1/2 go? Where does 1/3 go? Which number between 0 and 10 is closest to 1/3?

d) Find a whole number whose distance from $-1/3$ is less than 1. Find a whole number whose distance from $-1/3$ is less than 0.1. Can you find a whole number whose distance from $-1/3$ is less than 0.000001? Hint: earlier on in this sheet you might have got quite good at figuring out questions like what $1 + 4 + 4^4 + 4^3 + \cdots$ was in the 2-adic numbers. Now look at the most significant figures...

e) Let's check some fundamental properties of this distance function. If you have understood the definition correctly, you will know that $|a| = 0$ if, and only if, $a = 0$. This means that $\mathrm{dist}(x, y) = 0$ if, and only if, $x = y$. That's a pretty sensible property for a distance function to have!

f) Check for any rational numbers x and y, we have $\mathrm{dist}(x, y) = \mathrm{dist}(y, x)$. In other words, to measure the distance from x to y we can either start at x and move to y, or start at y and move to x. That's another sensible property! We will see the killer property in the next question though.

8. Let's look more at the distance function from the previous question. This question is quite subtle.

a) Here is a deeper property of a distance function – the shortest distance between two points is a straight line. Here's a mathematical way to express this: say x, y and z are three whole numbers. The direct distance from x to y is $\mathrm{dist}(x, y)$. The distance from x to y via z should be at least as long! This latter distance is $\mathrm{dist}(x, z) + \mathrm{dist}(z, y)$. So we should definitely check that for any whole numbers x, y and z, we have

$$\text{dist}(x, y) \leq \text{dist}(x, z) + \text{dist}(z, y).$$

If our distance function did not satify this, it would not really be worthy of being called a distance function. But it does! Can you check this? Here's a hint: if $z - x$ is divisible by a big power of 2, and $y - z$ is also divisible by a big power of 2, then their sum is also divisible by a big power of 2. . . .

b) In fact we can do better. Prove that if x, y and z are whole numbers then

$$\text{dist}(x, y) \leq \max\{\text{dist}(x, z), \text{dist}(z, y)\}.$$

In words, the distance from x to y is not just at most the sum of $\text{dist}(x, z)$ and $\text{dist}(z, y)$, it is actually at most the *maximum* of these numbers. This is a stronger assertion, because a distance is a non-negative real number. This is why part c of the previous question is so weird – the distance function we have defined is much weirder than the usual distance function – it has all the usual properties of the usual distance function, and more too.

c) Prove that if x, y and z are whole numbers and $\text{dist}(x, z) \neq \text{dist}(z, y)$, then

$$\text{dist}(x, y) = \max\{\text{dist}(x, z), \text{dist}(z, y)\}.$$

You can either prove this directly using the definitions, or you can try and find a way of logically deducing it from part (b).

d) If you're feeling brave, re-do parts (b) and (c) for x, y and z rational numbers.

9. Let's finally do some geometry using this funny distance.

You have to have tried Q6 and Q7 to make sense of this question, and will need something from Q8 (but just assume Q8 is right if you didn't do it, and use the conclusions).

Our set where we're going to do geometry is the set of whole numbers. This question would work fine using the rational numbers or the 2-adic numbers, but let's stick to the whole numbers. Let's prove that every triangle is isosceles. Also, let's prove that every point inside a circle is its centre! Notation for this question: let's define a "disc" to mean a circle, plus everything inside the circle. We will prove that every point in a disc can be thought of as a centre for the disc.

a) If x, y and z are three distinct whole numbers, we could consider the triangle with corners at x, y and z. [If you're thinking this is daft because x, y and z all lie on a straight line – the number line – then remember that that's not at all how we're thinking of numbers in this question, and a much better picture would be the picture you drew in Q7 part (c).] The three sides of the triangle are $\text{dist}(x, y)$, $\text{dist}(y, z)$ and $\text{dist}(x, z)$. Prove that every triangle is isosceles in this version of geometry. Hint: use part (c) of the previous question (whether or not you managed to solve it) and deal with the three possibilities $\text{dist}(x, z) < \text{dist}(z, y)$, $\text{dist}(x, z) = \text{dist}(z, y)$, $\text{dist}(x, z) > \text{dist}(z, y)$ separately.

b) Let's look at some discs. First let's look at discs centre zero. Let C be the set of whole numbers n such that $|n| \leq 0.5$. This is the disc centre 0 radius 0.5. Write down some numbers in this disc. What's an easy way to tell whether or not a number is in this disc? Which numbers are in the disc centre zero radius 0.4? The answers are the same! The size of any whole number is either 1 or 0.1 or 0.01 or… so we can't really tell the difference between 0.4 and 0.5. What about the disc centre zero radius 0.003? Which numbers are in that?

c) Now let's look at some discs with other centres. The definition is this: if x is a whole number, and r is a positive real number, then the disc centre x and radius r is the collection of all whole numbers n such that $\mathrm{dist}(x, n) \leq r$. Work out which whole numbers are in the disc centre 1 radius 0.5. Convince yourself that every number in the disc centre 0 radius 1 is either in the disc centre 0 radius 0.5 or in the disc centre 1 radius 0.5, but not both. Conclude that we can write a disc as the union of two smaller discs which don't overlap!

d) Now work out which whole numbers are in the disc centre 2 radius 0.5. Can you show that the disc centre 2 radius 0.5 contains exactly the same numbers as the disc centre 0 radius 0.5? That means those discs are *equal*! That means that a disc can have more than one centre.

e) Prove that if D is a disc with centre x and radius $r > 0$, and if y is in D, then D is also equal to the disc centre y radius $r > 0$. In particular every point in a disc is its centre. Use Q8 part (c).

f) Prove that if C and D are two discs, then either they are disjoint, or one is contained completely within the other.

10. Here is how a mathematician builds the real numbers and how they build the p-adic numbers. This is not actually a question.

To build mathematics, a mathematician starts with set theory. One of the axioms of set theory is that there exists a set. Another one says that you can take subsets of a set, and they're sets too, so that means there must exist an empty set. Let's define 0 to be the empty set.

Let's define 1 to be the set containing the empty set. So 1 is a set with 1 element, namely the empty set.

Let's define 2 to be the set containing 0 and 1. So 2 is a set with 2 elements.

Let's keep going, and before you know it we have defined every non-negative whole number.

Mathematicians can press on in this way, and then define negative numbers, and then define rational numbers as just pairs of whole numbers (we don't attempt to divide a set by another set, we just regard a rational number as a pair of sets, one for the numerator and one for the denominator).

That's how a mathematician might build the rational numbers.

The interesting part is what to do next. The problem with the rational numbers is that they have loads of holes in. For example, there are rational numbers whose square is just a little less than 2, and rational numbers whose square is just a little more than 2, but there is no rational number whose square is exactly 2 – there's a hole in the rational numbers at that point. So classically what mathematicians did was they gave rational numbers a size.

They said the size of x was x if $x \geq 0$, and $-x$ if $x < 0$. And now they filled in the holes, by which I mean they *defined* a real number to be a sequence of rational numbers such that, as you moved along the sequence, all of the rational numbers were closer and closer to each other (that is, their differences always had very small size). For example π is *defined* to be the sequence of rational numbers 3, 3.1, 3.14, 3.141,.... This is a complicated idea and you won't see it presented carefully unless you do maths at university (or read a book or watch a youtube video...). But this is how to rigorously build the real numbers – you build the rationals first, and then fill in the holes by looking at sequences which should tend to a limit but for which the limit isn't there.

Historically this sort of approach is well over 100 years old, but it was only around 100 years ago that it dawned on Kurt Hensel that if instead of the usual notion of distance, we used the version of distance in questions 6 to 9 above, then we could fill in the holes using just the same technique and get a new number system. He called the resulting number system the p-adic numbers, and that was how it was originally constructed. Just like the infinite decimal expansion of π is a way to explain how you're filling in the hole at π in the rationals, the infinite expansions in the other direction that we used for p-adic numbers are a formal way to fill in the holes that take us from rational numbers to p-adic numbers.

One might now get very excited and start wondering whether there are even more crazy notions of distance on the rational numbers so that when we fill in the holes we get even more crazy number systems. Unfortuately Ostrowski proved that the only distance functions on the rational numbers that satisfied that the shortest distance between two points was a straight line, are the usual notion of distance, and the p-adic distance for p a prime number. Even the 10-adic distance doesn't work – p has to be prime. The problem with the 10-adic distance is that perhaps you want 2 and 5 to be small – but then 2^{100} is really small, and 5^{100} is really small, but $2^{100} + 5^{100}$ is much much bigger, so if we want to walk from 0 to $2^{100} + 5^{100}$ it's better to go via 2^{100} than going directly, and you can't do geometry when that sort of thing happens – it becomes mayhem. This is why I started using p-adic numbers for $p = 2$ later on – if you want to make 10 small, you have to make either 2 small or 5 small but not both, so really you're doing the 2-adic numbers or the 5-adic numbers at the end of the day.

So in fact, when you think about things this way, Ostrowski's theorem says that the p-adic numbers and the real numbers are the only ways you can fill in the holes in the rational numbers. Traditionally we use the real numbers in real life, because we life in a universe that locally looks like it's a 3-dimensional real space. But if you're just interested in the rational numbers, like I am, then you should treat the real numbers and the p-adic numbers with equal weight. There are amazing theorems that say that certain equations can be solved in rational numbers if and only if they can be solved in the real numbers and the p-adic numbers, and these theorems really show you the symmetry between the real world and the p-adic world.

Kevin Buzzard, 2nd December 2015.