Paradox

Issue 4, 2012

The Magazine of the Melbourne University Mathematics and Statistics Society

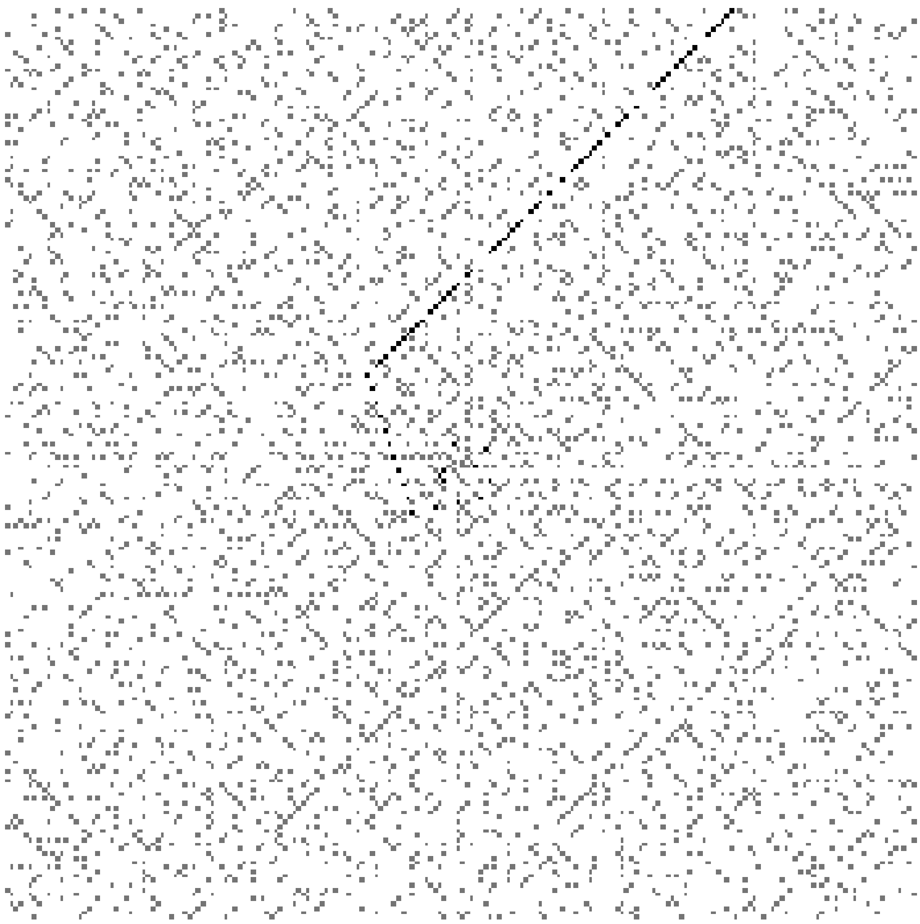

MUMS

President:	Giles David Adams gilesdavidadams@gmail.com
Vice-President:	Andrew Elvey Price andrewelveyprice@gmail.com
Secretary:	Aaron Chong awychong@hotmail.com
Treasurer:	Mel Chen m.chen3111@gmail.com
Editor of Paradox:	Kristijan Jovanoski paradox.editor@gmail.com
Education Officer:	Dougal Davis dougal.davis@gmail.com
Publicity Officer:	Franky Cheung fycheung@student.unimelb.edu.au
Undergrad rep:	David Batt davidbatt@optusnet.com.au
Undergrad rep:	Benjamin Hague bhague@student.unimelb.edu.au
Undergrad rep:	Patrick Kennedy p.kennedy4@student.unimelb.edu.au
Undergrad rep:	Julia Mccoey julia.mccoey@gmail.com
Undergrad rep:	Emma Yeap mintymelody@gmail.com

MUMS Website:	www.mums.org.au
MUMS Email:	mums@ms.unimelb.edu.au
MUMS Phone:	(03) 8344 4021
Paradox Website:	http://www.ms.unimelb.edu.au/~mums/paradox/
Paradox Email:	paradox.editor@gmail.com
Printed:	19 August, 2012
Cover:	The polynomial $4x^2 - 2x + 41$ generates a very interesting Ulam spiral. Find out more on page 6!

In This Edition of Paradox

Regular Features

Articles

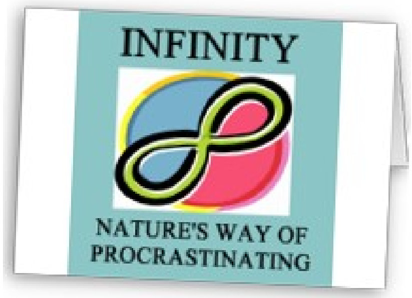

A physicist and a mathematician are sitting in a faculty lounge. Suddenly, the coffee machine catches on fire. The physicist grabs a bucket and leap towards the sink, filled the bucket with water and puts out the fire. The next day, the same two sit in the same lounge. Again, the coffee machine catches on fire. This time, the mathematician stands up, gets a bucket, hands it to the physicist, thus reducing the problem to a previously solved one.

Words from the Editor

Welcome to this year's fourth issue of Paradox, the magazine produced by the Melbourne University Mathematics and Statistics Society (MUMS). This is the first regular issue since the new committee has been elected, but the radical Editor remains!

Every issue this year has been radically different in some way, and this School Maths Olympics (SMO) and Open Day issue is no different. The regular features have all taken a break this time around, and three Presidents past and present have stepped in to fill the void instead. Hence, it is only fitting to tribute this issue to the Presidents of MUMS. Without your steady hands, MUMS (and Paradox) would be lucky to survive, right?

Anyway, in this edition you will discover how mathematics toys with chess and why drawing during boring presentations might have profound implications someday. Otherwise, exercise your mind trying to use the modern notion of statistical significance in a not-so-modern situation or go on a brief Pythagorean journey that might whet your appetite for more!

Paradox thanks all contributors for providing so much material over the past year and encourages everyone to continue submitting articles, puzzles, reviews, and all-important jokes. Just ask about Paradox in the MUMS room or contact us via email (see page 2). Be prepared for quite a few surprises and comebacks in record issue number 5!

Yours radically,
Kristijan Jovanoski
Paradox Editor

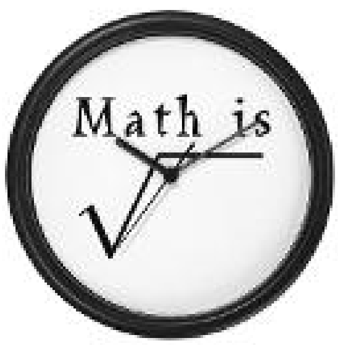

Words from the President

Welcome all to Paradox! As we wander through another Open Day, SMO, and into a second third of a second semester, we reminisce on our fellow MUMSians who have left us for brighter futures. Some, off into the brave world of actual careers, others off to universities abroad. One certainly can't say that maths doesn't get you places.

Back here at the MUMS room, the new MUMS year (AGM to AGM) is going mathemagically. Seminars have been going very well (many thanks to our Education Officer Dougal) with many new students joining our inner sanctum within the walls of the MUMS room. Everyone is welcome, especially you! In other news, we've reinstated the old tradition of having in-depth lecture series, and our overseas-bound Sam is kicking things off with three weeks on number theory. Look forward to more of these throughout the semester and into the new year.

For now, enjoy whatever it is you are currently doing. Whether that be Open Daying, Semestering or discovering this Paradox in an old bag long since discarded. And remember, maths is awesome (see further in this Paradox for examples thereof.)

Yours presidentially,
Giles "Da Prez" Adams
MUMS President

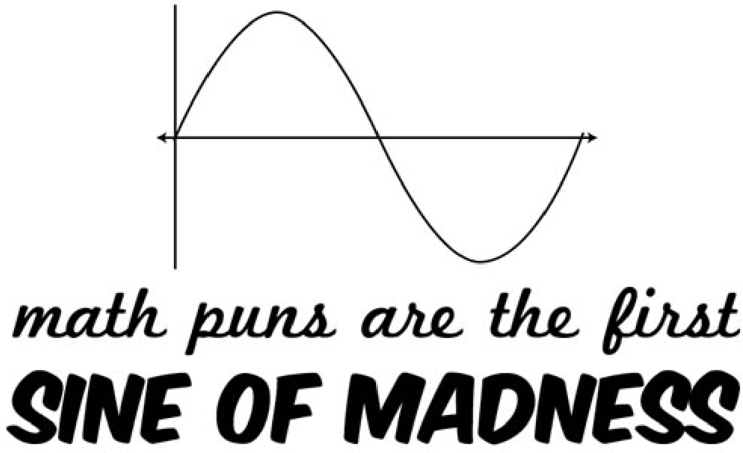

The Ulam Spiral

The Ulam spiral is a method of visualizing the prime numbers that shows the apparent tendency of certain quadratic polynomials to generate unusually large numbers of primes. It was discovered by Stanisław Ułam in 1963 while doodling during the presentation of a long and boring paper. He constructed the spiral by writing down a regular rectangular grid of numbers, starting with 1 at the centre, and spiralling out:

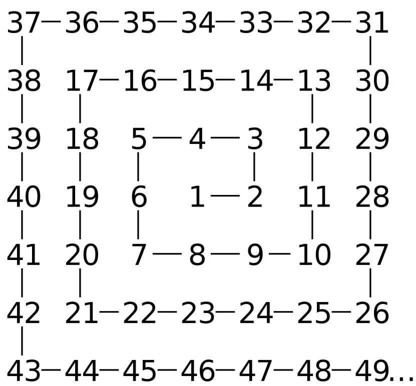

He then circled all of the primes to get this picture, discovering that the prime numbers tended to lie on some diagonals more than others:

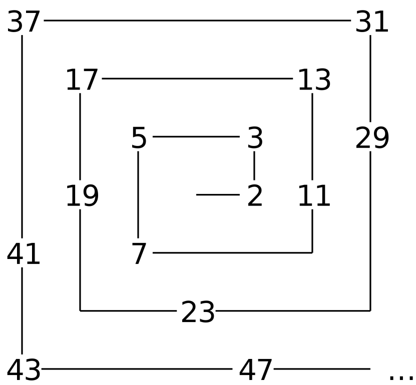

Later Ułam used the first-generation electronic computer MANIAC II at Los Alamos Scientific Laboratory with collaborators Myron Stein and Mark Wells to produce pictures of the spiral for numbers up to 65,000.

Almost all of the prime numbers lie on alternating diagonals, since all except the number 2 are odd. Note that the difference between rows or columns

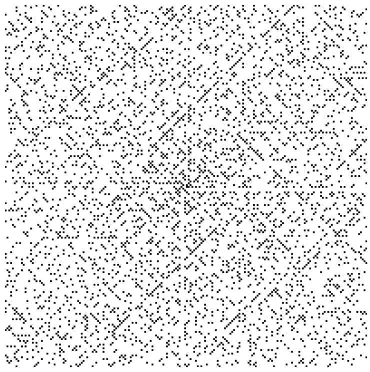

Figure 1: 200×200 Ulam spiral.

going towards the outside is always odd, but if we then move one along the spiral, we can see that the difference is always even. This corresponds to alternating diagonals giving only odd or even numbers. We can see some of these patterns in Figure **??**.

The pattern of lines imply that there are many quadratic polynomials of the form

$$f(n) = 4n^2 + bn + c$$

where $b, c \in \mathbb{Z}$ and $n \in \mathbb{N}$.[1] Since these are the equations for numbers along lines emanating from the central area in the Ulam spiral, with b odd for lines at a 45° angle to the horizontal, and b even for vertical and horizontal lines.[2]

Why is it interesting? First, not much has actually been discovered about the Ulam spiral. The reason why the prime numbers align diagonally while non-prime numbers align horizontally and vertically is not clear yet. Nonetheless, the spiral is important primarily because it shows a clear pattern among prime numbers, even when the number at the centre is not 1.

Hardy and Littlewood's Conjecture F

Many decades earlier in 1923, Hardy and Littlewood stated a conjecture, that, if true, may explain many of the striking properties of the Ulam spiral. Their Conjecture F concerns polynomials of the form $ax^2 + bx + c$ where a, b, c are integers and a is positive. It states that if a, b, c contain a common factor greater than 1 or if the discriminant $\Delta = b^2 - 4ac$ is a perfect square, then the polynomial factorises and produces composite numbers for almost all $x \in \mathbb{N}$, and if $a + b$ and c are both even, the polynomial will produce only even

[1] That is, b and c are whole numbers (integers) and n is a whole number greater than 0.

[2] This actually implies the result in footnote 1.

numbers. Otherwise, such a polynomial gives infinitely many prime values for $x \in \mathbb{N}$.

Furthermore, they conjecture that the number $P(n)$ of primes, asymptotically,[3] of the form $ax^2 + bx + c$ less than n is given by

$$P(n) \sim A \frac{1}{\sqrt{a}} \frac{\sqrt{n}}{\log n}$$

where

$$A = \varepsilon \prod_p \left(\frac{p}{p-1} \right) \prod_{\bar{\omega}} \left(1 - \frac{1}{\bar{\omega} - 1} \left(\frac{\Delta}{\bar{\omega}} \right) \right)$$

Where the first product ($\prod$ symbol) p is a prime number dividing both a and b.[4], in the second product, $\bar{\omega}$ is an odd prime number not dividing a, ε is defined to be 1 if $a + b$ is odd, and 2 if $a + b$ is even, Δ is the discriminant, and $\left(\frac{\Delta}{\bar{\omega}} \right)$ is a special multiplicative function known as the Legendre symbol.[5]

A basic explanation of this formula would be that in the case where $A = 1$, we get via the prime number theorem the asymptotic number of primes less than n expected in a random sample of numbers having the same density as those of the form $ax^2 + bx + c$. But A can take on values greater or smaller than 1, so the lines apparent in Ulam's spiral correspond to the polynomials where A is significantly higher than 1.

For example, the polynomial $4x^2 - 2x + 41$ has a value for A of about 6.6, meaning that it generates about 6.6 times as many prime numbers as you would get if you just looked at random numbers of comparable size. It is unusually rich in primes and generates a visible line in the Ulam spiral which graces the cover of this issue of Paradox.

Finally, it might seem that diagonal lines can be seen in the Ulam spiral simply because our eyes seek patterns and groups even among random clusters of dots. However, a quick comparison between Figures **??** and **??** on the next page will put this matter to rest:

[3] That is, the limit as n approaches infinity.

[4] E.g., if $a = 12$ and $b = 6$, then $\prod_p \left(\frac{p}{p-1} \right) = \frac{2}{2-1} \cdot \frac{3}{3-1} = 3$. That is, we multiply together the expression following $\prod$ for all possible values of p.

[5] A thorough explanation of this function has been omitted for the sake of simplicity.

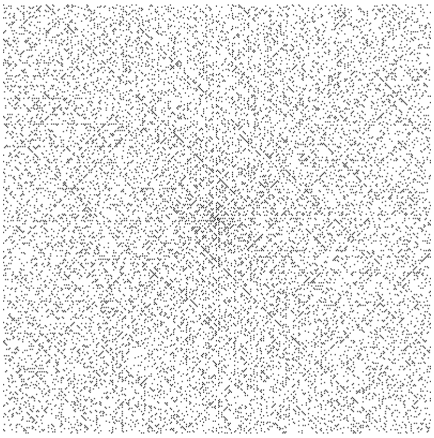

Figure 2: A spiral where the black dots denote prime numbers.

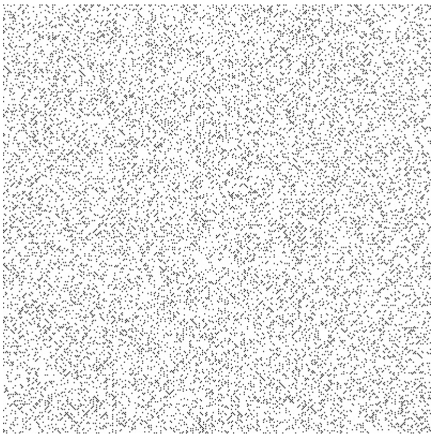

Figure 3: A spiral of random numbers.

If the Ulam spiral with prime numbers one day provides us with enough information for the discovery of a new polynomial that can generate even more prime numbers than the ones we already know, then we might have a better understanding of other mysterious conjectures involving prime numbers, such as the prime conjecture and Goldbach's conjecture, but that's another story…

— Mel Chen

The reason that every major university maintains a department of mathematics is that it is cheaper to do this than to institutionalize all those people.

How Mathematics Changed the Rules of Chess

Can a game of chess have infinitely many moves?

In 1935, Max Euwe defeated Alexander Alekhine to become World Chess Champion. Amazingly, he was only a part-time chess player.[1] Euwe completed a doctorate in mathematics at the University of Amsterdam in 1926.[2] In 1929, he published a paper answering the above question in the affirmative! The relevant rule at the time was as follows:

> A player may claim a draw if the same sequence of moves occurs twice in succession and is immediately followed by the first move of a third repetition.

Many people believed that this rule made the game of chess finite, but this was not the case. Euwe's idea was to use the following sequences of knight moves over and over in some order:[3]

[1]See *Max Euwe: The Biography*, by Alexander Münninghoff.

[2]His doctoral research was in differential geometry.

[3]This is a summary of the description given in a lecture by Jeffrey Shallit, *The Ubiquitous Thue-*

$$0 \mapsto \quad Nf3 \qquad Nf6$$
$$\qquad\quad Ng1 \qquad Ng8$$

and

$$1 \mapsto \quad Nc3 \qquad Nc6$$
$$\qquad\quad Nb1 \qquad Nb8.$$

Euwe wondered if he could find a sequence of 0s and 1s with the property of being *overlap-free*, i.e., there does not exist a subword (substring of 0s and 1s) of the form $axaxa$, where a is a letter (0 or 1) and x is a (possibly empty) word. If so, then it would correspond to an infinite chess game under the above map. Thus, Euwe independently discovered the Thue-Morse sequence.[4]

The Thue-Morse sequence

The *Thue-Morse sequence* is a sequence t such that if $n \in \mathbb{Z}_{\geq 0}$ then

- $t_0 = 0$,
- $t_{2n} = t_n$, and
- $t_{2n+1} = 1 - t_n$.

It looks like

$$t = 0110100110010110\ldots.$$

Some equivalent definitions are:

Morse Sequence, found at http://www.cs.uwaterloo.ca/~shallit/Talks/green3.pdf.

[4]It was originally discovered by Eugène Prouhet (who applied it to number theory) in 1851, but Prouhet's work went largely unnoticed for a long time. Consequently, several mathematicians independently discovered the sequence in different contexts. See http://homepages.fh-friedberg.de/boergens/english/problems/problem059engl.htm.

1. $t = \lim_{n\to\infty} X_n$, where $X_0 = 0$ and
$$X_{n+1} = X_n\overline{X_n},$$
where the bar denotes swapping the 0s and 1s in a word.
2. $t_n = s_2(n) \bmod 2$, where $s_2(n)$ is the sum of digits of n when n is written in base 2.

These are merely stated for curiosity, as we only need the initial definition. The first alternative definition tells you how to actually write down the sequence, while the second tells you how to get one term in the sequence without having to figure out all of the previous ones.

Proof[5]

So how do we show that the Thue-Morse sequence is overlap-free? Assume, for the sake of contradiction, that there is a subword of the form $axaxa$, where a is a letter and x is a word. Then we can write
$$t = uaxaxav,$$
where u and v are words (v is infinite). The overlap condition is then equivalent to
$$t_{k+j} = t_{k+j+m} \qquad \text{for } j = 0, 1, \ldots, m, \tag{1}$$
where $m = |ax|$ and $k = |u|$. Consider this for an overlap $axaxa$ such that m is minimal.[6] Note that $m > 1$, since $t_{2n} \neq t_{2n+1}$ for $n \in \mathbb{Z}_{\geq 0}$.

- Case 1: m is even. Let $m = 2m'$. We will find a smaller overlap, contradicting the minimality of m.
  - Case 1a: k is even. Let $k = 2k'$. We use the fact that $t_{2n} = t_n$ for $n \in \mathbb{Z}_{\geq 0}$:
  $$t_{k+j} = t_{k+j+m} \qquad \text{for } j = 0, 1, \ldots, m, \tag{2}$$
  so
  $$t_{2k'+2j'} = t_{2k'+2j'+2m'} \qquad \text{for } j' = 0, 1, \ldots, m', \tag{3}$$

[5] This is again from Shallit's lecture.

[6] By the well-ordering principle, if it is not obvious that such a thing exists. This is an example of a proof by *infinite descent*.

so

$$t_{k'+j'} = t_{k'+j'+m'} \qquad \text{for } j' = 0, 1, \ldots, m', \tag{4}$$

which is a smaller overlap as $m' < m$.

- Case 1b: k is odd. Let $k = 2k' + 1$, and proceed in a similar fashion.

- Case 2: m is odd. As $m > 1$, we either have $m \geq 5$ or $m = 3$.

  - Case 2a: $m \geq 5$. Choose j such that $1 \leq j \leq 4$ and $k + j \equiv 2 \pmod{4}$. Observe that

$$t_{k+j} + t_{k+j-1} = t_{k+j+m} + t_{k+j+m-1} = 1, \tag{5}$$

since

$$t_{k+j+m} = 1 - t_{(k+j+m-1)/2} = 1 - t_{k+j+m-1}, \tag{6}$$

so we get a contradiction once we show that $t_{k+j} = t_{k+j-1}$. One way to see this[7] is

$$t_{k+j-1} = 1 - t_{(k+j-2)/2} = 1 - t_{(k+j-2)/4} = t_{(k+j)/2} = t_{k+j}. \tag{7}$$

  - Case 2b: $m = 3$. Choose j such that $1 \leq j \leq 3$ and $k + j \equiv 2$ or $3 \pmod{4}$. If $k + j \equiv 2 \pmod{4}$ then we get a contradiction as in the previous case. If $k + j \equiv 3 \pmod{4}$ then consider

$$t_{k+j} + t_{k+j-1} = t_{k+j+3} + t_{k+j+2}. \tag{8}$$

The right hand side is even (the same as the left hand side in case 2a), while the left hand side is 1 since

$$t_{k+j} = 1 - t_{(k+j-1)/2} = 1 - t_{k+j-1}. \tag{9}$$

Contradiction.

Hence, the Thue-Morse sequence is overlap-free, so the corresponding game of chess was infinite, under the rules of the day.

[7] It is also immediate from the second alternative definition of the Thue-Morse sequence.

Subsequent developments

There is something unsettling about the possibility of a game having infinitely many moves, even if it is never going to happen in practice. The rules were reformulated several times during the 19th and 20th centuries, but the current statement of the threefold repetition rule is as follows:

> A player may claim a draw if the same position occurs for the third time.

This takes into account the eligibility status of castling and *en passant*. The number of possible positions is easy to bound, which now bounds the length of a game.

There are 13^{64} ways to put pieces on the board; here $13 = 2 * 6 + 1$ takes into account 6 types of pieces and 2 colours, as well as the possibility of no piece being on the square. There are 2^6 combinations of long-term castling eligibility (depending on whether or not each king or rook has been moved), and at most $2^{2(2+2*6)} = 2^{28}$ combinations of long-term *en passant* eligibility. We also multiply by two because it could be white to move or black to move. Thus, $C = 2^{35} * 13^{64}$ is an upper bound on the number of chess positions. By the pigeonhole principle, threefold repetition must be achieved within 2C ply (a ply is a move for white or black; the starting position is attained after 0 ply), which is C moves (for each side), provided that someone claims the draw when they are allowed to![8]

Euwe continued teaching mathematics at schools and universities, and became a professor at Tilburg University in 1964. He wrote over seventy chess books, presided over FIDE (the world chess federation) as President from 1970 to 1978, and was also involved in the development of computer science.

— Sam Chow

In a park people come across a man playing chess against a dog. Astonished, they say: “What a clever dog!"
But the man protests: "No, he isn’t that clever, I’m leading three games to one!"

[8]This bound is lowered significantly using the 50-move rule: a player may claim a draw if at least 50 consecutive moves have been made by each side without the capture of any piece or the movement of any pawn.

Pythagorean Triples

Triples of the Pythagorean kind are such well-known beasts that giving this article any other name besides the boring straightforward one just seems plain wrong. And yet, there are questions one can ask that lead to interesting answers; questions you will find below. So let's start at the natural start.

Fermat's Last Theorem states that there are no integer solutions (a, b, c) to $a^n + b^n = c^n$ for $n > 2$. But what about $n = 2$? Well...

Theorem
There is an integer solution (a, b, c) to $a^2 + b^2 = c^2$. That is, there exists a Pythagorean triple.

Proof
$3^2 + 4^2 = 5^2$.

So great, there is *one* solution. But are there more? Actually, yes.

Let (a, b, c) be a Pythagorean triple. Then

$$(ka)^2 + (kb)^2 = k^2a^2 + k^2b^2 = k^2(a^2 + b^2) = k^2c^2 = (kc)^2.$$

So if k is a positive integer, then (ka, kb, kc) is a Pythagorean triple.

Well that's all very good and interesting, but really all those Pythagorean triples are in some sense the same. At the very least, they can all be generated from the one original Pythagorean triple. What about when a and b have no common factors? That is, they are what's known as *coprime*?

In fact, the answer here is still yes, but now we have to introduce a formula of Plato's: Euclid's formula.

Euclid's Formula Theorem
Every Pythagorean triple can be constructed using positive integers n, m (and k) using the following formula:

$$a = k(m^2 - n^2),\ b = k(2mn),\ c = k(m^2 + n^2)$$

In particular if $k = 1$, n, and m are coprime and $m - n$ is odd, then the triple has no common factors and is called a *primitive triple*. If n and m are both odd, then a, b, and c are even and so the triple is not primitive (but dividing through by 2 will result in a primitive triple).

Proof
To show that this formula generates Pythagorean triples is easy; you just have to show that $a^2 + b^2 = c^2$. To show that any primitive Pythagorean triple can be written in that form for some n and m is also pretty easy, and involves writing the terms as a fraction in simplest form. I have omitted both proofs here, but they're easy enough that you should try them yourself. If you get stuck, just check out the helpful Wikipedia page.

But with $k = 1$, there are infinitely many n and m coprime with $m - n$ being odd (can you prove this?) so there are infinitely many primitive triples. Hence, does that mean that every number is in a primitive Pythagorean triple?

No, as 6 is not part of a primitive triple. We can prove this by showing that 6 cannot be element a or c (by bounding both of them) and then showing that when $6 = b$, the triple is not primitive. In fact, any number that can be written in the form $4d + 2$ for an integer d (that is, numbers congruent to $2 \, mod \, 4$) cannot be part of a primitive triple. Strikingly, the contrapositive is also true: every number that is not congruent to $2 \, mod \, 4$ is in a primitive triple.

This means that there are an infinite number of primitive triples, but it is a smaller infinity than the infinite number of integers. You might have to wait until beyond second year to fully appreciate this.

However, if we let $m = 2d + 1$ and $n = 1$, then $b = 2mn = 4d + 2$. So all those numbers we just excluded from the primitive triple club can in fact find their way into a non-primitive triple. This means that the set of all numbers that are in a Pythagorean triple is actually a big happy family. Every integer is in one! Except for 1 or 2...

— Giles David Adams

Some say that the Pope is the greatest cardinal. But others insist that this simply cannot be, as every Pope must have a successor.

Statistics on Graphs

Our civilization has become intimately familiar with statistics. We routinely accept likelihood estimates in everyday life, we don't blink an eye when our marks are scaled to a regular distribution, and we expect ordinary office software to perform least squares regression. We are used to seeing statistical analysis done on almost any numerical figure we can imagine, seemingly unaware of the implicit constraints required to fit the framework of classical statistics. What constraints, you ask? Perhaps this is a question best answered by example.

Example 1: Darwin's Finches. A quarter-century before *On the Origin of Species*, Charles Darwin was in the Galápagos Islands, and noticed that finches with similar beak shapes tended to occur on different islands. Since beak shape is directly related to diet, this suggests some species drove out others through competition for food. We all know where this train of thought eventually led, but what about the initial observation? How do we reconcile the statement "birds with similar beaks tend to occur on different islands" with our modern notion of statistical significance?

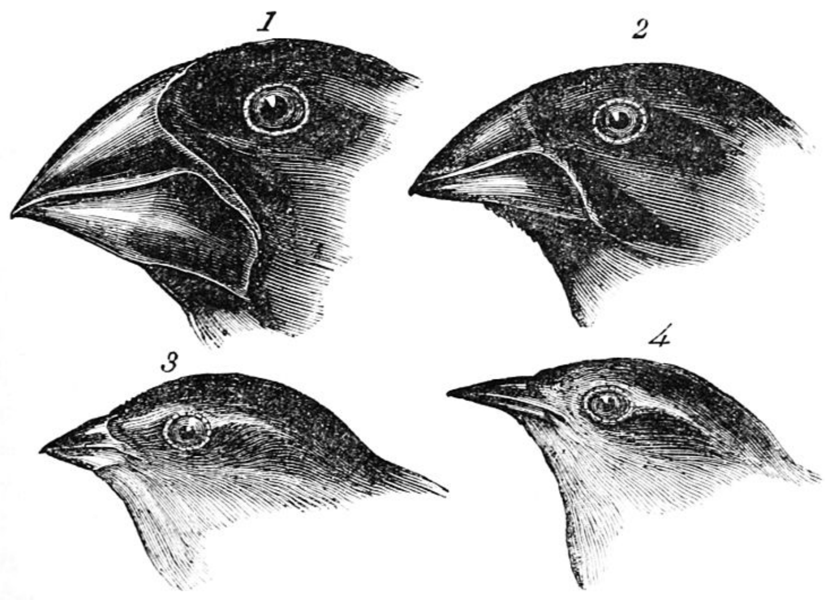

1. Geospiza magnirostris. 2. Geospiza fortis.
3. Geospiza parvula. 4. Certhidea olivasea.

Example 2: Six Degrees of Separation. In our increasingly connected world, the idea that two typical people are separated by about six friend-of-a-friend connections seems almost mundane. Recent studies have found the average distance on Twitter to be around 4, while Facebook lies slightly below 6. Whatever the figure might be, we still have a statistical conundrum: how do we decide whether or not this is surprising? More precisely, is this a mathematical phenomenon that naturally arises whenever we form links between things, or is it a sociological one that reflects the way humans form friendships?

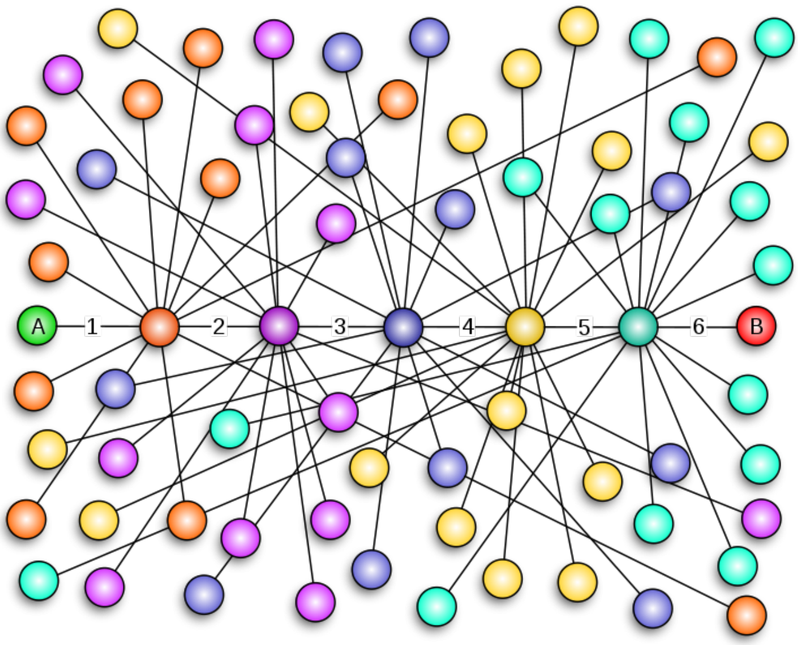

Finding statistics in these examples is not difficult. They both ultimately boil down to a single figure—for Darwin's finches, we are interested in the number of occurrences of similar-beaked birds on the same island, and for six degrees of separation, we care about the average separation distance. We can still naïvely ask our usual question: "What is the likelihood of observing this statistic?"

The problem arises when we try to define what we mean by "likelihood". Neither statistic in question readily fits the classical interpretation of a number sampled from a repeatable experiment. In both cases, there is only a single data point—there is exactly *one* distribution of finches in the Galápagos Islands, and there is exactly one friendship network of the world. Without a probabilistic backdrop against which to compare, the idea of likelihood makes no sense.

Yet within our single data point lies a multitude of information. The friendship network of the world is a graph—each pair of vertices (people) contains one bit of information, namely whether or not those two people are friends. These bits are far from independent (friends of friends are more likely to be friends), but we should still be able to extract plenty of useful information. Similarly, Darwin's observation is a bipartite graph—vertices are either finches or islands, and our bits of information are again the absence or presence of edges, that is, whether or not a particular finch lives on a particular island.

Now we can pull out some probabilistic machinery. If we inject randomness into our edges, we obtain a random source of graphs, which provides the missing backdrop against which to compare our real-world observations. Armed with a probability distribution of graphs, the question "What is the likelihood of observing this statistic?" suddenly makes sense.

But what distribution should we pick? We already concluded that the edges aren't independent, so picking edges randomly won't work. Looking at the vertices instead, we see that they are fixed—there are a fixed set of people in the world, and a fixed set of finches on a fixed set of islands. Going further, it makes sense to consider the number of connections at each vertex to be fixed, that is, to fix the friendmaking potential of each person, the prolificness of each finch, and the food abundance of each island. Fixing the *degree sequence*—the set of vertices along with the number of connections at each vertex—we obtain a natural probability distribution on graphs, where each graph with the given degree sequence occurs with equal probability.

Problem solved, right? Actually, not quite. Here statistics ends and mathematics begins. Statistics defines the problem—assume the underlying distribution is uniform among those with the observed degree sequence, and calculate the likelihood of the statistic in question. Mathematics is what we need to actually deal with this complicated yet somehow elegant underlying distribution.

Probability is perhaps unique as a field of mathematics where a few non-technical pages can bring us to the forefront of active research. We conclude by presenting a beautiful solution by Joseph Steger and Melbourne's very own Nick Wormald from the turn of the century, which was proved by Bayati, Kim and Saberi in 2010 (some minor tweaks omitted for simplicity):

1. Arrange all the people in the world around a circle (you'll need a very large circle, or a very abstract imagination).
2. Ask each person how many friends they would like to have. Give each person one end of that number of strings, and gather the loose ends of everyone's strings at the centre of the circle.
3. Pick two loose ends at random, and tie them together, as long as it doesn't connect a person to themselves or repeat a previous connection.
4. Repeat the previous step until there are no loose ends left.

As long as no one requests a huge number of friends, you won't get stuck with loose ends that can't be tied, and the distribution of the friendship graph you create is asymptotically close to the one we want for our statistical analysis.

Can you figure out the corresponding solution for Darwin's finches?

— James Zhao

Q: How many mathematicians does it take to screw in a light bulb?
A1: None. It's left to the reader as an exercise.
A2: None. A mathematician can't screw in a light bulb, but he can easily prove that the work can be done.
A3: The answer is intuitively obvious.
A4: Just one, once you've managed to present the problem in terms he or she is familiar with.
A5: In earlier work, Weiner (2004) has shown that one mathematician can change a light bulb.
A6. If k mathematicians can change a light bulb, and if one more simply watches them do it, then $k+1$ mathematicians will have changed the light bulb. Therefore, by induction, for all n in the positive integers, n mathematicians can change a light bulb.

Paradox would like to thank Giles David Adams, Mel Chen, Sam Chow, and James Zhao for their contributions to this issue.