

ANNEX

1. [to Chapter Three, page 66] The discovery of the largest twin-primes to date.

David Underbakke and Phil Carmody found these titanic twins on March 27, 2001, in Minnesota. The work involved sieving on several machines for many gigahertz-months—but the sieving was much more sophisticated than that of Eratosthenes, and the cleverness lay in knowing where to look, how to search economically, and how to test whether a candidate was prime. Their research, in contemporary style, is compounded with that of many others.

Here is the announcement they e-mailed out.

X-External-Networks: yes
Precedence: bulk
Approved-By: "Victor S. Miller" <victor@IDACCR.ORG>
Date: Thu, 29 Mar 2001 13:33:58 -0500
Reply-To: Phil Carmody <fatphil@altavista.com>
Sender: Number Theory List <NMBRTHRY@LISTSERV.NODAK.EDU>
From: Phil Carmody <fatphil@altavista.com>
Subject: Worlds largest twins found in Minnesota
To: NMBRTHRY@LISTSERV.NODAK.EDU

Four months ago, to the day, David Underbakke was pleased to announce[1] the discovery of the largest known twin primes[2], $665551035 \cdot 2^{80025} \pm 1$, with 24099 digits. The above was achieved after searching less than the expected search-space, so, feeling encouraged by that, we decided to collaborate again to attack a higher target.

The target has finally been reached: at 29603 digits each, $1807318575 \cdot 2^{98305} \pm 1$ are twin primes.

One of the lessons learnt from the former search was the importance of thorough pre-sieving of the ranges to remove all candidates with small ($<10^{14}$ or similar magnitude) prime factors. Pre-sieving reduced the number of lengthy probable-primality tests required. This is doubly effective when searching specifically for twin primes, as the density

improvement #twins/PrP-test is squared (e.g. sieving further to increase the single-prime PrP density by 10% increases the twin density by 21%).

After some discussion, we came to the conclusion that we could get the most from the sieving stage by again looking at primes of the form $K \cdot 2^N \pm 1$, and by choosing an exponent N optimised for the sieving stage. Phil had in the past used this technique, and we decided to simply use the same exponent, $N = 98305 = 0x18001$. As before, a dedicated twin sieve was used, using Phil's own code on his Alpha 21164 (which was hand-optimised specifically for this exponent) and using Paul Jobling's NewPGen[3], a general-purpose pre-siever.

An arbitrary odd number of this size (assuming the mean K tested is of order $K=10^9$) is prime with probability $P(\text{single}) = 2/\ln(10^9 \cdot 2^{98305}) = 1/34080$

Therefore the probability of finding a twin prime was $>P(\text{Twin}) = P(\text{single})^2 \cdot 1.32 = 1/880M$ where the adjustment factor 1.32 is twice the twin prime constant[4] (and M is million).

Therefore an odd K range of 1 .. 2.09G (i.e 1045M candidates) could be expected to yield one twin. (In retrospect, it appears our failure potential was a bit high perhaps!)

We sieved for about 2GHz-months on several machines, until we had sieved up to $p=50T$. Using Merten's theorem[5], we can predict that a simple sieve to $p=50T$ increases the density for arbitrary numbers by $1.781 \cdot \ln(50 \cdot 10^{12}) = 56.2$, but a twin sieve increases the density by $56.2^2/1.32 = 2390$. However, we were not checking even numbers anyway, by the construction of our expression, so the real density increases ought to be 1:28.1 and 1:1195, respectively.

This predicts a candidate count within .1% of what we ended up with - 875000 candidates.

After that it was 'simply' a matter of distributing the work across the various machines. The probable primality tests were done using George Woltmann's PRP program. In order to give Phil's Alpha something to do once the sieving was complete, we opted to do the +1 test first (for the Alpha to prove formally) and only perform the -1 test on PrPs from the +1 test.

Annex/The Art of the Infinite

As hinted at before, the bulk of the work began about four months ago (Phil started the sieve in advance). David was able to put between 10 and 15GHz onto the project on average, and Phil about half that. During the search we found many hundred solo primes, but a few weeks ago, as we started handing out the final blocks >to the various PCs, it looked as if the twin search would be fruitless. However, yesterday, with perfect thriller timing, PRP provided a positive result to the -1 test on one of David's machines. It was a very quick matter to then prove the -1 case formally[6] using Yves Gallot's Proth, and the $+1$ case[7] using Phil's own code.

Our thanks go to the authors of the software mentioned above and to Professor Caldwell for the Prime Pages website, without which we'd not have been equipped to make our predictions (and thus know where to aim).

David Underbakke,
Phil Carmody - 28/03/2001

References

[1] Underbakke; Largest Twin Primes Record Broken, Nov 28 2000
><http://listserv.nodak.edu/scripts/wa.exe?A2=ind0011&L=nbrthry&F=&S=&P=2877>

All of the following are from Professor C. Caldwell's paedagogical Prime Pages, and are given relative to and can be navigated to from both
<http://primepages.org/> and
<http://www.utm.edu/research/primes/>

[2] Twin Primes.
</lists/top20/twin.html>

[3] NewPGen.
</programs/NewPGen/index.html>

[4] Twin Prime Constant
</glossary/TwinPrimeConstant.html>

[5] Merten's Theorem
</glossary/MertensTheorem.html>

[6] $n - 1$ tests and Pepin's Test For Fermats.
/prove/prove3_1.html

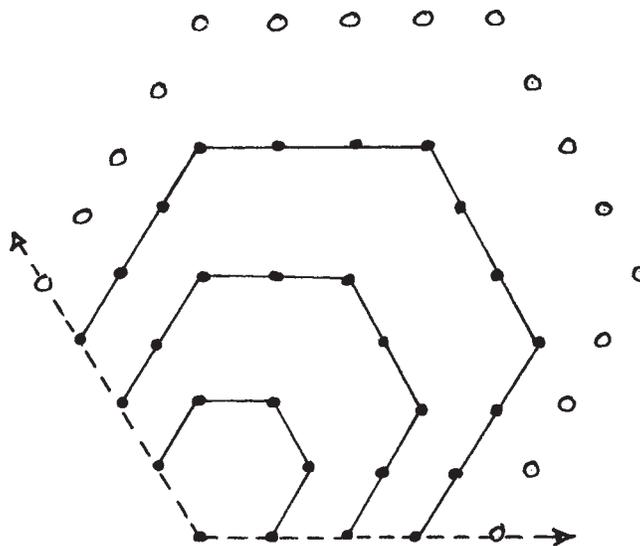
[7] $n+1$ tests and the Lucas Lehmer Test.
/prove/prove3_2.html

2. [to Chapter Four, page 85] The formula for the n th k -gonal number.

For our proof that the n th k -gonal number, P_n^k , is

$$\frac{n(nk - 2n - k + 4)}{2},$$

we'll go back to our study of how a pentagonal number grew from the one before it— P_{n+1}^5 from P_n^5 —and notice that what was structurally true there remains true for any k -gon: the number of dots on a side will increase by 1 in this growth. P_n^k has n dots on a side, P_{n+1}^k has $n + 1$. Since the new figure is made from the old by extending two old sides and then completing the circuit of sides to the total k , it must always be that $k - 2$ new sides are made for the new figure.



Annex/The Art of the Infinite

That means we are adding $(k - 2)(n + 1)$ new dots—but with the same qualification as in the text on page 81 about shared dots at corners—of which there will be $k - 3$ new ones. The number of dots in P^k therefore is:

$$\begin{array}{ccccccc}
 P_{n+1}^k & = & P_n^k & + & (k-2)(n+1) & - & (k-3) \\
 \uparrow & & \uparrow & & \uparrow & & \uparrow \\
 \text{new} & & \text{old} & & \begin{array}{c} \text{new sides} \\ \text{new dots} \\ \text{per side} \end{array} & & \begin{array}{c} \text{shared dots} \\ \text{at new} \\ \text{corners} \end{array}
 \end{array}$$

Too much is happening here too fast. We need any help we can get to make sense of it all. Let's simplify a bit by letting g stand for $k - 2$, as it did before. This gives us

$$P_{n+1}^k = P_n^k + g(n + 1) - (k - 3)$$

and since $k - 3 = k - 2 - 1$, which is $g - 1$, this artful flick of the wrist gives us

$$P_{n+1}^k = P_n^k + g(n + 1) - (g - 1)$$

or just

$$P_{n+1}^k = P_n^k + gn + 1.$$

This tells us how the $(n + 1)$ th k -gonal number grows from the next smaller. What we'd like to do, however, is watch this clamber up from the very first k -gonal number, P^k (which is always 1: the single dot \bullet).

Here is the tower of gymnasts:

$$\begin{array}{rcl}
 P_2^k & = & P_1^k + g + 1 \\
 P_3^k & = & P_2^k + 2g + 1 \\
 P_4^k & = & P_3^k + 3g + 1 \\
 P_5^k & = & P_4^k + 4g + 1 \\
 & \vdots & \\
 P_n^k & = & P_{n-1}^k + (n - 1)g + 1
 \end{array}$$

It is tempting now to add them all together, but one last bit of juggling will save us a lot of work. Let's just rearrange our tower by moving the first term on the right of each equation to the left, via subtraction. In this way all but two of the terms on the left will cancel out when we add (the canceling occurs "on the diagonal"):

$$\begin{aligned}
 P_2^k - P_1^k &= g + 1 \\
 P_3^k - P_2^k &= 2g + 1 \\
 P_4^k - P_3^k &= 3g + 1 \\
 P_5^k - P_4^k &= 4g + 1 \\
 &\vdots \\
 P_n^k - P_{n-1}^k &= (n-1)g + 1
 \end{aligned}$$

Now add:

$$P_n^k - P_1^k = g(1 + 2 + \dots + (n-1)) + (n-1).$$

Re-adding P_1^k to the right-hand side and undoing the last contortions, this is

$$P_n^k = P_1^k + \frac{g(n-1)n}{2} + (n-1).$$

Since P_1^k is 1, and g is $k-2$, we have

$$P_n^k = 1 + \frac{(k-2)(n-1)n}{2} + (n-1)$$

which, with a final trumpeting from the elephants, becomes

$$P_n^k = n \cdot \frac{(nk - 2n - k + 4)}{2},$$

just as we'd hoped. Our intuition may have leapt at the infinite from the corner of a table, but our proof has encompassed it smoothly through the flexible power of thought.

3. [to Chapter Five, page 106] Constructing a perpendicular to a line.

Euclid gives us (I. 11) an easy way to construct a perpendicular to a line ℓ at a point P on it. With your compass strike an arc from P intersecting ℓ at A to its left and B to its right; then with radius AB , arcs from A and B meeting at C . Draw PC . $\triangle APC \cong \triangle BPC$ (by SSS), hence $\angle CPA \cong \angle CPB$; and since their sum is a straight angle, each is a right angle: $CP \perp \ell$.

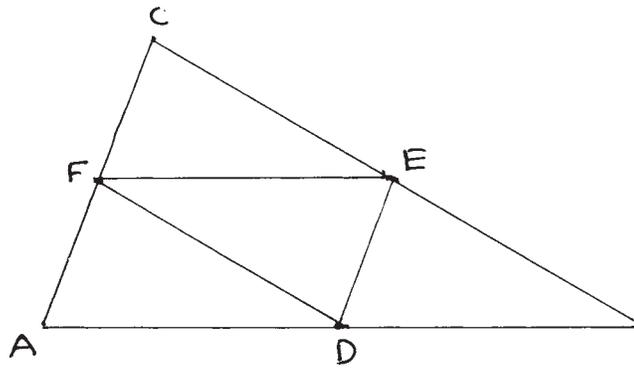
Behind this construction lies the question: why not use a right-angle, as carpenters do, to draw the line at P perpendicular to ℓ ? Why this Greek prejudice against tools other than the conceptual and, for that matter, against measuring lengths with a ruler and angles with a protractor?

Part of the answer is that it clarifies the mind and its view of the world to see the hierarchy of truths built up from the least foundations. If angles can be copied and perpendiculars and parallels drawn with compass and straight-edge alone, we then know that these constructions lie on a more fundamental level than where measuring takes place. The vast array of Euclid's results depend only on these two non-metric tools.

There is more, however. Measuring ever approximates to an ideal, mathematical construction is of the ideal itself. These objects of Euclid's, or of any part of mathematics, lie not in but behind, or below, or beyond the sensed world: *we invent them* as the limits toward which our strivings converge.

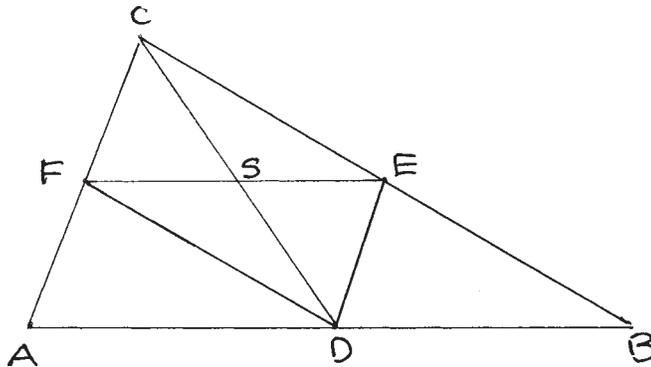
4. [to Chapter Five, page 116] A proof via infinite sequences that the medians of a triangle concur.

Let's take our $\triangle ABC$, put in the three midpoints D , E , and F of its sides, and connect them.



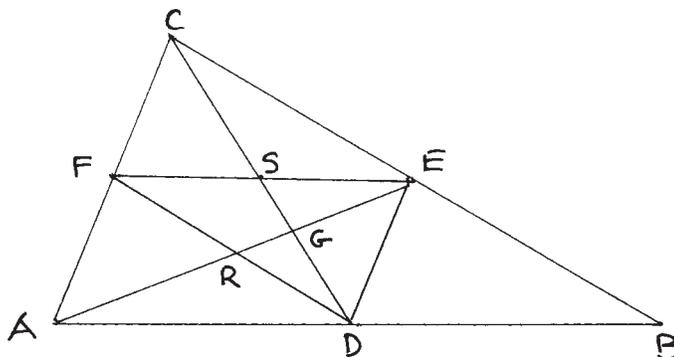
Because the line joining midpoints is parallel to the base (result (1) on page 114), $\triangle EFD$ is an upside-down, half-size copy of $\triangle ABC$: similar to it, that is (in symbols, $\triangle ABC \sim \triangle EFD$, where the order of letters represents the paired, similar parts). We'll call $\triangle EFD$ the "midpoints-triangle".

Construct the median CD of $\triangle ABC$, intersecting EF at S .



What's really nice is that DS is a median of the midpoints-triangle too. Why? Because on the diagram's right-hand side, $\triangle CSE \sim \triangle CDB$ (they share the top angle, and the parallel lines SE and DB make the angles formed with the transversal CSD the same). Since CE is half of CB , SE must be half of DB . On the diagram's left-hand side $\triangle CSF \sim \triangle CDA$ for the same reasons, so FS is half of AD . Since $AD = DB$ (D is a midpoint), $FS = SE$ —that is, S is the midpoint of FE , so DS is a median.

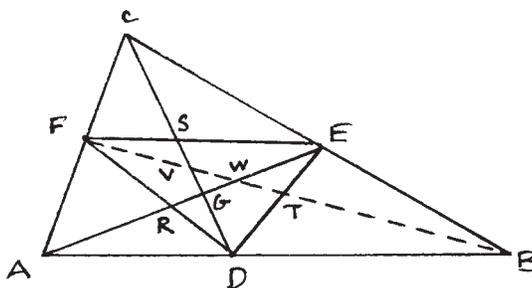
When we construct the median AE of the big triangle, intersecting DF at R , again ER will be the median of this midpoints-triangle.



Where is this proof going? We want to prove that the three medians of a triangle are concurrent at some point G , and so far have two of them meeting *inside* the midpoints-triangle—where parts of the medians turn out to be the full medians of that half-size triangle.

By the way, how do we know (other than by looking) that the medians AE and CD have to meet *inside* $\triangle DEF$? Basically because each must first pass through a different side of the midpoints-triangle.

Now when we draw the big triangle's median from B , going through DE at T , the worst case scenario would be that it miss G and instead intersect AE in some other point W , and CD in a third point V :

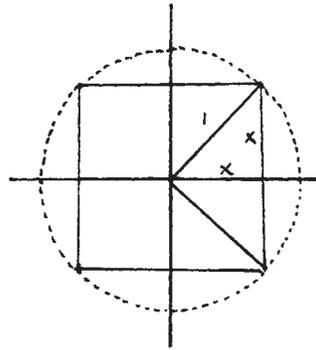


But even were this to happen, all three points— G , V , and W —would have to lie in the confined space of $\triangle DEF$ (a space one-fourth the area of the original triangle).

The touch of magic: repeat the argument exactly for the new, tiny midpoints-triangle $\triangle RST$ (with $\triangle DEF$ playing the role of the larger triangle). Again, the supposedly different points G , V , and W must now lie in the cramped quarters of this new $\triangle RST$ (one-sixteenth the area of $\triangle ABC$). Keep doing this (the midpoints-triangles flipping up and down as they diminish), getting an infinite sequence of nested triangles each a quarter the area of the preceding, with G , V , and W running around inside them. But this sequence of areas approaches zero—so that what might have been three different points collapses into one: the centroid G we sought.

5. [to Chapter Six, page 164] A Companion Miracle, which will appear at the end of a journey through thickets of square roots.

If we draw a square in a circle of radius 1 (a *unit circle*), the Pythagorean Theorem tells us that the square's sides will each be $\sqrt{2}$ long:



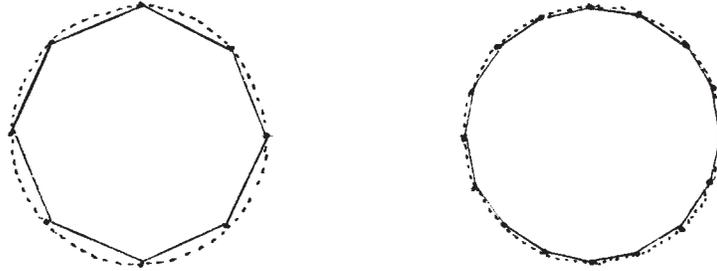
$$2x^2 = 1 \text{ so } x = \sqrt{\frac{1}{2}}, \text{ which is } \frac{\sqrt{2}}{2}$$

$$\text{and } \frac{\sqrt{2}}{2} + \frac{\sqrt{2}}{2} = \sqrt{2}.$$

The perimeter of the square will therefore be $4\sqrt{2} \approx 5.657$ —well short of the circle's circumference ($2\pi r$, which with $r = 1$ is $2\pi \approx 6.283$) as we would expect, since the square fits inside the circle with space left over.

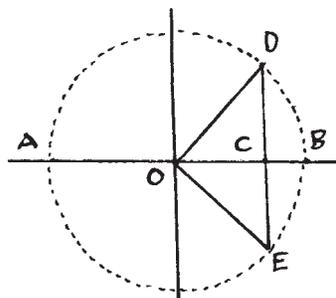
Annex/The Art of the Infinite

We know how to build from the square to an octagon, and doing this will, we know, increase its perimeter; and were we to double the octagon to a 16-gon, its perimeter would creep even closer to the circle's circumference as the polygon fits ever more snugly inside it.



Clearly the length of the circle's circumference will be the limit which these increasing perimeters approach. Let's find a way of expressing what the side-length at the k th stage of this doubling will be (since from side-length and number of sides we can quickly find perimeter). As always, we will bring back what we need from a remoter kingdom: we shall ask how to express the side-lengths of *any* $2n$ -gon in terms of the n -gon's side-length. The safari begins.

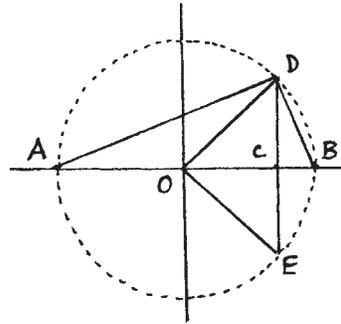
Let S_n stand for the side-length of an n -gon inscribed in the unit circle. Because it *is* a unit circle, all its radii are 1: $OA = OB = OD = OE = 1$, and DE is the n -gon's side, of length S_n .



DE crosses OB at C and $DE = 2DC$ ($\triangle OCD \cong \triangle OCE$ by SAS, so $CD = CE$).
Hence $CD = \frac{DE}{2}$, which is $\frac{S_n}{2}$.

Annex/The Art of the Infinite

Now let's draw in a side of the $2n$ -gon; it will be BD , since OB is the bisector of $\angle DOE$ (from those congruent triangles, $\angle DOB = \angle BOE$). So the length of BD is S_{2n} ; we also know that $AB = OA + OB = 2$. Now construct AD .



Reaching back to Thales, we know that $\angle ADB$ is a right angle, and—fetching from afar—we decide to make use of the fact that the *area* of $\triangle ADB = \frac{(\text{base} \cdot \text{height})}{2}$, which here would be $\frac{BD \cdot AD}{2}$.

On the other hand, the area of $\triangle ADB$ is *also* (using another choice of base and height) $\frac{AB \cdot CD}{2}$.

We have found two different ways of expressing the same quantity—always a promising step, because now we have an equality:

$$\frac{BD \cdot AD}{2} = \frac{AB \cdot CD}{2},$$

or,

$$BD \cdot AD = AB \cdot CD.$$

But we already know that $AB = 2$, $BD = S_{2n}$ and $CD = \frac{S_n}{2}$; so

$$S_{2n} \cdot AD = 2 \cdot \frac{S_n}{2}, \text{ that is, } S_{2n} \cdot AD = S_n.$$

Now, by the ever-valuable Pythagorean Theorem,

$$AD = \sqrt{AB^2 - BD^2} = \sqrt{2^2 - (S_{2n})^2} = \sqrt{4 - (S_{2n})^2}.$$

Annex/The Art of the Infinite

We said above that $S_{2n} \cdot AD = S_n$, or writing it backwards,

$$S_n = S_{2n} \cdot AD,$$

and our new expression for AD transforms that into

$$S_n = S_{2n} \cdot \sqrt{4 - (S_{2n})^2}.$$

If we lift the lid of the radical by squaring both sides, we get:

$$(S_n)^2 = (S_{2n})^2 \cdot (4 - [S_{2n}]^2).$$

This looks a little—but only a little—better. Remember that what we are looking for is a way of expressing S_{2n} in terms of S_n , and we have it almost in our grasp. For when we multiply out the last equation we get

$$(S_n)^2 = 4 (S_{2n})^2 - (S_{2n})^4.$$

The notation that helped us this far is now standing in our way, with all those subscripts and exponents. Let's simplify by taking a daring step, and letting x stand for $(S_{2n})^2$. Then the last equation looks like this:

$$(S_n)^2 = 4x - x^2$$

or

$$x^2 - 4x + (S_n)^2 = 0,$$

which is a quadratic equation of the form $ax^2 + bx + c$, and we can solve it by the Quadratic Formula:

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}.$$

In our equation we have $a = 1$, $b = -4$, and $c = (S_n)^2$, so we get

$$x = \frac{4 \pm \sqrt{16 - 4(S_n)^2}}{2}$$

which reduces to

$$x = 2 \pm \sqrt{4 - (S_n)^2}.$$

Algebra has produced two possibilities for x :

$$2 + \sqrt{4 - (S_n)^2} \text{ and } 2 - \sqrt{4 - (S_n)^2}.$$

But the geometry of our diagram shows us that x , which is $(S_{2n})^2$ (that is, the length $(BD)^2$), must be significantly less than 2, since BD is less than 1. So the only possibility is

$$x = 2 - \sqrt{4 - (S_n)^2}.$$

Now replace x by the $(S_{2n})^2$ it masked:

$$(S_{2n})^2 = 2 - \sqrt{4 - (S_n)^2}$$

so

$$S_{2n} = \sqrt{2 - \sqrt{4 - (S_n)^2}}.$$

We *have* come back with a way of expressing S_{2n} (the side-length of the doubled polygon) in terms of S_n , the side-length of the polygon it doubled.

Was it worth the voyage, as Michelin might ask? Watch.

For a square, as we saw, $S_4 = \sqrt{2}$.

$$\text{So } S_8 = \sqrt{2 - \sqrt{4 - (\sqrt{2})^2}} = \sqrt{2 - \sqrt{4 - 2}} = \sqrt{2 - \sqrt{2}}.$$

$$\text{Then } S_{16} = \sqrt{2 - \sqrt{4 - (\sqrt{2 - \sqrt{2}})^2}} = \sqrt{2 - \sqrt{4 - 2 + \sqrt{2}}} = \sqrt{2 - \sqrt{2 + \sqrt{2}}}.$$

$$\text{Next, } S_{32} = \sqrt{2 - \sqrt{2 + \sqrt{2 + \sqrt{2}}}}$$

and so (strangely) on.

This means that the side-length of S_{2^n} , when we start with a square, is

$$S_{2^n} = \sqrt{2 - \sqrt{2 + \sqrt{2 + \dots \sqrt{2}}}}.$$

($n - 1$ nested square roots)

To get the total perimeter of the 2^n -gon we have to multiply S_2 by 2^n , so that its perimeter is

$$P = 2^n \sqrt{2 - \sqrt{2 + \sqrt{2 + \dots + \sqrt{2}}}}$$

(for the 32-gon we get $P \approx 6.273$)

and P , as we said, approaches 2π as a limit:

$$\lim_{n \rightarrow \infty} 2^n \sqrt{2 - \sqrt{2 + \sqrt{2 + \dots + \sqrt{2}}}} = 2\pi$$

or, dividing both sides by 2,

$$\lim_{n \rightarrow \infty} 2^{n-1} \sqrt{2 - \sqrt{2 + \sqrt{2 + \dots + \sqrt{2}}}} = \pi.$$

We have an infinite expression made up of two of the most fundamental and elusive building blocks of mathematics: $\sqrt{2}$ and π .

6. [to Chapter Six, page 165] On Quadratic Reciprocity.

Gauss's great discovery of the criterion for polygon constructibility had an even greater offspring: it led to his Golden Theorem (*Theorema Aureum*), and two of its proofs. This theorem concerns small universes of natural numbers, like those on a clock-face, where 1 through 12 are all you have and all you need (13 o'clock is 1 o'clock again, 28 o'clock is 4 o'clock, and any natural number n appears in this world as the remainder you get on dividing n by 12). Any natural number, not just the clock cycle of 12, or the week cycle of 7, can set the limits to such a universe. In a 5-based world, for example, $2 + 2 = 4$, but $2 + 11 = 3$: the remainder on dividing 13 by 5.

We can have equations in such a world just as in the world of all the naturals: in the 5-based world, $7x = 2$ has the lowest solution $x = 1$, because $7 \cdot 1 = 7$ and 7 leaves a remainder of 2 when divided by 5. Of course $x = 6$ is also a solution: $7 \cdot 6 = 42$ which also leaves a remainder of 2 when divided by 5. Indeed, 11, 16, and any number of the form $1 + 5k$ will do the trick. If $7x = 2$ in the 12-based world, however, x can be 2 and its relatives, 14, 26—any number of the form $2 + 12k$. In general, once you have *one* solution to $ax = b$ in an m -based world, an infinite number of other solutions roll out—namely, the solution you found with any multiple of m added to it.

In these miniature worlds, as in the large one, some equations are harder to solve than others: so with some effort we can solve $x^2 = 13$ in a 17-based world ($x = 8$ will work, since $8^2 = 64$, which leaves a remainder of 13 when divided by 17). Some, however, are not just harder but in fact insoluble: $x^2 = 5$ has no solution in the world of 17.

The most interesting worlds are those with primes as their bases; since all the natural numbers can be expressed as the product of primes, worlds based on composites atomize to worlds based on primes. A major question, then, is when will there be a solution for $x^2 = q$ in a p -based world (where both p and q are prime)? What Gauss was able to prove—only four and a half months after coming up with the criterion for polygon constructibility—was that if $x^2 = q$ has a solution in a p -based world, then $x^2 = p$ will have a solution in a q -based world—or if one hasn't, the other won't either—*unless* p and q each leave a remainder of 3 when divided by 4 (such as $p = 11$ and $q = 19$). In this case, one of the two equations will have a solution and the other won't. This is the famous Law of Quadratic Reciprocity.

7. [to Chapter Eight, page 225] Another proof of Pappus's Theorem.

Magical as was the proof on pages 221–25 of Pappus's famous theorem, it leaned a little out of the projective plane into the Euclidean. Here is a proof that enters fully into the projective spirit.

You saw in the first section of the appendix to this chapter how an arbitrary *triple* of points ABC on a line ℓ could be sent to an equally arbitrary triple $A'B'C'$ on a line m by a projection—a chain of perspectivities (in this case, two):

$$\ell \xrightarrow{\alpha} n \text{ and } n \xrightarrow{\beta} m,$$

with

$$O(ABC) = A'B''C$$

and

$$P(A'B''C) = A'B'C',$$

so that altogether (writing the first perspectivity closest to what it worked on):

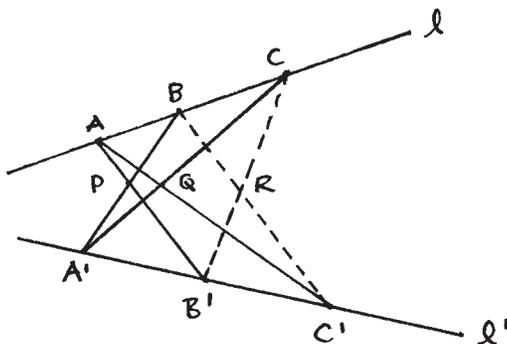
$$PO(ABC) = A'B'C'.$$

You also saw, thanks to the cross ratio, that *four* points on a line can't be sent to an arbitrary four on another, no matter how many or few perspectivities are in your chain. The segments made by the four points resulting from a projectivity must always have the same cross ratio as those made by the points they came from.

These observations led to the *Fundamental Theorem for Projectivity on a Line*, which we will put this way: If ABC is a triple of points on line ℓ , and $A'B'C'$ a triple on line ℓ' , then there is one, but only one, projectivity which sends A to A' , B to B' , and C to C' .

Rather than tracing out the proof of this key theorem, we will simply use it as an axiom in our new proof of

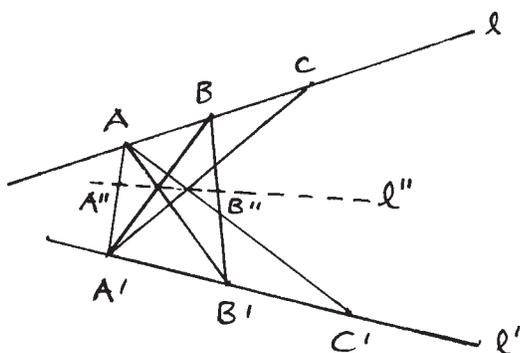
Pappus's Theorem: If A, B, C are on line ℓ and A', B', C' on line ℓ' , and AB' meets $A'B$ at P , AC' meets $A'C$ at Q and BC' meets $B'C$ at R , then P, Q , and R are collinear.



The strategy of this proof will be to draw the line ℓ'' through PQ and prove that R is on it. The tactics involve a series of perspectivities between the three lines, and proving (by means of the Fundamental Theorem) that a point R' that arises on ℓ'' from these perspectivities, is in fact R .

Proof:

1. Construct the line ℓ'' through P and Q , and lines AA' and BB' , meeting ℓ'' at A'' and B'' respectively.
2. Look first at $\ell \xrightarrow{A'} \ell''$: $A'(ABC) = A''PQ$.
Then look at $\ell'' \xrightarrow{A} \ell'$: $A(A''PQ) = A'B'C'$.
Doing A' followed by A therefore gives us: $AA'(ABC) = A'B'C'$.

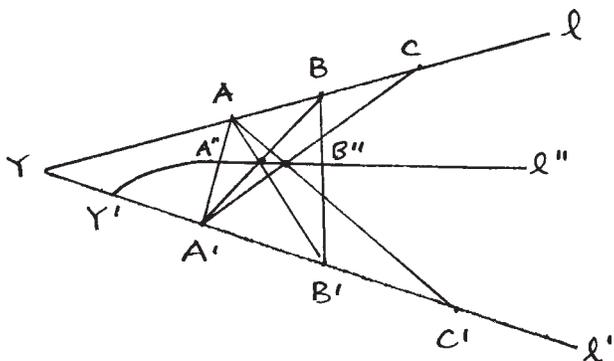


3. ℓ and ℓ' meet at a point—call it Y ; and ℓ'' meets ℓ' somewhere—call it Y' .

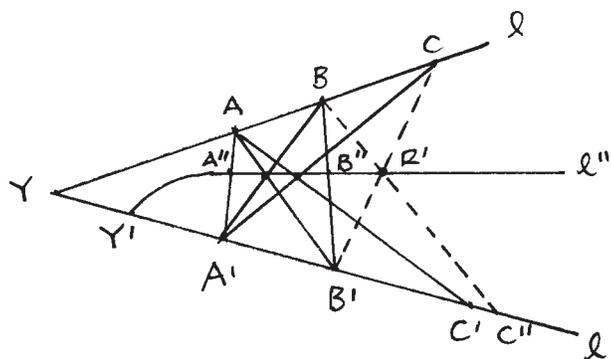
Let's see what our projectivity AA' does to Y : $A'(Y) = Y'$, and $A(Y') = Y'$.

Putting this all together: $A'(ABCY) = A''PQY'$, $A(A''PQY') = A'B'C'Y'$.

That is, $AA'(ABCY) = A'B'C'Y'$.



4. The time has come to construct $B'C$. Let it intersect ℓ'' at R' . Now construct BR' and extend it to meet ℓ' at C'' (we are given that $B'C$ and BC' meet at R but don't know that R is on ℓ'' . Our hope is to show that R' is R in disguise, by proving that C'' is really C').



5. Look now at the perspectivity $\ell \xrightarrow{B'} \ell''$: $B'(ABCY) = PB''R'Y'$,
and now at $\ell'' \xrightarrow{B} \ell'$: $B(PB''R'Y') = A'B'C''Y'$.

Doing first B' , then B , to $ABCY$ therefore gives us

$$BB'(ABCY) = A'B'C''Y'.$$

6. Comparing steps 3 and 5, you see that two different projectivities— AA' and BB' —have identical results on the three points A , B , and Y . By the Fundamental Theorem, these projectivities must be the same; hence, they must agree on the remaining point, which is C' in one and C'' in the other.

This can only mean that $C'' = C'$, hence R' must be R , and therefore R is on line l'' :

P , Q , and R are collinear, as desired!

NOTES

- 2: [Proof of general formula] Ross Honsberger, *Ingenuity in Mathematics* (Mathematical Association of America, 1970), pp. 117–19.
- 4: [Proof via diminishing midpoint-triangles] We first heard this proof from one of our Math Circle students, Jenny Chen. It may have been original with her.
- 5: [A companion miracle] Our appendix is based on the clear exposition given in Richard Courant and Herbert Robbins, *What Is Mathematics?* (Oxford University Press, 1941), pp. 124–25.
- 6: [Quadratic reciprocity] Euler and Lagrange first came up with this conjecture, but Gauss was the first to prove it (Goldman, p.184).