Solutions to October 2010 Problems

Problem 1. Find (with proof) all primes p such that $p+6$, $p+12$, $p+18$, and $p+24$ are all prime.

Solution. It is easy to check that $p = 2$ and $p = 3$ don't work. Let $p = 5$. We note that 5, 11, 17, 23, and 29 are all prime. So $p = 5$ works. We can now try $p = 7$, $p = 11$, $p = 13$, perhaps a few more. Nothing we tried except $p = 5$ works. We next *prove* that the only prime p such that $p+6$, $p+12$, $p+18$, and $p+24$ are all prime is given by $p = 5$.

The experimentation we did in testing various p may point the way. Test for example $p = 7$. We get the 5-term sequence 7, 13, 19, 25, 31. But 25 isn't prime. Try $p = 11$. We are looking at the 5-term sequence 11, 17, 23, 29, 35. The problem is at 35. Now try $p = 13$, $p = 17$, and so on. We may observe that we always seem to bump into a number greater than 5 whose decimal expansion ends in a 5. Any such number is divisible by 5, so cannot be prime. We will show that the phenomenon we observed with $p = 7$, 11, and 13 actually happens with any prime $p > 5$.

Suppose that p is an odd prime and $p \neq 5$. There are 4 possibilities: (i) the last digit in the decimal representation of p is 1; (ii) the last digit is 3; (iii) the last digit is 7; (iv) the last digit is 9.

In case (i), $p+24$ has last digit 5, so is obviously not prime. In case (ii), $p+12$ has last digit 5, so is not prime. In case (iii), $p+18$ has last digit 5, so is not prime. Finally, in case (iv), $p+6$ has last digit 5, so is not prime. We have shown that if p is a prime greater than 5, at least one of $p+6$, $p+12$, $p+18$, or $p+24$ is not prime. This completes the argument.

Comment. We worked with the decimal representation because it is concrete and familiar. But a minor modification is useful if we are interested in looking at similar phenomena. Let p be a prime other than 5. There are 4 possibilities: (i) the remainder when p is divided by 5 is equal to 1; (ii) the remainder is 2; (iii) the remainder is 3; (iv) the remainder is 4. In case (i), $p+24$ is divisible by 5. In case (ii), $p+18$ is divisible by 5. In case (iii), $p+12$ is divisible by 5. Finally, in case (iv), $p+6$ is divisible by 5. So, in all cases, one of $p+6$, $p+12$, $p+18$, or $p+24$ is divisible by 5 (and not equal to 5). Thus if p is a prime other than 5, at least one of $p+6$, $p+12$, $p+18$, or $p+24$ is not prime.

Comment. Many years ago, it was conjectured that for any positive integer n, there is an n-term arithmetic progression consisting entirely of primes. (It is easy to show that there cannot be an *infinite* arithmetic progression of primes.) Much work was done on the conjecture. By 2004, using a mixture of theory and heavy-duty computation, a 23-term arithmetic progression of primes had been found, but hopes for substantial progress beyond that seemed dim. However, in that year, using a difficult and beautiful argument, Ben Green and Terence Tao together proved that indeed, for every n, there is an n-term arithmetic progression consisting entirely of primes. Local note: Ben Green was a postdoctoral fellow at PIMS UBC in 2003-2004.

Comment. What is a prime? We consider only the integers, though the concept of prime has been defined in far more general settings. Fairly often, one sees the question "Why isn't 1 a prime?" A not entirely satisfactory answer is that the meaning of a mathematical term is defined by the mathematical community, and the mathematical community does not call 1 a prime. Why not?

Because it would be inconvenient. There are many theorems that go like this: "Let p be a prime. Then ...." If 1 were called a prime, almost all such theorems would have to start with "Let p be a prime other than 1. Then ...." A nuisance, more paper used, more forests cut down.

The above is not a full answer. The standard meaning of a mathematical term can change over time. (In Classical Greece, 1 was not even considered to be a *number:* the numbers were 2, 3, 4, ...!) Some people in the mathematical community *have* called 1 a prime, though as far as I know no one serious has done so for more than a century. I believe that the only prominent mathematician who called 1 a prime is Adrien-Marie Legendre (1752–1833).

However, there is good reason to consider -2, -3, -5, -7, and so on to be primes! In certain contexts (though not in school level discussions) *many* (indeed most) mathematicians do so. If we allow negative primes, there is one other answer to our problem, $p = -29$.

Problem 2. Consider the system of equations $x + y + z = a$, $xy + bz = 6$. For what integer values of a and b does the system have infinitely many integer solutions (x, y, z)?

Solution. It is natural to use the equation $x + y + z = a$ to eliminate one of the variables. And z looks like the best candidate, since getting rid of it leaves a more symmetrical expression than getting rid of x or y. So we look at the equation

$$xy + b(a - x - y) = 6. \tag{1}$$

The original system has infinitely many integer solutions (x, y, z) if and only if Equation 1 has infinitely many integer solutions (x, y). Rewrite that equation as

$$(x - b)(y - b) = b^2 - ba + 6 \tag{2}$$

(the idea is a mild variant of the usual "completing the square").

On the left-hand side of Equation 2 we have a product of two integers. If the integer $b^2 - ba + 6$ is non-zero, there are only finitely many integer solutions (x, y), since then $b^2 - ba + 6$ has only finitely many divisors. And if $b^2 - ba + 6 = 0$, there are infinitely many solutions, namely $x = b$, y anything, and x anything, $y = b$.

So we are trying to find the integer values of a and b such that $b^2 - ba + 6 = 0$. Rewrite this equation as $b(a - b) = 6$. The possibilities for b are ± 1, ± 2, ± 3, and ± 6. The rest is routine, and uninteresting. When $b = 1$, we need $a - b = 6$, so $a = 7$, and we get the pair $(a, b) = (7, 1)$. When $b = -1$, we need $a - b = -6$, giving the pair $(-7, -1)$. Continue. We get 6 other solutions, $(\pm 5, \pm 2)$, $(\pm 5, \pm 3)$, and $(\pm 7, \pm 6)$.

Another Way. There are other reasonable ways to continue after Equation 1. For example, we can attempt to solve for y. Equation 1 can be rewritten as

$$y(x - b) = bx - ba + 6. \tag{3}$$

Possibly $x = b$. That forces $b^2 - ba + 6 = 0$, and gives us infinitely many solutions, since y can be anything. So we end up examining the equation $b^2 - ba + 6 = 0$. This is done exactly as in the first solution.

If $x \neq b$, we can divide both sides of Equation 3 by $x - b$, obtaining

$$y = \frac{bx - ba + 6}{x - b}. \tag{4}$$

Rewrite this as

$$y = b + \frac{b^2 - ba + 6}{x - b}, \tag{5}$$

maybe by doing the standard "polynomial division" of $bx - ba + 6$ by $x - b$. To get infinitely many solutions (x, y) in integers, we must use infinitely many x, since y is determined by x. So $b^2 - ab + 6$ must be divisible by $x - b$ for infinitely many x. This is only possible if $b^2 - ab + 6 = 0$, an equation we have already dealt with.

Comment. We described two approaches partly because they introduce two different potentially useful ideas for dealing with an equation like Equation 1.

The second solution is much more awkward than the first. The first solution did not "break" the symmetry between x and y, while the second approach did. Often the best solution to a problem is the one that fully exploits any symmetry. One may be able to grind out a solution by breaking symmetry, but it is useful aferwards to look for a more symmetrical approach.

Problem 3. (a) Among all triangles that contain a 1×1 square, what is the smallest possible area? (b) And among these triangles that contain a 1×1 square and have smallest possible area, what is the smallest possible perimeter?

Solution. (a) It is reasonable to believe that in any triangle of smallest area, one side of the triangle must contain one of the sides of the square. For completeness one should prove this. The proof is not hard, but the full details involve a fair amount of picture drawing and writing. We omit this, but refer to Ivan Niven's highly recommended *Maxima and Minima Without Calculus.* (Local note: Ivan Niven was a UBC graduate.) Once we have decided that one side of the triangle must contain a side of the square, it is obvious that the other sides of the triangle must pass through the remaining vertices of the square. So the best configuration must look roughly like the figure below.

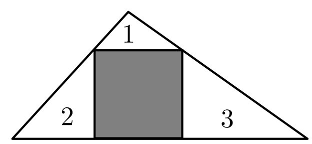

Note that Triangle 1 is similar to the full triangle. Imagine removing the square, and sliding Triangle 2 to the right until it meets Triangle 3. The result is a triangle (which we call Triangle 4, not drawn). Triangle 4 is also similar to the full triangle, and hence to Triangle 1.

To minimize the area of the full triangle, we need to minimize the sum of the areas of Triangles 1 and 4. Let x be the height of Triangle 1, and let y be the base of Triangle 4. Since the base of Triangle 1 is 1, and the height of Triangle 4 is 1, the sum of their areas is $(1/2)(x + y)$. And since the triangles are similar, we have

$$\frac{x}{1} = \frac{1}{y},$$

or equivalently $xy = 1$. So we want to minimize $x + y$ subject to the condition $xy = 1$. This is a problem that can be settled by the ever useful Arithmetic Mean—Geometric Mean Inequality. If that is not familiar, note that

$$(x + y)^2 = (x - y)^2 + 4xy = (x - y)^2 + 4.$$

To minimize $(x + y)^2$ (and hence $x + y$) we should make $(x - y)^2$ as small as possible. The best we can do is to make $x = y$. That forces $x = y = 1$.

Can we actually *achieve* this geometrically? Yes, and in many ways. Pick any line segment of length 2 that contains the bottom edge of the square. Draw the line that goes through the left end of this line segment and the top left corner of the square. Do the same with the right end of the line segment and the top right corner of the square. The triangle thus produced has the desired properties. Any such triangle has area 2.

Comment. Minimizing $x + y$ given that $xy = 1$ is equivalent to minimizing the perimeter of a rectangle given that it has area 1. It is not hard to produce a purely geometric argument that the square has the minimum perimeter among all these rectangles. However, the idea behind the algebraic argument above is far more widely applicable.

(b) In part (a), we saw that the "winning" triangles are precisely the triangles for which the triangle we called Triangle 1 has height equal to 1. And since Triangle 1 and the full triangle are similar, to minimize the perimeter of the full triangle it is enough to minimize the perimeter of Triangle 1.

It seems very plausible, some might even say obvious, that to minimize this perimeter we should make the triangle isosceles. If that is true, then the calculation of the smallest possible perimeter is easy. For then the winning full triangle is isosceles with base 2 and height 2. The two non-horizontal sides of the triangle each have length $\sqrt{1^2 + 2^2}$, so the triangle has perimeter $2 + 2\sqrt{5}$.

If we set up the problem algebraically, the minimization problem is very difficult to do without calculus, and would be considered kind of hard in a first year calculus exam. It is easier to look at the closely related problem of maximizing the area given the base and the perimeter. You may want to try. Heron's formula for the area of a triangle turns out to be the key.

But here is a little story that settles matters easily . Friends X and Y live 1 mile apart at the ends of the top edge of the square. A river runs 1 mile to the north, and parallel to this edge. X wants to go to the river, get some water, and take it to Y. What is the shortest possible path? Draw a possible path, like the one in the picture.

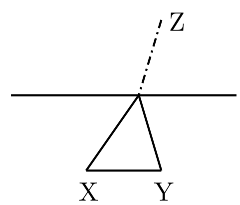

It so happens that X has another friend, Z, who lives exactly opposite Y, just as far to the north of the river as Y is to the south. The river is narrow and shallow, easy to cross, so X thinks about visiting Z instead. If so, the second part of his trip would be the dashed line in the picture—the reflection of the second part of the trip to Y. If the path south of the river was the minimal path to Y, then the part to the river, followed by the dashed line, is the minimal path to Z. The minimal path to Z, however, must be a straight line. Draw it. The line obviously meets the river due north of the halfway point between X's house and Y's house. Reflect the second part of this straight line path back so that it goes to Y (X decided not to visit Z). The two parts of the trip have the same length, the triangle is isosceles.

Comment. Another way to solve part (b) uses additional geometric knowledge. Look at the isosceles triangle $\mathcal{T}$ with (horizontal) base 1 and height 1. Call the endpoints of the base by the names F_1 and F_2. The sum of the distances from the top vertex to F_1 and F_2 is a certain number, which happens to be $\sqrt{5}$.

The points with sum of distances from the *foci* F_1 and F_2 equal to $\sqrt{5}$ trace out an ellipse. Points inside the ellipse have sum of distances from the foci less than $\sqrt{5}$. Points outside the ellipse have sum of distances from the foci greater than $\sqrt{5}$.

Now look at the top vertex of $\mathcal{T}$, and imagine moving it, parallel to the base, either to the left or to the right. Then this top vertex moves *outside* the ellipse (we need a Java applet here!), so the sum of distances becomes greater than $\sqrt{5}$.

Problem 4. Without using a calculator, express $2 \times 10^5 - \sqrt{10^{10}-1} - \sqrt{10^{10}+1}$ in "scientific" notation, correct to 2 significant figures. (You may want to use a calculator to help you guess what the answer might be.)

Solution. There is a high precision calculator bundled with Microsoft Windows. Or else we could use the very interesting (and free) Wolfram Alpha. There are many other free programs that do high precision calculations. And if one insists on paying for what is available for free, there is Mathematica, or (O Canada) Maple. The Microsoft Windows calculator gives something like

$$2.499999 \times 10^{-16}$$

(there are a bunch more 9's, then other stuff). So, to two significant figures, the Windows calculator gives 2.5×10^{-16}. This is correct. Maybe our analysis will come up with something similar. The Windows calculator actually shows about 32 digits. Interestingly enough, the answer it gives is already wrong in the 18th significant digit.

Now let us attack the calculation with a simple scientific calculator, whatever that means. Mine, which are simple, or at least cheap, all claim that the answer is 0. The answer is indeed *close* to 0. But it is not equal to 0, so none of my cheap calculators gives the result correct to even 1 significant figure! In fact they already fail disastrously when calculating

$$2000 - \sqrt{999999} - \sqrt{1000001}.$$

The problem is that 2000 and $\sqrt{999999} + \sqrt{10000001}$ are large and nearly equal, and a simple calculator does not carry enough digits to correctly find the difference. The difficulty we are in may seem artificial. But there are many real scientific computations in which the quantity we are interested in is the difference between nearly equal large numbers, so analogous issues come up surprisingly often in real world scientific computation.

Comment. Should we give up on the calculator? It may be a good idea to work with smaller numbers that the calculator really can handle. For example, one of my calculators gives

$$20 - \sqrt{99} - \sqrt{101} \approx 2.500079 \times 10^{-4}, \quad \text{and}$$

$$200 - \sqrt{9999} - \sqrt{10001} \approx 2.51 \times 10^{-7}$$

(here the third significant digit is already wrong). We could calculate related things, like $50 - \sqrt{24} - \sqrt{26}$. After some calculator experimentation, we could form a plausible conjecture about what happens in general. We do not pursue this further, but this kind of investigation can be very productive.

Enough preliminary comments! Let's tackle the problem directly, by finding excellent approximations to $\sqrt{10^{10}-1}$ and $\sqrt{10^{10}+1}$. Look first at $\sqrt{10^{10}+1}$. We want to solve the equation $u^2 = 10^{10}+1$. It is clear that u is close to 10^5. So let $u = 10^5 + v$ Note that v will be close to 0.

Substitute 10^5+v for u in the equation $u^2 = 10^{10}1$. Expand and simplify. We get $v^2+(2\times10^5)v = 1$. Since v is close to 0, the term v^2 is negligible compared to the other two terms. So $(2\times10^5)v \approx 1$. Thus $v \approx 5\times10^{-6}$. Let $v = (5\times10^{-6}) + w$. Substitute in the equation $v^2 + (2\times10^5)v = 1$ and simplify a bit. We get

$$w^2 + (2\times10^5 + 10^{-5})w + 1/(2\times10^{-5})^2 = 0.$$

The term w^2 is negligibly small in comparison with the other terms. And the coefficient of w is nearly 2×10^5. So w is approximately equal to $-1/(2\times10^{-5})^3$. We conclude that

$$\sqrt{10^{10}+1} \approx 10^5 + \frac{1}{2\times10^5} - \frac{1}{(2\times10^5)^3}.$$

Use the same technique to approximate $\sqrt{10^{10}-1}$. We get

$$\sqrt{10^{10}-1} \approx 10^5 - \frac{1}{2\times10^5} - \frac{1}{(2\times10^5)^3}.$$

Add the two estimates, subtract the result from 2×10^5. We get

$$x \approx 2(2\times10^5)^3 = 2.5\times10^{-16}.$$

Our approximation is in fact excellent, it is correct to many significant figures. That will be shown later. Note that we did not bother to write down explicitly the approximate decimal expansions for our square roots, but instead expressed each approximation as a sum.

Another Way. It is convenient, but not necessary, to notice that $\sqrt{10^{10}+1} = 10^5\sqrt{1+10^{-10}}$. To save typing, let $t = 10^{-10}$. We want to find an excellent estimate for $\sqrt{1+t}$.

Note that $(1+t/2)^2 = 1+t+t^2/4$. So the square of $(1+t/2)$ is a tiny bit bigger than $1+t$, meaning that the square root of $1+t$ is a tiny bit smaller than $1+t/2$. How can we adjust $1+t/2$ so that on squaring the $t^2/4$ term disappears? Look at $1+t/2-t^2/8$. When we square this, we get $1+t-t^3/8+t^4/64$. That's *awfully close* to $1+t$, a very tiny bit less. So our square root is a very tiny bit bigger than $1+t/2-t^2/8$.

More or less the same reasoning works for $\sqrt{10^{10}-1}$, which is equal to $10^5\sqrt{1-t}$. The square of $1-t/2-t^2/8$ is $1-t+t^3/8+t^4/64$, a very tiny bit too big. So $\sqrt{1-t}$ is a very tiny bit smaller than $1-t/2-t^2/8$. Note that the errors in these new estimates for $\sqrt{1+t}$ and $\sqrt{1-t}$ are in opposite directions, and will at least partially cancel. So we have good reason to think that our number is approximately

$$10^5(2-(1-t/2-t^2/8)-(1+t/2-t^2/8)), \quad \text{that is,} \quad 10^5t^2/4.$$

But $t = 10^{-10}$, so our estimate is 2.5×10^{-16}.

We have *almost* solved the problem (in two ways). Our estimates for the square roots are fantastically good, but a formal proof has not been given that they are indeed fantastically good. Maybe one should relax, this is an applied numerical problem, maybe being morally sure is good enough. But we can easily check that, for example with $t = 10^{-10}$, the number $1+t/2-t^2/8$ is close enough to $\sqrt{1+t}$. We had already noted it is a tiny bit too small. Look at the slightly bigger number $1+t/2-t^2/9$. Its square is $1+t+t^2/36-t^3/9+t^4/81$, which is bigger than $1+t$. So our estimate for $\sqrt{1+t}$ has error less than $t^2/8-t^2/9$, which is $t^2/72$, about 1.4×10^{-24}, good enough, by a lot.

Another Way. To save typing, and for better reasons, let $a = 10^5$ and let $e = 1$. We are trying to estimate

$$2a - \sqrt{a^2 - e} - \sqrt{a^2 + e}.$$

But

$$\sqrt{a^2 - e} - \sqrt{a^2 + e} = a\left(\sqrt{1-t} + \sqrt{1+t}\right),$$

where $t = e/a^2$. So we are estimating

$$a\left(2 - \sqrt{1-t} - \sqrt{1+t}\right).$$

(In our case, $t = 10^{-10}$.) It is fairly natural to rewrite our expression as

$$a\left((1 - \sqrt{1-t}) - (\sqrt{1+t} - 1)\right).$$

We find a better expression for $1-\sqrt{1-t}$ by "rationalizing the numerator," in this case by multiplying "top" and "bottom" (which is 1) by $1 + \sqrt{1-t}$. We get

$$\frac{t}{1 + \sqrt{1-t}}.$$

Do the same sort of thing for $\sqrt{1+t} - 1$. So we arrive at

$$\frac{at}{1 + \sqrt{1-t}} - \frac{at}{\sqrt{1+t} + 1}.$$

Bring to a common denominator. We have reached

$$\frac{at(\sqrt{1+t} - \sqrt{1-t}}{)}(1 + \sqrt{1+t})(1 + \sqrt{1-t}).$$

Again, rationalize the numerator. We reach

$$\frac{2at^2}{(1 + \sqrt{1+t})(1 + \sqrt{1-t})(\sqrt{1+t} + \sqrt{1-t})}.$$

Finally, we have something that looks nice. It can be easily approximated with any scientific calculator, or, more easily, without a calculator. Since t is incredibly tiny, the denominator is very close to 8. So our expression is nearly equal to $at^2/4$, or, since $t = e/a^2$, our expression is nearly equal to $e^2/(4a^3)$. (By *nearly equal* we mean that the *ratio* of our expression to $e^2/(4a^3)$ is nearly equal to 1. Note that although 0.0002 and 0.0003 are not far from each other, they are not "nearly equal" in this sense.) Now recall that $a = 10^5$ and $e = 1$. Our estimate is therefore 2.5×10^{-16}.

There is nothing particularly special about $a = 10^5$ and $e = 1$. The only thing that matters is that e is much smaller than a. So the approximation we have obtained is in fact *general.* If a is positive and $|e/a|$ is close to 0, then $2a - \sqrt{a^2 - e} - \sqrt{a^2 + e}$ is nearly equal to $e^2/(4a^3)$.