Solutions to December 2005 Problems

This document is quite long. That is partly because solutions are very detailed—and not edited tightly enough. Several solutions are given for most problems, together with additional questions and comments.

There is not much point to *reading* solutions to problems that you have not attempted. If you have attempted a problem without full success, glance at the solution for a hint on how to proceed, and then solve the problem. After you have solved the problem, you can read at leisure the solution given here.

Problem 1. Find the solutions of the equation $x^2 - 123456789\,x + 10 = 0$, correct to 4 significant figures.

Solution. We use the Quadratic Formula, which says that if $a \neq 0$, then the solutions of the equation $ax^2 + bx + c = 0$ are:

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}. \tag{1}$$

The Quadratic Formula is derived at the end of the solution.

In Formula 1, put $a = 1$, $b = -123456789$, and $c = 10$. Most scientific calculators have no problem finding $(-b + \sqrt{b^2 - 4ac})/2$ correct to 4 significant figures. My Sharp calculator (bought on sale, \$8.98) gives 123456789 as the answer. Thus, to 4 significant figures, the answer is 1.235×10^8.

(The standard TI only allows entry of 8 digits. With the TI, we can during the calculation enter the rounded value 1.2345679×10^8 instead of 123456789. This small lie will not affect the first four significant figures.)

Note that using the calculator was a waste of time. Since 40 is negligible compared to $(123456789)^2$, we are basically adding 123456789 to itself and dividing the result by 2.

Unfortunately, the second root gives trouble. My Sharp calculator claims that the answer is 0. This is definitely *not* correct to 4 significant figures, indeed not even to 1 significant figure. What can we do?

We could use a more powerful tool. For example, the scientific calculator program bundled with Microsoft Windows can handle 32 digits. The powerful— and (central) Canadian—program Maple can handle thousands of digits. There are many other programs for high precision arithmetic, including free ones: search the Internet. With any such tool, we find that the second root, to 4 significant figures, is 8.100×10^{-8}.

Comment. We ran into trouble because when we computed the second root directly from the Quadratic Formula, we needed the difference between two large nearly equal numbers.

Calculators store numbers internally to a limited precision, usually no more than a dozen digits. The numbers 123456789 and $\sqrt{(123456789)^2 - 40}$ are, from the calculator's point of view, equal, so it reports that their difference is 0.

Subtracting nearly equal large numbers can lead to catastrophic loss of precision. This issue comes up surprisingly often in scientific work. It has to be guarded against when we use (or design) automated computing procedures.

Another Way. If $c = 0$, the solutions of $ax^2 + bx + c = 0$ are easy to find. So we can assume that $c \neq 0$.

In Formula 1, multiply "top" and "bottom" by $(-b \mp \sqrt{b^2 - 4ac})/2a$. After some algebraic manipulation, we find that an alternate expression for the roots of the quadratic is

$$x = \frac{2c}{-b \mp \sqrt{b^2 - 4ac}}. \tag{2}$$

The expression

$$\frac{123456789 - \sqrt{(123456789)^2 - 40}}{2},$$

which a simple scientific calculator can not handle, turns into the harmless

$$\frac{20}{123456789 + \sqrt{(123456789)^2 - 40}},$$

which we can evaluate quite precisely with an ordinary calculator. Unfortunately, the new version of the Quadratic Formula has trouble with the big root—we get a division by 0 error. Nobody is perfect.

I have seen Formula 2 referred to as the *Citardauq Formula*. (Really! I am not imaginative enough to make it up.)

There is another way of obtaining Formula 2, which introduces an occasionally useful idea. Look at the equation $ax^2 + bx + c = 0$, where $c \neq 0$. Since $x = 0$ can not be a solution, we can divide both sides by x^2 and obtain the equivalent equation $a + b/x + c/x^2 = 0$. Now let $y = 1/x$. We obtain $cy^2 + by + a = 0$. By the ordinary Quadratic Formula,

$$y = \frac{-b \pm \sqrt{b^2 - 4ca}}{2c}. \tag{3}$$

But $x = 1/y$; flip Formula 3 over to obtain Formula 2.

Comment. Formula 2 has a square root in the denominator. But for computing the small root it was far more useful than the conventional Quadratic Formula.

Many a high-school student has been penalized for leaving square roots in the denominator. That may once have made some sense: paper and pencil division by an ugly number is more painful than multiplication by an ugly number.

However, as our successful computation of the small root shows, it can be a very poor idea to insist that square roots should not occur in the denominator. A similar comment can be made about the command "simplify." What is simple for one purpose may not be simple for another.

Another Way. If r and s are the roots of $ax^2 + bx + c = 0$, then $rs = c/a$. One way to show this is to multiply together the expressions for r and s given by the Quadratic Formula: a little manipulation yields c/a.

There is a better way. Our equation is equivalent to $x^2 - (b/a)x + c/a = 0$. It turns out that the polynomial $x^2 - (b/a)x + c/a$ is identically equal to

$(x - r)(x - s)$, that is, to $x^2 - (r + s)x + rs$. Thus we can conclude that $r + s = -b/a$ and $rs = c/a$. This idea can be generalized to higher degree polynomials, and is *very* useful.

In our situation, a simple calculation shows that the large root is nearly equal to 123456789. But the product of the two roots is 10, and therefore the other root is very nearly $10/123456789$.

Another Way. Note that

$$\sqrt{(123456789)^2 - 40} = 123456789\sqrt{1 - \epsilon},$$

where $\epsilon = 40/(123456789)^2$.

For any ϵ close to 0, $\sqrt{1 - \epsilon}$ is nearly equal to $1 - \epsilon/2$. We can verify this informally by squaring $1 - \epsilon/2$, and noting that the result is $1 - \epsilon + \epsilon^2/4$. The number $\epsilon^2/4$ is utterly negligible in comparison with ϵ when the latter is close to 0.

Using this approximation to the square root, we find that the small root is about $(1/2)(123456789 - 123456789(1 - \epsilon/2))$, that is, $123456789\epsilon/4$. This simplifies to $10/123456789$.

Another Way (Informal). If we think about the equation of the problem, we can see that one root is very large, and the other is close to 0.

Insofar as the small root is concerned, the equation is really not a quadratic at all! The x^2 term only perturbs things a little bit, so for the small root we are (almost) looking at the linear equation $-123456789\,x + 10 = 0$, which is trivial to solve. Similarly, insofar as the large root is concerned, the 10 at the end makes little difference, so we are almost looking at $x^2 - 123456789\,x = 0$, which again reduces to a linear equation.

Sophisticated versions of this kind of thinking occur all over mathematics, in particular applied mathematics. Nonlinear problems are in general difficult to handle. It is often useful to replace a nonlinear problem by a nearly equivalent linear problem.

Comment. The idea above can be viewed in a more formal way. Rewrite the original quadratic equation as

$$x = \frac{10 + x^2}{123456789}. \tag{4}$$

The small root is nearly 0. Let $r_0 = 0$. Substitute this estimate r_0 for x into the right-hand side of Formula 4 to get an improved estimate r_1. We get $r_1 = 10/123456789$.

Now substitute the improved estimate r_1 for x into the right-hand side of Formula 4 to get a new estimate r_2. There is virtually no change.

But in other less extreme situations there might be. In that case, we can substitute r_2 for x into the right-hand side of Formula 4 to get an improved estimate r_3, and so on. This kind of *iterative* procedure can be highly efficient.

Comment (The Quadratic Formula). We find an expression for the root(s) of the equation $ax^2 + bx + c = 0$, where $a \neq 0$. Note that

$$ax^2 + bx + c = 0 \qquad \text{if and only if} \qquad 4a^2 + 4abx + 4ac = 0.$$

3

(We multiplied both sides by $4a$. The usual argument, which involves dividing both sides by a, is more awkward.) Note that

$$4a^2x^2 + 4abx + 4ac = (2ax + b)^2 - b^2 + 4ac,$$

and therefore

$$ax^2 + bx + c = 0 \qquad \text{if and only if} \qquad (2ax + b)^2 = b^2 - 4ac,$$

or equivalently if and only if

$$2ax + b = \pm\sqrt{b^2 - 4ac} \qquad \text{or equivalently} \qquad x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}.$$

We make some final comments about the Quadratic Formula.

- The key step in the calculation was "completing the square," that is, recognizing that $4a^2x^2 + 4abx$ is almost equal to the square of the linear expression $2ax + b$, except that we must subtract the correction term b^2. Variants of completing the square, some exceedingly sophisticated, occur often in more advanced mathematics.

- What do we mean by $\sqrt{b^2 - 4ac}$, or more concretely by $\sqrt{9}$? The expression "$\sqrt{9}$" is a *word* in the language of mathematics. Its meaning, like the meaning of any word, is determined by agreement among the speakers of the language.

  Unfortunately, speakers of mathematics do not all agree! The meaning of $\sqrt{9}$ has varied with time. And even now different subcommunities give the word different meanings. So we have to make sure that we use the word in the way our community—in particular the grader of our exam—uses it. In elementary work, there is now general agreement that $\sqrt{9} = 3$: if x is positive then $\sqrt{x}$ is the positive number whose square is x. But people working in the theory of complex variables find it useful to agree that $\sqrt{9} = \pm 3$.

- What about $\sqrt{-1}$? Is there such a thing? It is easy to see that *if* there is an object whose square is -1, then there must be two. If we decide that $\sqrt{x}$ has only one value, which of the two possible values do we pick for $\sqrt{-1}$? In elementary work with complex numbers, such questions are seldom addressed, and blind manipulation is all too frequent.

Problem 2. The figure below represents a network of one-way streets. How many different ways are there to drive from A to B?
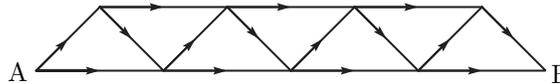
Solution. We could carefully list all possible paths, and then count. To make the list, it is useful to have names for the paths. Label all points where roads meet, as in the figure below.

4
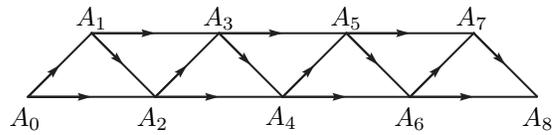
Now begin to list. The problem with this approach is that it is tedious, and unless we are systematic and careful, it is easy to miss a path, or to count a path twice.

One systematic way of listing is alphabetical. Think of A_0, A_1, and so on as letters of an alphabet, where A_0 is the first letter of the alphabet, A_1 is the second, and so on. Then the paths are *words* in that alphabet. We can list the words alphabetically. (The first word is $A_0 A_1 A_2 \ldots A_7 A_8$, the second is $A_0 A_1 A_2 \ldots A_6 A_8$.) We can continue. The list is not long—it turns out that there are 34 paths.

The problem with the above approach is that it is time-consuming, and becomes completely impractical if the network is extended on the right in the natural way.

Another Way. We first solve a (much) simpler problem: how many paths are there from A_0 to A_1? Obviously there is only 1. How many paths from A_0 to A_2? There are 2, namely $A_0 A_1 A_2$ and $A_0 A_2$.

How many paths from A_0 to A_3? Again, the problem is easy, we can make an explicit list: there are 3 paths. So to A_1 there is 1, to A_2 there are 2, to A_3 there are 3. We must not jump to the conclusion that there are n paths from A_0 to A_n. And even if the answer was n, we would need to *prove* it. Calculations with small numbers can exhibit apparent "patterns" that turn out to be wrong.

And indeed it turns out that there are 5 paths from A_0 to A_4. They can be counted directly, but a nice idea lets us recycle our earlier results.

There are two ways of getting to A_4:

1. Somehow we get to A_3, and then go to A_4, or

2. Somehow we get to A_2, and then go directly to A_4 without going through A_3.

By our earlier work, there are 3 ways to get to A_3, and hence 3 ways of getting to A_4 via A_3. And there are 2 ways to get to A_2, and hence 2 ways of getting to A_4 without going through A_3. Thus there is a total of $3 + 2$ ways of getting to A_4.

We do a similar analysis for paths from A_0 to A_5. There are two kinds of path to A_5, namely:

1. The paths that include A_4, and

2. The paths that miss A_4, that is, go directly from A_3 to A_5.

By our previous work, there are 5 paths of the first kind, and 3 of the second kind, for a total of 8.

5

A similar analysis shows that there are $8 + 5$, that is, 13, paths from A_0 to A_6, and 21 paths from A_0 to A_7, and finally 34 paths from A_0 to A_8. If the diagram is extended in the obvious way, there are 55 paths to A_9, 89 paths to A_{10}, "and so on."

Another Way. Well, this way is not really different, but it *looks* different.

For any path, for example $A_0 A_2 A_3 A_5 A_6 A_8$, concentrate on the indices (subscripts). The path could be more efficiently described as $0\,2\,3\,5\,6\,8$. Note that the indices jump by 1 or by 2. Any sequence of numbers that starts at 0, ends at 8, and jumps by 1 or 2, describes a valid path.

Let's simplify even further, and concentrate on the *jumps.* In the example we are looking at, the jumps are $2 - 0$, $3 - 2$, $5 - 3$, $6 - 5$, and $8 - 6$. So the path can be very briefly described by $(2, 1, 2, 1, 2)$. Note that this is a sequence of 1's and/or 2's that adds up to 8.

Now we ask: how many sequences of 1's and/or 2's are there that add up to 8? This is precisely the number of paths from A_0 to A_8.

More generally, imagine continuing the street pattern indefinitely to the right. Then the number of sequences of 1's and/or 2's that add up to n is equal to the number of paths from A_0 to A_n. So our original problem can be reworded as follows.

Let s_n be the number of sequences of 1's and/or 2's that add up to n. What is s_8?

Like in the first solution, we solve first some simpler problems, by finding s_1, s_2, s_3, and so on. It is obvious that $s_1 = 1$. And there are two sequences that add up to 2, namely $(1, 1)$ and (2).

Now we imitate the analysis of the first solution. Let $n > 2$. There are two kinds of sequence that add up to n, namely:

1. The sequences that end (on the right) with 1, and

2. The sequences that end with 2.

Any sequence that ends with 1 and has sum n can be obtained by appending a 1 to a sequence that has sum $n - 1$. By definition there are s_{n-1} of these.

Any sequence that ends with a 2 and has sum n can be obtained by appending a 2 to a sequence that has sum $n - 2$. By definition there are s_{n-2} of these.

Thus we arrive at the important recurrence:

$$s_n = s_{n-1} + s_{n-2}. \tag{5}$$

Now we can compute: $s_3 = s_2 + s_1 = 3$, $s_4 = s_3 + s_2 = 5$, $s_5 = 8$, $s_6 = 13$, $s_7 = 21$, and $s_8 = 34$. Recurrence 5 can be used to compute s_n easily whenever n is not too large.

Comment. The sequence above is almost the same as the famous *Fibonacci sequence*, which is defined as follows: $F_0 = 0$, $F_1 = 1$, and for any $n > 1$,

$$F_n = F_{n-1} + F_{n-2}.$$

6

The first few terms of the Fibonacci sequence are 0, 1, 1, 2, 3, 5, 8, 13, 21, 34, 55, 89, 144.

The sequence seems to have been first discussed indirectly in Indian technical analysis of poetry (Piṅgala, around 450, Virahāṅka, around 650), and then in considerable detail by a number of Indian mathematicians from the thirteenth century on.

The first published European problem that (in hindsight) involves this sequence is a problem about rabbits in Leonardo of Pisa's *Liber Abaci* (1202). This book was the first important European mathematical work since the ancient Greeks. It was largely cribbed from Arabic sources.

In the 19th century, a historian of mathematics decided, probably wrongly, that Leonardo of Pisa should be called Fibonacci. A 13th century document refers to him as Leonardo Bigollo.

The Fibonacci sequence shows up in amazingly many areas of mathematics. It is likely that there will be more problems in the next few months that involve this sequence. You may want to look at the following question.

A path is 2 feet wide and 18 feet long. We want to pave the path with 1×2 concrete paving blocks. How many ways are there to do this? (The answer is 4181.)

Problem 3. Alphonse and Beti have identical cone-shaped paper drinking cups. At the beginning, the tea in Alphonse's cup (left) is 8 cm deep, and Beti's cup (right) is empty. Alphonse pours half his tea into Beti's cup. How deep is the tea in Beti's cup?
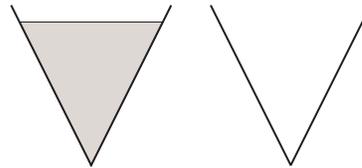
Solution. We first give a solution which is a bit formula-heavy. In the diagram below, the picture on the left represents Alphonse's cone of tea before he poured half into Beti's cup, and the picture on the right represents Beti's cone of tea after the gift. For clarity, we are looking only at the tea, not the cups.
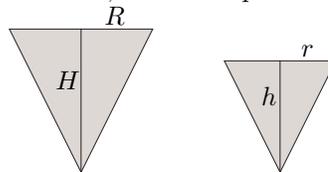
Let R be the radius of the top rim of Alphonse's original cone of tea, and let H be the height of this cone. Let r and h be the corresponding numbers for Beti's cone of tea. Let V be the original volume of Alphonse's cone of tea, and let v be the volume of Beti's cone after the gift.

We were told that $H = 8$ and that $v/V = 1/2$, but we might as well be general. Tomorrow Alphonse may want to give Beti two-thirds of his tea. And our approach allows us to preserve symmetry. Symmetry is good.

7

By the standard formula for the volume of a cone,

$$V = \frac{\pi R^2 H}{3} \qquad \text{and} \qquad v = \frac{\pi r^2 h}{3}. \tag{6}$$

The cone on the left and the cone on the right are *similar*: the cone on the right is just a scaled down version of the one on the left. It follows that

$$\frac{r}{R} = \frac{h}{H}. \tag{7}$$

Alternately, we can think of the pictures above as cross-sections through the central axis of the two cones. Then the pictures represent similar triangles.

Finally, we can compute. From the Equations 6 we obtain

$$vR^2 H = V r^2 h.$$

But then by Equation 7

$$vH^3 = Vh^3, \qquad \text{or equivalently} \qquad \frac{h}{H} = \sqrt[3]{\frac{v}{V}}.$$

Since $v/V = 1/2$ and $H = 8$, we obtain $h = 8\sqrt[3]{1/2}$. Beti's tea is about $6.35\,\mathrm{cm}$ deep—maybe deeper than we would have guessed.

Another Way. The "small" cone of tea is just a scaled down version of the big cone. If we multiply the dimensions of a solid object by the linear scaling factor t, then the volume scales by the factor t^3. Someone whose body dimensions are all twice yours will weigh eight times as much as you do.

In our case, the volume of the small cone is $1/2$ times the volume of the big one, so $t^3 = 1/2$, and therefore $t = \sqrt[3]{1/2}$. It follows that the height of the small cone is $8\sqrt[3]{1/2}$.

Comment. The second argument is much simpler, much more structural, than the first. And we didn't even have to know the formula for the volume of a cone. The argument also generalizes easily. First we make a definition.

Definition. A *generalized cone* is obtained by taking a closed curve $\mathcal{C}$ in a plane, and a point P not in the plane, and drawing all line segments that join P to a point on $\mathcal{C}$.

In our situation, the point P is the bottom of the paper cup, $\mathcal{C}$ is the rim of the cup, and the line joining P to the center of $\mathcal{C}$ is perpendicular to the plane of $\mathcal{C}$. The cup is called a *right circular cone*.

We can look at the tea problem for a generalized cone. If the rim is a circle, the bottom point of the cup need not be directly "below" the center of the circle. And the rim of a generalized conical cup could be an ellipse, or a square, or a more exotic curve.

The scaling argument with which we found the depth of Beti's tea works, word for word, for generalized cones!

8

Problem 4. Alicia is looking through binoculars as a distant ship sails away. Her eyes are 10 metres above sea level, and the tip of the ship's smokestack is 40 m above sea level. How far away is the tip of the smokestack at the instant that it disappears from Alicia's view? Assume that the Earth is a sphere of radius 6400 kilometres.

Solution. Let A be the location of Alicia's eye, let S be the tip of the smokestack, and let C be the centre of the Earth. Imagine slicing through the Earth along the plane through C, S, and A. We get a picture somewhat like the one below. Note that the scale is grossly incorrect: if draw the picture to scale, the circular arc is visually indistinguishable from a straight line.
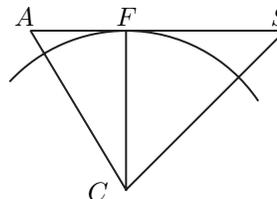
At the instant that the ship's smokestack disappears from view, the line segment AS is tangent to the circular arc. Let F be the point of tangency. Then $\angle CFA$ and $\angle CFS$ are right angles.

To save space, and to make the argument more general, let a be Alicia's height (in kilometres) above sea-level, and let s be the height of the smokestack. Let R be the radius of the Earth. (As much as possible, one should work with letters rather than numbers. Using numbers from the beginning can hide vital structural relationships.)

We have $CA = R + a$, $CS = R + s$, and $CF = R$. By the Pythagorean Theorem, $(FA)^2 + (CF)^2 = (CA)^2$, and therefore

$$FA = \sqrt{(R+a)^2 - R^2} = \sqrt{2Ra + a^2}.$$

Similarly, $FS = \sqrt{2Rs + s^2}$, and therefore the straight-line distance between A and S is equal to

$$\sqrt{2Ra + a^2} + \sqrt{2Rs + s^2}.$$

Finally, it is time to compute. We have $a = 0.01$ and $s = 0.04$. So the required distance is approximately 33.941 km.

Comment. Note that we "simplified" $\sqrt{(R+a)^2 - R^2}$ to $\sqrt{2Ra + a^2}$. Actually, I think that the first expression is simpler. But in this case there was a good reason to do this "simplification." Recall from the solution of Problem 1 that it can be dangerous to subtract two very large nearly equal numbers. The first expression asks us to do that, while the second expression does not.

We can take advantage of the fact that R is much larger than a or s to produce a simpler *approximate* formula for the distance AS. If x is much smaller than R, then $\sqrt{2Rx + x^2}$ is nearly equal to $\sqrt{2Rx}$, so the distance AF is roughly equal to $\sqrt{2R}(\sqrt{a} + \sqrt{s})$.

For our particular numbers, the correct distance, to 6 decimal places, is 33.941165 km, and the approximate formula yields 33.941125 km. The approximation differs from the truth by roughly 4 centimetres!

You may want to consider the following question. An old lighthouse stands by the shore and is 50 meters high. We want to replace it with a new lighthouse that can be seen from three times as far away as the old one. How high should the new lighthouse be?

Comment. Let A_0 be the point on the Earth's surface directly below Alicia, and S_0 the point on the Earth's surface directly below the smokestack. We can ask for:

a) the minimum distance between A_0 and S_0, if we travel along the Earth's surface, that is, on a great circle, and

b) the straight line distance between A_0 and S_0.

It turns out that the answer to a) is about 33.941059, and the answer to b) is about 33.941019. (The calculation requires material that is not met until Grade 12.)

Note how astonishingly close these numbers are to each other, and to the true distance AF. The difference between the answers to a) and b) is about 4 cm! This shows how very flat the Earth is when we look at a region only a few miles long.

You may want to look at the following question. You are on the surface of a deep part of the ocean, like the Marianas Trench, and want to swim to a point 20 km away as measured along the surface of the ocean.

- How much distance do you save by swimming in a straight line, underwater, towards your destination?

- If you do swim underwater in a straight line, what is the maximum depth that you reach?

Problem 5. The large triangle below has been decomposed into two triangles and a parallelogram as shown. Given that the two triangles have area α and β, what is the area of the large triangle?
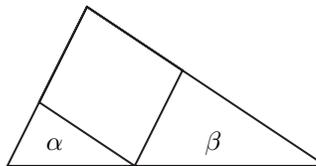
Solution. Label the diagram as in the picture below, and draw the line segment PW. We will show that $\triangle PAW$ has area $\sqrt{\alpha}\sqrt{\beta}$. From this we can then easily conclude that $\triangle PQR$ has area $\alpha + 2\sqrt{\alpha}\sqrt{\beta} + \beta$.
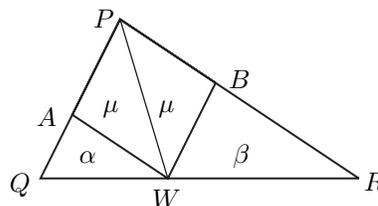
Let μ be the area of $\triangle PAW$. Think of PA as a base of $\triangle PAW$, and think of AQ as a base of $\triangle AQW$. With respect to these bases, the two triangles have the same height, and therefore

$$\frac{\mu}{\alpha} = \frac{PA}{AQ} = \frac{BW}{AQ}. \tag{8}$$

But μ is also the area of $\triangle PBW$. So by similar reasoning we get

$$\frac{\mu}{\beta} = \frac{PB}{BR} = \frac{AW}{BR}. \tag{9}$$

But $\triangle AQW$ is similar to $\triangle BWR$. It follows that

$$\frac{BW}{AQ} = \frac{BR}{AW}, \qquad \text{and therefore} \qquad \frac{BW}{AQ}\frac{AW}{BR} = 1. \tag{10}$$

Now multiply the left-hand sides of Equations 8 and 9, and the right-hand sides. With the aid of Equation 10 we obtain

$$\frac{\mu^2}{\alpha\beta} = \frac{BW}{AQ}\frac{AW}{BR} = 1.$$

It follows that $\mu^2 = \alpha\beta$ and therefore $\mu = \sqrt{\alpha}\sqrt{\beta}$. Not pleasant—too many letters. The next three solutions are more attractive.

Another Way. This is based on a suggestion by Mr. Owen Ren, Grade 10, Magee Secondary School. We first prove a result about parallelograms that is more general than necessary. The parallelogram in the picture below and to the left has been divided into four parts by two lines parallel to the sides.
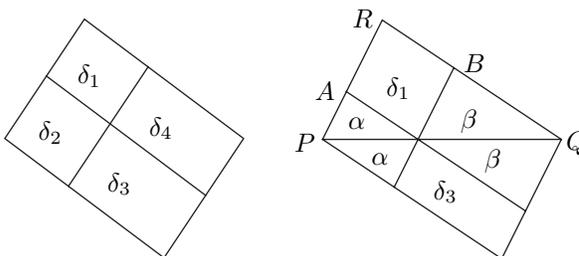
Suppose that the areas of these parts are δ_1, δ_2, δ_3, and δ_4 as in the picture. We will show that

$$\delta_1\delta_3 = \delta_2\delta_4.$$

11

The proof is easy. View the parallelograms with areas δ_1 and δ_2 as having their common line as base. The ratio of their areas is equal to the ratio of their heights.

Similarly, the parallelograms with areas δ_4 and δ_3 can be viewed as having their common line as base. So the ratio of their areas is equal to the ratio of their heights.

But the parallelograms with areas δ_1 and δ_4 have the same height. So do the ones with area δ_2 and δ_3. Thus $\delta_1/\delta_2 = \delta_4/\delta_3$, and therefore $\delta_1\delta_3 = \delta_2\delta_4$.

Now look at the picture on the right. We have completed the original triangle PQR to a parallelogram, and extended the sides of the old small triangles into the new part.

The two triangles whose areas are labelled α do have the same area, since they are congruent. The two triangles whose areas are labelled β are also congruent.

It follows that $\delta_1\delta_3 = 4\alpha\beta$. But in this case, we have $\delta_1 = \delta_3$, because the part of the picture above PQ has the same area as the part below PQ. Thus $\delta_1^2 = 4\alpha\beta$, and therefore $\triangle PQR$ has area $\alpha + \beta + 2\sqrt{\alpha}\sqrt{\beta}$.

Another Way (Informal). When we start thinking about the problem, it is not obvious that the area of the large triangle is completely determined by α and β. So it is tempting to assume that the large triangle has a special shape, say equilateral, or right-angled isosceles, anything that makes the problem easier.

That would be cheating: we would still have to show that the answer is *indeed* independent of shape. We now proceed to do precisely that.

Imagine pushing each point of the original figure straight left (or right) by an amount proportional to the distance of that point from the base. (This kind of transformation is called a *shear*.) It is intuitively reasonable, and indeed true, that a shear transforms lines into lines.

Push until the apex of the original triangle is directly over one of the ends of the base. The bases and heights of all our triangles remain unchanged, so all areas remain unchanged, and we arrive at a shape like the left half of the picture below.
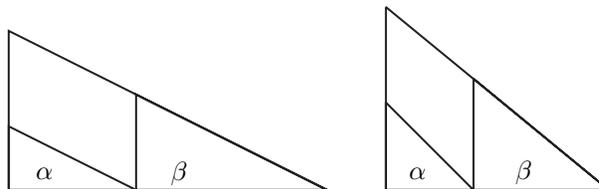
Let's improve the picture some more, by transforming the picture on the left into an *isosceles* right-angled triangle.

We do this by scaling in the vertical direction by a certain scaling factor t, and in the horizontal direction by the scaling factor $1/t$. The vertical scaling multiplies areas by t. Then the horizontal scaling multiplies areas by $1/t$. The net effect is that areas are unchanged.

If we know the base and height of the triangle on the left, it is easy to find the scaling factor t that makes base and height equal.

So *whatever* the shape of the original triangle, we can, without changing any areas, transform it into a right-angled isosceles triangle as in the picture on the right. Now all we need to do is to find the area of that triangle.

The right-angled isosceles triangle with area α has "legs" of length $\sqrt{2}\sqrt{\alpha}$; the right-angled isosceles triangle with area β has legs $\sqrt{2}\sqrt{\beta}$. So the full triangle on the right has legs $\sqrt{2}(\sqrt{\alpha}+\sqrt{\beta})$, and therefore area $(\sqrt{\alpha}+\sqrt{\beta})^2$.

We could have transformed to a standard shape other than right-angled isosceles. For example, if we transform to a right-angled triangle which is whose base is twice the height, we can remove the slightly ugly $\sqrt{2}$ from the argument.

Another Way. The original large triangle, and the two smaller triangles, are all *similar.* They differ in size but not in basic shape: the two triangles of area α and β are scaled down versions of the large triangle. For any of these three triangles, call the side that goes in a left-right direction the *base* of the triangle.

Let k be the area of a triangle which is similar to all of these triangles, but has base length equal to 1. Let a and b be the base lengths of the triangles of area α and β.

Recall that if we scale all lengths of a figure by a factor of t, then areas scale by a factor t^2. It follows that

$$\alpha = ka^2, \quad \beta = kb^2, \qquad \text{and therefore} \qquad a = \sqrt{\alpha/k}, \quad b = \sqrt{\beta/k}.$$

The original triangle has base $a+b$. Its area is therefore $k(\alpha+\beta)^2$. But

$$k(\alpha+\beta)^2 = k(\sqrt{\alpha/k}+\sqrt{\beta/k})^2 = (\sqrt{\alpha}+\sqrt{\beta})^2.$$

Comment. Note the efficiency of the above argument. We can be a little more tricky. Choose the unit of length in such a way that a triangle similar to all of our triangles, but with base of length 1, has area 1. This can be done: there is no rule saying that if the unit of length is the *ertem*, then the unit of area has to be the *ertem*^{2}. For example, in measuring fields in the United States, the usual unit of length is the foot or yard, but the usual unit of area is the acre; an acre is 4840 square yards.

By choosing the unit of area suitably, we can, without loss of generality, take $k = 1$. Thus $a = \sqrt{\alpha}$, $b = \sqrt{\beta}$, so the original triangle has base $\sqrt{\alpha}+\sqrt{\beta}$, and therefore its area is the square of that.

Comment. Pick a point in the interior of a large triangle, and through this point draw lines parallel to the three sides of the triangle, as in the figure below.
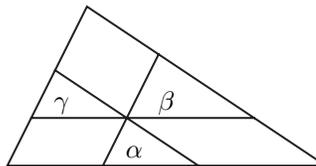
The lines divide the triangle into three parallelograms and three triangles. Suppose that the triangles have area α, β, and γ. It is not hard to show, by a similarity argument of the type used above, that the large triangle has area

$$(\sqrt{\alpha}+\sqrt{\beta}+\sqrt{\gamma})^2.$$

Problem 6. Find all integers n such that $(28.5)^n + (99.5)^n$ is an integer.

Solution. Let $x_n = (28.5)^n + (99.5)^n$. If $n \le -1$, then $(28.5)^n$ and $(99.5)^n$ are both "small." In particular, $0 < x_n < 1$, so x_n can not be an integer.

If $n = 0$, then $x_n = 2$. From here on we look at positive n.

We could (and should) explore with a calculator for a while. The results will be quite calculator-dependent.

My Sharp calculator is typical. It thinks that x_2 and x_4 are not integers, and that x_1 and x_3 are. Because it only displays 10 digits, it says nothing useful about x_5.

As a first step, we should probably work with integers, not fractions. As a general rule of thumb, number-theoretic problems that involve fractions should be transformed to problems about integers.

Note that

$$x_n = (57/2)^n + (199/2)^n,$$

so $57^n + 199^n = 2^n x_n$. Thus we want to find all (positive) integers n such that $57^n + 199^n$ is an integer multiple of 2^n.

First we deal with the case when n is positive and *even*. If n is even, let $m = 2n$. Then $57^n = 57^{2m} = (57^m)^2$. Similarly, $199^n = (199^m)^2$. Thus 57^n and 199^n are both the squares of *odd* integers.

The square of an odd integer $2k + 1$ is $4(k^2 + k) + 1$, that is, 1 more than a multiple of 4. Thus $57^n + 199^n$ is 2 more than a multiple of 4, and in particular is not divisible by 4.

But 2^n *is* divisible by 4 for $n \ge 2$, so if n is even and positive then x_n is not an integer. Reasoning is much more efficient than the calculator! A relatively simple argument has taken care of infinitely many n.

Next we deal with the case when n is positive and *odd*. Recall, maybe, the general identity

$$a^n + b^n = (a + b)(a^{n-1} - a^{n-2}b + \cdots - ab^{n-2} + b^{n-1}). \tag{11}$$

(This useful formula can be verified easily by multiplying out the right-hand side. Note the wholesale cancellation. The formula is only valid when n is odd.)

Now let $a = 57$ and $b = 199$. By Identity 11, $57^n + 199^n$ is divisible by $a + b$, that is, by 256. The second term in the product above is a "sum" of an odd number (namely n) of odd terms, so it is odd. Thus $x_n = (256/2^n)q$, where q is odd.

So if n is positive and odd, then x_n is an integer precisely if 2^n divides 256, that is, if $n = 1$, 3, 5, or 7.

We conclude that the only integers n for which x_n is an integer are 0, 1, 3, 5, and 7.

More generally, let a and b be *any* positive odd integers, and let $x_n = (a/2)^n + (b/2)^n$. Then 0 is the only even integer n for which x_n is an integer. And the only odd n for which x_n is an integer are the n such that 2^n divides $a + b$.

Another Way. The algebraic factorization of $a^n + b^n$ for odd n (Identity 11) is probably the simplest approach to the problem, if we happen to think of it.

We used the identity to show that if $a = 57$ and $b = 199$, then for any positive odd n the largest power of 2 that divides $a^n + b^n$ is 256. But we can establish this fact in other ways.

Note for example the following sequence of identities, which are easy to verify by multiplying out the right-hand sides.

$$a^3 + b^3 = (a^2 + b^2)(a + b) - ab(a + b),$$
$$a^5 + b^5 = (a^4 + b^4)(a + b) - ab(a^3 + b^3),$$
$$a^7 + b^7 = (a^6 + b^6)(a + b) - ab(a^5 + b^5),$$
$$a^9 + b^9 = (a^8 + b^8)(a + b) - ab(a^7 + b^7),$$

and so on.

Now let $a = 57$ and $b = 199$. Look at the right-hand side of the first identity. Note that $a^2 + b^2$ is even, and $a + b$ is a multiple of 256. (In fact $a + b = 256$.) So $(a^2 + b^2)$ is a multiple of 512. Also, $ab(a + b)$ is 256 times an odd number. It follows that the difference $(a^2 + b^2)(a + b) - ab(a + b)$ is 256 times an odd number. So the highest power of 2 that divides $a^3 + b^3$ is 256.

Now, using the second identity, we can show that the highest power of 2 that divides $a^5 + b^5$ is 256. For $a^4 + b^4$ is even, so $(a^4 + b^4)(a + b)$ is a multiple of 512, while by the result for $a^3 + b^3$, the number $ab(a^3 + b^3)$ is 256 times an odd number. It follows that $a^5 + b^5$ is 256 times an odd number.

We can now use the just established fact that 256 is the highest power of 2 that divides $a^5 + b^5$, and the third identity, to show that the highest power of 2 that divides $a^7 + b^7$ is 256.

And now the fourth identity can be used to deal with $a^9 + b^9$. Clearly we can continue in this way forever.

Other identities could be used in the proof, such as the sequence that starts

$$a^3 + b^3 = (a + b)(a^2 + b^2) - ab(a + b),$$
$$a^5 + b^5 = (a^2 + b^2)(a^3 + b^3) - a^2 b^2(a + b),$$
$$a^7 + b^7 = (a^3 + b^3)(a^4 + b^4) - a^3 b^3(a + b).$$