HOW LONG IS THE COAST OF AUSTRALIA ?

Suppose that our Government has decided to improve the capabilities of the Royal Australian Navy. It has resolved to buy one new destroyer for each 1000 km of the Australian coast-line. The only problem left for the Navy is to measure the length of the coast-line.

Method 0. Buy a map of Australia, cut along the line separating the blue bit from the green bit and measure the coast-line with a tape measure. (This method goes against the great Australian tradition established by such great Australian men as Ludwig van Leichhardt, the ill-fated Burke N. Wills and Captain Mathews Flinders, who at least had the good sense to go by sea. We are appalled that you would even consider such a solution. Anyway, the map we bought was drawn by the Army and obviously unreliable.) Besides that, the actual coast will be much longer than our tape measure can measure; think of all those lovely little bays round Sydney Harbour!

Method 1. Take a handy straight bit of wood of length η and use it to mark segments of length η round the coast so that each new step starts where the previous step leave off. The number of steps multiplied by η gives an approximate length $L(\eta)$ for the coast-line. If we make the length η smaller and smaller and repeat the operation, we expect this $L(\eta)$ to approach a well-defined value which we can call the true length of the coast-line. Figure 1 shows how we can use this procedure to measure the circumference of a circle of radius 5×10^3m. One step of length η turns the radius through an angle $\theta°$ given by $\sin \frac{1}{2}\theta = \eta \times 10^{-4}$ and the number of steps needed to turn the radius through $360°$ and bring us back to our starting point is $360/\theta$. So $L(\eta) = 360\eta/\theta$. As you can see, $L(\eta)$ converges very rapidly to $\pi \times 10^4$. Indeed, $L(\eta) = 360\eta/\theta = 360 \times 10^4 \sin \frac{1}{2}\theta/\theta$. Now, if θ is a small angle measured in degrees, $\sin \theta \approx \pi\theta/180$. We can use this approximation when η is small and we find

$$L(\eta) \approx 360 \times 10^4 \times \frac{\pi(2\theta)}{180} \times \frac{1}{\theta} = \pi \times 10^4,$$

This idea is the basis of Archimedes' method for determining π. (See Parabola, Volume 15, Number 2.) Measurements of the Australian coast have been made for various values of η and $L(\eta)$ does not seem to stabilise. Indeed, $L(\eta)$ appears

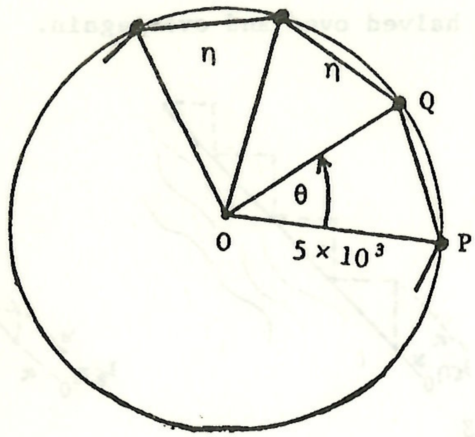

η	L(η)
5×10^3	3×10^4
4×10^3	3.05×10^4
3×10^3	3.09×10^4
2×10^3	3.12×10^4
10^3	3.14×10^4
10^2	3.14×10^4
10	3.14×10^4

Figure 1

to increase without limit as η is made smaller and smaller. The behaviour of L(η) for the circle and for our coast-line is shown in Figure 2.

Here, to see the pattern more clearly, we plotted the data on log-log scales, which is a common device to change awkward curves on ordinary graph paper into straight lines on log-log paper. Our data for the coast-line, represented by the straight line in Figure 2 was obtained from actual measurements.

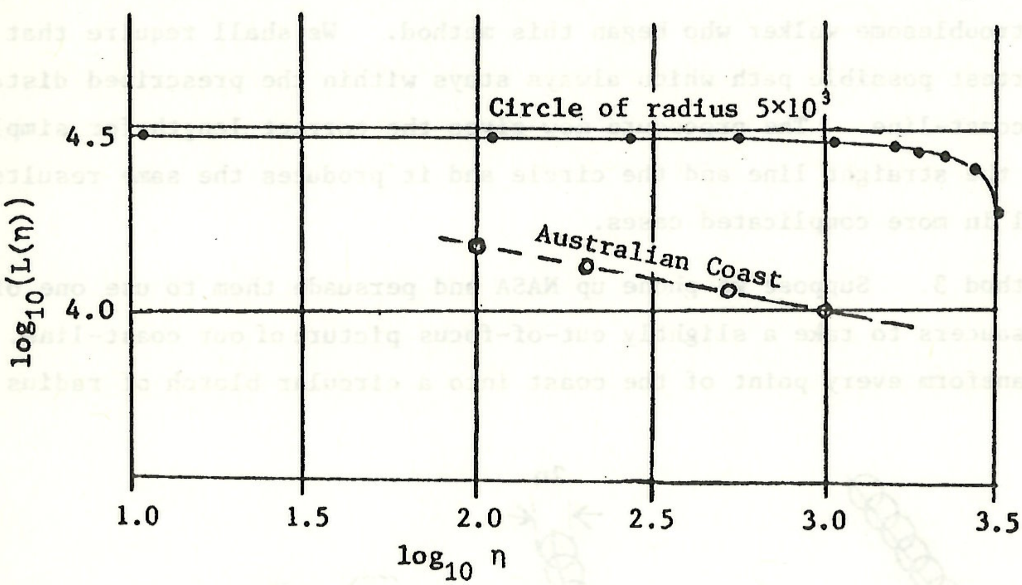

Figure 2

Method 2. Imagine a man walking round the coast, following whatever route takes his fancy, but always staying within a prescribed distance η of the coast-line. The experiment can be repeated with smaller and smaller values of η to force our walker to follow finer and finer details in the coast-line. This seems much the same as method 1, but there is a parodox lurking around the corner. Suppose the walker insists on walking in a north-south or east-west direction and consider a hypothetical straight coast-line running in a north-easterly direction.

Figure 3 shows the walker's path as η is halved over and over again. If the

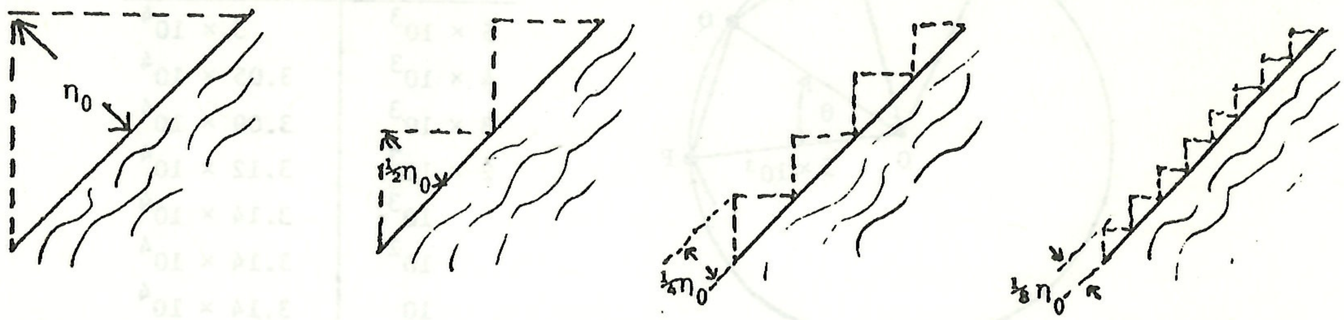

Figure 3

piece of coast-line is the diagonal of a square of side 100m, then we know that its true length is $100\sqrt{2}$m. However, in each of the diagrams, the walker goes altogether 100m north and 100m east, so in every case the "approximate length" measured is 200m. Notwithstanding all appearances to the contrary, the path formed by the steps does not approach the diagonal line. Are you still convinced that the polygons in figure 1 must give the correct value for the circumference of the circle? To avoid the difficulty, we have to take away some of the freedom of the troublesome walker who began this method. We shall require that he takes the shortest possible path which always stays within the prescribed distance η of the coast-line. The procedure now gives the correct length for simple curves such as the straight line and the circle and it produces the same results as method 1 in more complicated cases.

Method 3. Suppose we phone up NASA and persuade them to use one of their flying saucers to take a slightly out-of-focus picture of our coast-line. This will transform every point of the coast into a circular blotch of radius η,

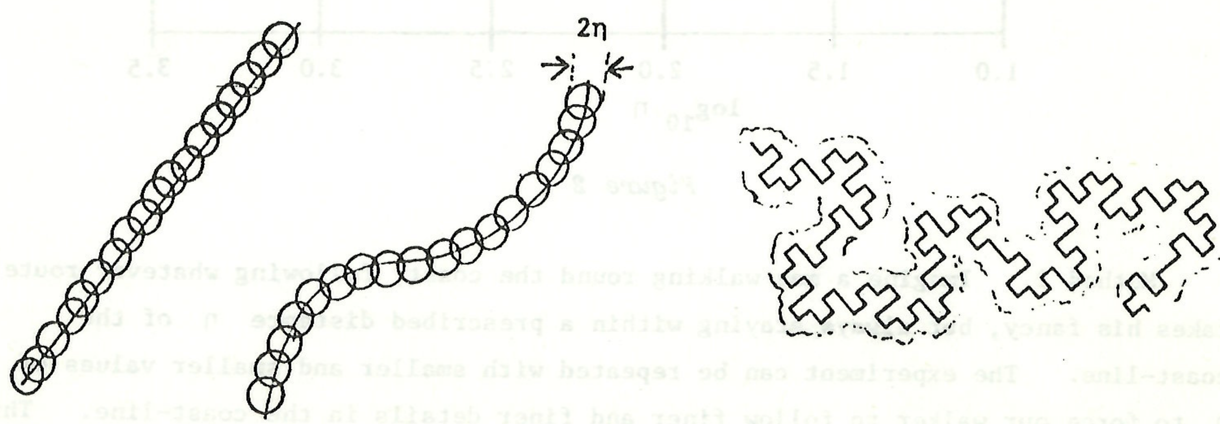

"Minkowsky Sausage"

Figure 4

forming a kind of sausage of width 2η which covers the coast-line. We then measure the area of the sausage and divide by 2η to get an estimate of the length of the coast-line. For example, for a straight-line coast, the sausage is a rectangle of width 2η and its area divided by 2η is the length of the line. For an actual coast-line, the sausage smooths out the wiggles whose scale is small compared to η, just as in methods 1 and 2. As η decreases, the sausage has to slink round more and more bays and headlands and, again, the estimated length becomes larger and larger, apparently without bound. This trend has been observed for a number of rugged coast-lines and cannot be explained away by coffee stains on the chart table or too much splice on the mainbrace.

Now let us examine the results of all this methodology in more detail.

We have observed, to our amazement, that the data for the Australian coast fall on a straight line of negative slope. In fact, any set of data can be made to fall on a straight line in order to make it look more scientific. That is why graph paper was invented. None of this, of course, should detract from the extremely impressive straight line in figure 2. A little figuring gives its equation, namely $\log_{10} L(\eta) = 4.4 - 0.14 \log_{10}\eta$, or after exponentiating, $L(\eta) = 25{,}000\ \eta^{-0.14}$. This means that, at the scale-length η, our coast-line is approximated by a polygon with $25{,}000\ \eta^{-1.14}$ sides, each of length η, giving the estimate $25\text{-}000\ \eta^{1-1.14}$ for the total length. How can we understand the funny numbers in the formula? Let us remind ourselves first that to measure the length of a straight line of length 25,000m, say, with a short ruler of length η, we divide it into $25{,}000\ \eta^{-1}$ segments each of length η and then simply add the lengths of the segments to get the total length $25{,}000\ \eta^{-1}\eta = 25{,}000$m. To measure the area of a square of side 158m, say, we can pave the area with lots of little squares of side η, raise η to the power 2 to get the area of each little square, and then add the areas together. Since there are $(158\ \eta^{-1})^2 \approx 25{,}000\ \eta^{-2}$ little squares, the total area is $25{,}000\ \eta^{-2}\eta^{2} = 25{,}000$ sq.m. In the first case, we are measuring a length which has dimension 1; in the second case, we are measuring an area which has dimension 2. In both cases, our final measure is independent of η and this is what makes it useful. To measure our coast-line in a similar way, we can approximate it by small segments of length η, raise the length of each segment to the power 1.14 and add them together. Since there are $25{,}000\ \eta^{-1.14}$ segments, we get the measure $25{,}000\ \eta^{-1.14}\eta^{1.14} = 25{,}000$ and again this is independent of η. We might say that our coast-line has dimension 1.14 and that its 1.14-dimensional measure is 25,000.

All this is very strange. How can a line have a dimension bigger than 1? And how can a dimension possibly be a fraction? It becomes a little more acceptable, perhaps, when we see that curves with these peculiar properties were considered by a number of highly respectable mathematicians around 1900. These arose in attempts to clarify the intuitive ideas about curves and surfaces which everyone understood but which no-one could define satisfactorily. It is strange, too, that these products of a debate in the foundations of the calculus should find an application to the real world, but that seems to be the way in which mathematics work.

Our first example of a funny curve is the snow-flake curve given by Helge von Koch in 1904. We start with an equilateral triangle, as in figure 5.

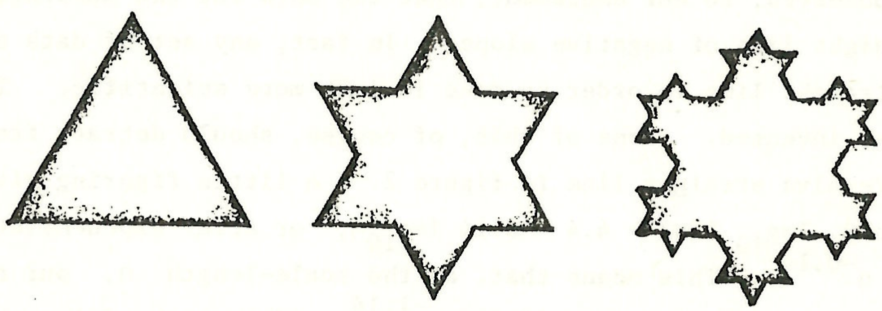

Figure 5

On the middle third of each side, we construct an equilateral triangle, giving the star of David, Then the middle third of each side of the star we construct an equilateral trianlge and so on, ad infinitum. This sequence of curves converges to a limit which is illustrated in figure 6. You can imagine that the stages in the construction of the snow-flake curve show the detail which becomes

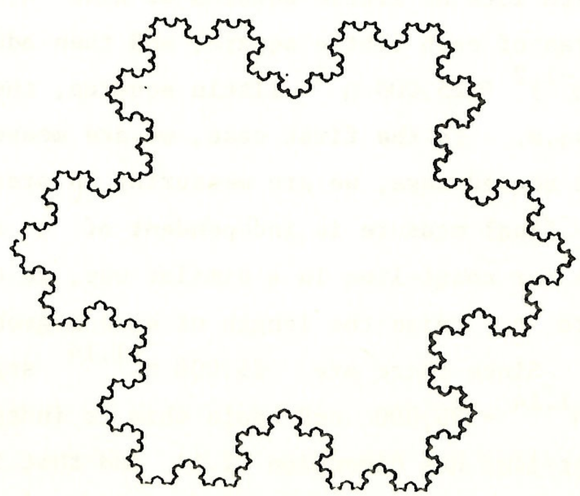

Figure 6

visible if we examine the curve at finer and finer resolution. In the same way if we look at a piece of coast-line at a scale of 1/100,000 we may see just a smooth bay. At a scale of 1/10,000 innumerable sub-bays become visible. At a scale of 1/1,000, sub-sub-bays appear, and so on. At each stage in the construction of the snow-flake, the total length is increased by a factor of 4/3. Thus, just like a coast-line, the snow-flake curve has infinite length. The analogy extends further, as we shall see by applying method 1 to measure the snow-flake at various scales. Suppose the construction begins, as in figure 5, with an equilateral triangle of side 1. If we approximate the snow-flake using a yardstick of length 1, we get the first figure in figure 5, so $L(1) = 3$. If we approximate the snow-flake using a yardstick of length 1/3, we get the middle figure in figure 5, so $L(3/4) = 4$. The approximation with a yardstick of length 1/9 is the third figure in figure 5, so $L(1/9) = 16/3$. Continuing in this way we get $L(3^{-k}) = 3(4/3)^k$. Indeed, at the scale $\eta = 3^{-k}$, the approximate snow-flake consists of 3×4^k segments, each of length 3^{-k}. Can we find a "dimension" in which the measure of the approximate snow-flake comes out independent of k? According to the procedure above, we raise each side to the power D and then add them together, giving $3 \times 4^k \times 3^{-kD} = 3 \times 10^{k(\log 4 - D \log 3)}$ and this is just 3 if we choose $D = \log_{10}4/\log_{10}3 = 1.26$. So the snow-flake curve has dimension 1.26 and its 1.26-dimensional measure is 3. The dimension measures the degree of wiggliness of the curve, so that the snow-flake is apparently more wiggly than our coast-line.

The next curve was discovered by Giuseppe Peano in 1890. We begin with a square of side 1, say, as in figure 7. To get the next stage, the typical side

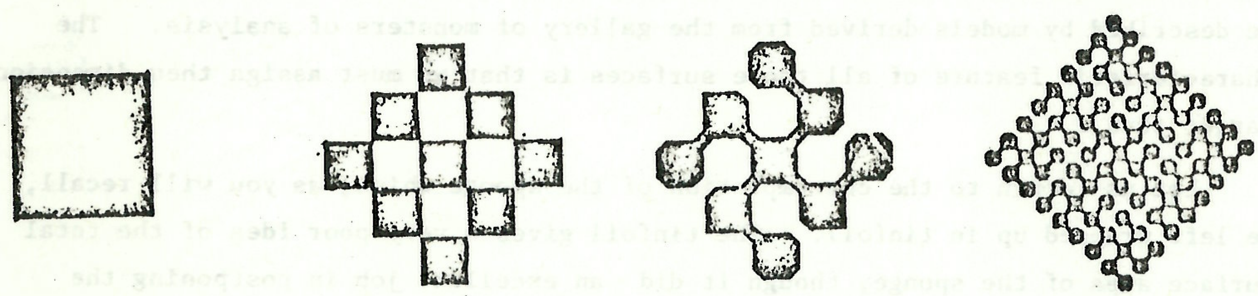

Figure 7

AB of the square is replaced by 9 segments of length 1/3; the corners are cut off so that the curve can be followed unambiguously from A to B. For the next stage, each side of this figure is replaced by 9 segments of length 1/9 in the same pattern, and do on. The limit curve fills the whole of a square of side $\sqrt{2}$.

This, of course, is impossible. Yet, since it has been done, it gives new meaning to the well-known boast that the impossible merely takes a little longer. After k steps in the construction, Peano's curve consists of 4×9^k segments, each of length 3^{-k}. If we raise each side to the power 2 and add, we get $4 \times 9^k \times 3^{-2k} = 4$, independent of k. So we have confirmed that the Peano curve has dimension 2, which should be the case since it fills an area in the plane. As though this monster was not bad enough, it is possible to construct curves along these lines which fill a cube in three-dimensional space. This, of course, is impossible. All such constructions were branded as monsters at the turn of the century, but it is not too hard to see their counterparts in real life. Consider, for example, the theory of the circulation of the blood as propounded in "The Merchant of Venice". The downfall of Shylock requires that there should be an artery and a vein infinitely near every point of the flesh. That is, every point of the flesh lies on the boundary between the two blood networks. On the other hand, the volume of all the arteries and veins is only a small fraction of the body volume, leaving the bulk to flesh. Thus the flesh is a surface which fills most of the volume of the body. This, of course, is impossible.

There are many other natural surfaces which are so irregular that they cannot be assigned an area in the usual way which enables us to measure the surface area of a sphere or a cylinder. The craters on the moon form cascades in the same way as the bays and headlands along a coast-line. Big craters have little craters and little craters have lesser craters, not to mention the holes made by all those marauding astronauts with their flag-poles. Again, you cannot measure the surface area of a sponge by wrapping a sheet of tin-foil around it. (This might be a good way of dealing with a sponge cake that didn't rise.) Many of these phenomena can be described by models derived from the gallery of monsters of analysis. The characteristic feature of all these surfaces is that we must assign them dimensions larger than 2.

Let us return to the contemplation of the sponge which, as you will recall, we left wrapped up in tinfoil. The tinfoil gives a very poor idea of the total surface area of the sponge, though it did an excellent job in postponing the window-washing last week-end. A more realistic way to tame the sponge is to apply method 3. Imagine each point of the sponge replaced by a drop of water of radius η. Instead of a sponge, we now have dozens of sheets of lasagne, each of thickness 2η, and hopelessly stuck together. Now, we can measure the volume of the lasagne, by squeezing the water out of the sponge, and divide by 2η to get our estimate of the surface area. Do you think this is a practical method? Let us return even further to method 1. The idea here is to replace the surface by something smooth enough to measure, which follows the meanderings of the surface down to the scale η. We might try to replace the surface by a surface made up from trillions of little triangles in the same way as we replaced our coast-line by lots of little line segments. This is rather hard to visualise, so instead of triangulating the sponge, let us try to triangulate a cylinder.

Our modest aim, for the moment, is to measure the curved surface of the right circular cylinder

$$S = \{(x,y,z) : x^2 + y^2 = 1, \quad 0 \leqslant z \leqslant 1\},$$

shown in figure 8. The idea is to approximate the surface by a network of little triangles which we construct as follows.

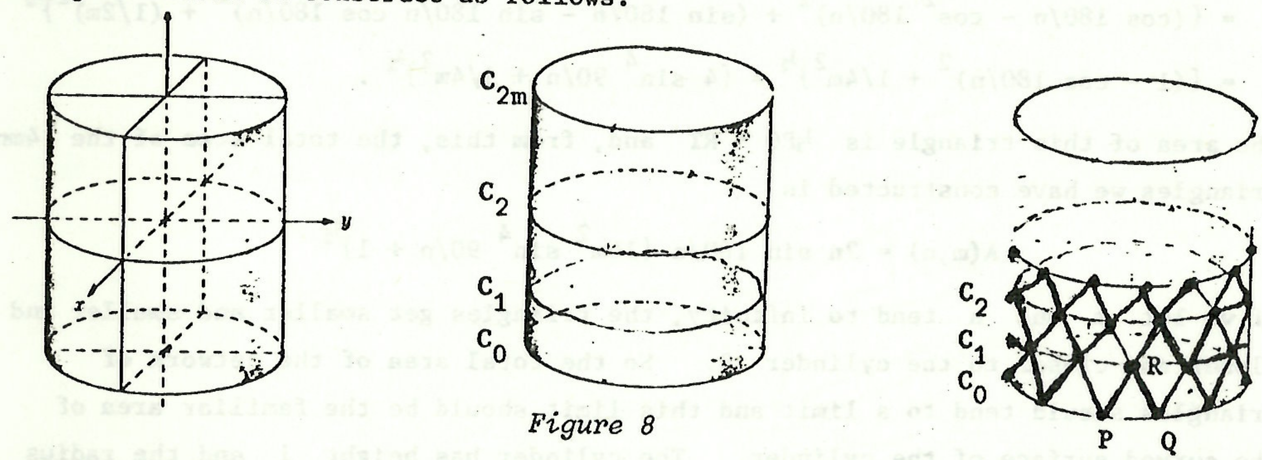

Figure 8

First draw $2m + 1$ circles $C_0, C_1, \ldots, C_{2m}$ round the cylinder, spaced equally in the z-direction. The equation of the j-th circle is

$$C_j = \{(x,y,z) : x^2 + y^2 = 1, \quad z = j/2m\}.$$

Next, mark n equally spaced points round each of these circles. To get the n points on the circle C_0, we start at $(1,0,0)$ and move round the circumference through successive angles of $360°/n$; the points we get have coordinates $(\cos(360k/n), \sin(360k/n), 0)$ with $k = 0, 1, 2, \ldots, n - 1$. The starting point on C_1 lies above the midpoint of the arc joining the first two points on C_0; it is $(\cos 360/2n, \text{win } 360/2n, 1/2m)$. Then we continue round C_1 through successive angles of $360/n$ as before. The starting point on C_2 lies above the starting point on C_0, that on C_3 lies above the starting point on C_1, and so on. So the points on the even-numbered circles are

$(\cos(360k/n), \sin(360k/n), j/2m)$, $j = 0, 2, 4, \ldots, 2m$, $k = 0, 1, 2, \ldots, n - 1$;

the points on the odd-numbered circles are

$(\cos 360(k + \tfrac{1}{2})/n, \sin 360(k + \tfrac{1}{2})/n, j/2m)$, $j = 1, 3, 5, \ldots, 2m-1$, $k = 0, 1, 2, \ldots, n-1$.

We join the points on neighbouring circles by straight lines in a zig-zag fashion as shown in the figure. In this way, we get a network of triangles with their vertices lying on the cylinder S. There are 2n triangles between the circles C_0 and C_1, so there are 4mn triangles altogether. These triangles are all congruent, so we only need to calculate the area of the one with vertices $P = (1,0,0)$, $Q = (\cos(360/n), \sin(360/n,0)$, $R = (\cos(180/n), \sin(180/n), 1/2m)$. This triangle has base

$$PQ = \{(1 - \cos 360/n)^2 + \sin^2 360/n\}^{\frac{1}{2}} = \{2 - 2\cos 360/n\}^{\frac{1}{2}} = 2 \sin 180/n.$$

The altitude of the triangle joins R to the midpoint T of PQ, and its length is

$$RT = \{(\cos 180/n - \tfrac{1}{2}(1 + \cos 360/n))^2 + (\sin 180/n - \tfrac{1}{2}\sin 360/n)^2 + (1/2m)^2\}^{\frac{1}{2}}$$

$$= \{(\cos 180/n - \cos^2 180/n)^2 + (\sin 180/n - \sin 180/n \cos 180/n)^2 + (1/2m)^2\}^{\frac{1}{2}}$$

$$= \{(1 - \cos 180/n)^2 + 1/4m^2\}^{\frac{1}{2}} = \{4 \sin^4 90/n + 1/4m^2\}^{\frac{1}{2}}.$$

The area of this triangle is $\frac{1}{2}PQ \times RT$ and, from this, the total area of the 4mn triangles we have constructed is

$$A(m,n) = 2n \sin 180/n \{16m^2 \sin^4 90/n + 1\}^{\frac{1}{2}}$$

If we let m and n tend to infinity, the triangles get smaller and smaller and closer and closer to the cylinder S. So the total area of the network of triangles should tend to a limit and this limit should be the familiar area of the curved surface of the cylinder. The cylinder has height 1 and the radius of its base is 1, so the area in question is 2π. That seems fair enough, but since we have gone to all this trouble to calculate A(m,n), we may as well carry on. Recall that if θ is a small angle measured in degrees, then

$$\sin \theta \approx \pi\theta/180$$

and the relative error in this approximation gets smaller and smaller as θ tends to 0. (You can see this from figure 1. If θ is small, then $\sin\theta$ is approximately equal to the segment PQ divided by OP and this is approximately equal to the arc PQ divided by OP. Now the arc PQ is proportional to the angle θ and it is equal to the circumference of the circle when $\theta = 360°$, so we get $\sin\theta \approx (\text{arc } PQ)/OP = 2\pi\,\theta/360$.) If we use this approximation in the expression for A(m,n), we get

$$A(m,n) \approx 2\pi\{(\pi^4 m^2/n^4) + 1\}^{\frac{1}{2}}$$

and this approximation becomes better and better as m and n get larger and larger. Now we shall make some special choices for m and n.

(i) Suppose $m = n$. Then $A(m,n) \approx 2\pi\{1 + \pi^4/n^2\}^{\frac{1}{2}}$, so $A(n,n) \to 2\pi$ as $n \to \infty$. There, told you so! The total area of the triangles approaches the area of the cylinder. All that ghastly trigonometry was obviously a waste of time.

(ii) Suppose $m = xn^2$ with $x > 0$. Then $A(m,n) \approx 2\pi\{1 + \pi^4 x^2\}^{\frac{1}{2}}$, so $A(xn^2,n) \to 2\pi\{1 + \pi^4 x^2\}^{\frac{1}{2}}$ as $n \to \infty$. But that's ridiculous. If we choose a suitable value for x, we can make the total area of the triangles tend to anything we like, so long as its bigger than 2π. All that ghastly trigonometry certainly was a waste of time.

(iii) Suppose $m = n^3$. Then $A(m,n) \approx 2\pi\{\pi^4 n^2 + 1\}^{\frac{1}{2}}$, so $A(n^3,n) \to \infty$. No comment.

This construction was suggested by Herman Schwarz in 1880. It is an advanced case of the paradox discussed in method 2 above. What happens in (ii) and (iii) is that the rings of triangles become more and more pleated, like a Chinese lantern, and their area can therefore add up to something much larger than the area of the cylinder. (See Parabola, Volume 17, Number 1.) The moral is that neither method 1 nor method 2 is much use in measuring a complicated surface because we will probably not be able to tell if the approximating surfaces are too pleated.

Here, to conclude our story, is a positive note. The answer to the question in the title is 19,540km (including Tasmania).

<u>Reference</u>. This article was inspired by the fascinating book by Benoit Mandelbrot called "Fractals : form, chance, and dimension" (Freeman, 1977). For example, the data for the Australian coast-line was taken from Chapter 2 of this book.