





The importance of being peripheral Professor John D Barrow FRS 25 November 2008

Today I want to talk about one simple idea, and it is the relation between a boundary and the area or the volume that it encloses. But this simple mathematical idea has all sorts of consequences in the physical and biological world, and even in the esoteric world of black holes.

There is a curious history about anything to do with surfaces and boundaries and perimeters. If you look at old maps, then you will often find cities are built around rivers. If you are going to build a walled city, it is a good idea to save on building materials perhaps by creating a boundary that is semicircular - you will see why that might be in a moment - and you are saving perhaps a good deal of space, compared with having a boundary that is square.

There is a famous surface problem that dates from antiquity; mathematicians even know it as Queen Dido's problem. Dido was the Queen of Carthage, perhaps the founder of Carthage, and when she arrived with her followers on the North African coast and decided that she would ask for some land on which to create her new city state, she was clearly treated with some impunity by the rulers that be in the region. They told her that she would be entitled to have all the land that she can encompass by one ox-hide. So clearly, the person making the offer had in mind that perhaps she would have a kingdom about the size of a couple of tabletops; but she was rather more ingenious, and set about cutting the ox-hide into the narrowest of strips that she could possibly create. Eventually, you could make a very remarkable length of tiny strips of ox-hide, and to use that as the perimeter to enclose a very wide area of the coastline, around which to establish her city state. So this is the beginnings in literature of this problem that approaches the relationship between the length of a boundary and the area that is enclosed within it.

If you do a few little experiments, you take perhaps a circle of string, starting with it in a circular configuration and then wiggle it a little, you are not changing the length of the circumference - the perimeter - of your string, but you can get some sense as to what happens to the area that it encloses. As you get more and more oval-like, flatter and flatter, the area which you are enclosing becomes smaller and smaller. If you think the other way round, if you want to start by enclosing a particular area, then you can see that you could make your boundary more and more and more wiggly, on a finer and finer and finer scale, and you would not really change the enclosed area very much. The intuitions you get from these sorts of examples persuade you that perhaps the circle is the situation where the most area is enclosed by a given perimeter. Conversely, if you change the shape of the region from being a circle, the area enclosed is going to get smaller. If you stick with enclosing a given area, then there is no limit to how large you can make the perimeter which is going to enclose that area, but the smallest perimeter is going to arise when it is a circle.

So mathematicians enshrine these simple, intuitive ideas in the form of some theorems that have become known as isoperimetric theorems. They really originated long ago through many people, but through Steiner in particular. What an isoperimetric theorem will tell you, for a flat curve on a flat surface, that if you enclose an area by a closed curve, the square of the perimeter distance must always be bigger than or equal to 4 pi times the area that is enclosed:

(Perimeter) $2 \ge 4p x Area$

The case of equality is the case of the circle, where the perimeter is two pir:



(2pir)2 = 4pi x pir2

So if you deviate from the shape of the perfect circle, then you enter the region where this is an inequality, and the perimeter will always be bigger than 4 pi times the area.

A similar situation holds if you are not dealing with a flat surface, but you are dealing with a solid closed object: the maximum volume for a given circumference that you can enclose is the case of the sphere. So we have an analogous inequality, so the cube of the area is always bigger or equal to 36 pi times the square of the volume; and the equal case, the best that you can do, is the situation where it is a sphere, the area is 4 pi R squared, and the volume is four-thirds pi R cubed:

(Area)3 ≥ 36pi x (vol)2

(4pir2) 3 = 36pi x (4pir3/3)2

The proof is not entirely straightforward - you need to know different ways to do it - but you can see various thought experiments that you could carry out in order to arrive at a proof. You could start with the situation of the circle or the sphere, and perturb it slightly, and show that the area enclosed always went in one direction and always got smaller.

Or you could take a situation where you have a curve with an indentation:

You could draw a line across, where the indentation goes into the figure. You could then rotate your indentation about that line, so the length of that line still stays the same and the perimeter of the object is staying the same. But, obviously, now more area is being enclosed because you are moving from a concave figure to a convex one.

So what you see from these theorems that is intuitively obvious is that the optimal cases, the maximum area enclosed, is the case of the circle or of the sphere. If I were to pick up a little metal ring, dip it in some washing-up liquid, and blow some bubbles in the air, all the bubbles would be spherical, because they want to minimise the surface area which they expose to the area around them. In this way they exploit the second one of these theorems to produce the minimum possible area.

So this idea of the importance and uniqueness of the circular or the spherical configuration has all sorts of interesting consequences in the world around us. One example of this is an animal is cold and wants to retain as much of its own body heat as possible. Here the animals will curl up into as small a ball as possible in order to minimise its boundary and to try and be as spherical as it possibly can. So in situations where you want to keep to warm, you want to minimise your surface exposure to the atmosphere around you and you exploit the theorem by becoming as spherical as you possibly can. Perhaps also, if you want to avoid being detected, you might think you want to minimise the interface between your edge and anything that is trying to search you out.

Let us look at it just dimensionally and a little more simply at this business of keeping warm. If you have a device - it could be a computer, or it could be a human body - that has a source of heating within it so it generates heat energy, and it has to radiate heat away to remain in equilibrium or else it will just heat up and melt. If it does not keep generating heat, it will just cool off and lose all its energy and freeze, down to the temperature of the surroundings. So in practice, the heat generation rate is proportional to the volume, so your heat generation rate is proportional to your bodily volume, which is also proportional to your body mass, because your density is essentially constant.

Heat generation \propto volume \propto L3

But what about the heat loss? If you haven't got any clothes on, your heat loss is exactly proportional to your surface area. If you put clothes on, you can reduce that rate of cooling in different ways, but ultimately, it will still be proportional to the area that you are exposing to the cooler environment:

Heat loss \propto surface \propto L2

If you divide the heat generation rate by the heat loss, the surface area is proportional to some length squared, characterising your volume, to some length cubed, then your heating over your cooling is proportional to your size:

Heating/Cooling \propto L

What does that mean? It means if you want to keep warm, it pays to be big! So your ability to generate heat grows with you more quickly than the rate at which you lose heat. So if you were just trying to build a

bigger and bigger computer, just copying the same design of your little laptop computer, ultimately, it would melt, because it would not be able to keep cooling itself at a rate that would keep pace with the processor heating power that it would generate to keep warm.

So there are a number of lessons that you can learn here: if you want to be an animal that is going to live in a very cold climate, then you want to be big. So if you look in the Arctic, you do not find lots of tiny little birds and mice and so forth: they are not thermodynamically viable; they are not able to keep as warm as they need to in order to maintain constant body temperature. Similarly, if you want to keep cool, if you want to live in the Tropics, spend your life in the Caribbean in a tax haven, it pays to be small. If you look at the average size of birds as you move from the equator to the poles, then the gradually become, on the average, larger. So you find lots and lots of small creatures in hot climates and increasingly larger average sized creatures in cold climates.

The other thing that you can do of course if you want to keep warm is you can stick together. What is going on here is a reflection of what we have been looking at. If animals cluster together in tights packs or herds, what they are doing is minimising their external surface area. So if they were standing separately, with space between them, they would be able to cool from the entire boundary of their body, but by gathering together in a huddle, they reduce the exposed surfaced area, and they behave as though they were one larger creature, which is well set-up to remain warm and in equilibrium in a cold situation.

The other situation where we see a boundary effect that is a little more complicated is the problem of being intercepted. If you look at large collections of wildebeest or other creatures in the savannah regions of Africa and elsewhere, the herd behaviour is an interesting mathematical problem. There are lots of things going on at the same time, and if you were going to build a computer simulation of this, you would ask what the rules of the game that you would want to tell the computer would be. Well, one sort of rule is that you do not want to be an outlier; you do not want to be on the edge of the collection, because the first one that the lion or the cheetah or the leopard is going to grab is the one on the outside. So they keep moving around all the time, almost like a mixing process, and the idea of each wildebeest is to make sure that there is somebody else along the line of sight in between you and the lion that is sitting nearby. So once a predator appears, this movement begins, to try and make sure that there is somebody else in between you and the predator. This has a particular mathematical consequence for the structure of the herd, that you produce what mathematicians call a Voronoi tessellation structure for the pattern of the herd - but that is another story.

But there are different sorts of predator problems that are much more human-based that have a close link to our problem. If we go back to World War II, there was a really critical problem, particularly at a stage of the War when Britain was suffering badly from having its supply lines from the United States decimated by U-Boat attacks. So the challenge was: what is the optimal strategy to try to make sure that cargo boats coming across from the USA and from Canada should adopt? Should you stick together in one big convoy, or should you split up into lots of small convoys?

This was a problem that was put to mathematical physicists of the day, and it was Lord Blackett, Nobel Prize winner in Physics and also a commander of a ship at some stage, he who produced the convincing argument that was adopted. Intuitively, you tend to think that if you all split up into lots of little groups, surely you can sneak by and no one will notice you, but if you are going along in one big convoy you appear to be much easier to see, so surely that is the worst thing to do. In fact, it turns out that it is not, and that the optimal strategy is to stick together as one big convoy.

So why is that? It is another feature of our problem of boundaries and areas. If you have a large region, like a piece of paper, and you split it up, because it is rectangular, it is very easy to see that when you split it, that you have got a lot more edge than you had before. So as a general rule, if you split big things up, and you conserve the total area in this case, you always create more boundary. The same is true if your convoys are not square or rectangular, but if they are circular. Therefore, what Blackett pointed out was that if you have one large convoy and you divide it up into three convoys whose total area equals the same area - because there is a limit to how close together the ships can be, so they are always as close together as they can be - this is the worst scenario. There are two reasons.

When you send the convoy across the Atlantic, you have to protect it with destroyers, and what the destroyers do is to orbit around in circles, and of course the distance that they go is equal to two pi times the radius of the convoy. If you split it up into three pieces, the sum of these perimeters is bigger than what it is for the single convoy, so you need more guard destroyers - they have to spend more time and fuel patrolling the outskirts of your convoy.

G

But the second is that you are actually more likely to be seen by the U-Boat, because what the U-Boat does is you look through a periscope, and what you see is basically the diameter of the convoy, because they look at the convoy's circular shape from one direction only. So you see only the diameter, and the sum of the three diameters of these convoys is bigger than the diameter of the one big convoy. So you are more likely to be seen by the U-Boat periscope if you were three small convoys of the same total size than if you were one.

To go into the mathematics of this, you start with your big convoy that is covering an area of the sea, A, and you just split it into two, each of size A over 2, so half A. The perimeter of the single big convoy is just two pi r, so it is twice the square root of pi times the area:

2pi√A

That is what the boat would have to steam around to protect it, and it is what the U-Boats would see. But when you split it into two of area A/2, the total perimeter of the two, of A/2, is two times two times the square root of pi times half-A:

2x2xpi√½A

This is bigger, by the square root of two, which is roughly 1.41, than the perimeter of the single convoy. So you are 41% more likely to be seen, you are 41% harder to protect, than if you stayed as one group.

This was an example of what we discovered earlier through the idea of cutting up a piece of paper: when I cut up a piece of paper you increase the amount of surface, because cutting up a rectangular piece of paper in half creates these two new edges. This is a simple fact that has potentially catastrophic consequences.

For example, if you have a cube something and then you divide it up into quarter cubes, and then divide those quarter cubes into quarter cubes again, then you find that your surface area increases very quickly. If the initial surface area is A, then when you have sliced it for the first time, the total surface is doubled, and if you slice it again, it is multiplied by four, and so on. If this thing that we are dividing up is dust, junk, crud, the sort of stuff that sits around in old warehouses or under your stairs or under the escalators at underground stations in the old days, then you begin to see why this is potentially catastrophic. What determines the rate of combustion of material is the area of their surface which is in contact with the oxygen in the air; so the rate of combustion is determined by the total area that is exposed. So if you cut things up into lots of little pieces, if they fragment into grains of dust, the area grows massively, and the potential for a spontaneous flash-fire grows dramatically. So this is why dust is such a fire hazard. Ordinarily, you do not think of dust as being terribly inflammable, but it is especially dangerous because of the huge surface area that is exposed. So you hear cases of people going into an old warehouse where everything is covered in dust, they see a fire has started, they pick up the fire extinguisher, they point it at the fire and squirt it, and the fire gets ten times worse immediately, because the fire extinguisher has blown dust all up into the air, and it is that that has combusted.

There is another curious example of this: that surface area also affects the scattering of things. So if you look at droplets of water, or pollutants in the Earth's atmosphere, then those pollutants become the points around which liquid condenses. When sunlight comes into the earth's atmosphere, it scatters off particles in the atmosphere. If the particles are very small, you will get a form of scattering that produces the blue of the sky, because the scattering depends upon the colour and the energy of the light that comes in. When the particles become a bit bigger, then all the wavelengths scatter pretty much the same, and that is why the sky is blue during the day, but at sunset, when you are looking at the light that has not been scattered very much from the other side of the sky, it is red.

Some years ago, after the 9/11 attacks, there was a rather interesting phenomena that was observed. What we see, in general, is that sunlight scatters back off little water droplets in the atmosphere that coagulate around pollutants. Lots of these pollutants are caused by aircraft exhaust, from jet aircraft. If you have lots more flights, you have more pollutants, you have more particles and you get lots of coagulation around them. But even if there are more particles, there is still the same amount of water vapour, so each particle ends up with less water vapour coagulating around it, so you get lots of smaller droplets. What this means, as we have seen, if you divide something up into lots of smaller bits, you get relatively more surface area than if they all stayed as one big drop or a small number of fairly big drops. So the more pollutants you had, the more surface area that you had exposed to the sunlight, the more back scattering you got, and the less radiation made it to the Earth's surface, and so the Earth remained slightly cooler than if there had been fewer droplets.

This effect, which has become known as global dimming, became particularly obvious after the 9/11 events. You will remember that there was a week or more where there was a complete ban on jet flight internationally. Even in that short period of time, the lack of new small pollutants in the atmosphere led to a slight and detectable warming, because the droplets became larger, there were a smaller number of them, there was less back scattering of sunlight and so more sunlight reached the Earth, and in the evaporation experiments, there was less evaporation. So it was an effect that had been hidden because there was so much over-flying, but it meant that global warming was really worse than we had thought it had been, because this effect had been going on to produce a little bit of cooling for you through the back scattering.

All these cases we have looked at so far are some of the advantages of having rather small boundaries, but there are situations where it is good to have large boundaries. So if you want to keep cool, it is good to have a large boundary. So, if you are an exhibitionist and you want to be seen, you would want to have a large surface area.

If you make a towel, you want it to soak up lots of moisture, so you make the surface of the towel as unsmooth as possible, so it is rather crenulated and knotted. The reason for that is to increase the surface area of the towelling which is going to be in contact with your body and absorb moisture. If you need to ingest nutrients, if you are a tree or a plant, it pays to have a lot of surface area, because that is going to be the interface with the outside world, where your food and drink is floating around waiting to be absorbed.

If you are having to take some medicine at the moment, you will notice that the little tablets are never spherical. They are always lozenge shaped so they are flat, and they have a larger surface area for their volume and for their weight than they would if they were spherical. Why is that? Because their job in life is to dissolve, so if you want something to dissolve quickly, you want it to have as large a surface area as you possibly can.

If you have a swimming pool and you have those cleansing tablets that keep the pool clean and chlorinated, you will find that they are never spherical. They are really as flat as they can be without them falling to pieces, so that they will dissolve rapidly and efficiently.

There is another, rather seasonal, example I thought that it was important to put in, where surfaces and areas play a role, and it is in cooking turkeys. If you look through cookery books, and you look for the cooking time for a given weight of turkey, then it is really very perplexing if you have a scientific bent of mind. For instance, there are rules like you add an extra ten minutes for every pound of the turkey, and then add another half an hour for luck, and you do something else, and then somebody else just says it is this fixed time, and then you add a little bit if it does not seem to be done after that time, but somebody else will say, if you have stuffing, it is this time, and if you don't, it is the other.

When you have analysed all these and put them together and turn them into formulae, I seem to remember it was only Mrs Beeton's ancient book that really got it right. Only that one, at some level, had a scientific basis for what it was proposing, because if you are cooking something, then you are engaged in a problem of heat diffusion. What is happening is that heat has to find its way from outside right to the centre of the turkey.

Heat spreads in a rather random way: one molecule gets a bit more energetic, bounces into the next one, which deflects in a random way, and that one deflects in a random way and so on. It is not at all like a straight line walk - it is a bit slower than that. So whereas it might take you just n steps to walk from here to there, if you are proceeding like a heat flow, it will take you n2 steps to do that, because it is as if after each step you toss a coin to decide which direction you should go in next. So if it takes you n steps to walk home in a straight line when you are sober, it will take you n2 steps to walk home if you are completely inebriated. So that is why this random walk is sometimes called the drunkard's walk. This is how the heat spreads from the outside to the inside. What this means is that the time is proportional to a distance squared - the square of the number of steps. The distance squared, in the turkey, is proportional to the surface area of the turkey. Its actual shape will tell you what the constant of proportionality is, whether it is four pi, two pie, three or one or whatever, but in terms of proportionality, the time for the heat to get in goes like the surface area, it goes like a size squared. The weight of a turkey is proportional: it is just its density times its volume. Its volume looks like its size cubed. Density is a constant thing, so size squared is proportional to the two-thirds power. So here's the golden turkey rule:

Time \propto area \propto (size)2 \propto (weight)2/3

Weight \propto density x (size)3



The cooking time is proportional to the two-thirds power of the weight, or if you do not like two-thirds, the cube of the cooking time is proportional to the square of the weight.

To get to the mathematics underlying all of this, we should look at this so-called heat conduction equation, or diffusion equation:

T/t=k ∇ 2T so T/t \propto T/d2 and d2 \propto t

Let us use T say for temperature, delta T by delta t, the rate of change in time of the quantity is equal to some constant, which tells you how good the thing is at actually spreading the heat - if it is made of metal or wood or something, it is rather different than if it is made of turkey meat - and it is proportional to grad squared, or the second spatial derivative of T. So dimensionally, this means our big T divided by a time is proportional to big T divided by a distance squared. Therefore, the big T's cancel out, so Delta squared is proportional to T; the time is proportional to size squared.

There is a rather analogous thing: the strength of things is also proportional to your surface area. So the strength of your arm is not proportional to your volume, because if you were to break it, it would break just on a slice, on a cross-sectional area, and you would break all the atomic bonds through that slice. So there would be an argument like this that tells you that your strength is proportional to your area. Of course, your weight is proportional to your volume, so as you grow, your strength does not keep pace with your weight, so your strength divided by your weight is like one over your size, and if you tried to build people just by inflating them, making them bigger and bigger, eventually they would not be strong enough to stand up and they would collapse. So your strength is proportional to your area, and thus your strength is proportional to your weight to the two-thirds.

As a little test of this, if you look in my 100 Essential Things You Didn't Know You Didn't Know book, you will find this plotted and worked out. An interesting example to think about: look at the world weightlifting records. They are rather meticulously kept. The weightlifters are in particular weight limits dictated by their weight. If you plot the weight lifted against the weight of the weightlifter for all the world records, it beautifully follows this law, so the weight lifted and the strength of the lifter, grows like the two-thirds power of the weight of the weightlifter across all the world records from the flyweights up to the super-heavyweights.

But, to get back on track with our boundaries, one of the things that we have seen is that you can do things to boundaries to make the boundary get bigger and bigger without changing the area that it is enclosing. This is something that might be rather useful to you if you are one of these things, like a tree or a leaf, that wants to have as big a boundary as you possibly can. Mathematicians have gradually latched upon the recipe and the idea for making structures of that sort. Nature figured it out a much longer time ago.

For example, if we have a triangle, then we can make an object which has a bigger and bigger area by, say, taking out the middle third of each side and erecting another triangle, and then doing that again on every side:

You get something that looks increasing crenulated. Eventually, it is possible, if you go on forever, to make the boundary as large as you like. So, if you wanted to, it could become infinitely long, if you made an infinite number of those little crenulations on it, but the area it encloses is completely finite. So this is completely in agreement with our isoperimetric theorem 'the square of the perimeter has to be bigger than something times the area, so there is no limit on how large the perimeter can be that encloses a particular area. This is one of the recipes for making it as large as you want.

There is also a three-dimensional example of this:

This is a mathematical construction that is usually called Menger's Sponge, after Karl Menger, who first constructed it in the 1930s. So again, there is a recipe for removing volumes from a cube, and you can remove an infinite number of these volumes, of smaller and smaller size, so you end up with an object which has a fixed, finite volume, but has an infinitely large surface area.

This leads us to think more carefully about what we mean by the length of surfaces or boundaries. A classic, rather specific, problem was to ask how long the British coastline is. At first, you might think that this was a pretty obvious answer - you look it up on your map, you drive around it - but If I sent you out with a ruler which was 50 metres long to measure the length of the coastline, and sent someone else out with a ruler that was one centimetre long to measure the coastline, when you came back, you would provide very

different answers. The person with the shorter ruler would measure a rather different length of coastline than the person with the longer one, and it is clear why that would be the case. Our coastline us rather wiggly, unstructured, and if you have a long ruler, you would miss lots of lengths around those wiggles; you really rather underestimate the length of the coastline. If you have a shorter ruler, you capture a lot more of the wiggliness, and you reckon the coastline is much longer. If you took an even shorter ruler, you would get an even longer measure of the coastline, and so you might begin to worry that, if you pick a short enough ruler, you can make the coastline as long as you wish, and, indeed, you can. So there is no sort of invariant definition of what the length of our coastline is, but the key piece of information is going to be how does your estimate of the length change as you change the length of your ruler, and what that is telling you is how wiggly and complicated the coastline is.

Let us suppose that we have our ruler of length d and we are going to use that to cover the boundary of the coastline, and we ask how many times we have to put that ruler down to cover the whole of the coastline; that is going to be a measure of the length. For particular types of wiggliness, the answer you get depends on some number - it will depend on the particular coastline - divided by the length of your ruler to some power:

N(d) = M/dD

If the line of the coast you were measuring was just a straight line - which would be a rather weird coastline - that number would be one, but the more wiggly and complicated the coastline becomes, so the larger that number becomes and it tends increasingly towards two. As examples, the west coast of Britain is about 1.25; Australia is 1.13; South Africa, which is rather even is very close to one.

So the quality of complicated curves, that they have a structure that when you try and measure their length with a ruler of length d, you get an answer of this sort, is called a fractal curve. So if the answer did not just depend on a power of the size of the ruler - perhaps there was another term here like d squared or the cosine of d - then that would not be a fractal. So the key feature here is that if you change the length of your ruler, by multiplying it by some constant, say doubling it or tripling it, then this proportionality would still remain. So if you change d to 2d, N would still be proportional to d to the D. So we call that property self-similar, so the structure has the property that, if you keep looking at it under the magnifying glass, on smaller and smaller scales, it has the same geometrical structure than if you looked at it magnified. So it takes the same pattern, the same design plan, and copies it again and again on smaller and smaller scales. Examples of this found in the natural world include trees, flowers, human lungs, metabolic systems and much else besides. So such structures have been called fractals; they have been known and studied at some level since about 1904. The word 'fractals' was given by Mandelbrot in 1972. You can see that this type of wiggliness that a fractal structure creates is a basic recipe that could be used if you want to evolve in such a way that you maximise your surface while keeping your interior volume or area essentially the same.

All over the universe, as we will see later, and certainly around the Earth, we see many situations where these structures have been exploited in the natural world, and then sometimes even in artistic creation, to create a form of complexity that we like. Many of Jackson Pollock's paintings make use of this type of fractal invariance, and that is one of the reasons why they look rather nice when you reproduce them in a book, even though the real thing might be eight feet by four feet, or bigger, on the gallery wall. Most works of modern art do not have that scaling invariance property, so when you see them in a book, you get a completely different impression than when you see them actual size on the gallery wall.

A very good visual example is cauliflower, broccoli:

If you look at it under the magnifying glass, the same basic structure is repeated on smaller and smaller scales. Genetically, it is a simple recipe to programme in, but biologically, it is very useful in practice because you are increasing your surface area, but not increasing your volume and your mass at the same rate.

Inside your chest, the same design plan is exploited in the lungs, so you have lots of surface boundary, which is available to absorb oxygen, with very little increase in mass. But there is another interesting feature of the lungs, and it rests on the fact that, if you perturb it, if you cough or breath heavily, you do not want lungs that are liable to break or are badly perturbed. This type of structure is very good at absorbing noise and vibration, and so it is an ideal design plan for something that is going to be coughing and heaving and breathing and panting: it absorbs those vibrations very efficiently. The vibrations are sort of carried, as it were, to all these little extreme limits of the surface, where they dampen out.

G

This is not unrelated to the coastline problem. So having a coastline where we saw there was the problem of the length, this type of fractal structure, with many small inlets and divisions, is also useful if you want a certain type of robustness against the perturbations of the incoming and outgoing tides and winds. So this type of structure is, in some ways, the most stable against erosion and being washed away by heavy winds and tides.

There is one other rather curious feature of this use amongst living things of fractal structure. In recent years, biophysicists have come to appreciate that many organisms spread nutrients through their bodies, whether they are trees or animals, in a way that is characteristic of the distribution of nutrients through a fractal network, a network that works as efficiently as possible from the smallest scale of a single capillary right up to the largest scales for nutrient transport.

Kleiber's law is a famous law of biology, which had, until recently, remained something of a mystery. What it plots here is the metabolic rate in some weird and strange units that biologists like to use, kilocalories per hour - think of it as joules per second or watts, something like that - against a measure of the mass of the organism (they are log-log scales):

You can see all sorts of creatures are on this picture, from tiny scales, unicellular organisms, all the way up to things like frogs and mice and hyenas and elephants. What is curious is that there really is a systematic pattern here; you might even think that there was a single, rather fixed, straight line through all this, but maybe there is just several straight lines, according to the different body plans of the organisms, but they all have essentially the same slope. What the slope is is that the metabolic rate is proportional to the three-quarters' power of the mass of the organism, so that is the three-quarter power law, and that is Kleiber's law: log(metabolic rate) μ log((mass)3/4).

This is the sort of thing that you should challenge a physicist with: 'Can you, within twenty seconds, explain the origin of this law, just on the back of an envelope, from dimensional analysis?' It might be the sort of thing that you could ask at interview of a prospective student. But when you try to do it for this law, you always got a very strange result, because if we just use the sort of reasoning we used before, with the cooling and heating, so all these organisms are in equilibrium, so the rate at which they generate heat metabolically is the same as the rate at which they cool, otherwise they would heat up and melt. So, if you want, you know that the metabolic rate of heating is equal to the heat loss of the body, and that is proportional to its area, and that is proportional to its size squared. As we saw earlier, the mass, the weight is equal to the mass times acceleration due to gravity, is proportional to the volume, so the metabolic rate is proportional to the two-thirds' power - there's our two-thirds' power again. So this is what we would predict on the back of our envelope, that when we plotted all the animals on our pictures, we should find that the metabolic rate is proportional to their mass to the two-thirds' power, so the log-log plot would have a slope of two-thirds. That is the mystery - that this is not the case. It has a slope of three-quarters.

So people started looking at how these organisms might be modelled by some of these fractal networks which are producing a very large surface area of capillaries and nutrient transport tubes inside them. What this means is that, although a tube might be a line in some sense, because it is a very wiggly line, what it is doing is covering an area, and it is behaving as though it were an area even though it is a line. So these organisms, by having interior wiring that is a big wiggly muddle, have a wiring that behaves as though it has got one dimension greater dimension than it truly has.

Let us see if that really holds good. Suppose that we imagine that our organism is D-dimensional and it behaves in the same way. Then, just as before, in a D-dimensional world, the area will be proportional to size to the D minus one, the mass will be proportional to the volume, which will be size to the D, and so when you put this together here, you will not have 2 over 3 - you will have D minus 1 over D, where D is the dimension of the creature.

Rate ∝ (Mass)(D-1)/D

So if it is simply three dimensions, D minus 1 over D is 2 over 3; but our world has creatures that behave as though the world for them is 4 dimensional, D minus 1 over D is 3 over 4. So that is, in effect, what is happening in the internal structure of these living things: although things are communicated along linear tubes, they are fractal and complicated in structure, so they behave as though there is one extra dimension of packing and complexity inside the animal, and that explains why everything you see has this rate proportional to three-quarters mass structure.

That is the last I was going to tell you about such things. I will just say a few more things to show you how the whole issue of surface is on the edge of very important speculations and research in fundamental physics at the moment. The source of that, of all things, are black holes.

A black hole, as you probably learnt in the Gresham Astronomy lectures, is a part of the universe whose gravitational field is so strong that not even light can escape from its surface. If you looked at it in the language of Newton, it would be like a planet whose escape velocity was equal to the speed of light. So its size, the radius of its event horizon, is twice Newton's gravitation constant times its mass divided by the speed of light squared.

R = 2GM/c2

For a simple black hole, its area is just that of a sphere, 4 pi r squared, and so you notice its area is proportional to its mass squared.

Area = 4piR2 ∝ M2

Density $\propto 1/M2$

In 1974, Stephen Hawking discovered that when you include quantum mechanics in the description of a black hole, black holes are actually thermodynamically black bodies. So quantum mechanically, particles can tunnel through the horizon of the black hole and are radiated away as heat. What was mysterious about this process, was that when it was explored in much more detail, it was discovered that the rules that determine how it behaves mirror exactly the laws of thermodynamics, and that the black hole has a temperature given by its surface gravity, and it has an entropy given by its surface area. In anything you do to a black hole - if you throw something at it, if you collide two black holes - that surface area, the total surface area of all the black holes involved can never go down, just like the entropy never goes down in a physical system that is closed. So physicists began to see that the surface area of structures had a sort of hidden gravitational, thermodynamic complexion, and this has led the way to people's study of how you might apply quantum theory to gravitation itself. Where one has reached, then one understands that there is a type of gravitational analogue of entropy, which is proportional to the surface area of the horizon of a black hole, and if you look at it using information theory, you discover that the information content, the number of bits of information that would have to be given to completely specify the internal structure of the black hole, is given not by its volume, as you might have expected, but by its surface area, its entropy and its area. So, rather remarkably, black holes are completely described by the structure of their surface and the magnitude of their surface area.

In recent years, physicists have moved towards investigating what they have called a holographic principle. This is the idea that, much more generally in physics, the measure of the information content of any volume of space is characterised by its surface area. One can prove a theorem of a sort, that if you have any particular region in space, then the maximum amount of information that it can contain is that given when you turn that region into a black hole, and it has the so called Bekenstein-Hawking information content. Now people are exploring whether one can characterise the information content and complexity of the universe in terms simply of its surface area. So if you believe there are other dimensions of space of course, then our three-dimensional universe may simply be the surface of another four-dimensional structure, and so by looking at our three-dimensional surface of volume, you may be looking at the surface area of a larger dimensional structure.

So that is just an aside to whet your appetite, to show you that the issue of surfaces and boundaries is very much a live one, both in physics as well as in mathematics. You might even say that it is the edge of something that really provokes us to look into it much more deeply.

© Professor John D Barrow FRS, 2008