A Tutorial on Vector Areas
(for Natural Sciences / Computer Sciences Tripos Part IA Maths)

Corrections to Dr Ian Rudy (http://www.robinson.cam.ac.uk/iar1/contact.html) please.

The topic of vector areas traditionally causes considerable confusion for students. Vector areas are important in their own right, and they are needed in Lent Term when you cover surface integrals. Somehow, when they resurface in Lent Term, they are much less confusing, so if you understand the basics now, you should be OK.

The starting point for vector areas is the idea that we can use a vector to describe some of the properties of a surface. It's easiest to begin with a small (let's say rectangular) planar surface. The vector area $\underline{dS}$ we use to describe this surface is defined as being perpendicular to the surface and having a magnitude equal to the scalar area of the surface:
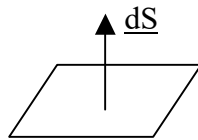

Figure 1

Note that a vector area is a *vector*, and not (as an "area" normally would be) a scalar.

The vector has a very useful property:

The component of the vector in any direction   =   The scalar area of the surface we see when viewing the surface from that direction   [Property 1]

To see that this is true, take a small rectangular piece of paper or card like the one in Figure 1, and view it:
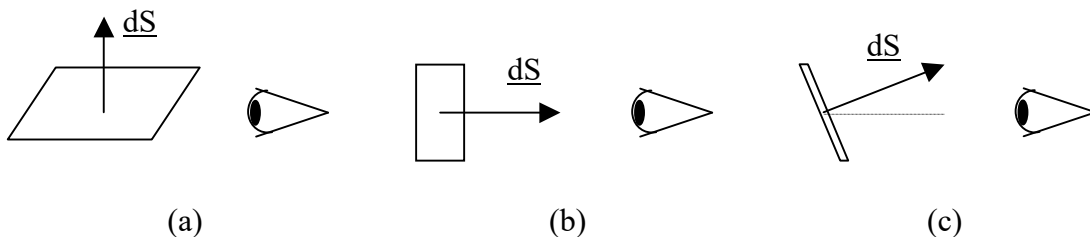

(a)   (b)   (c)

Figure 2

(a) From the side, so that you see just the edge of the paper, in other words, zero area. The component of the vector area in your viewing direction is also zero, because the vector is perpendicular to your viewing direction. So Property 1 holds.

(b) So that the surface is full-on towards you. You see the whole scalar area of the surface, and the component of the vector in your viewing direction is the magnitude of the vector. Because the magnitude of the vector was defined to be equal to the whole scalar area of the surface, Property 1 still holds.

(c) So that the surface is at an angle to the viewing direction. Because both the component of the vector and the scalar area you see vary as the *cosine* of the angle, Property 1 still holds.

One detail of this idea needs attention. In Figure 1, the vector was drawn going *upwards*, but we could equally well have drawn it going *downwards*. Which is correct? By convention, this issue is determined by specifying the direction of traversal of the *rim* of the surface:

Figure 3

Once the direction of traversal of the rim is specified, the direction of the vector is given by the right hand thumb rule: your fingers point in the direction of traversal of the rim, and your thumb points in the direction of the vector.

This means you need to take note of how the surface is defined. If a question describes "a rectangle OABCO", that tells you the traversal direction of the rim. The opposite traversal direction would be described as "OCBAO".

So far we have considered only a simple, planar, surface. Let's make things a bit more sophisticated by putting two of them together, like the surface of a roof of a house:
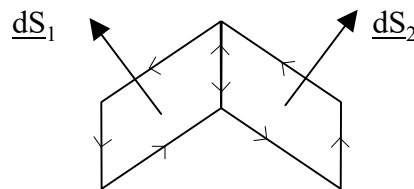

Figure 4

Each of the two planar surfaces has its own vector. Does it mean anything to add the vectors? Well, the resultant is vertically upwards (towards the sky, in the house analogy). If we viewed the roof from the sky, looking down, the scalar area we would see is given by Property 1, because the component of the resultant vector is just the sum of the vertical components of $\underline{dS}_1$ and $\underline{dS}_2$. What if we viewed the roof from the side, horizontally:
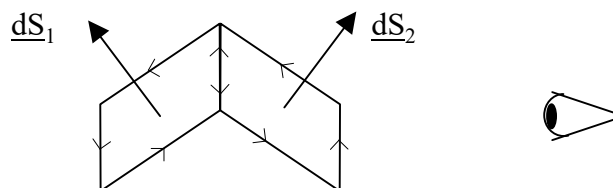

Figure 5

Then the component of the resultant vector in our viewing direction is zero. This still ties in with Property 1, as long as we regard the back face of the roof as making a *negative* contribution to the total area. Then the negative area of the back face exactly cancels the positive area of the front face, giving a total viewable scalar area of zero. Notice that the rim traversal arrows "cancel" along the common edge of the two planar surfaces. This means the resultant rim arrows simply go all around the overall rim, and the idea of the right-hand thumb rule still holds.

Now in principle, we can make up *any* surface, including curved surfaces, from a collection of infinitesimally small planar elements like the ones above. As with the roof in Figures 4 and 5, Property 1 continues to hold, as long as you count back faces as showing a negative scalar area. This makes Property 1 a very general and powerful property.

Here is a check to see if you understand the idea of back faces counting negative. What is the vector area of a sphere? It is zero, because regardless of what direction you view it from, you see a front face with a positive scalar area, but also a back face (on the other side of the sphere) which makes an equivalent negative contribution. In fact, more generally, the vector area of any *closed* surface is zero.

The idea that back faces count negatively to the overall scalar area we see when viewing the surface from some direction has an interesting implication. *The vector area of a surface depends only on its rim and not on the detailed nature of the surface*. Imagine we could take the roof in Figures 4 and 5 and pull the ridge (the common edge of the two planar surfaces) of the roof upwards towards the sky. The scalar area you see when viewing from vertically above or horizontally from the side would not change. In fact the scalar area you see would not change regardless of which direction you view the surface from. The only thing that matters is the rim, not the detailed nature of the surface.

How do we actually find the vector area of a complicated surface in practice? Here are two ways of doing it:

(a) Divide the surface up into planar elements. Find the vector for each element and then just add up all the vectors. A useful trick here is to note that the vector area of a parallelogram with edges defined by the vectors a and b is just a×b. Similarly, the vector area of a triangle with edges defined by the vectors a and b is just ½ a×b. Make sure you get the direction of the vector area to be consistent with the right hand thumb rule!

Figure 6

(b) View the surface from the $\underline{i}$, $\underline{j}$ and $\underline{k}$ directions and estimate the net scalar area you see (don't forget to count back faces as negative). Property 1 then says you've found the $\underline{i}$, $\underline{j}$ and $\underline{k}$ components of the vector area, which completely specifies the vector area itself.

After building all these foundations, can we do anything useful? Here is an example:

Imagine a flat circular loop of wire in a river, fully immersed in the river. What volume of water flows through the loop each second? We'll assume for simplicity that the water in the river is flowing steadily in a horizontal direction. First consider the extreme case where the normal to the plane of the loop is vertically upwards, in other words, the loop is edge on to the flow of the water. Then the volume of water flowing through the loop per second is zero. At the other extreme, consider the case where the normal to the loop is horizontal, in other words, the loop is presenting its maximum area to the flow of the water. Then the volume of water flowing through the loop per second is the area of the loop multiplied by the speed of the water. And what if the loop is tilted at some intermediate angle to the flow, between these two extremes? Then the volume of water flowing through the loop per second depends on the cosine of the angle of tilt. But this all means we can say:

$$\text{Volume of water flowing through the loop per second} \quad = \quad \underline{A}.\underline{v}$$

Where $\underline{A}$ is the vector area of the surface, and $\underline{v}$ is the velocity vector of the water. The dot product takes care of the cosine issue, and at the maximum, the volume flowing per second is just the magnitude of $\underline{A}$ (ie the scalar area of the loop) multiplied by the magnitude of $\underline{v}$ (ie the speed of the water). Furthermore, this formula works even if the loop is not flat and not circular.