Lecture on Solid Angle

Ben Kravitz

1 Introduction

The idea of a solid angle comes up a lot in remote sensing applications, so it's important that you know what it is. As such, I've made this lecture for you. This lecture was designed to be a take-home lecture, so the explanations are a bit more thorough than one would normally expect from a classroom lecture. It would be a good idea for you to read through all of it to get a better understanding. I've also put some homework problems in here, which I will collect. If there's anything in here you don't understand, feel free to ask.

2 Projections

In order to talk about solid angles, we first need to talk about projections. There are lots of different kinds of projections, some of which you already know. We should begin with a very simple definition of what a projection is, and then we'll get more complex as we go along. A **projection** is the transformation of points and lines in one plane onto another plane by connecting corresponding points on the two planes with parallel lines.[1] A better way to illustrate what this means is to look at figure 1. Basically, think of a thin sheet of paper with some points and lines drawn on it. Some of the lines may make shapes, and some may not. The sheet of paper is thin enough that you can see through it (like tracing paper). Now you hold it up to a very strong light, and you look at the shadow that the points and lines make on the floor or table below the light. That's a projection. If you like, you can think of a movie projector (as in when you go to the movie theatre). You take a square of film with some pictures on them, and you shine a very bright light behind the film. The light *projects* whatever is on the film onto the wall.

You will notice, both from figure 1 and from your own personal experience, that projections can distort the original picture. In figure 1, the triangle and line look different in the top plane from the bottom plane. When you go to a movie theatre, what you see on the screen is obviously much, much larger than what is on the film. This distortion can depend on the orientation of the plane the light is shining through, as well as how far away the light is, etc.

[1] Weisstein, Eric W. "Projection." From MathWorld–A Wolfram Web Resource. http://mathworld.wolfram.com/Projection.html

Now, this definition is actually a bit restrictive, since there are other kinds of projections. Really, we could say that we want to project from one surface to another. A surface is something that, locally (i.e. when you zoom in to a small enough area), looks flat. An example is the Earth. We know that the Earth is a sphere (or pretty close to it). However, if you just look out your window, the Earth looks pretty flat. This will actually work for any sphere, and in fact, a sphere is an example of a surface.

So let's think of examples of projecting a sphere onto a plane. Well, we know of one very easily - it's called a map. You see maps all the time. Those maps are just projections of the Earth (a sphere - the surface we're talking about) and a plane (the piece of paper on which the map is printed).

This is where it starts to get a little more tricky, since not all map projections are the same. The most common map projection that you see is called the Mercator projection. It makes all lines of longitude perfectly up-and-down, and it makes all lines of latitude perfectly left-to-right. However, this won't give you an accurate picture, because the real lines of longitude are not parallel. They actually all intersect at the poles. So what happens is that you get a lot of distortion toward the poles. Figure 2 shows an example of a Mercator projection. You can clearly see that Greenland appears to be much, much bigger than all of Africa, which we know is not correct. Similarly, Antarctica, despite being a very small continent, appears to be the largest land mass on the globe. For more information on Mercator projections, I encourage you to read *http://www.math.ubc.ca/~israel/m103/mercator/mercator.html*.

So let's take the idea of projection one step further. We know that we can project a plane onto a plane. We can also project a sphere onto a plane. Well, how about projecting a plane onto a sphere? So let's start with a cube inside a sphere, as in figure 3. When we project, as in the red lines, the faces of the cube get pushed out onto the surface of the sphere. You'll notice that the entire square gets bowed out, as do the edges. You'll also notice that the corners of each face used to be 90 degrees (since each face of a cube is a square), but after the projection, the corners aren't 90 degrees anymore. Figure 4 shows the same idea, but now subdivided. (If any of you want to be climate modelers in the future, you'll eventually learn why the subdivisions are important.)

More carefully described, the way this is done is the following (which you can actually see in figure 3): You take your plane, and you draw lines from its corners to the center of the sphere. Then you extend those lines until they intersect the sphere. Those are now the corners of your projection. An example I'm going to use is an icosahedron, which is a polyhedron made of 20 triangles. You can see what it looks like in figure 5. Then we're going to project one of the faces (which is a triangle) onto a sphere, as in figure 6. Looking at figure 6, we draw lines from the center of the sphere through the corners of the face (these lines are yellow). Then we extend these lines until they intersect the sphere. Once they do that, these become the corners of the projection on the sphere. You can see this in figure 6: the face of the icosahedron is made of white lines, and the projection onto the sphere is in purple. (You can ignore the part on the right of figure 6. It's just for fun.)

There's one more combination we haven't covered, and that is projecting an arbitrary surface onto a sphere. However, it's really the same idea. You take your surface, and you draw lines from the center of the sphere to the edges of your surface. Above, we said that you could draw it to the corners of your surface. However, that won't always work, since your surface might not have corners. For example, you might be projecting a circle onto the sphere. In cases like that, you can think about drawing as many lines as necessary to get a clear idea of what the shape of the projection will be.

So why are we going through all of this? Well, as it turns out, the idea of a solid angle depends very heavily on this last idea of projecting onto a sphere. So, if you understood everything I just talked about, you're well on your way to understanding solid angles.

3 Solid Angle

The **solid angle** Ω subtended by a surface S is defined as the surface area Ω of a unit sphere covered by the surface's projection onto the sphere.[2] That's a very complicated way of saying the following: You take a surface. Then you project it onto a unit sphere (a sphere of radius 1). Then you calculate the surface area of your projection. That's it.

Solid angle is, effectively, a measure of how big an object looks to an observer. For example, let's say you go outside on a sunny day. To shield the sun from your eyes, you might hold up your hand. Obviously, the sun is much bigger than your hand. However, it is far enough away that the sun takes up a smaller portion of the sky than your hand. So we would say that the sun *subtends a smaller solid angle* than your hand. (You can think of the sky as part of a sphere which completely surrounds you. The ground below you takes up half of the sphere. Buildings or other things that obstruct the horizon take up even more of the sphere.)

Before we get into why solid angle is useful (which I will do in the next section), we should establish a little bit more background. So let's go back to our definition of solid angle and try to calculate some solid angles by projecting some surfaces onto the unit sphere. Let's start with the easiest one: projecting a sphere onto the unit sphere. The way you can think of this is to take the unit sphere and put a sphere of smaller radius (say 0.5) inside of it. The projection is simply blowing up the sphere (doubling the radius). Or we can take a sphere that completely surrounds the unit sphere (say of radius 2). The projection is shrinking the sphere (cutting the radius in half). So when we project a sphere onto the unit sphere, we get the unit sphere. Now, according to the definition of solid angle, we need to calculate the surface area of the unit sphere. Well, we know the formula for surface area of a sphere: $4\pi r^2$. Since this is the unit sphere, $r = 1$, so **the solid angle subtended by a whole sphere is** 4π. The

[2]Weisstein, Eric W. "Solid Angle." From MathWorld–A Wolfram Web Resource. http://mathworld.wolfram.com/SolidAngle.html

units of solid angle are called **steradians**, abbreviated *sr*. So we would say that a sphere subtends a solid angle of 4π *sr*.

As I'm sure you can guess, calculating surface area on a sphere can get a bit tricky. However, we don't always need to do it. For example, let's calculate the solid angle that is subtended by the face of a cube, as in figure 3. We know that a cube has 6 sides, all of them equal in size. We also know that a cube, if projected onto the sphere, will cover the entire sphere, as in figure 4. So if we project the cube onto the sphere, the surface area covered by one of the faces will be 1/6 of the total surface area. Then the solid angle subtended by one face of the cube will be 1/6 of the solid angle subtended by the entire sphere. So that makes the math really easy. The solid angle subtended by one face of the cube is

$$\frac{1}{6} \cdot 4\pi = \frac{4\pi}{6} = \frac{2\pi}{3}$$

Now let's put forth a more mathematically rigorous way of calculating solid angle. The easiest way to define a way to calculate solid angle is to use spherical coordinates.[3] So the definition of solid angle is

$$\Omega = \int \int_S \sin(\phi)\, d\theta\, d\phi$$

where S denotes an integral over the surface (the projection of your object onto the sphere), ϕ is the colatitude, and θ is the longitude. Dr. Wilkin used different terminology in his lecture (specifically, right ascension and declination), but it's basically the same idea. For a more clear definition of longitude and colatitude, refer to figure 7. Note that colatitude is 0 at the North Pole and π at the South Pole. (Yes, we use radians. Sorry if you don't like them, but they are important.)

This seems like a complicated equation, so let's talk about how you use it practically. Really, it's just the same as any integral in two dimensions. (All of you have taken calculus 3 and remember it all, right? Of course you do....) So, effectively, the integral is all set up for us, if we can figure out the boundaries for S. Unfortunately, this is the hard part. Since it's so difficult, I'm going to briefly go over what might be the most important solid angle calculation with you: a cone. (I'll talk about why it's important in the next section.) I don't want to get too bogged down in the math, so this is going to be a rather cursory calculation. If you have further questions on it, I'd be happy to go over them with you.

Figure 8 shows a cone embedded inside the unit sphere. You can see that its projection onto the sphere is a bowed-out circle. Let's orient this so the cone opens directly upward. Then the circle lies perfectly on a line of latitude. This makes our calculation much easier, since that means we can integrate longitude (θ) from 0 to 2π without having to worry about any details. Let's say the angle

[3] Note that there is another way to define it, using a surface integral and the divergence theorem, which is most often the way you'll see solid angle defined. However, I don't think that way of defining it is particularly instructive, and it makes calculations very difficult.

at which the cone opens is ζ. (ζ is fixed, since the cone doesn't change.) As it happens, if you draw out the picture of what you're doing, this is also the colatitude ϕ. So we can integrate ϕ from 0 to ζ. So our integral becomes

$$\Omega = \int_0^{\zeta} \int_0^{2\pi} \sin(\phi)\, d\theta\, d\phi$$

When we integrate this (I'll skip the details), we get $\Omega = 2\pi(1 - \cos(\zeta))$.

4 Why we care about solid angles

Let's use an example that a lot of people are familiar with: weather radar. Dr. Miller is going to talk in great detail about weather radar later, but I can at least use it as a brief example. You've often seen radar weather figures that show a line sweeping around in a big circle, as in figure 9. This is, effectively, sending out an electromagnetic signal and measuring how much of that signal comes back. (As you learned, this is called an *active* sensor.) As I'm sure you can guess, there are limits to weather radar. For example, it cannot tell you about precipitation happening 2000 miles away, i.e. the signal is attenuated (reduced) with distance from the radar device. Also, the sensor only has so much power available to it (for various reasons, including both electrical costs and the amount of noise that is generated with increased amounts of power). Therefore, it makes sense to concentrate all of that power in a narrow beam as opposed to spreading it in all directions. Put in terms of what we've been learning, the amount of signal that can be sent out in any one direction is attenuated with increased solid angle.

So now we can get more specific with our definitions. There is a standard language for describing energy flux.

1. The **radiant flux** Φ is the rate at which energy is transported toward or away from a surface. It has units of Watts (W). For example, total radiant flux emitted by the sun is $\Phi_s = 3.9 \times 10^{26}$ W. Radiant flux is called **power** in radar meteorology.

2. The **radiant intensity** $I = d\Phi/d\Omega$ is the radiant flux per unit solid angle. It has units of Watts per steradian (W sr^{-1}). It is used in the description of radiation propagating from a point source.

3. The **flux density** $E = d\Phi/dA$ has units of Watts per square meter (W m^{-2}). It is the radiant flux per unit area that is either incident upon or emitted from a unit surface area A.

4. The **radiance** L has units of Watts per square meter per steradian (W m^{-2} sr^{-1}). It is the radiant flux propagating toward or away from a surface in a specified direction with solid angle $d\Omega$. The flux is emitted from or incident upon a differential unit area dA inclined at an angle θ to

the direction of energy propagation and is written

$$L = \frac{d^2\Phi}{d\Omega dA\cos\theta}$$

So what do all of these mean? Well, let's take the sun as an example. The sun is putting out a lot of energy all the time. Photons are streaming forth from it continually. As I'm sure you all remember (or do now that I've reminded you), energy is measured in Joules. Now, since the sun is continually putting out energy, it doesn't make any sense to talk about how many joules the sun has put out or will put out. Instead, a more useful measurement is how many Joules per second the sun puts out. Joules per second is also known as Watts and is a measure of power. (Atmospheric scientists call this radiant flux.)

Now we know that the sun is putting out so many Watts of power. However, it's putting that much power out in all directions. Well, the Earth isn't in all directions. It's only in one direction. When viewed from the sun, the Earth only subtends a very small solid angle. Therefore, if we want to know how intense solar radiation is in any direction (for example, as received by the Earth), we need to divide by the solid angle, giving us radiant intensity.

So that describes the energy from the sun's perspective, but now we should look at it from the Earth's perspective. If we take a beam of energy from the sun, we know how intense it is, but that doesn't tell the whole story. So let's take the energy that is being emitted from the sun and calculate quantities for it at various distances away from the sun. If we stop the radiation very close to the sun (for example, as far away from the sun as Mercury), the radiation hasn't dispersed over a very wide area. (Remember that area of a sphere equals $4\pi R^2$, so the smaller the radius of the sphere, the smaller the area of the sphere.) If we stop the radiation very far from the sun (for example, as far away as Pluto), the area covered is much larger, so that same power is now spread over a very large area. Thus, we have come up with an important quantity that measures the amount of power per unit area. This is in units of Watts per square meter and is called either flux density or *irradiance*. Note that the amount of irradiance decreases as we move farther from the source, but the amount of radiant intensity does not.

Before we talk about radiance, we should talk about cones and why they are important to solid angles. As an experiment, take a flashlight and point it straight downward at a table. The light is coming out of the flashlight in a cone from a single point (the lightbulb). The lightbulb puts out light in all directions, but the flashlight (a cylinder) only allows the light to travel in certain directions. The collection of all of those directions forms a cone. You can get a better idea of this in figure 10. You'll notice that when you point the flashlight straight downward (at a colatitude of 0), the spot is a bright, perfect circle. Now take the flashlight and start moving the beam across the table (increasing the colatitude). You'll notice that the spot gets dimmer, which happens for the reasons we discussed in the previous part on irradiance. However, you'll also notice that the spot starts to elongate. Thus, the area covered by the beam has increased, but the solid angle taken up by the beam has not. If you imagine a remote sensor in this fashion, as you change the viewing angle (colatitude) of the remote sensor, it can view a larger area for the same solid angle, but what it can view is not as strong of a signal.

And now we get to radiance. Another definition of radiance is the amount of power emitted by a surface that will be received by a sensor looking at the surface from some angle of view (which we have called θ above). Basically, take a surface with surface area ΔA. This surface is emitting power in the amount of $\Delta\Phi$. Now you have your sensor which is off somewhere. Your sensor sees the surface as subtending a certain solid angle $\Delta\Omega$ of its field of view. You can think of a camera with a fixed aperture (if you like that example). Also, your sensor may not be directly overhead of the surface. Your sensor could be off to the side. That means the cone that makes the solid angle might not be pointed directly upward. It could be angled away from the vertical. We call that angle θ. (If the cone is pointed directly upward, $\theta = 0$.) You can see this illustrated in figure 11. So we collect all of these terms and, using a bit of geometry and guided by what we want to calculate, we get the quantity of *radiance*, which is given by the formula

$$L = \frac{\Delta\Phi}{\Delta\Omega\Delta A\cos\theta}$$

given in units Watts per square meter per steradian.

Now that we've established some definitions, let's do a quick calculation. Let's say that we have a *Lambert surface*. That means that incident radiation (coming in at any angle) is emitted equally in every direction. We wish to calculate the irradiance E that is being emitted from the Lambert surface. We know that, by definition $E = d\Phi/dA$. From the formula for L, we can substitute in to get $E = L\cos\theta d\Omega$. Also by definition of solid angle, $d\Omega = \sin\theta d\theta d\phi$. So $E = Ld\phi\cos\theta\sin\theta d\theta$. Then we integrate over the entire sphere to get

$$E = L\int_0^{2\pi} d\phi \int_0^{\pi/2} \cos\theta\sin\theta \; d\theta$$

Performing the integral, we get $E = \pi L$. Since L is now independent of angle, L, and hence E, is constant. This means that for a Lambert surface, emitted irradiance is independent of direction.

5 Conclusion

So why do we care about all of these quantities? Well, the reason we care is that radiance is the variable directly measured by remote sensing instruments. You can think of radiance as the amount of light (or other electromagnetic radiation) that the instrument "sees" from the object it is observing. The atmosphere can emit radiation, resulting in an increased amount of radiance measured by the sensor. It can also absorb radiation, resulting in a decreased amount of radiance measured by the sensor.

Since you're going to be working with remote sensors throughout the entire semester, it's important that you know what radiance is. And you can't really understand radiance without understanding solid angle. So that's it for my lecture! I hope you understood everything, and if you didn't please feel free to ask.

6 Homework

This homework will be due one week from the day it is assigned. Since I'm handing this lecture to you on September 10, 2009, the homework will be due on September 17, 2009 at the beginning of class. In any problem where I have asked you to perform a computation, show all of your work. Your work is much more important than the answer you get.

1. List one map projection that is used, other than a Mercator projection. (Note: Do NOT choose another cylindrical projection. For a definition of what that means, see the website below.) A fun example of a projection you can choose can be seen in figure 12. Describe what the projection you choose is called and how it looks. (You should also submit pictures. This is why Rutgers gives you a printing allotment. Greyscale pictures are fine.) Why does it look the way it does, and why is it different from the Mercator projection? (Note that I said WHY are they different from each other, and not HOW. This means that you need to actually talk about the method of projection that is used and how using that method distorts the final picture.) Discuss advantages and disadvantages to both methods. As a starting point, you may wish to look at this website: *http://www.colorado.edu/geography/gcraft/notes/mapproj/mapproj_f.html*. Please remember that if you use any reference sources, you MUST cite them.

2. Now consider a satellite orbiting the Earth. You may choose which type of orbit you wish to consider, but you must indicate which one. Be specific. Take both of your maps that you used in problem 1 and trace out the ground track of the satellite as it completes an orbit. Be sure to indicate which direction the satellite is moving on this path. (Note that this is also a projection onto a sphere. It's just a traveling projection of a point instead of a line or surface.)

3. Calculate the solid angle subtended by a cone whose circular base has a radius of 0.5.

4. Calculate the solid angle subtended by a cone whose circular base has a radius of 1. What does this look like? Describe (in enough detail that I understand you know what you're talking about) another way of obtaining this quantity.

5. Calculate the solid angle subtended by one face of an icosahedron. (See figure 5.)

6. According to http://www.satobs.org/seesat/Jun-2001/0063.html, the sun is seen from Earth as subtending a solid angle of 6×10^{-5} *sr*. Calculate what percent of the sky it occupies. (Remember that the sky, as you see it, is only half of the total sphere.)

7. I stated earlier that irradiance decreases as you move farther away from the source. Let's say you have an object that is a distance d from the source of radiation. Calculate how much less the irradiance is for an object that is a distance $3d$ from the source. Note: you do not need to know what d is, how strong the radiation from the source is, or any other information than what I have given you to do this problem. Work out a mathematical relationship between E and d. (An example of such a mathematical relationship would be saying that E is proportional to d^3. I'll tell you right now that this is not the correct answer, but it is an example of what I'm looking for.)

8. Calculate the solar constant, i.e. the irradiance that Earth receives from the sun. Hint: Given that the sun puts out a certain amount of power (which I have given to you earlier in the lecture), think about the area over which that power is distributed. You will need to know that the distance between the sun and the Earth is 1.496×10^{11} meters.

7 Figures

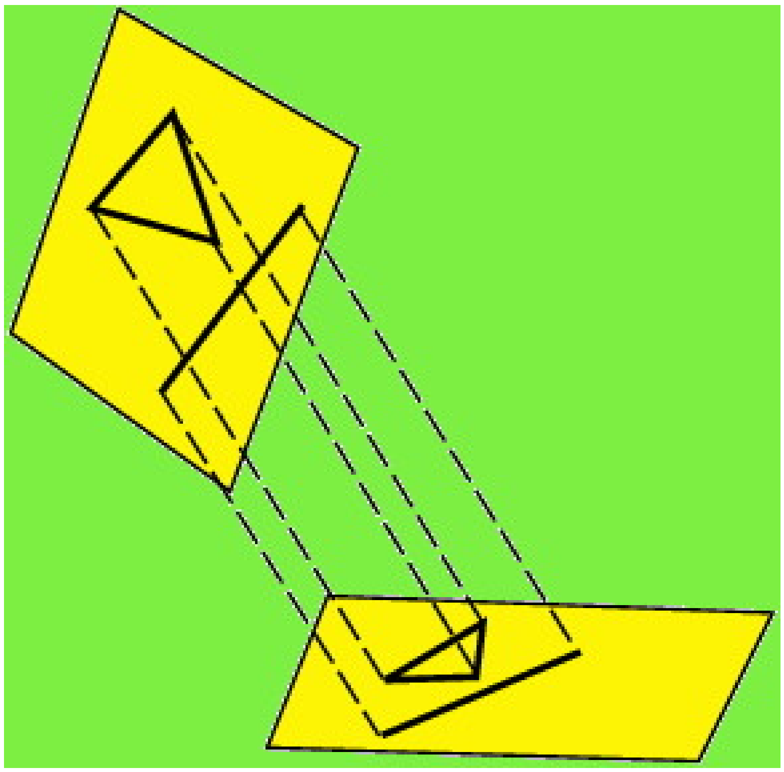

Figure 1: An example of how projection works. The triangle and line are being *projected* from the plane on the top to the plane on the bottom. From Weisstein, Eric W. “Projection.” From MathWorld–A Wolfram Web Resource. http://mathworld.wolfram.com/Projection.html

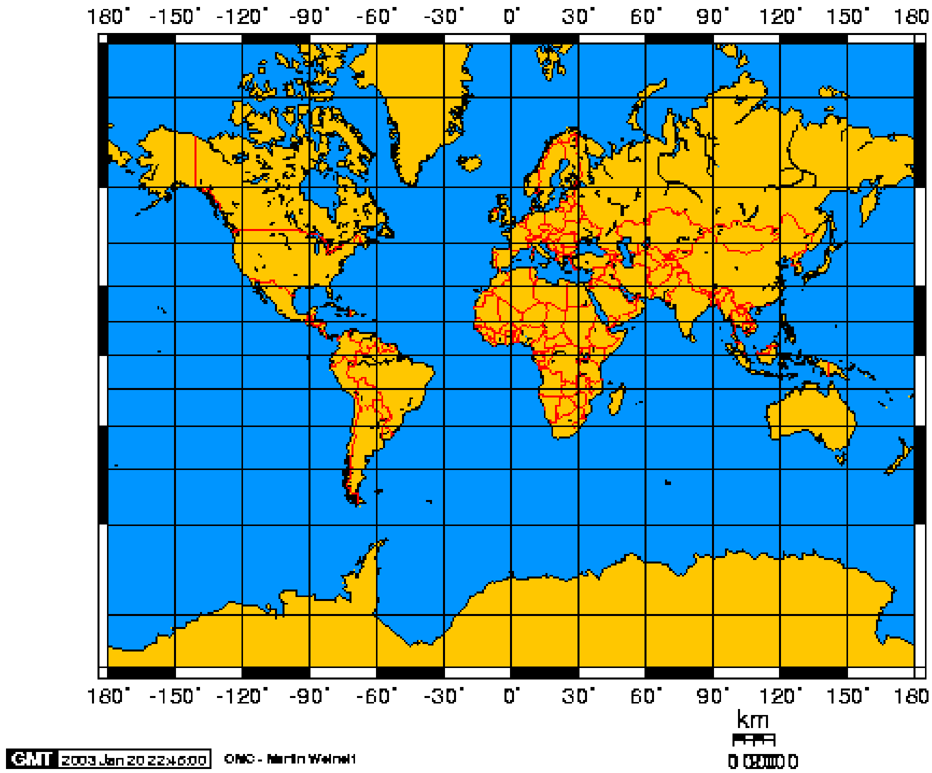

Figure 2: An example of a Mercator projection of the globe. From http://www.math.ubc.ca/~israel/m103/mercator/mercator.html

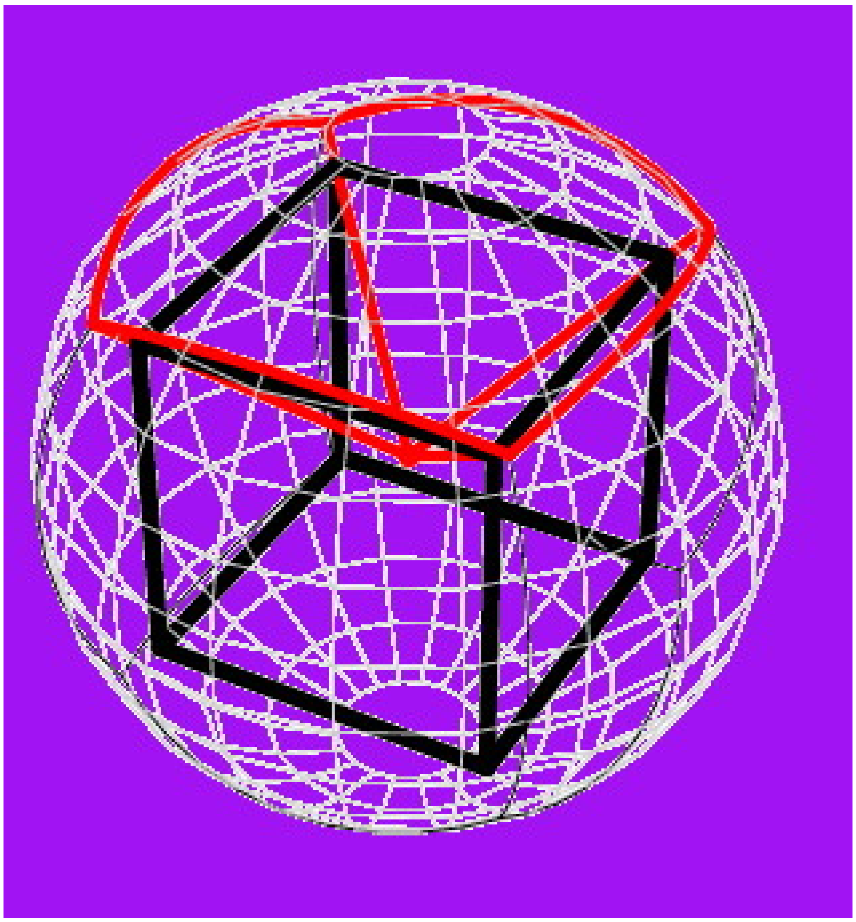

Figure 3: A cube inside a sphere, being projected onto the sphere. From Weisstein, Eric W. “Solid Angle.” From MathWorld–A Wolfram Web Resource. http://mathworld.wolfram.com/SolidAngle.html

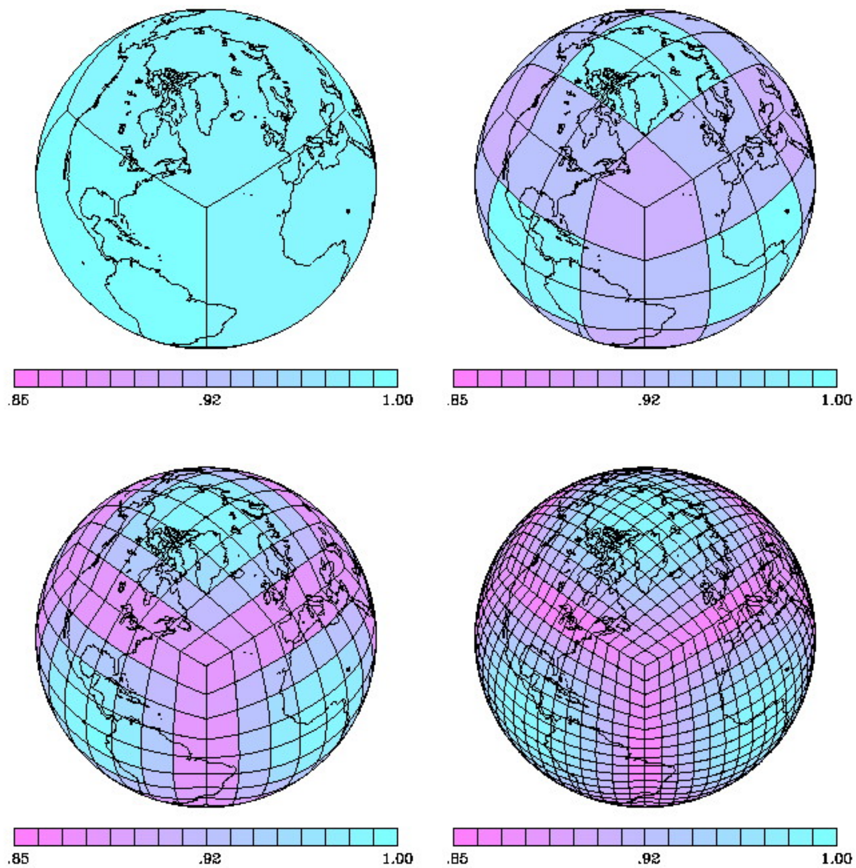

Figure 4: A cubed sphere (the result of projection of a cube onto a sphere). From http://www.cgd.ucar.edu/gds/wanghj/taylor/sphere2.gif

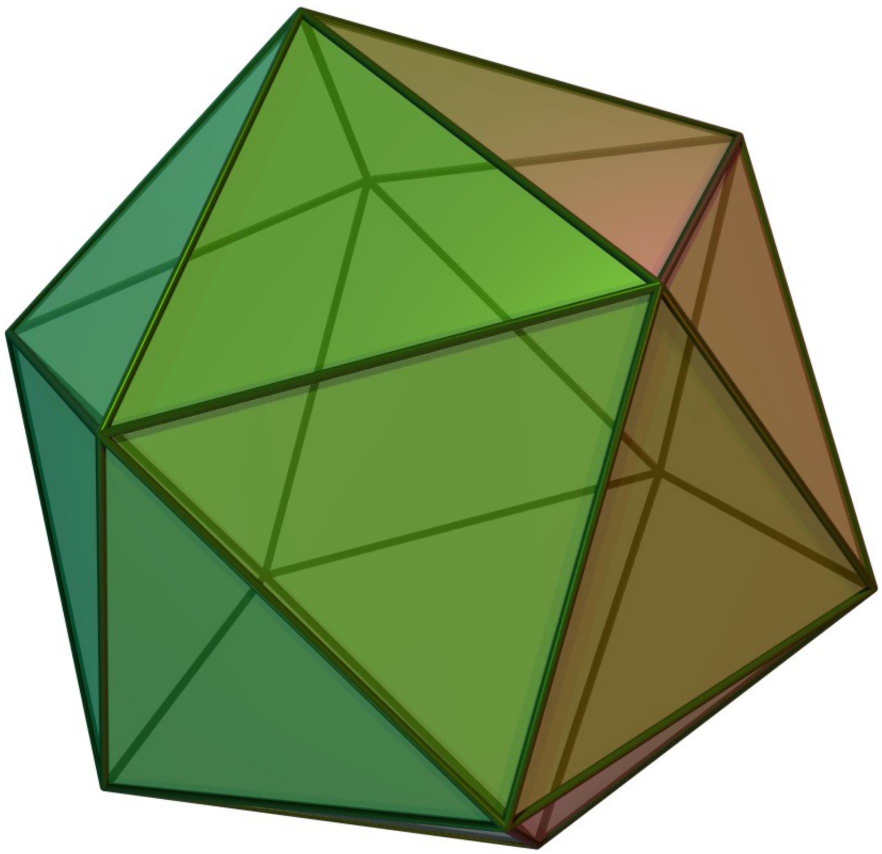

Figure 5: An icosahedron. From
http://upload.wikimedia.org/wikipedia/commons/e/eb/Icosahedron.jpg

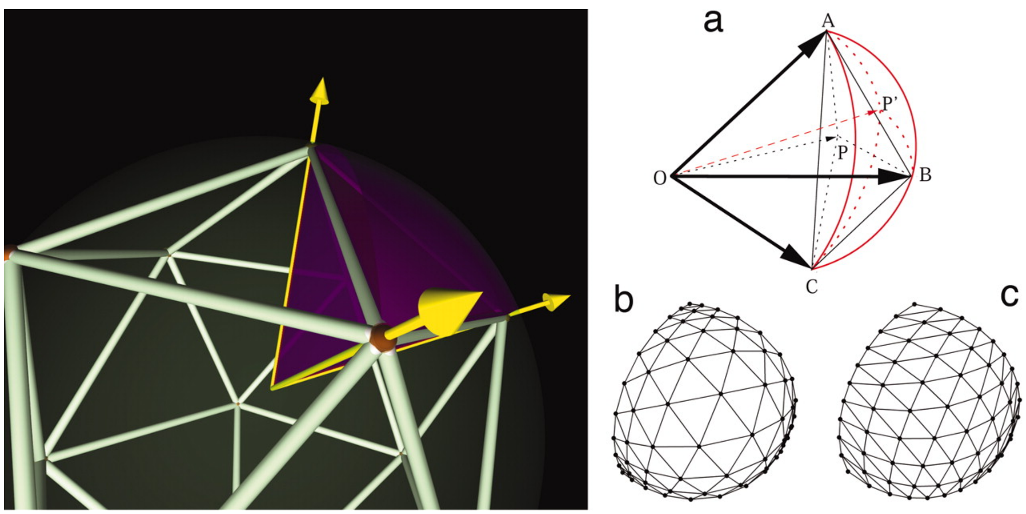

Figure 6: An icosahedron face (a triangle) projected onto a sphere. From http://www.pnas.org/content/104/47/18382/F3.large.jpg

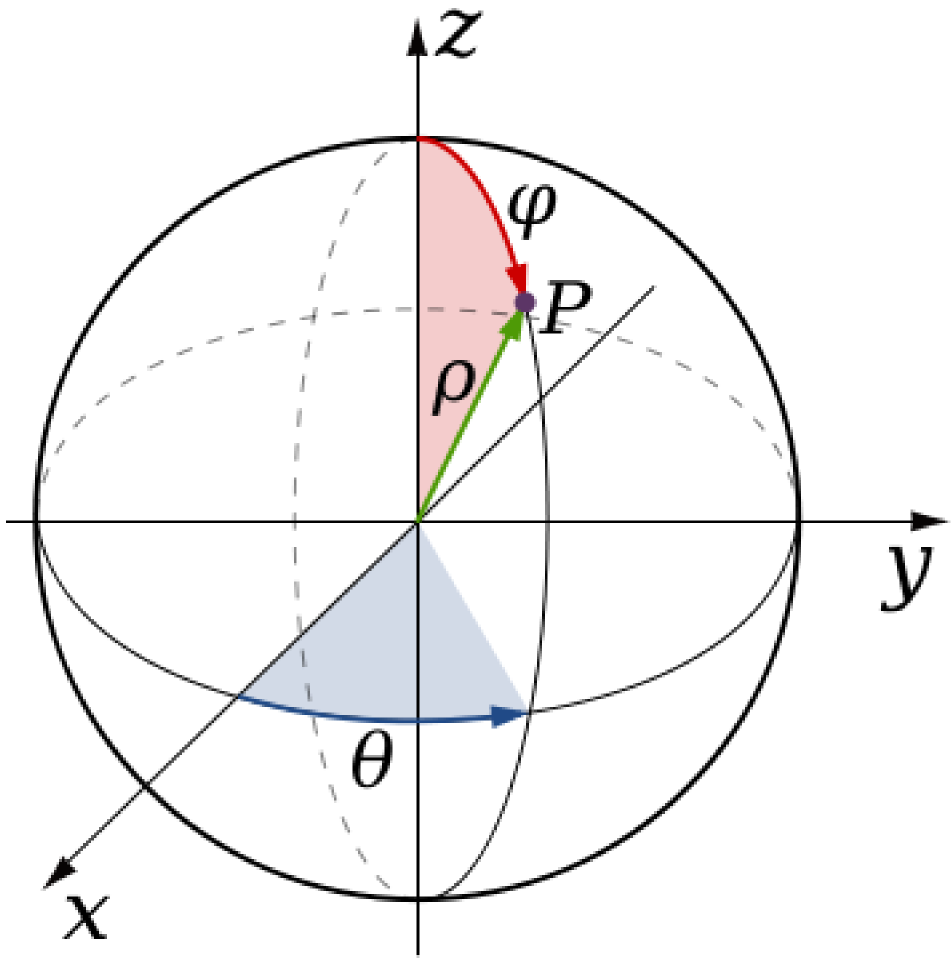

Figure 7: Showing definitions for spherical coordinates in more detail. From http://upload.wikimedia.org/wikipedia/commons/thumb/5/51/Spherical_Coordinates_(Colatitude,_Longitude).svg/360px-Spherical_Coordinates_(Colatitude,_Longitude).svg.png

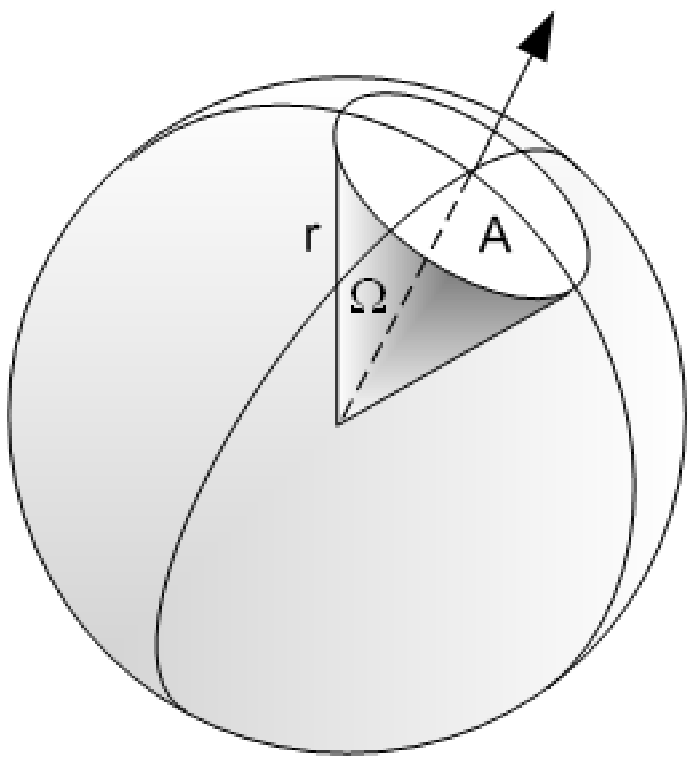

Figure 8: The projection of a cone onto a sphere. From
http://upload.wikimedia.org/wikipedia/commons/3/3e/Solid_Angle.png

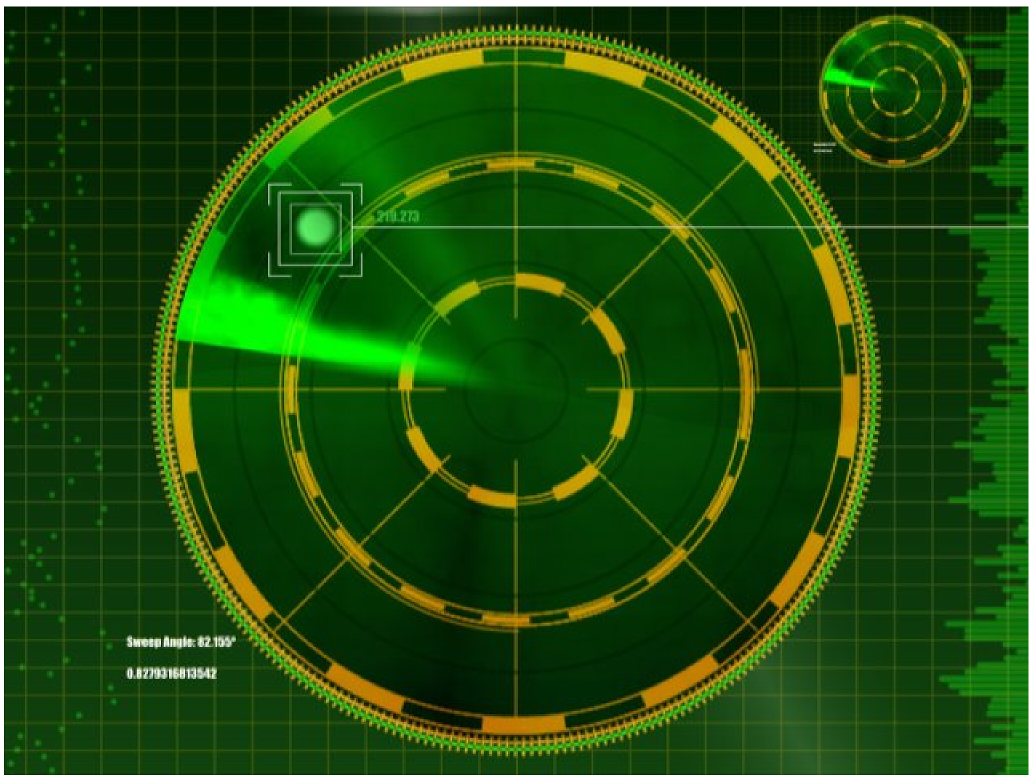

Figure 9: A sample radar screen. From
http://www.aefreemart.com/uploaded_images/radarScreen-719485.jpg

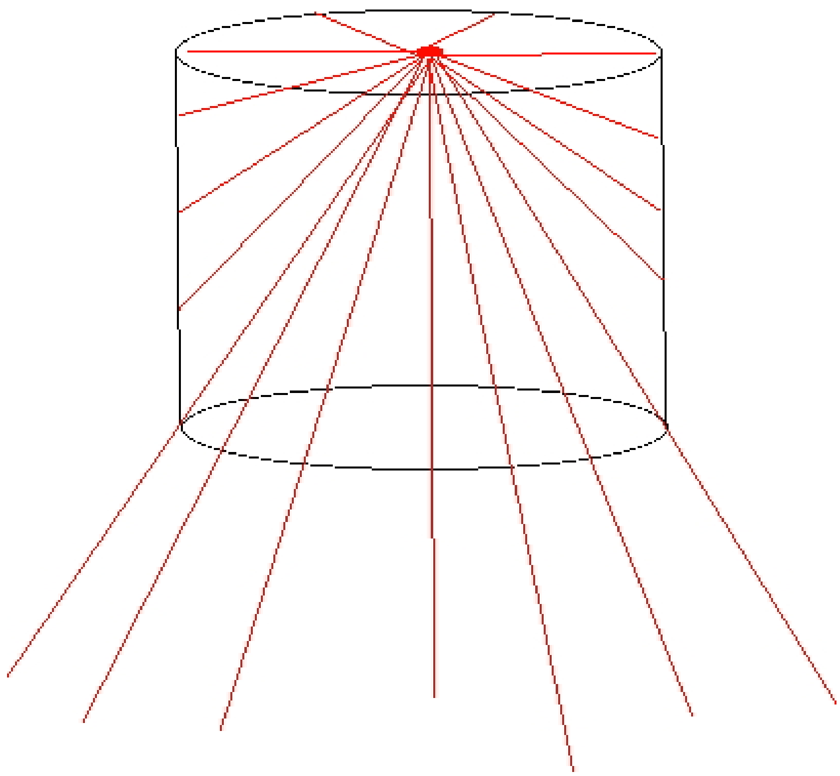

Figure 10: Solid angle actually describes how rays are allowed to be oriented. Note that the result is a cone.

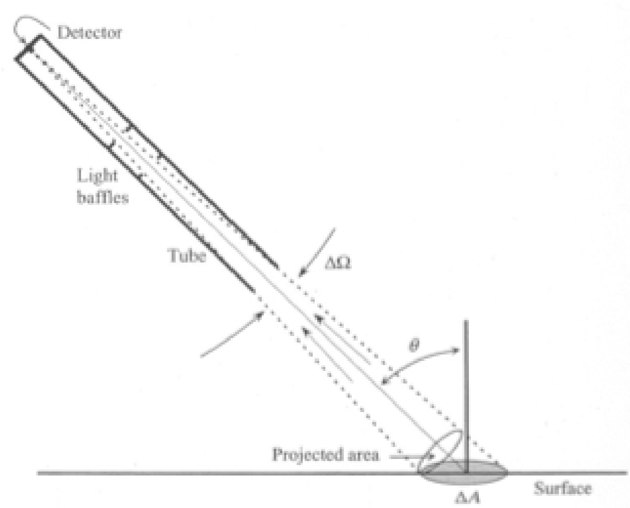

Figure 3.8. Schematic of a radiance meter viewing the surface (Radiance meter adapted from Figure 5.6 of Kirk, 1996).

Figure 11: A radiance meter whose field of view subtends a certain angle and has a certain orientation.

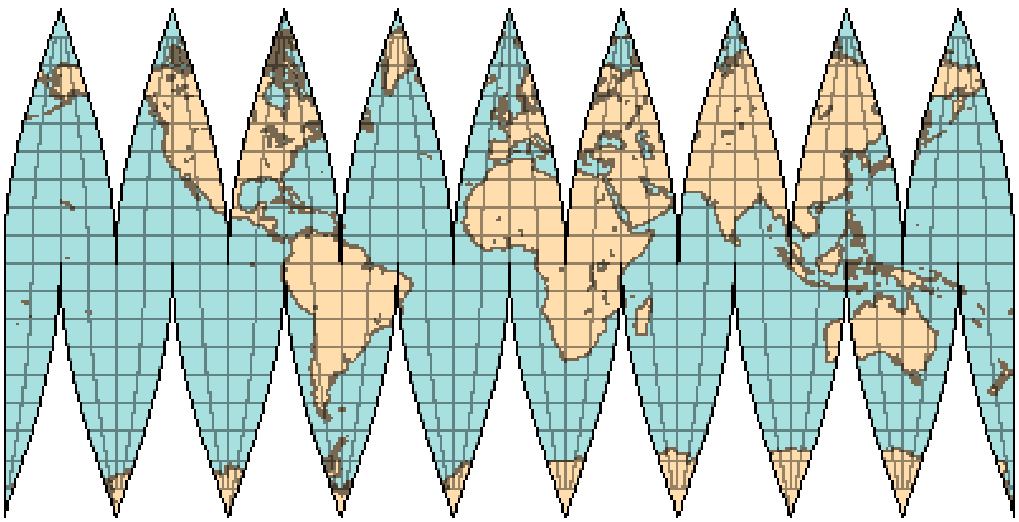

Figure 12: An example of a map projection.
From http://www.cs.ucla.edu/ kohler/z/icsilogo/isf-n9.png