

Weather Observation and Analysis

John Nielsen-Gammon

Course Notes

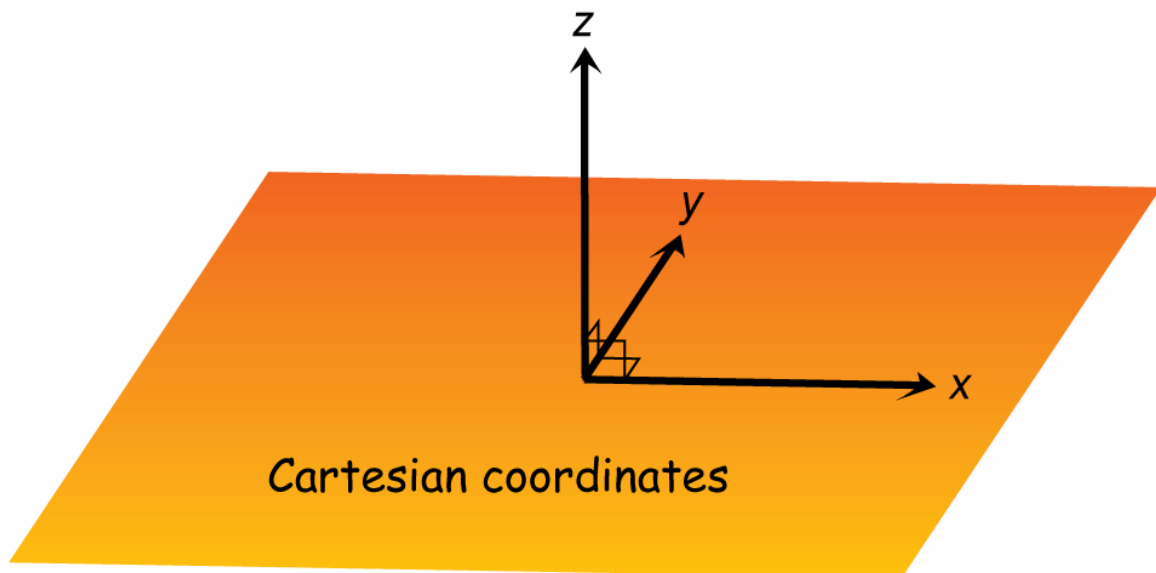
These course notes are copyrighted. If you are presently registered for ATMO 251 at Texas A&M University, permission is hereby granted to download and print these course notes for your personal use. If you are not registered for ATMO 251, you may view these course notes, but you may not download or print them without the permission of the author. Redistribution of these course notes, whether done freely or for profit, is explicitly prohibited without the written permission of the author.

Chapter 8. COORDINATE SYSTEMS

8.1 Conventional Coordinate Systems

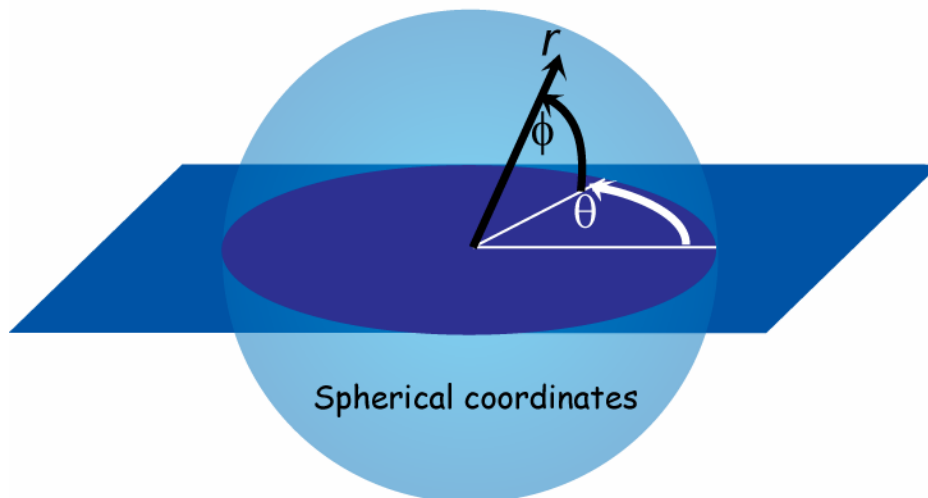
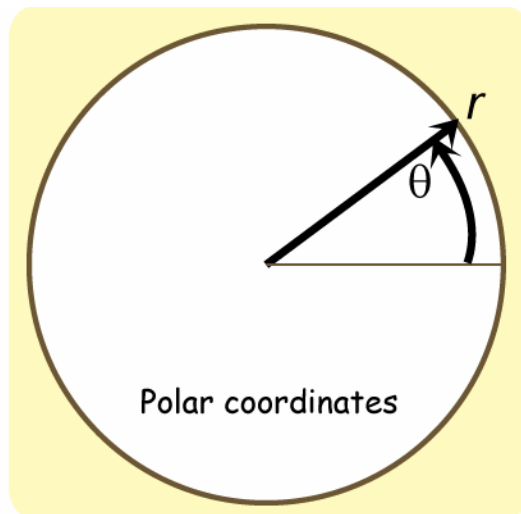
In all of your math classes up to this point, you have worked mostly in what are called orthogonal coordinate systems. Orthogonal here means that each coordinate axis is perpendicular to every other coordinate axis.

The simplest such example is Cartesian coordinates. In three-dimensional Cartesian coordinates, the three coordinate axes perfectly

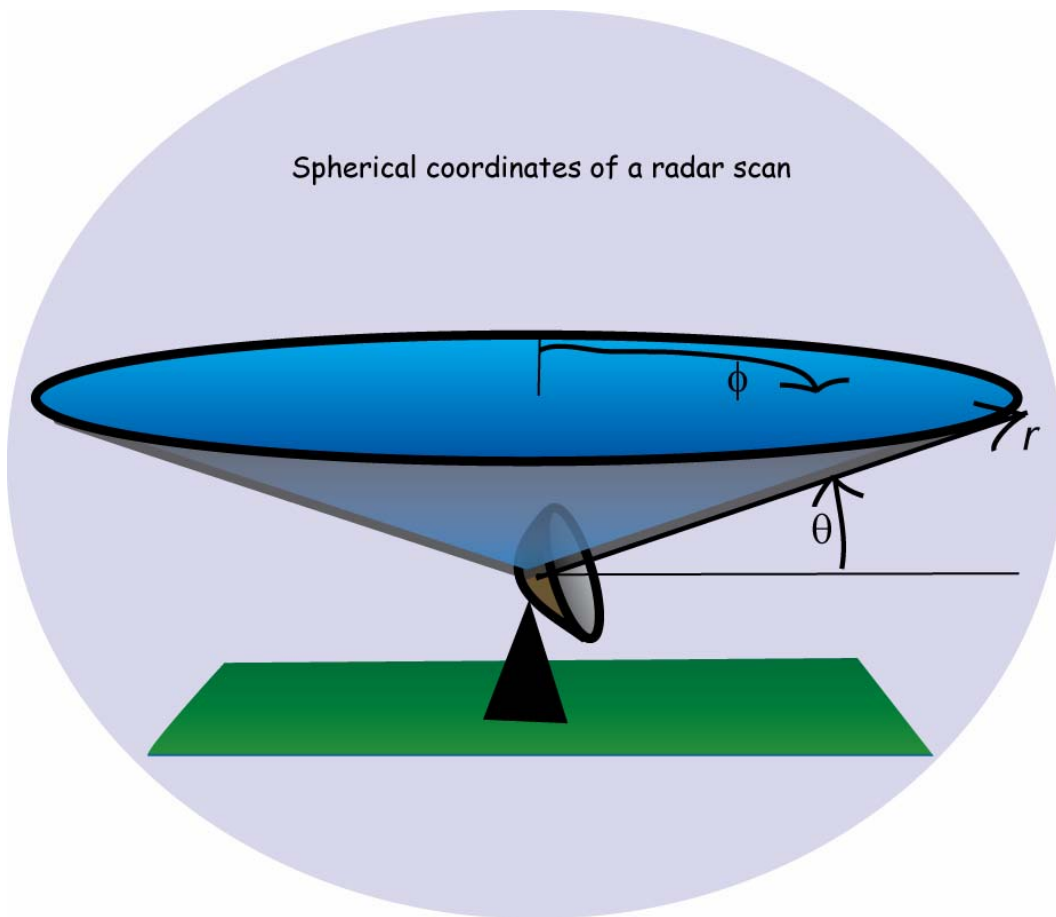
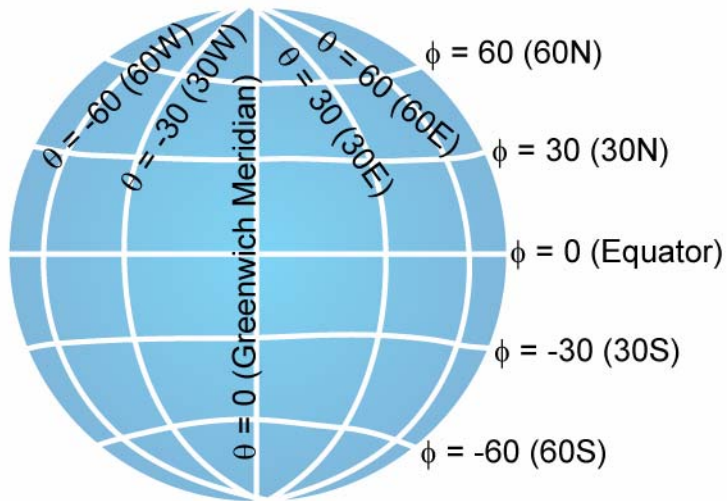


straight and exactly perpendicular to each other. Each axis extends to infinity. The horizontal coordinates are labeled x and y , and the vertical coordinate is labeled z .

The second most common coordinate system is polar coordinates (in two dimensions) or spherical coordinates (in three dimensions). These coordinates are orthogonal everywhere except at the origin, where directions are undefined. In polar or spherical coordinates, only the radial axis is straight. Unlike Cartesian coordinates, where the orientation of the axes is the same throughout the known universe, polar and spherical coordinates are oriented differently at different locations.



The most common use of spherical coordinates in meteorology is with the Earth itself. Longitude corresponds to the azimuthal angle in spherical coordinates, with the origin defined arbitrarily as the Greenwich Meridian. Latitude corresponds to the elevation angle, with the origin defined as the Equator. Longitudes West are equivalent to negative azimuthal angles.



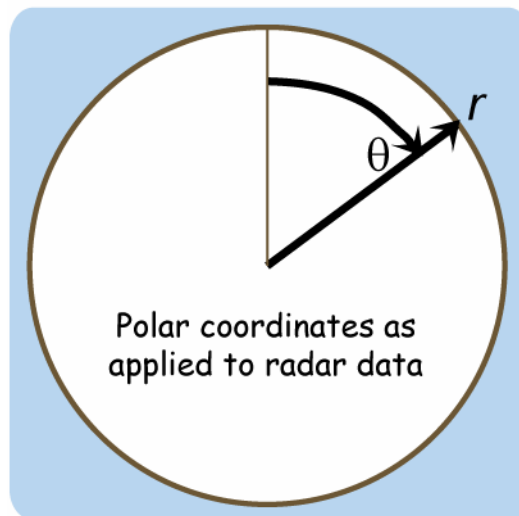
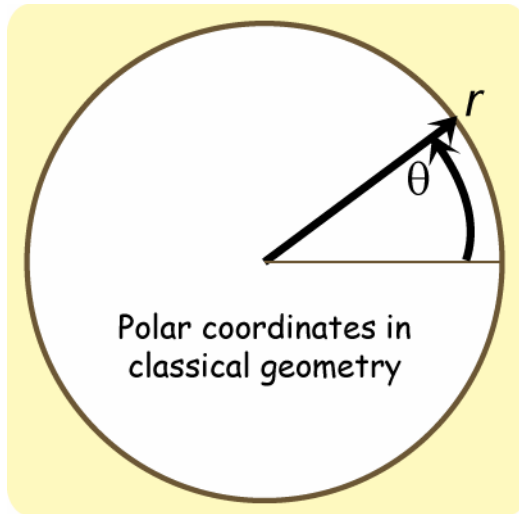
The next most common use of spherical coordinates in meteorology is with radar data. The coordinate system origin is defined as the radar transmitter location, the elevation angle is zero along a direction parallel to the horizontal, and the azimuthal angle is defined to be similar to compass headings, with a zero azimuth corresponding to a beam pointing toward the north, a 90 degree azimuth corresponding to a beam pointing toward the east, etc.

The two different applications of spherical coordinates, Earth and radar, illustrate the arbitrary nature of coordinate systems in general. When applied to the Earth, the center of the Earth is the origin of the coordinate system. When applied to radar data, the radar location is the origin of the coordinate system. In practice, one can define the origin of a coordinate system to be anywhere.

Another important distinction between Earth spherical coordinates and radar spherical coordinates is the orientation of the coordinate system. In Earth coordinates, the “up” direction (the direction where the elevation angle/latitude is 90 degrees) is the direction of the North Pole. In radar coordinates, the “up” direction is up, directly outward from the center of the Earth.

The final distinction between the two coordinate systems is subtler, but it is the most important of all. Looking down from the “up” direction, the azimuth increases in the counterclockwise direction in Earth coordinates and in the clockwise direction in radar coordinates. The Earth coordinate convention is the one that corresponds to the normal mathematical definition of polar or spherical coordinates, the one you’re probably familiar with. In polar coordinates, the azimuth is normally defined as zero in what in Cartesian coordinates would be x direction, and increases counterclockwise. But in meteorology, we speak of directions relative to North, like on a compass, with zero toward the y direction and the azimuth angle increasing clockwise. This may take you quite a bit of getting used to.

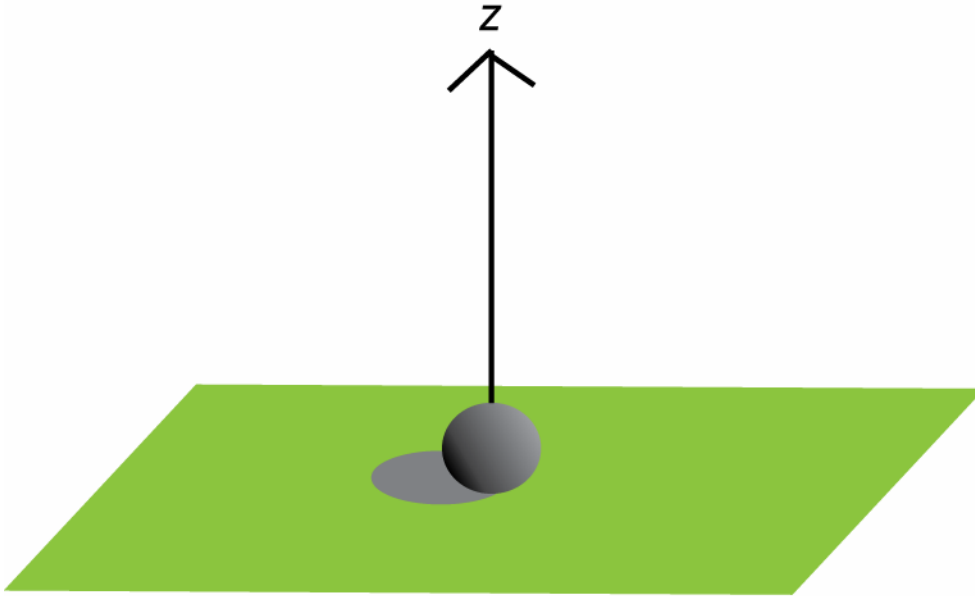
Besides that specific, important information, the general lesson is that even the direction in which particular coordinates increase is arbitrary. Indeed, there are an infinite variety of possible coordinates, and the only rule is that each coordinate value be unique. They need not even be orthogonal, although orthogonal coordinate systems should be used whenever possible because the math is easier. Sometimes, such as a few pages from now, we will even encounter coordinate systems that are *almost* orthogonal, so that mathematical manipulations can be performed as though the coordinates are orthogonal, but they aren’t really.



8.2 Earthbound Coordinates

On Earth, spherical coordinates are the most natural coordinates, but most of what happens meteorologically happens over a relatively small segment of the atmosphere, small enough that Cartesian coordinates work well. Because the math is simplest in Cartesian or almost Cartesian coordinates, we'll spend almost all our time working in local (almost) Cartesian coordinates. In fact, we have already done so, by defining x to be toward the east, y to be toward the north, and z to be up.

On Earth, the orientation of z is commonly defined as the direction opposite to a "flat" plane. If you have a table that's perfectly flat rather than sloping, z is straight upward from that table. You can tell if something is flat and level by placing a ball on it. If the ball doesn't roll off, it's flat. In physical terms, you might say that the gravitational force is exactly perpendicular to the surface of the table; if it wasn't, gravity would cause the ball to roll off.



The z direction is defined as straight up from a perfectly flat and level surface.

You might say that, but it wouldn't quite be right. The correct statement is that the overall sum of the real and apparent body forces is exactly perpendicular to the level surface. In other words, every force acting on the ball, when added together, is pulling it straight down. Gravity is by far the strongest such force, and if the Earth wasn't rotating on its axis, it would be the only force. But the thing you feel, the one you normally call "gravity", is actually the sum of two forces: the real gravity and the centrifugal force.

The centrifugal force, you may remember from physics, is an apparent force that, in a rotating reference frame, pulls objects outward, away from the axis of rotation. That happens on Earth, because the Earth is rotating. If you were standing on the Equator, and the gravitational pull of the Earth were suddenly turned off, you'd spin right off the Earth and fly off into outer space. (Please don't try this at home.)

This outward pull at the Equator, while much weaker than gravity, is still strong enough to cause the Earth, including the oceans, to bulge outward at the Equator by about 20 km compared to the poles, about 1/3 of 1% of its total radius. The next time someone tries to tell you that the world isn't flat, you tell them that it isn't round, either. The technical term for the basic shape of the Earth is an "oblate spheroid". Say it out loud a few times...it's great exercise for your lips!



What the Earth would look like if it spun really, really fast.

From now on, then, while we'll simply call g the gravitational acceleration, remember that a small part of it is actually a centrifugal acceleration.

While we're on technical terms, the technical term for a surface (real or imaginary) that is perfectly level in the sense described above is a "geopotential" surface. "Sea level" is a geopotential surface, with the geopotential height of that surface arbitrarily defined as 0. In meteorology, altitudes, or distances along the z axis, are often specified as geopotential heights (or heights, for short) above sea level.

The departure of the Earth from sphericity is so small that spherical coordinates work pretty well for large-scale phenomena. Similarly, the effect of the curvature of the Earth is small enough for more local phenomena that we can work in ordinary Cartesian coordinates, with up (z) defined as normal to the geopotential. You (and the air) can do things on Earth that you can't do in Cartesian space, such as head off in a particular direction and end up back where you started. But, for most purposes, they're close enough. It's sort of like the laws of relativity versus Newton's laws. Relativity is exact, but in most circumstances, you can apply Newton's laws and not notice the difference.

Meteorologists don't stop there. Rarely do they actually deal with the three-dimensional atmosphere in Cartesian coordinates. Instead, two alternative coordinate systems are employed: pressure coordinates and potential temperature (or isentropic) coordinates. Furthermore, both of these coordinate systems are non-orthogonal: the three coordinate axes are not quite at right angles to each other.

Just like we get away with Cartesian coordinates because the Earth is almost flat, we get away with non-orthogonal coordinates because the coordinates are almost orthogonal. A typical pressure surface might vary

in altitude by as much as a kilometer over a thousand kilometers, making it almost flat.

Why bother? For one thing, pressure and potential temperature are easier to measure than height. A cheap barometer and thermometer on board a radiosonde are sufficient to determine exactly the pressure and potential temperature of the sonde at any point. Until the advent of accurate GPS systems, the only ways to determine the height of the rawinsonde were to inflate the balloon to a precise ascent rate (and hope that the air around the balloon wasn't going up or coming down) or compute the height by integrating the hydrostatic equation upward from the ground using the pressure and temperature observations from the rawinsonde.

But even with GPS, pressure and isentropic coordinates are here to stay. Isentropic coordinates are nice because, if there's no latent heating or other diabatic heating going on, potential temperature is conserved. Thus if air starts out at a particular potential temperature, it stays at that particular potential temperature. All the air on a particular isentropic surface stays on that isentropic surface. So any other conserved quantity, such as mixing ratio, is simply moved around by the horizontal wind on that isentropic surface. This only works in isentropic coordinates (or coordinates based on other conserved quantities).

Pressure coordinates have their own advantages. One is the fact that the geostrophic wind equation is simpler in pressure coordinates. As we'll see later, lines of constant height on a pressure surface are parallel to the geostrophic wind just like lines of constant pressure at a particular height. Stronger winds are associated with closer contours in both cases too.

Another nice feature is that the vertical distance between pressure surfaces is proportional to the average temperature between those surfaces. We'll talk about this now.

8.3 The Hydrostatic Equation Revisited

The equation used to determine the height of a rawinsonde, and therefore the height of individual pressure surfaces, is the hydrostatic equation. This is the equation that quantifies the approximate balance between the gravitational acceleration and the acceleration due to the vertical pressure gradient force.

In Chapter 6, we saw that the vertical accelerations were given by

$$\frac{Dw}{Dt} = -g - \frac{1}{\rho} \frac{\partial p}{\partial z}$$

The partial derivative of pressure with respect to height is the measured at a particular instant (holding time fixed) in a particular vertical column (holding the horizontal location fixed). The strange notation on the left-hand side is a ‘total derivative’, and refers to the derivative with respect to time of some characteristic of a particular air parcel. In this case, it’s the rate of change of the parcel’s upward velocity, or, in other words, the upward acceleration.

Most of the time, the vertical acceleration of air is almost zero, so the two accelerations are (approximately) equal and opposite:

$$0 = -g - \frac{1}{\rho} \frac{\partial p}{\partial z}$$

or, rearranging it into the conventional form of the *hydrostatic equation*,

$$\frac{\partial p}{\partial z} = -\rho g$$

In addition to its application to buoyancy analysis, another interpretation of this equation is that it provides the conversion between height and pressure. Next, substitute for density using the ideal gas law,

$$\frac{\partial p}{\partial z} = -\frac{p}{R_d T} g$$

stick everything on the left-hand side,

$$\frac{R_d T}{g} \frac{1}{p} \frac{\partial p}{\partial z} = -1$$

and use the rule that $dy/y = d(\ln y)$:

$$\frac{R_d T}{g} \frac{\partial \ln(p)}{\partial z} = -1$$

8.4 The Hypsometric Equation

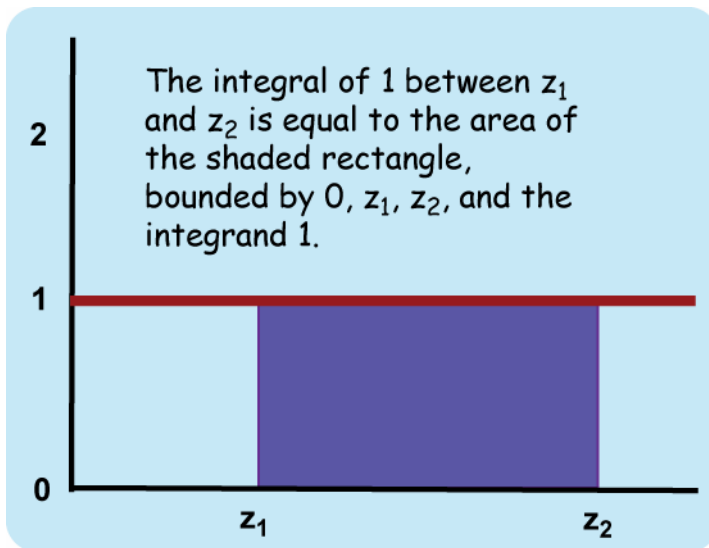
Next, integrate from some low level (p_1, z_1) to some higher level (p_2, z_2):

$$\int_{\ln(p_1)}^{\ln(p_2)} \frac{R_d T}{g} \partial \ln(p) = - \int_{z_1}^{z_2} 1 \partial z$$

To keep the results simple, multiply both sides by -1 and reverse the limits on the left-hand side so that the upper bound has a higher value on both sides of the equation. We can bring the constants outside the integral. And we may as well actually carry out the (trivial) integration on the right hand side.

$$\frac{R_d}{g} \int_{\ln(p_2)}^{\ln(p_1)} T \partial \ln(p) = z_2 - z_1$$

Now, integration is an operation that is graphically expressed as “computing the area under the curve”. The integration we just performed on the right-hand side computed the area of a rectangle (since the integrand was constant) whose height (the integrand) was 1 and whose length (the bounds of integration) was $z_2 - z_1$. The integral on the left does the same sort of thing, but the temperature throughout an atmospheric layer generally does not have a single value so the area is not a rectangle. Instead, in general, temperature will change as you go to higher or lower pressures. Otherwise, plots of soundings would be very boring.

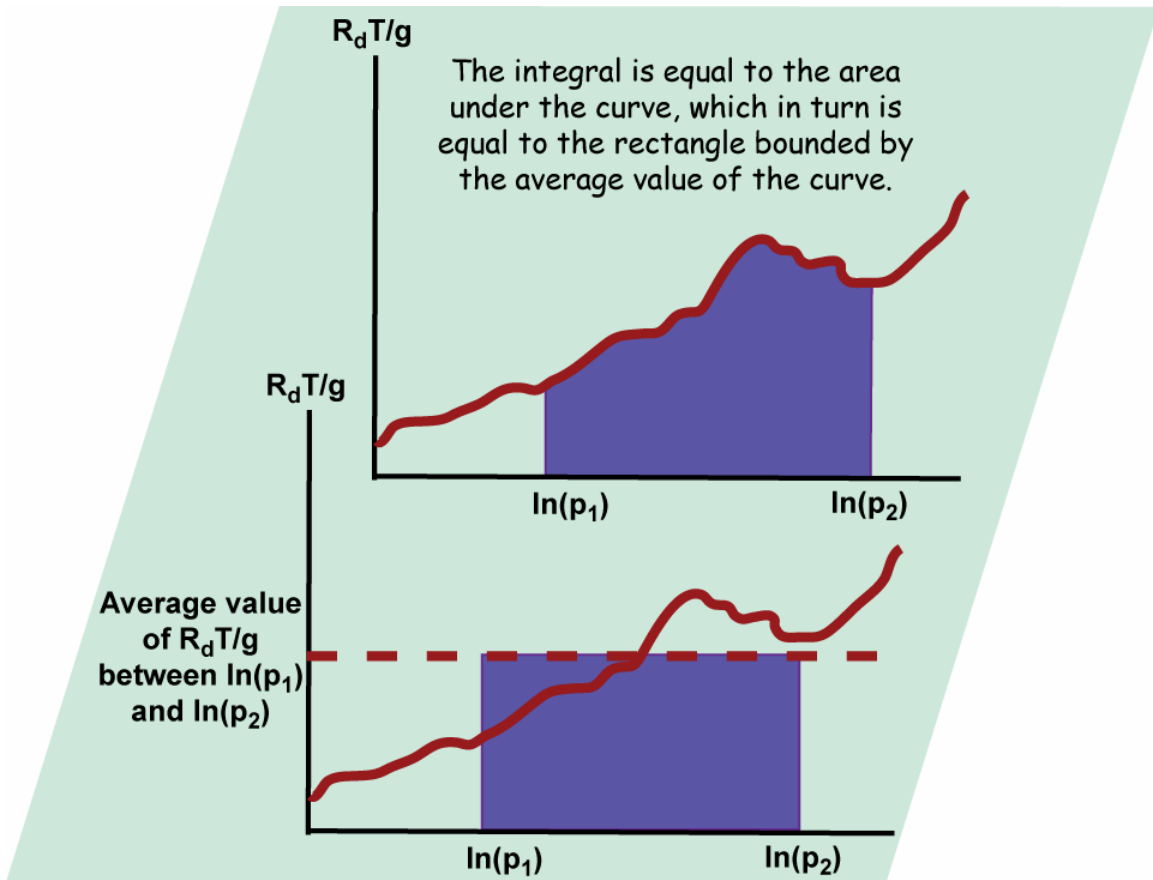


In calculus, you learn all about doing integration and determining the answer based on mathematical theorems and principles. Here, we’re just going to close our eyes, snap our fingers, and get “an” answer. For the value of the integral, whatever it is, is equal to the area of a rectangle whose length is $\ln(p_1) - \ln(p_2)$ and whose height is the average value of T over that interval (as long as you average over equal increments of log pressure). So rather than compute the integral, which can only be done

numerically and which changes for each different location, we will simply record the answer as

$$\frac{R_d}{g} [\ln(p_1) - \ln(p_2)] \bar{T} = z_2 - z_1$$

where the overbar indicates an average value for T .



Strictly speaking, the gravitational acceleration is not a constant. Instead, it depends on distance from the center of the Earth, making it a function of height. We can make that problem go away by using the true definition of a quantity, the geopotential height Z , as

$$Z = z (g/g_o)$$

where g_o is the gravitational acceleration at the Earth's surface.

Also strictly speaking, we've used the gas constant for dry air R_d when air is not necessarily dry. Instead, it generally has some nonzero fraction (mixing ratio) of water vapor, and the gas constant dry air is only 0.611 times the gas constant for water vapor. The intelligent thing to do might be to determine the correct value for the gas constant in each

circumstance based on the water vapor content of the air, but then another constant will have turned into a variable. Meteorologists instead do something rather remarkable in their audacity: they define a *virtual temperature* T_v as the temperature that a mixture of dry air and water vapor would have, given its density and pressure, if the ideal gas law worked using the gas constant for dry air alone. Fortunately, the formula for computing it is simple:

$$T_v = T (1 + 0.611 w)$$

where w is the mixing ratio of water vapor, expressed in kg/kg. Since a typical value of w is 0.010 kg/kg, the virtual temperature is usually within about a couple degrees of (and slightly warmer than) the actual temperature.

Okay, so at this point in your education those differences are inconsequential, but we may as well be strictly correct here:

$$\frac{R_d}{g_o} [\ln(p_1) - \ln(p_2)] \overline{T}_v = Z_2 - Z_1$$

This equation, which is merely an integrated form of the hydrostatic equation, is sufficiently important to get a name of its own: the *hypsometric equation*. The hypsometric equation states that the vertical distance between two pressures in the atmosphere is proportional to the average temperature between those two pressure levels.

If you just have pressure and temperature information, this lets you compute the height of each pressure level. You also need a starting point, say Z_1 , so that there's only one unknown in the equation. Usually the starting point will be the ground, and level 2 will be the height/pressure of the first data point above the ground. From there, the height of the next data point can be computed, and so forth. So it is that the height of the balloon at all times can be computed with standard rawinsonde data.

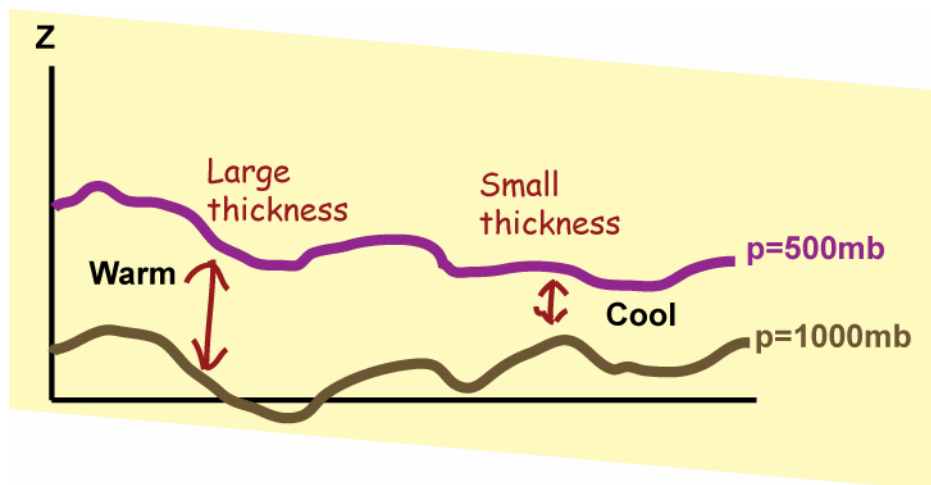
8.5 Thickness

Once meteorologists figured out which pressure levels they wanted to use for plotting maps in the troposphere, it didn't take long for international convention to dictate that the values of temperature, wind, and even height at those pressure levels be automatically included in each rawinsonde report. In the context of the upper-air reports, these levels are called *mandatory levels*, because reporting data at those levels are mandatory. From the ground to 100 mb, the mandatory pressure levels are the following: 1000 mb, 925 mb, 850 mb, 700 mb, 500 mb, 400 mb, 300

mb, 250 mb, 200 mb, 150 mb, and 100 mb. Of these, the levels most commonly used for pressure maps are 850 mb, 700 mb, 500 mb, and 250 mb.

Now, suppose you are looking at geopotential height on a 500 mb map and you are curious about what the 1000 mb map would look like. You also have a general idea about the pattern of temperatures in the lower troposphere, between 500 mb and 1000 mb. According to the hypsometric equation, you should already know what the 1000 mb map would look like. Once you've chosen two particular pressure levels, the only unknowns left in the hypsometric equation are the average (virtual) temperature between the two levels and the geopotential heights of those two levels. If you know the height of one pressure surface and the temperatures between them, you can compute the height of the other surface.

Rarely, though, will you do that computation, except for homework. Instead, think about the quantity $Z_2 - Z_1$. This difference is the vertical distance between the two pressure surfaces. In very real terms, it is the thickness of the layer of air bounded by those two pressure surfaces, and indeed that's the name attached to that difference: the *thickness*. Of course, there are as many thicknesses as there are combinations of pressure levels, so to be specific, one has to specify the two bounding pressure surfaces. In our example, we would be working with the 1000 to 500 mb thickness.



Now, wherever the average temperature in that layer is cold, the thickness must be small. Wherever it's warm, the thickness must be large. So suppose you had a perfectly flat 500 mb surface: no (geopotential) height variations at all. Under a cold spot, the thickness is small, and under a warm spot, the thickness is large. So the height of the 1000 mb surface must be large under the cold spot and small under the warm spot.

Let's try it the other way: suppose you know the 1000 mb heights, and they're almost constant. Yet to the north the temperatures are cold, and therefore thicknesses are low to the north too. Thus, the 500 mb geopotential heights must tend to be low to the north and higher to the south.

The importance of the concept of thickness lies in the close connection between heights and temperatures. Since there's also a close connection between heights and winds, the concept of thickness ties together the wind and temperature fields in the atmosphere. We'll circle back on this concept again and again.

8.6 Critical Thickness

Thickness also has a simple but practical application: snow forecasting. To get snow, you need the temperature at just about all levels of the atmosphere to be at or below freezing. Looking at constant pressure maps gives you the temperature at a few levels, but what's happening in between? You could examine a sounding, but if you're forecasting snow chances for a large area, you'd have to examine soundings (either observed or from a forecast model) throughout that area.

The solution is thickness. Since thickness is proportional to the layer average of temperature, it gives you information throughout a deep column of the atmosphere. Hard-earned experience has shown that thickness is a good predictor of snow possibilities.

At most snow-vulnerable locations, forecasters have determined the "critical" thickness: the thickness value below which snow is likely. Originally, the 1000-500 mb layer was used, and typical critical thicknesses would be 534 dam or 540 dam. More recently, since the air close to 500 mb is always below freezing, layers closer to the ground have been used, such as 1000-700 mb or 850-700 mb.

Critical thickness rules of thumb work because in the precipitation-producing layers of the atmosphere during winter storms, temperature profiles tend to follow a very consistent pattern, being slightly more stable than moist adiabatic. The rules can be adjusted for particular situations; once rain and snow have started, the forecaster can check which value of thickness most closely corresponds to the rain/snow line and apply that value to forecasts of precipitation type.

The 1000-500 mb thickness map is still widely used even when snow is not an issue. Since it represents the average temperature of the lower troposphere, it's a good measure of the intensity and depth of warm and cold air masses. Where there's a strong thickness gradient, there's

going to almost certainly be a strong surface front. Because such a deep layer is used, the thicknesses are not subject to the large daytime-nighttime variations that beset surface temperatures.

Questions

1. A thermodynamic diagram is based on a grid of pressures and temperatures. In what sense(s) can pressure and temperature be regarded as coordinates? In what sense(s) can they not be regarded as coordinates?

2. Describe an experiment that might be used to detect and measure the centrifugal force caused by the Earth's rotation.

3. In a sounding, the temperature at 850 mb is 18 C and the dewpoint is 15 C. Plot the dewpoint on a sounding diagram and use the saturation mixing ratio lines to read off the mixing ratio of the air parcel. Then, compute the virtual temperature.

4. The 1000-500 mb thickness of the atmosphere typically ranges from 480 dekameters (dam, 4800 m) to 580 dam (5800 m). Create a table giving the values of mean layer virtual temperature corresponding to various values of 1000-500 mb thickness.

5. The height of the 850 mb surface at a particular point is 146 dam, the temperature at 850 mb is 15 C, and the temperature at 700 mb is 5 C. Estimate the height of the 700 mb surface.

6. The sea level pressure at a particular point is 1008 mb, and the temperature is 48 F. What is the approximate height of the 1000 mb surface?

7. (a) Estimate the largest possible 850mb-700mb thickness in which the air is below freezing throughout the layer. (b) Estimate the largest possible 850mb-700mb thickness in which the air is below freezing throughout the layer and the lapse rate is moist adiabatic. A rough estimate with the help of a sounding diagram is okay.