Thermodynamic Diagrams

SkewT diagrams can look pretty forbidding but in the following discussion I will provide sufficient guidance to enable anyone to take a quick look at them and draw some pretty valid conclusions about what kind of a soaring day might be expected. With a little diligence expect to be able to glance at a once forbidding skewT, and conclude in this instance that it's going to be decent blue day with lift to 6,000 ft.

SkewT diagrams are at once a compact way to present a lot of data, and calculating devices. I will draw clear distinctions between these two properties. Balloon soundings, and the output of numerical models, including the temperature and dewpoint of the atmosphere from the surface to about 80,000 feet, are conveniently displayed on skewTs so it pays to understand them.

Those spoiled by Bill Moninger's great FSL website should be aware that nothing that happens on that site happens without the underlying thermodynamics captured in the skewT diagram. The same holds for BLIPMAPS so be advised that pretty though they be, if the underlying numerical model is wrong, and it sometimes is, the forecast will be wrong. When you understand what the skewT is showing, you are in a position to critically assess the BLIPMAP forecast.

Parenthetically I note that although these diagrams are ideal for producing thermal soaring forecasts, they were developed as an aid in forecasting convective storms. The enormous effort which goes into gathering and disseminating the data we use, and the comparable effort which goes into numerical modeling is driven by the destructive potential of deep cyclones, hurricanes, and thunderstorms, not by our desire to have good soaring forecasts.

To reiterate: SkewT's are graphs which display temperature and dewpoint data vertically in the earth's atmosphere, and lines superimposed upon those graphs which do certain kinds of calculations. Each of these calculating lines (dry adiabats, saturated adiabats, lines of constant mixing ratio) is backed up by thermodynamic equations. Their utility lies in the fact that we don't need to worry about the equations, and don't need to do any math. We do however need to understand what they represent.

Temperature and Pressure Lines

Stripped of all the calculating lines, and devoid of any data, the skewT looks a lot like an empty graph and except for the fact that the red lines are sloping and that is exactly what it is.

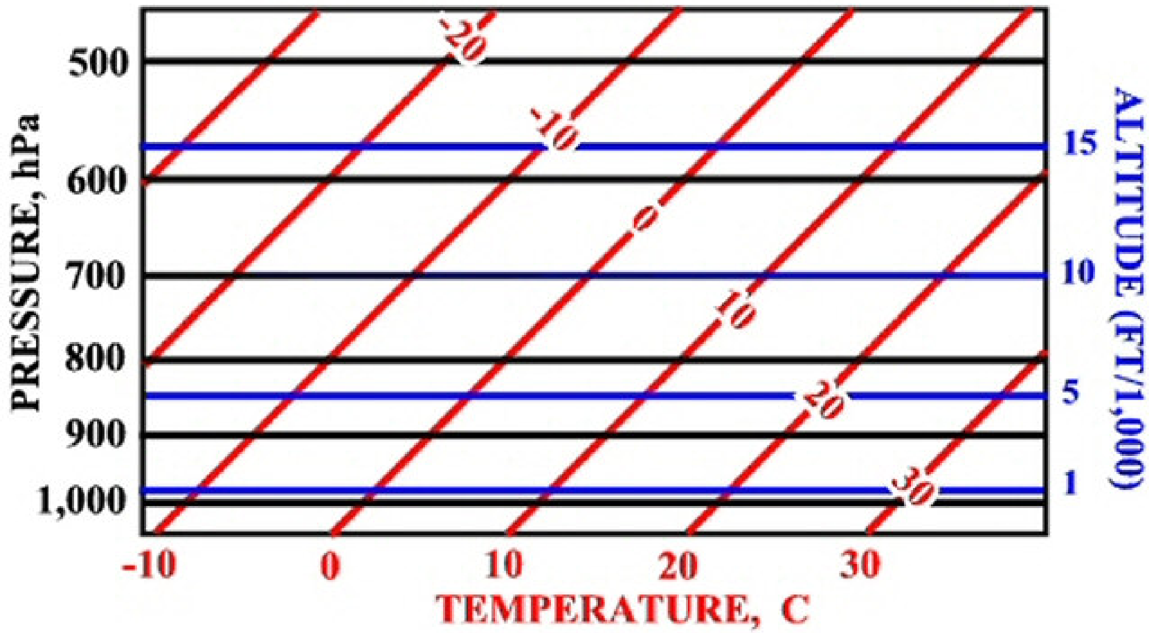

The temperature lines are at an angle ("skewed"), and the pressure scale is logarithmic because these choices make quite a few of the variables we will want to examine more-or-less linear. That is convenient when making extrapolations on a graph. For clarity I had added pressure altitude lines - these are a helpful aid, not an integral part of the diagram.

Dry Adiabatic Lines

The dry adiabatic lines tell us what happens to the temperature of a parcel of air as it rises or falls in the absence of condensation. We need to know this because the density, and hence the buoyancy of air, depends in part upon its temperature. These are the first example of what I am calling "calculating lines". They are depicted in cyan in the simplified skewT to which I have added only the dry adiabats.

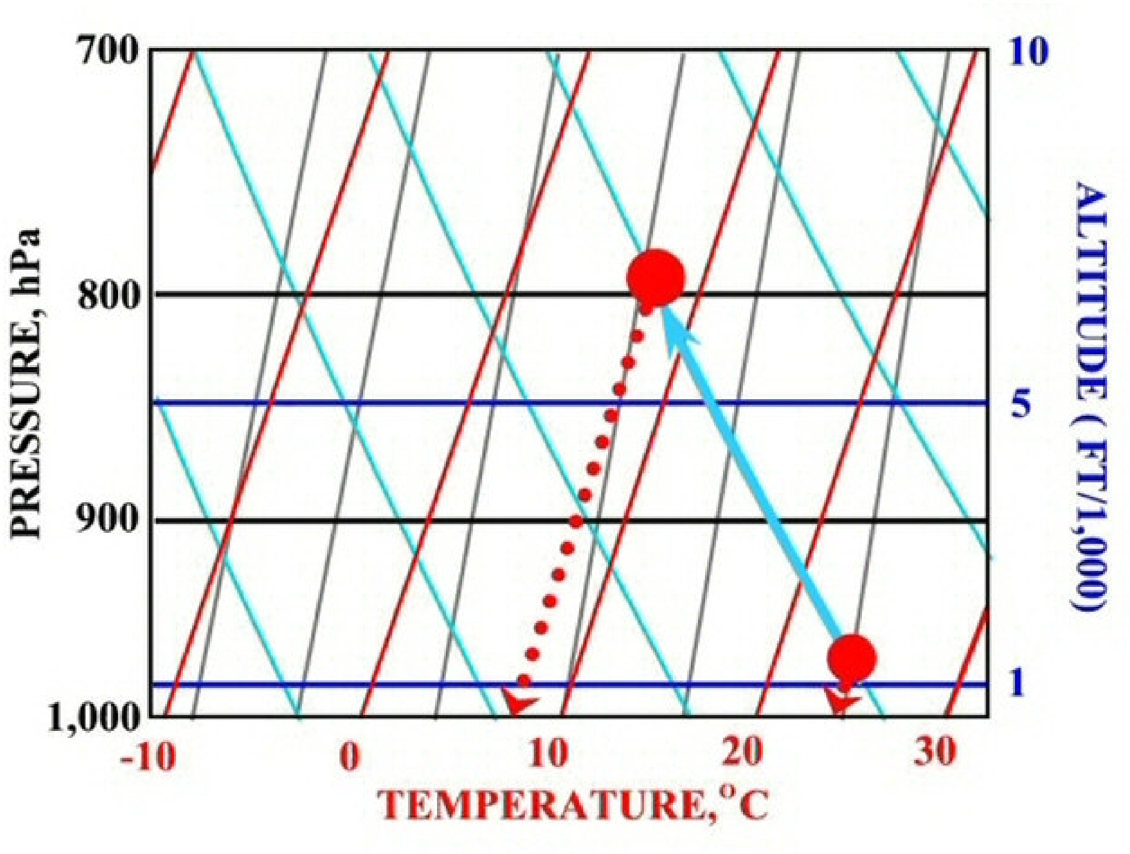

It depicts a spherical bubble of air at the surface (in this case at 1,200 ft msl and at a temperature of 23ºC) and tells us what happens when the bubble of air rises to about 6,000 ft msl. This spherical bubble is a simple model for a thermal. As it rises it cools, and the precise amount of cooling is determined by the dry adiabatic lines. Note that because the temperature axis is skewed, care must be taken in choosing the correct temperature. I have added the red arrows to make it clear how the temperature is to be read. Note too that often there is no dry adiabat coinciding with the surface temperature and when this happens it is easy enough to construct it.

By the time the bubble is at 6,000 ft it has cooled to 8ºC and has expanded by about 20%. The cooling is driven by the expansion of the bubble. As it rises the pressure of the surrounding air decreases and the bubble expands and does work pushing away the enveloping air. It is this work, achieved by tapping the internal energy of the bubble, which is responsible for the cooling.

An adiabatic process is one which takes place without exchange of heat with the surroundings so, in assuming an adiabatic process, we are assuming that the parcel of air maintains its identity and does not mix with the air through which it is moving. This is an acceptable assumption, at least in the sense that it yields useful and consistent results.

"Dry" in DALR does not mean quite what it says, since air is never dry. Specifically it means that no condensation occurs. When condensation does occur, the Saturation Adiabatic Lapse Rate (SALR) takes over. I will consider the SALR lines later when considering what goes on above cloudbase.

Lines of Constant Mixing Ratio

We started with the empty skewT logP graph to which we added dry adiabats. To describe what happens to the dewpoint of a rising parcel of air we need to add more lines - the lines of constant mixing ratio.

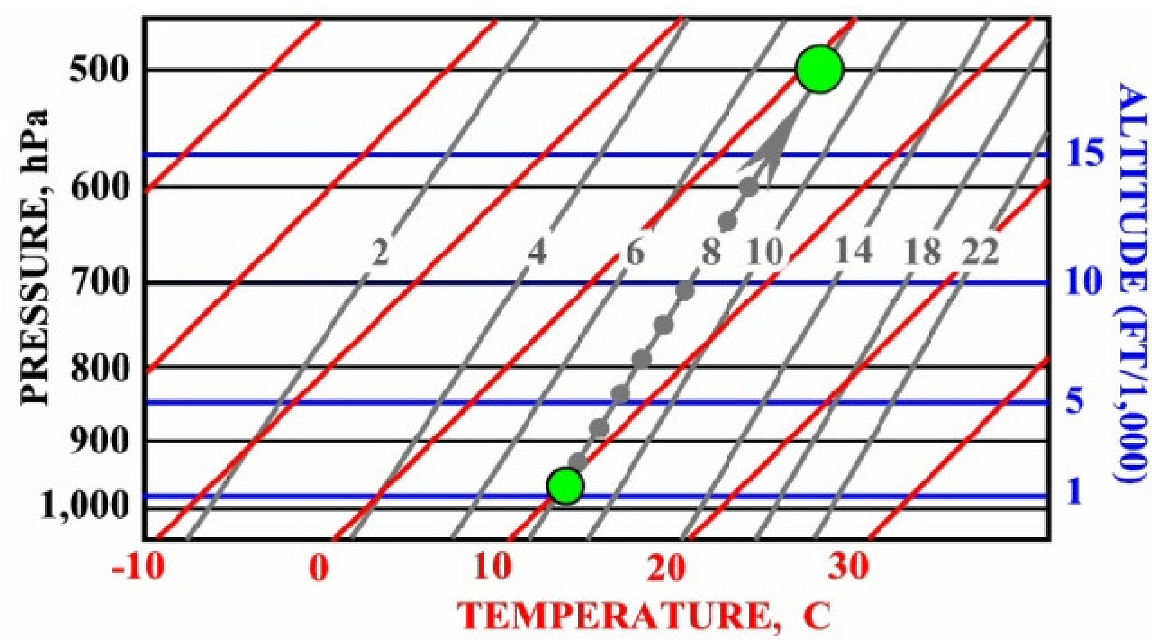

The change in the dewpoint air as it rises in a thermal concerns us because if the temperature and dewpoint converge until they are equal condensation occurs and a cloud will begin to form. We know from the dry adiabats how quickly the temperature cools with ascent of the bubble. To predict whether or not clouds will form we need to know how the dewpoint of the bubble changes during its ascent. Lines of constant mixing ratio tell us. The "mixing ratio" is the concentration of water vapor in the air, expressed in grams of water per kilogram of air. Since this is a mass ratio (as opposed to a volume ratio), it does not change as the bubble expands. However, the dewpoint does change. A surface bubble with a dewpoint of about 10ºC has a dewpoint of about 2ºC at 18,000 feet.

To re-iterate: Lines of constant mixing ratio tell us how the dewpoint of a bubble changes with altitude. The mixing ratio line passing through the surface dewpoint tells us what the dewpoint of the lifted surface parcel is at any height. Since the temperature of rising air falls faster than does the dewpoint, unchecked ascent always results in cloud formation.

The origin of the decrease in dewpoint as a parcel ascends and cools is the associated expansion as the pressure drops. This expansion causes cooling, as we saw when considering the dry adiabats, but it also increases the average distance between water vapor molecules. Since condensation *ipso facto* involves an agglomeration of molecules, it is to be expected that having fewer molecules in a given volume will result in a lower dewpoint - i.e. the air will have to be colder to force condensation to occur.

Making Cumulus Clouds

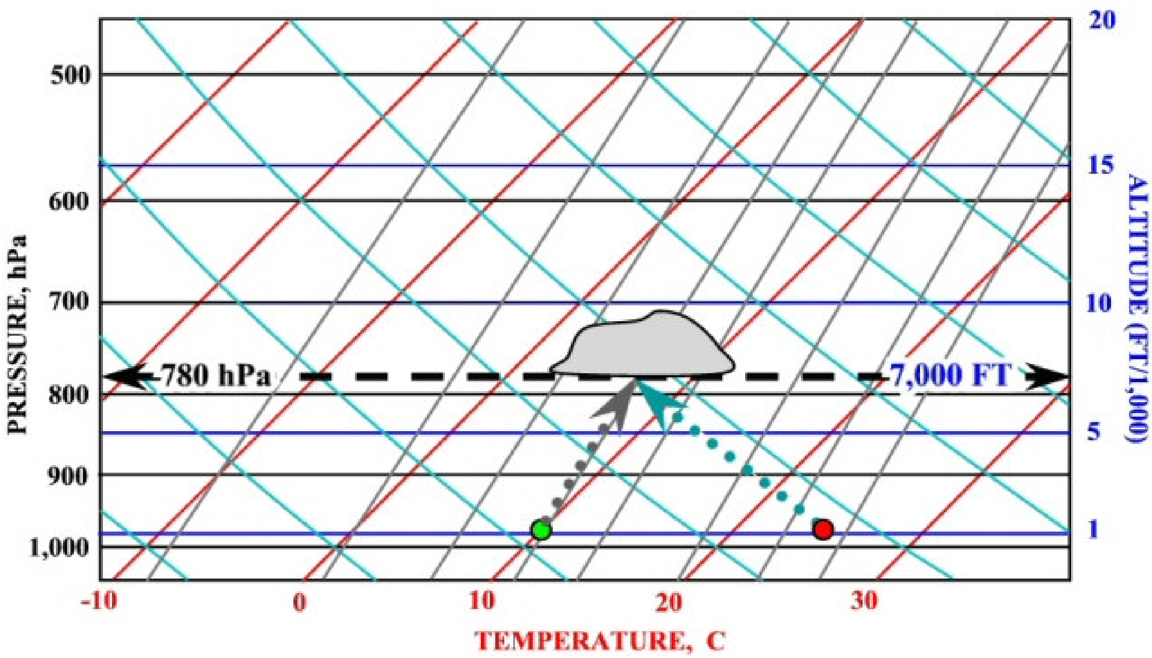

We are now in a position to understand how the skewT handles clouds. Suppose we start with a surface parcel of air having a temperature of 25ºC and a dewpoint of 10ºC as depicted here . If the bubble is lifted from the surface to 7,000 ft its temperature drops to about 8ºC, and its dewpoint will drop to the same temperature: Since the dewpoint and temperature become equal at 7,000 feet condensation must occur, and a cloud forms.

The widely used formula for cloudbase is a consequence of the differing rates of change of the temperature and dewpoint of a lifted parcel of air.

TEMPERATURE LAPSE RATE = 5.3 °F / 1,000 FT. (approximately)

DEWPOINT LAPSE RATE = 0.9 °F / 1,000 FT (approximately)

CLOUDBASE = ((T - DP) / 4.4°F) * 1,000 FT

TEMPERATURE LAPSE RATE = 2.94 °C / 1,000 FT. (approximately)

DEWPOINT LAPSE RATE = 0.5 °C / 1,000 FT (approximately)

CLOUDBASE = ((T - DP) / 2.44°C) * 1,000 FT

It is comforting to see that our first use of the skewT diagram predicts a familiar result.

Those who regularly apply this rule will have noticed that it occasionally fails since cumulus clouds never appear. The reason for this will soon become clear, as will the folly of relying too heavily upon the blind application of rules.

Lapse Rates

As I noted at the beginning of this discussion the skewT diagram does double duty: It calculates things (the change in the temperature and dewpoint of a rising bubble in the absence of condensation, and the change in the temperature of a rising bubble where condensation is occurring) and it presents data. Perhaps it is this dual role which gives rise to some of the confusion which SkewT's occasion. Further confusion arises because the term "lapse rate" also does double duty.

There are four lapse rates with which we will be dealing: The dry adiabatic, the saturated adiabatic, the temperature, and the dewpoint. Since the term "lapse rate" tends to get used imprecisely, I will make the distinction clear: The first two, the dry and saturated adiabatic lapse rates are calculated (using the adiabatic approximation) and that is why both always appear unchanged on every skewT plot. The second two, the temperature and dewpoint lapse rate are measured (or in the case of numerical model soundings, predicted values at the valid time of the forecast).

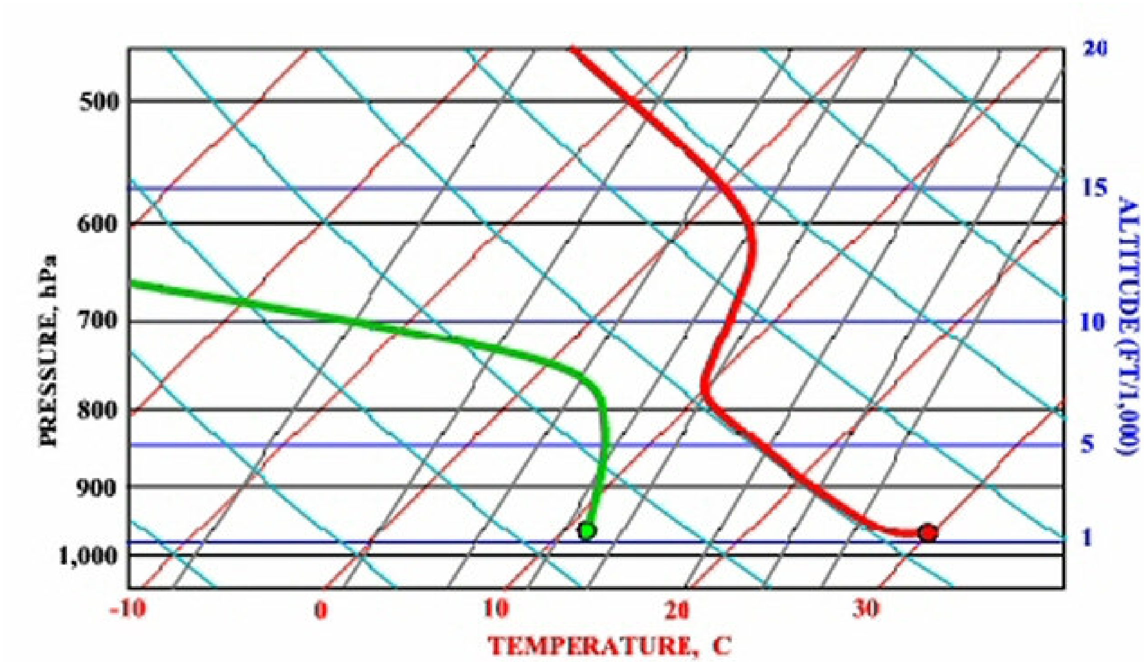

In the next illustration I have superimposed data lines upon the calculating lines, and it will now become apparent why the skewT diagram is so useful. The solid red line is the temperature. It's important to appreciate that every point on this line represents the actual or forecast temperature of the air at a given height.

There are three regions of interest, three different lapse rates, in the red line: Close to the surface the temperature decreases rather dramatically with increasing height. Then, from about a few hundred feet above the ground to about 6,000 feet msl, the temperature decreases at the DALR. From 6,000 ft the temperature decreases much more slowly. We need to understand why each region is the way it is, and what this implies for thermal forecasting.

The lowest layer, referred to as the "super adiabatic layer" exists courtesy of the sun. Its persistence is of kinetic, not thermodynamic origin. As soon as the forcing insolation is cut off it decays because any vertical displacement will result in the air finding itself in colder air. It is "unstable" in exactly the same sense that a rock on the edge of a cliff is unstable - a small push and it's gone. Although in the East super adiabatic layers are generally thin, they can grow to many hundreds of feet under sufficient sun, and in the absence of anything doing the equivalent of rolling the rock off the cliff, often do. This is origin of dust devils which can break loose with great energy and rise to 15,000 feet or more.

From a few hundred feet above the surface to about 6,000 ft the temperature tracks the DALR - this is no accident: It is the very thermals glider pilots seek, and the downdrafts they avoid, which mix up the air so that its actual lapse rate inevitably becomes approximately dry adiabatic when the sun is shining.

The next layer is generally referred to as "the inversion". In this example, as is often the case, the temperature continues to drop with increasing height, so strictly speaking there is no inversion, however, the air above 6,000 feet is getting colder with increasing height a lot more slowly. On a blue day, it is the lowest lying inversion which caps the lift, and although this seems to place inversions in a poor light, their absence can cause problems, as I will show later on.

The solid green line is the dewpoint. The shape is typical of a thermal soaring day, particularly in the tendency for the dewpoint and temperature to converge in the vicinity of the inversion and this too is no coincidence and has consequences to which I shall return.

Can We Soar?

We might as well start with a good day, so here's one I made up.

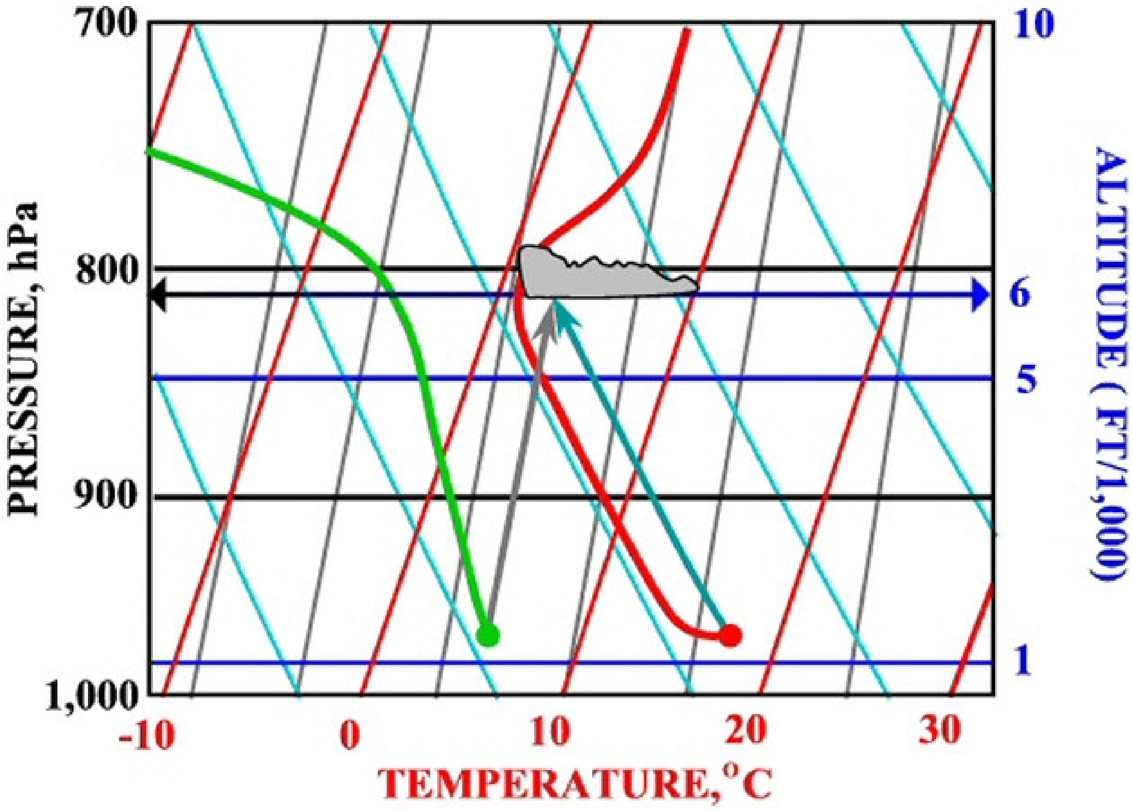

It's typical of a good spring day in the East. The red line is what the model predicts the temperature lapse will be. The red dot is what the model predicts the average surface temperature will be. The cyan arrow is the surface adiabat and because it is displaced from the temperature by about 2.5 C° to over 6,000 ft. this is a pretty unstable day. The greater the displacement, the greater the instability.

The green line and green dot are what the model predicts the dewpoint lapse rate will be. I have constructed the constant mixing ratio line passing through this value. That line intersects the surface adiabat below the inversion, telling me that on this day lift will be marked by cumulus cloud. As an aside, I note that at cloudbase, the air is still forecast to be 2.5 C° warmer than its surroundings, so on this day I would not expect lift to decrease with increasing height, as sometimes happens.

The dewpoint is well behaved, so I would not anticipate spreadout, and the inversion is more than enough to cap cloud growth. Neither of these predictions can be made on the basis of what we have so far considered, and I will return to these two critical elements of the soaring forecast.

I am going to repeat much of what I said about the good day with cu, since I want to be sure everyone gets the picture, so hear is what a good day without cu looks like.

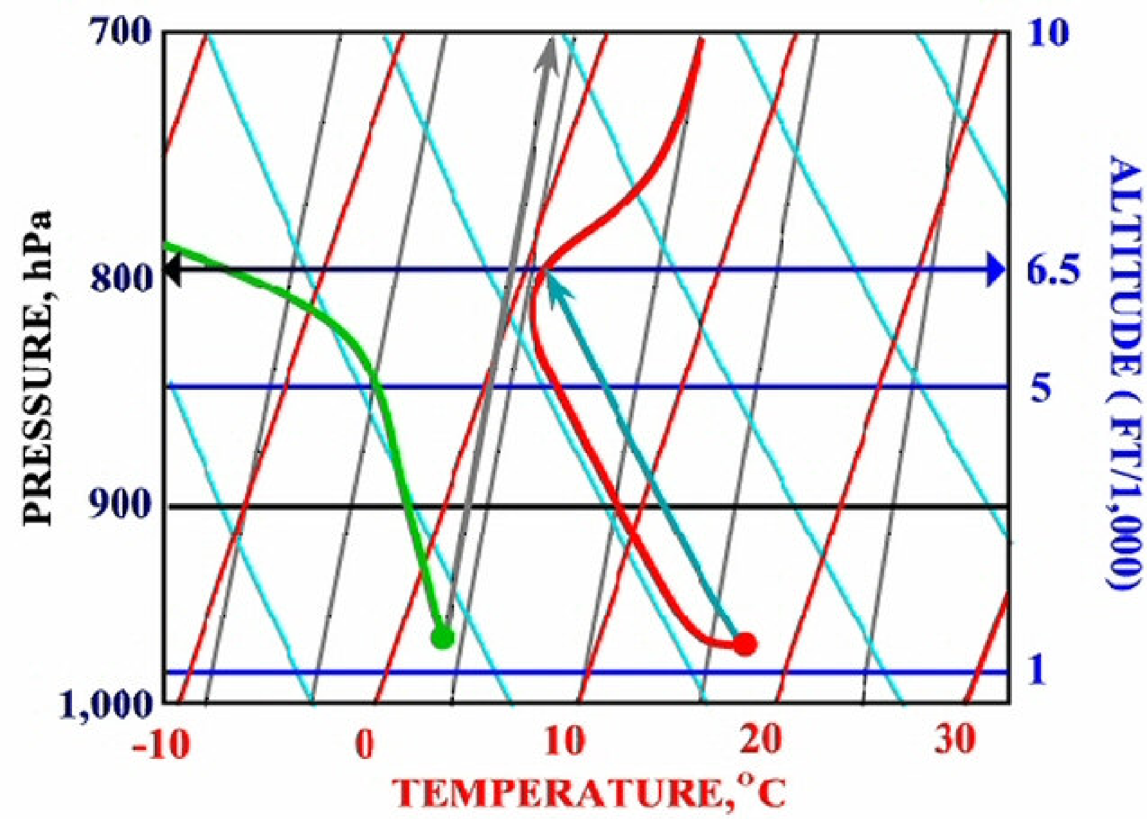

The red line is what the model predicts the temperature lapse rate will be. The red dot is what the model predicts the average surface temperature will be.

The cyan arrow is the surface adiabat and is displaced from the temperature by about 2.5C to over 6,000 ft. so again it's a pretty unstable day.

The green line and green dot are what the model predicts the dewpoint will be. I have constructed the constant mixing ratio line passing through this value. That line intersects the surface adiabat above the inversion, telling me that on this day lift will not be marked by cumulus cloud, and as a result, lift will be capped by the inversion.

What Goes on Above Cloudbase

In discussing the dry adiabats, I noted that when condensation occurs, a different adiabat comes into play. Condensation involves the release of heat - in the case of water vapor condensing to liquid water, large quantities of heat. Enough heat to drive convective storms and hurricanes.

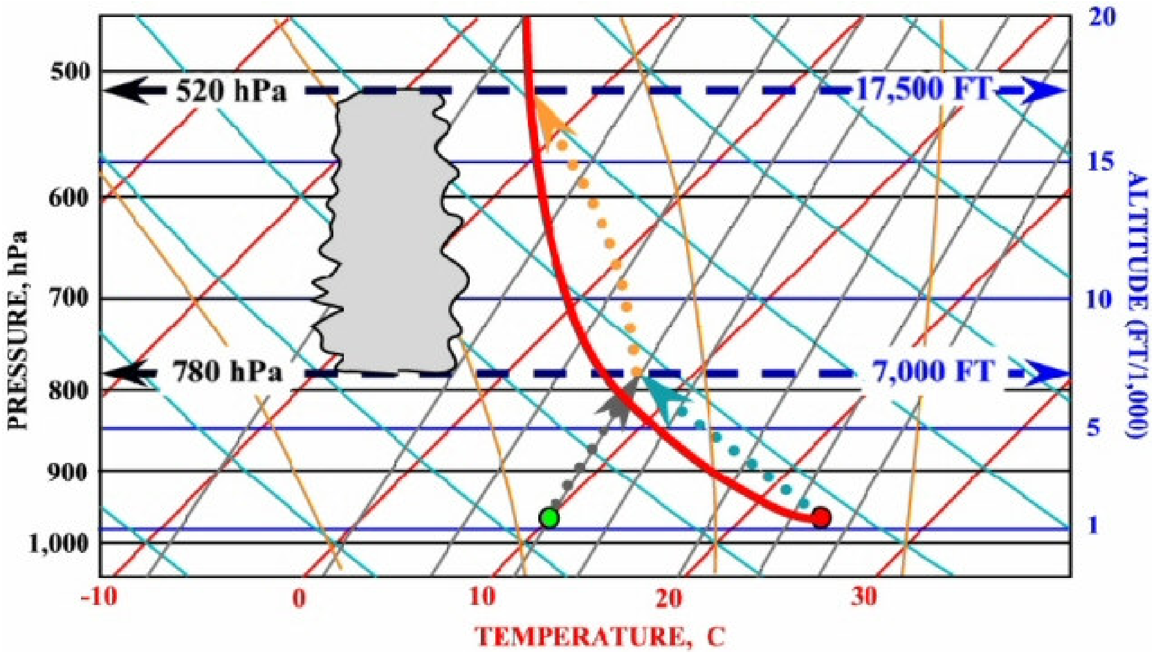

I have now added (orange) saturated adiabats. Observe that these lines and the dry adiabats (cyan) have a dramatically different slope at lower altitudes where the atmosphere is able to hold more water than it can at higher levels. Above about 500 hPa the saturated adiabats have almost the same slope as the dry adiabats.

Air undergoing condensation as it ascends cools more slowly than does non-condensing air because of the release of latent heat from the vapor/liquid phase transformation.

The area of the figure below 780 hPa ought by now to look pretty familiar: Unstable surface air at about 24°C ascends until, at 7,000 ft its temperature and dewpoint are the same, at which point condensation occurs. So far so good, after all we want cu don't we?

The problem here is that there is no inversion in the vicinity of cloudbase. As soon as clouds form, the dotted orange adiabat takes over and the cloud will grow to about 17,500 ft. If a cloud grows to above the freezing level (which this one easily does - the temperature lapse rate crosses the 0 degree dry adiabat at 700 hPa) rain can form. Not obvious from my illustration is that when convection extends to heights where winds are much stronger, those winds can mix down to the surface. My figure was adapted from the forecast sounding for a contest day at Mifflin County a few years ago. Although it was a generally good day, we saw isolated heavy rain showers and a 60 mph gust at the field. This is why, at least on days with cu, we need an inversion.

Spreadout ("Over Development")

Many good soaring days are ruined by overdevelopment, or as I prefer to call it, spreadout. The origin of spreadout is immediately evident on inspection of this skewT:

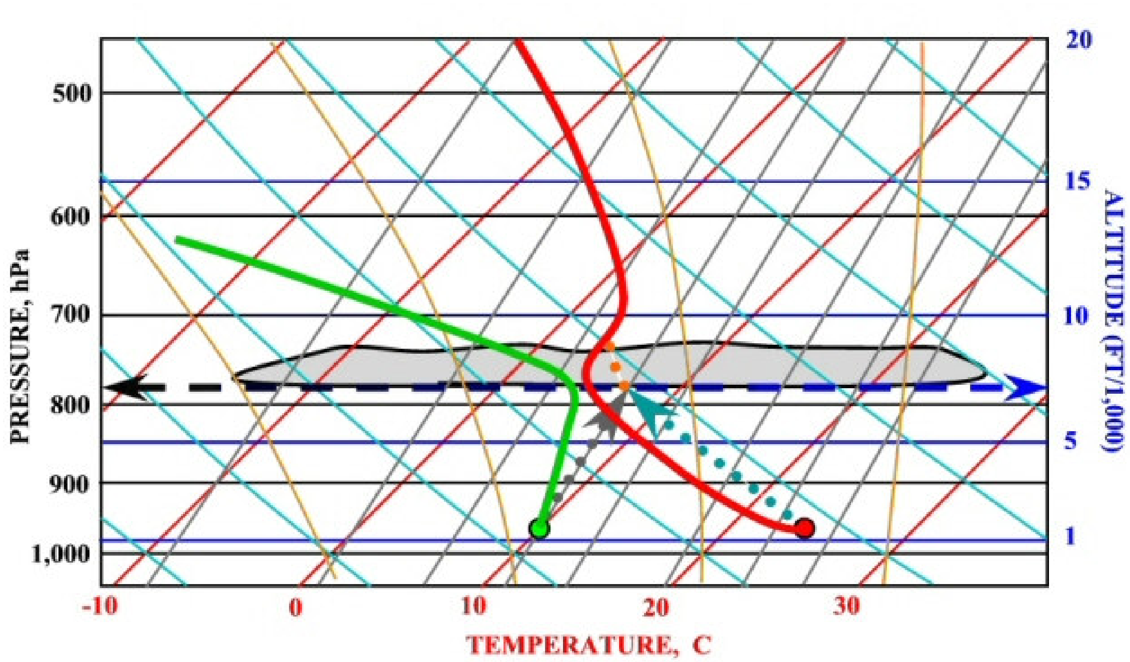

In the vicinity of cloudbase the dewpoint and temperature of the general airmass (i.e. all of the air NOT in thermals) are rather close to one another. Humidity is thus high, and dispersal of clouds through evaporation is slow.

If I see a temperature difference of less than 3°C at or in the vicinity of cloudbase, I will advise of the possibility of spreadout. The smaller the difference, the greater the chance. It may seem perverse that the dewpoint so frequently seems to approach the temperature at the inversion but this is no accident: The moisture is transported there by convection, and prevented from rising by the inversion.

The Effect of Moisture on Buoyancy

Air is never completely dry, and since water vapor is less dense than air, addition of water vapor decreases the density. Since buoyancy depends only on the density difference between the thermal and the surrounding air, we should in principle take account of water vapor.

The effect can be significant. It is accounted for by calculating the temperature the air would be at if it were dry, and of the same density. That temperature is known as the virtual temperature. On a typical East Coast soaring day the virtual temperature is about 2 degrees higher than the 2 meter surface temperature. The formula for calculating virtual temperature is:

$T_v = T + W/6$ (approximately).

where T is the surface temperature in degrees C and W is the mixing ratio in grams of water/kg of air.

Perversely, what is displayed on every skewT diagram that I am familiar with is the uncorrected temperature, not virtual temperature so I generally make the assumption (not completely true) that a correction in the surface temperature can be made. This mostly affects the predicted height of clouds.

Cirrus and other Clouds

General weather forecasts will usually include terms such as "partly" or "mostly" cloudy but they never say what kind of cloud. TAF and MOS (Model Output Statistics) forecasts are more helpful since they give the heights of cloud layers. Any skewT will make it clear the height and thickness of any cloud layer because the dewpoint and temperature lines will be close whenever cloud is likely.

Summary

The skewT diagram captures for a point on the surface the profile of the atmosphere above that point. It presents a complete picture of the temperature, the dewpoint and the wind speed and direction from the surface to about 80,000 ft. In addition to presenting a great deal of data in a highly compact form, the skewT also makes it simple to perform a variety of "what if?" calculations through the overlay of dry and wet adiabats, and lines of constant mixing ratio.

These diagrams are ideally suited to producing thermal soaring forecasts. With a little effort anyone interested in cross-country soaring (or even just staying up for that matter) can get a very good feel for the day with just a few minutes spent studying them. Even those pilots relying upon others for their forecasts would do well to take a look at the underlying data, and there is no better way to do that than the skewT.