Synoptic Meteorology I: Skew-T Diagrams and Thermodynamic Properties

For Further Reading

Most information contained within these lecture notes is drawn from Chapters 1, 2, 4, and 6 of "The Use of the Skew T, Log P Diagram in Analysis and Forecasting" by the Air Force Weather Agency, a PDF copy of which is available from the course website. Chapter 5 of *Weather Analysis* by D. Djurić provides further details about the utility of skew-T/ln-p diagrams.

Thermodynamic Diagrams and the Skew-T Chart

Thermodynamic diagrams serve three general purposes. First, they provide a means to plot and analyze observations – namely temperature, dew point temperature (or mixing ratio), and wind – taken above the Earth's surface at a given location. Second, thermodynamic diagrams enable us to graphically assess or compute stability-related fields that are functions of the observed variables. Finally, thermodynamic diagrams allow for atmospheric processes such as lifting an air parcel to be simulated, enabling an estimate of the effects of such processes on an air parcel's properties.

On any thermodynamic diagram, we require that there be at least five different kinds of isolines:

- **Isotherms**. Isotherms allow for observations of temperature and dew point temperature to be plotted. They allow for the assessment of derived fields such as potential temperature, equivalent potential temperature, and convective temperature.

- **Isobars**. Isobars allow for the vertical variability in temperature, dew point temperature, and wind speed and direction to be plotted. They serve as reference lines by which lifting levels and atmospheric layers may be identified.

- **Dry adiabats**. An *adiabatic process* is one that occurs without the transfer of heat or matter between an air parcel and its surroundings. The ascent of an unsaturated air parcel is one example of an adiabatic process. A parcel undergoing an adiabatic process follows a dry adiabat as it does so. A dry adiabat, given its relationship to the isotherms, is a measure of the rate at which temperature cools or warms (9.8°C per 1 km) upon ascent or descent, with this temperature change due to the adiabatic expansion (cooling) or compression (warming) of an air parcel as it changes its pressure.

- **Mixing ratio lines**. Mixing ratio is a measure of the mass of water vapor per kilogram of dry air. For an unsaturated air parcel, mixing ratio is conserved (does not change) as an air parcel changes altitude. Mixing ratio lines provide a means of representing this property.

- **Saturated adiabats**. The primary atmospheric exchange of heat between an air parcel and its surroundings (a *diabatic process*) is that associated with phase changes of water. For example, as an air parcel ascends, condensation or freezing may result, and the associated change of phase transfers heat from water to air. A saturated adiabat is a measure of the rate at which temperature cools upon ascent after an air parcel has become saturated. Note that the precise value depends on moisture content, with higher moisture content associated with a smaller amount of cooling, and assumptions made regarding what happens to water in the air parcel immediately upon condensation, deposition, or freezing. In the lower-to-middle troposphere, saturated adiabats indicate a cooling rate of ~6-7°C per 1 km of ascent.

There exist many different types of thermodynamic diagrams, though a discussion of the benefits and drawbacks of each is beyond the scope of this class. Instead, we focus on one specific type of thermodynamic diagram: the *skew-T/ln-p*, or skew-T for short, and its benefits and applications.

The skew-T diagram takes its name from the isotherm and isobar orientations on the diagram. The isotherms on a skew-T diagram are skewed 45° from the vertical. The natural logarithm of pressure (ln *p*) is the vertical coordinate. Isobars are given by horizontal lines.

A sample skew-T diagram is provided in Figure 1. Temperature and dew point temperature observations are plotted on the chart upward from the bottom, with the former (latter) typically given by a solid red (green) line. Wind speed and direction on a skew-T diagram are typically not plotted directly on the chart but, rather, to one side of the chart (typically the right side).

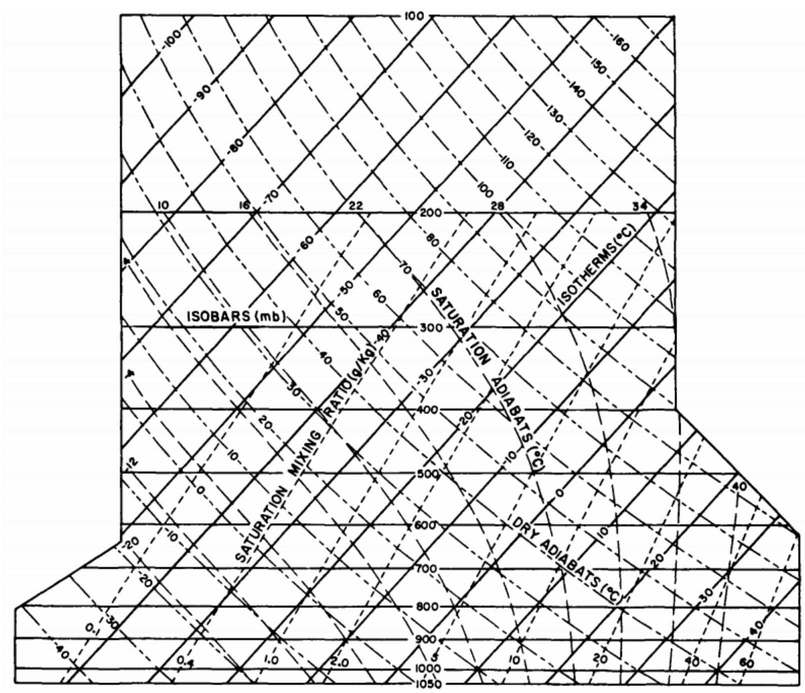

Figure 1. A sample skew-T/ln-p diagram. Isotherms are depicted by the solid skewed lines, isobars by the solid horizontal lines, mixing ratio lines by the short dashed lines, dry adiabats by the alternating short-long dashed lines, and saturated adiabats by the long dashed lines. Reproduced from "The Use of the Skew T, Log P Diagram in Analysis and Forecasting", their Figure 1d.

There are five desired qualities about thermodynamic diagrams that a skew-T diagram meets:

- The important isopleths are straight, or nearly so, rather than curved. See Figure 1 above.
- A large angle between isotherms and adiabats. This helps facilitate estimating atmospheric stability, as we will illustrate in our next lecture.
- There exists a direct proportionality between thermodynamic energy and the area between two lines on the chart, and this proportionality is approximately equal over the entire chart. We will also illustrate this in our next lecture.
- Vertical profiles of temperature, dew point temperature, and wind speed and direction can be plotted through the entirety of the troposphere and through the tropopause.
- The vertical coordinate of the diagram, ln p, approximates the vertical coordinate of the atmosphere. Indeed, ln p is approximately inversely proportional to –z.

Definitions

Before we can describe how we can use a skew-T diagram to draw atmospheric inferences, whether at a given location or over some larger region, we must first consider a bit of terminology that will help us along the way.

Lapse Rate

Lapse rate refers to the change in some quantity with respect to height. Thus, the temperature lapse rate, or Γ, refers to the change of temperature with height,

$$\Gamma = -\frac{\partial T}{\partial z} \tag{1}$$

Temperature typically decreases with increasing height above sea level. Given the leading negative sign on the right-hand side of (1), this situation is characterized by $\Gamma > 0$. In the case where temperature increases with increasing height above sea level, $\Gamma < 0$.

The slopes of the dry and saturated adiabats on a skew-T diagram are given by the dry adiabatic lapse rate (Γ_d) and saturated adiabatic lapse rate (Γ_s), respectively. The dry adiabatic lapse rate Γ_d is equal to g/c_p = 9.81 m s^{-2}/1005.7 J kg^{-1} K^{-1} = 9.75 K km^{-1}. Given that 1 K and 1°C changes in temperature are equivalent, Γ_d is often expressed as 9.75°C km^{-1} or, rounding to two significant figures, 9.8°C km^{-1}. The saturated adiabatic lapse rate Γ_s is smaller than the dry adiabatic lapse rate given the release of latent heat to the environment as water vapor condenses or deposits, or as liquid water freezes. The precise value of Γ_s varies as a function of water vapor content, pressure, and air temperature but is typically between 6-7°C km^{-1} in the lower-to-middle troposphere. Note how the slope of the saturated adiabats in Figure 1 is steeper than that of the dry adiabats.

If one lifts an air parcel along an adiabat, the dry adiabatic lapse rate describes how that air parcel's temperature will change upon ascent *so long as it remains unsaturated*. It also describes how any air parcel's temperature will change upon descent. The saturated adiabatic lapse rate describes how that air parcel's temperature will change upon ascent after it has become saturated.

In most situations, the temperature (or environmental) lapse rate is smaller than or equal to the dry adiabatic lapse rate. When the environmental lapse rate is larger than the dry adiabatic lapse rate, the environmental lapse rate is said to be *superadiabatic* (i.e., larger than adiabatic). This is not common; it typically occurs only at and immediately above a strongly heated surface, particularly during the warm-season (i.e., late spring through early fall).

Layers in the Atmosphere

There exist multiple types of layers bounded by two isobaric surfaces that can be identified using a skew-T diagram. For example, an *adiabatic layer* is one in which $\Gamma \approx \Gamma_d$. This is usually caused by turbulent vertical mixing and is most commonly found within the planetary boundary layer (i.e., near the surface). Turbulent vertical mixing homogenizes, or makes uniform, certain variables over the vertical layer in which the mixing occurs – namely potential temperature and mixing ratio. It is most commonly found near the surface since strong sensible heating (leading to dry convection and localized thermals) and vertical wind shear are the primary drivers of the requisite turbulence.

An *isothermal layer* is one in which $\Gamma = 0$, or one where the observed temperature is constant with increasing height above sea level. In such a case, the observed temperature trace is parallel to the isotherms. Similarly, an *inversion layer* is a layer in which the environmental lapse rate is negative ($\Gamma < 0$), describing a situation where the observed temperature increases with increasing height.

Applications of Skew-T/Ln-p Diagrams

Inversions

There are four different types of inversion layers, or inversions. The first is known as a *subsidence inversion*. When an air parcel sinks, no matter whether it is initially unsaturated or saturated, it warms at the dry adiabatic lapse rate (because no heat is exchanged between the air parcel and its surroundings upon descent). It also becomes drier. Thus, a subsidence inversion typically separates a dry air mass above from a moist air mass below. Subsidence inversions are often found in the subtropics, particularly over water, in conjunction with the descending branch of the Hadley cell circulation. An example of a subsidence inversion is given in Figure 2.

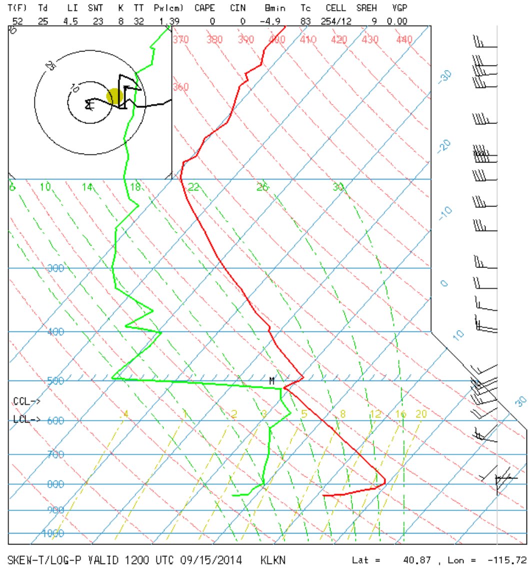

Figure 2. Skew-T diagram valid at 1200 UTC 15 September 2014 at Elko, NV (KLKN). In this diagram, a subsidence inversion is located just below 500 hPa. Note the distinct spreading apart of the temperature (thick solid red) and dew point temperature (thick solid green) traces at this altitude. Dry air, with $T_d << T$, is found above the inversion while moister air, with $T_d \leq T$, is found below the inversion. Image obtained from http://weather.ral.ucar.edu/upper/.

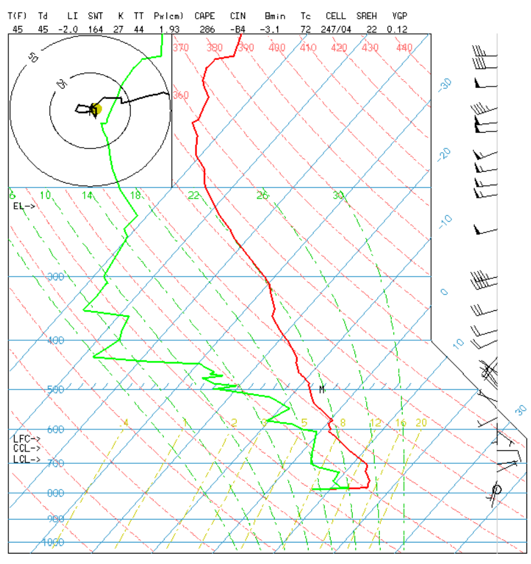

Figure 3. Skew-T diagram valid at 1200 UTC 15 September 2014 at Flagstaff, AZ (KFGZ). In this diagram, a radiation inversion is found right at the surface just above 800 hPa. Note how temperature rapidly increases with height through the radiation inversion, which is confined to a very shallow layer near the surface. Image obtained from http://weather.ral.ucar.edu/upper/.

The second type of inversion is known as a *radiation inversion*. At night, the Earth's surface loses heat to outer space by means of outgoing longwave radiation. When near-surface winds are calm, or nearly so; there is little to no cloud cover overhead; and particularly when the nights are long, a radiation inversion may form due to this heat loss. Radiation inversions are shallow, confined near the surface, and are typically found primarily in the observed temperature trace. An example of a radiation inversion is given in Figure 3.

The third type of inversion is known as a *frontal inversion*. Recall that a) fronts are not lines but are transition zones between two distinct air masses and b) fronts slope upward over relatively cold air. Frontal inversions separate relatively cold, dry air masses below from relatively warm, moist air masses above. The vertical extent of a frontal inversion is limited to the vertical extent of the frontal zone. A schematic cross-section through a frontal zone and accompanying frontal inversion is given in Figure 4.

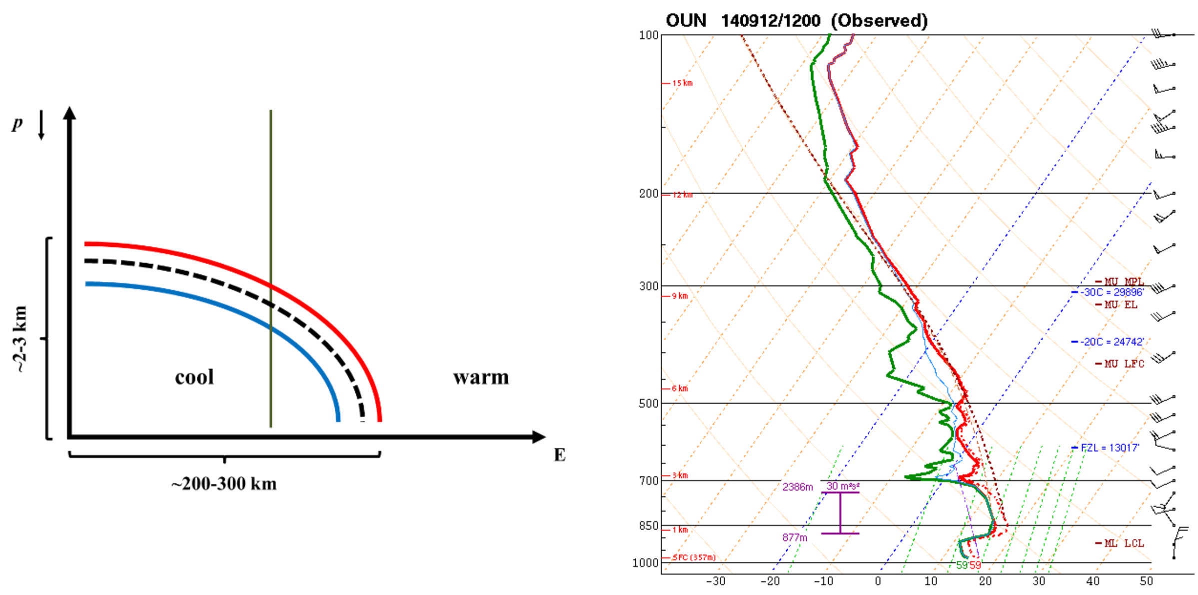

Figure 4. (left) Idealized schematic of a cold frontal zone that separates cool, dry air to the west from warm, moist air to the east. The cold front itself is denoted by the black dashed line while the vertical line denotes the location of the skew-T/ln-p diagram presented at right. (right) Skew-T diagram valid at 1200 UTC 12 September 2014 at Norman, OK (KOUN). In this diagram, a frontal inversion is centered at 900 hPa. Note how both temperature and dew point temperature rapidly increase with height through the frontal inversion, which separates cooler, drier air below the inversion from warmer, moister air above. Image from http://www.spc.noaa.gov/exper/soundings/.

Finally, the *tropopause*, or layer that separates the troposphere from the stratosphere, represents the fourth type of inversion. The tropopause can be identified from the observed temperature trace as the layer over which Γ is $< 2°C\ km^{-1}$ for at least 2 km of depth. Oftentimes, $\Gamma < 0$ through the tropopause. An example tropopause on a skew-T diagram is given in Figure 5.

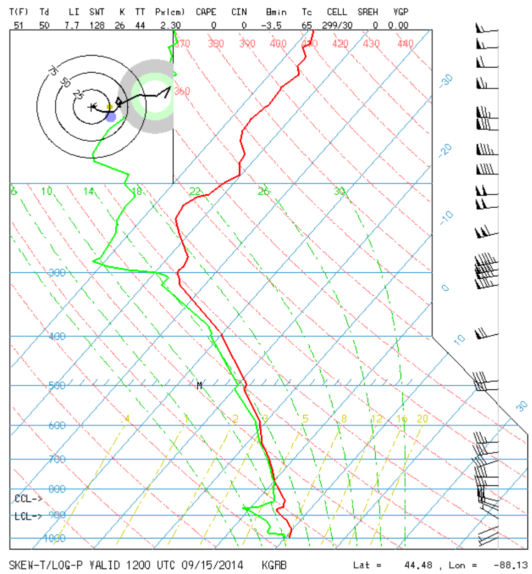

Figure 5. Skew-T diagram valid at 1200 UTC 15 September 2014 at Green Bay, WI (KGRB). In this diagram, the tropopause is found beginning just below 200 hPa. Note how both the temperature and dew point temperature lapse rates are small over a large vertical extent at and above this level. Image obtained from http://weather.ral.ucar.edu/upper/.

Cloud Layers

Cloud layers may be inferred from a skew-T diagram via consideration of the spacing between the temperature and dew point temperature curves. Clouds may be present if the dew point depression, defined as $T - T_d$, is $\leq 5°C$. (Why do we not require the dew point depression to be approximately zero? Although a full discussion of the underlying physics is beyond the scope of this class, note that the temperature at which an air parcel becomes saturated depends on if the water substance is liquid or frozen. For liquid water, saturation does not occur until the dew point temperature equals the temperature; for frozen water, however, saturation can occur at a dew point temperature slightly less than the temperature. More discussion on this point is given in the next paragraph.) A small dew point depression is no guarantee for cloud cover, however! When $T_d \approx T$, such that the air is saturated (or nearly so), clouds are likely present.

Note that on a skew-T diagram, dew point temperature is plotted at all altitudes and temperatures. We do not change from dew point temperature to frost point temperature when the temperature is $\leq 0°C$. Consequently, while the temperature and dew point temperature traces must overlap for there to be saturation when the temperature is above 0°C, the dew point temperature trace may be offset slightly to the left of the temperature trace for saturation to exist when the temperature is $\leq$

0°C. We define the former situation as saturation with respect to liquid water and the latter situation as saturation with respect to ice.

A representative example of cloud layer identification is provided by Figure 5. The temperature and dew point temperature traces overlap above 800 hPa and below 300 hPa, with minor offset at temperatures below 0°C. We can thus reasonably infer that clouds are present over a deep vertical layer bounded by 800 hPa and 300 hPa. Indeed, infrared satellite imagery from approximately the same time as the skew-T profile (Figure 6) indicates the presence of cloud cover at Green Bay with relatively cold cloud tops between -40°C and -50°C, corresponding well to the observed 300 hPa temperature (-46°C). Note, however, that we are unable to infer cloud depth from this image.

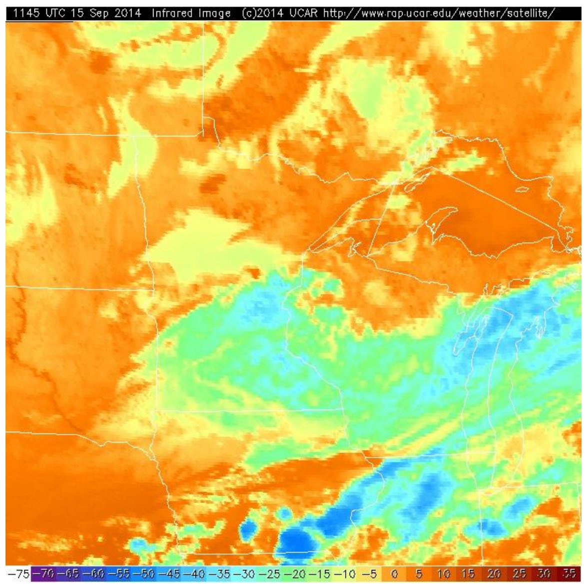

Figure 6. Infrared satellite image from the GOES-EAST geostationary satellite valid at 1145 UTC 15 September 2014. The display window is centered over Duluth, MN. Shaded is cloud-top brightness temperature in °C per the color scale at the bottom of the image. Image obtained from http://weather.ral.ucar.edu/satellite/.

Fronts

The presence of a frontal inversion on a skew-T/ln-p diagram enables us to determine whether a front is present at a given location and, if so, at what altitude it may be found. Please refer to the discussion accompanying Figure 4 for more information on this concept.

Temperature Advection

Wind observations at multiple levels plotted along the side of a skew-T diagram can be used to infer layer-mean temperature advection following principles of the thermal wind. Please refer to our lecture on thermal wind balance and its accompanying examples for more information.

Precipitation Type

Information plotted on a skew-T diagram can be used to estimate the precipitation type that is present or forecast to occur. To illustrate this, let us consider two thought exercises:

- **Saturated**. Here, we assume that the troposphere is saturated between the surface and the level at which precipitation originates. Examples of each situation are given in Figure 7.
    - **Snow**: Snow is the most likely precipitation type if the temperature is below 0°C throughout the entire layer between where snow formed and the surface.
        - **Exception**: An exception to this rule occurs if the level or layer from which precipitation falls is characterized by temperatures no colder than -10°C. Snowflake growth is most efficient when the air temperature is -10°C to -20°C. At temperatures between 0°C and -10°C, supercooled water and small ice pellets dominate. Thus, even if the temperature is < 0°C between the precipitation formation layer and the surface, sleet, graupel, or freezing rain/drizzle are more likely than snow in this scenario.
    - **Rain**: Rain is the most likely precipitation type when there is a sufficiently deep (≥ 1 km) and/or sufficiently warm (≥ 3°C) layer ending at the surface over which the temperature is above 0°C.
    - **Sleet/Graupel**: Sleet or graupel is the most likely precipitation type if precipitation that begins as snow falls through a layer in the lower troposphere over which the temperature is > 0°C and subsequently falls through a deep (~1 km) layer ending at the surface in which the temperature is below 0°C.
    - **Freezing rain**: Freezing rain is the most likely precipitation type if precipitation that begins as snow falls through a layer in the lower troposphere over which the temperature is > 0°C and subsequently falls through a shallow (~100-500 m) layer ending at the surface in which the temperature is below 0°C.
- **Unsaturated**. Here, we assume that the troposphere is unsaturated over one or more layers between the surface and the level at which precipitation originates.
    - To assess precipitation type, we must first consider that some of the precipitation that falls into the sub-saturated layer will evaporate (liquid to vapor) or sublimate (solid to vapor) before reaching the ground. Both phase changes require heat input into the water substance and thus *cool* the surrounding air. The increased water vapor content results in a higher dew point temperature and, eventually, saturation may be reached – but with a colder temperature than initially observed.

- When this is the case, one must first identify the temperature to which the air will be cooled by evaporation and/or sublimation, after which the saturated rules above may be used to identify precipitation type. We will discuss how to identify this temperature, known as the *wet bulb temperature*, shortly.

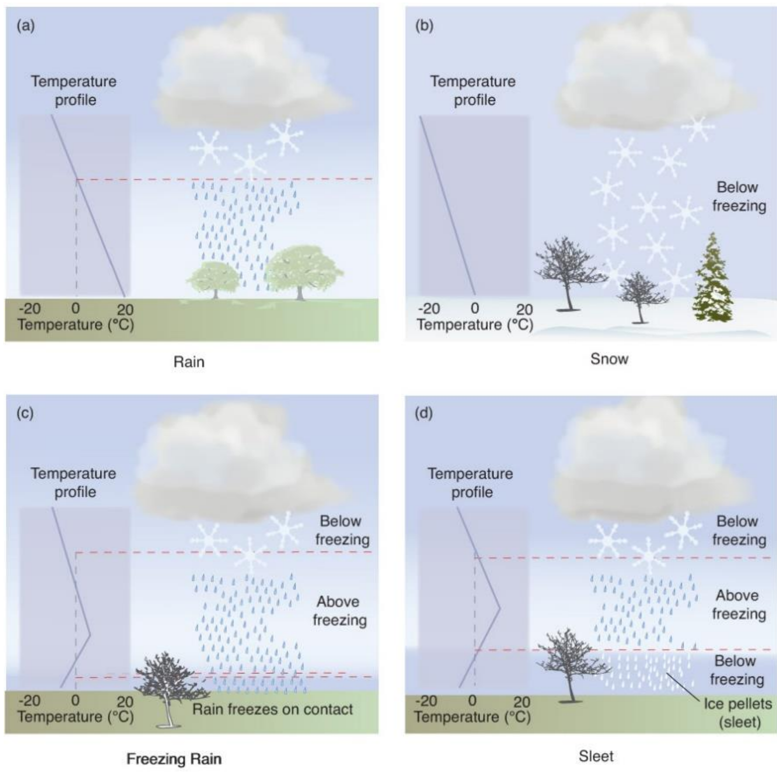

Figure 7. Idealized temperature profiles, assuming a saturated troposphere between surface and the level from which precipitation originates, leading to (a) rain, (b) snow, (c) freezing rain, and (d) sleet or graupel. Figure reproduced from *Meteorology: Understanding the Atmosphere* (4th Ed.) by S. Ackerman and J. Knox, their Figure 4-36.

Derived Thermodynamic Variables and the Skew-T Diagram

Plotted on a skew-T diagram are four elements: temperature, dew point temperature, wind speed, and wind direction. We are particularly interested in the first two of these elements. However, we can readily compute many other thermodynamic variables given these two fields and a properly constructed skew-T diagram. In this discussion, we wish to identify many of the most commonly used derived thermodynamic variables, describe how they can be obtained from a skew-T diagram, and, in many instances, state why we are interested in that given field. Note that there are many more derived thermodynamic variables that may be computed using a skew-T diagram; please refer to Chapter 4 of "The Use of the Skew T, Log P Diagram in Analysis and Forecasting" for details on these variables and their computation.

Mixing Ratio and Saturation Mixing Ratio

Mixing ratio (w) is defined as the ratio of the mass of water vapor contained within a given sample of air to the mass of dry air contained within that sample of air. It is one of several measures by which moisture content in the air may be quantified. To find the mixing ratio at a desired isobaric level, read or interpolate the value of the mixing ratio line that intersects the observed dew point temperature curve at the isobaric level of interest.

Saturation mixing ratio (w_s) is defined as the ratio of the mass of water vapor contained within a given sample of air *if it is saturated* to the mass of dry air contained within that sample of air. The units of mixing ratio and saturation mixing ratio are g kg^{-1}. The same procedure used to find mixing ratio is used to find saturation mixing ratio, except substituting the observed temperature curve for the observed dew point temperature curve.

Consider Figure 10 on page 4-2 of "The Use of the Skew T, Log P Diagram in Analysis and Forecasting." In this example, at 800 hPa, the mixing ratio is ~3.4 g kg^{-1} and the saturation mixing ratio is 6 g kg^{-1}.

Relative Humidity

Relative humidity (RH) is defined as the ratio, in percent, of the mixing ratio to the saturation mixing ratio. In other words, it represents the fraction of water vapor present in the air compared to the water vapor that would be present if the air were saturated. To find the relative humidity, first find the mixing ratio and saturation mixing ratio as described above, then divide w by w_s and multiply the result by 100.

Vapor Pressure and Saturation Vapor Pressure

Vapor pressure (e) is defined as the portion of the total atmospheric pressure (at a given isobaric level) that is contributed by water vapor molecules. It, like mixing ratio, is one of several measures by which moisture content in the air may be quantified. To find vapor pressure, first identify the dew point temperature at the desired isobaric level. Next, follow an isotherm up (or down) from this observation until you reach 622 hPa. Read the value of the mixing ratio line that intersects this isotherm at 622 hPa. This value, in hPa, is your vapor pressure.

Saturation vapor pressure (e_s) is defined as the portion of the total atmospheric pressure (at a given isobaric level) that would be contributed by water vapor molecules if the air sample were saturated. The procedure to find the saturation vapor pressure is identical to that for vapor pressure, except starting by identifying the temperature at the desired isobaric level. This process is illustrated in Figure 12 on page 4-4 of "The Use of the Skew T, Log P Diagram in Analysis and Forecasting."

Why do we read up to 622 hPa to determine vapor pressure and saturation vapor pressure? We can quantify this utilizing the relationship between mixing ratio and vapor pressure (or, equivalently, saturation mixing ratio and saturation vapor pressure). To wit, the saturation mixing ratio is related to the saturation vapor pressure by the following relationship:

$$w_s = 0.622 \frac{e_s}{p - e_s} \tag{2}$$

In (2), w_s has units of g g^{-1} rather than g kg^{-1}. By dividing the left-hand side by 1000, or equivalently multiplying the right-hand side by 1000, we can get w_s in units of g kg^{-1}. Further, typically p is much larger than e_s such that $p - e_s \approx p$. Substituting, (2) becomes:

$$w_s = 622 \frac{e_s}{p} \tag{3}$$

When we described how to obtain the saturation vapor pressure, we stated to read up (or down) an isotherm until reaching 622 hPa. This means that saturation vapor pressure *does not* change as pressure changes but that it *does* change as temperature changes (i.e., reading up or down a different isotherm). Under this condition, we can plug in any value for p – say, 622 hPa. If we do so, then we find that $w_s = e_s$, noting the different units (w_s in g kg^{-1}, e_s in hPa), the conversions for which we have neglected to include with our 0.622 value.

Potential Temperature

Potential temperature (θ) is the temperature that a sample of air would have if it were brought dry adiabatically (i.e., without heat exchange between an air parcel and its surroundings) to 1000 hPa. To determine potential temperature, identify the temperature at a desired isobaric level and proceed (typically, downward) along the dry adiabat that intersects the temperature curve at that isobaric level until reaching 1000 hPa. The isotherm at 1000 hPa defines the potential temperature. This process is illustrated in Figure 13 on page 4-6 of "The Use of the Skew T, Log P Diagram in Analysis and Forecasting." The units of potential temperature are K.

When diabatic heating is not occurring, potential temperature is conserved following the motion in both the horizontal and vertical directions. This attribute is incredibly beneficial for identifying synoptic-scale areas of ascent and descent, as we will demonstrate next semester. Note that so long as the lapse rate is less than the dry adiabatic lapse rate ($\Gamma < \Gamma_d$), potential temperature increases with increasing altitude above sea level. This implies that potential temperature can be used as an alternative vertical coordinate, the applications of which we will explore next semester.

Wet Bulb Temperature

The wet bulb temperature (T_w) is the lowest temperature to which a sample of air (at constant pressure) can be cooled by evaporating water into it. If the sample of air is saturated, such that no more water can be evaporated into it, then $T_w = T$. Otherwise, since evaporation requires heat input from the surrounding environment (the environment cools), $T_w < T$. However, since evaporation increases water vapor content in the air sample, $T_w > T_d$. Thus, $T_d < T_w < T$. The units of wet bulb temperature are °C or K.

To obtain the wet bulb temperature, identify the dew point temperature at the desired isobaric level. Draw a line upward along the mixing ratio line that intersects this dew point temperature reading. Next, identify the temperature at the desired isobaric level. Draw a line upward along the dry adiabat that intersects this temperature reading until you intersect the line drawn upward along the mixing ratio line. Follow the saturated adiabat that intersects this intersection point down until you reach the isobaric level at which you started. The value of the isotherm at this level gives you the wet bulb temperature. This process is illustrated in Figure 14 on page 4-7 of "The Use of the Skew T, Log P Diagram in Analysis and Forecasting."

Equivalent Temperature and Equivalent Potential Temperature

Equivalent temperature (T_e) is the temperature that a sample of air would have if all of its moisture were condensed out by a pseudoadiabatic process (one in which all liquid or solid water substance falls out of an air parcel immediately after it forms) and then is brought back dry adiabatically to its original pressure. Equivalent potential temperature (θ_e) is identical to T_e, except the sample of air is brought dry adiabatically to 1000 hPa rather than its original pressure. The units of both equivalent temperature and equivalent potential temperature are K.

To obtain T_e, identify the dew point temperature at the desired isobaric level. Draw a line upward along the mixing ratio line that intersects this dew point temperature reading. Next, identify the temperature at the desired isobaric level. Draw a line upward along the dry adiabat that intersects this temperature until you intersect the line drawn upward along the mixing ratio line. From here, follow the saturated adiabat that intersects this intersection point upward until you reach the isobaric level at which the dry and saturated adiabats become parallel to each other (i.e., $\Gamma_d \approx \Gamma_s$,). Finally, follow the dry adiabat that intersects this point downward until you reach the isobaric level at which you started. The value of the isotherm at this level gives you the equivalent temperature. You obtain the equivalent potential temperature if you continue downward to 1000 hPa and read the value of the intersecting isotherm. This process is illustrated in Figure 15 on page 4-8 of "The Use of the Skew T, Log P Diagram in Analysis and Forecasting."

Note that $T_e > T$ and, by extension, $\theta_e > \theta$. Why would we expect this to be the case? Both T_e and θ_e involve condensing moisture out of an air parcel. Condensation, freezing, and deposition cause the water substance to lose heat to its surrounding environment. This release of latent heat to the environment causes the environment to warm, thus resulting in $T_e > T$ and, by extension, $\theta_e > \theta$. Only in the case where there is no water vapor in the atmosphere ($w = 0$) does $T_e = T$ and $\theta_e = \theta$.

Equivalent potential temperature is approximately conserved following the motion, whether or not an air parcel is initially saturated, under pseudoadiabatic conditions. It provides a measure of both the potential warmth and moistness of a sample of air. While we will not make extensive use of equivalent potential temperature in this class it is beneficial to keep this dialogue in mind for other situations in which it might be beneficial.

Virtual Temperature

Virtual temperature (T_v) is the temperature at which dry air (w = 0) would have the same density as a observed air sample (w > 0). Virtual temperature is related to temperature T and mixing ratio w by the following equation:

$$T_v = T(1 + 0.6w) \quad (4)$$

Note that w here is expressed in g g^{-1} and not g kg^{-1}. The units of virtual temperature are K.

For the case where $w = 0$ (no water vapor present in the atmosphere), $T_v = T$. Otherwise, where $w > 0$, $T_v > T$. Why? Warm air and moist air are both less dense than cool air and dry air. Since our observed air sample contains some water vapor, it is less dense than dry air independent of the temperature. For the two to have equal density, the temperature of the dry air must be warmer such that it is less dense than the cool air. In other words, the effects of warm vs. cool and moist vs. dry cancel each other out. Thus, the virtual temperature is always greater than the temperature, albeit by a small amount (1-5 K) given typically observed values of w ($w < 0.04$ g g^{-1}). The temperature and virtual temperature are nearly identical at and above 500 hPa, where w is typically ~ 0 g g^{-1}.

Precipitable Water

Precipitable water (PW) is the total water vapor contained within a vertical column of air (over a unit area of 1 m^2) bounded by two isobaric levels. The total precipitable water (TPW) is the special case where the two isobaric levels are the ground and the top of the atmosphere, or the tropopause. The units of precipitable water and total precipitable water are kg m^{-2} or, in equivalent notation, mm. Mathematically, total precipitable water can be expressed as:

$$TPW = \frac{1}{g} \int_{p_{sfc}}^{p_{trop}} w\,dp \quad (5)$$

In (5), w is the mixing ratio, g is the gravitational constant (9.81 m s^{-2}), p_{trop} is the pressure at the tropopause, and p_{sfc} is the pressure at the surface. High values of TPW are associated with greater moisture content and, in precipitating regions, greater liquid equivalent values of accumulated precipitation. However, because it represents only an instantaneous vertical snapshot (and thus does not consider processes such as advection that can influence local moisture amounts), the TPW does not provide a minimum or maximum bound on the amount of precipitation that can or does fall during a precipitation event.