Mesoscale Meteorology: Sounding and Stability Analysis

7, 9 February 2017

Introduction

This is not intended to be a comprehensive overview of atmospheric thermodynamics, including concepts such as the thermal wind, or sounding and stability analysis. Primers on several of these topics are available at:

- Thermal Wind: http://derecho.math.uwm.edu/classes/SynI/ThermalWind.pdf
- Sounding Analysis: http://derecho.math.uwm.edu/classes/SynI/SoundingAnalysis.pdf
- Stability Analysis: http://derecho.math.uwm.edu/classes/SynI/Stability.pdf

Important Definitions

The *dry adiabatic lapse rate* defines the rate at which the temperature of an air parcel changes due to the effects of changing pressure – i.e., compression upon descent or expansion upon ascent. It explicitly assumes no heat exchange between the air parcel and its surrounding environment. The dry adiabatic lapse rate is equal to g (9.81 m s^{-2}) divided by the specific heat for a constant pressure process (1005.7 J kg^{-1} K^{-1}), or approximately 9.8 K km^{-1}.

The *potential temperature* defines the temperature that an air parcel would have if it were brought dry adiabatically (i.e., with no heat exchange between it and the surrounding environment) to some reference pressure p_0, typically taken to be 1000 hPa. It is conserved following the motion for dry adiabatic motions. Potential temperature is mathematically related to temperature and pressure by *Poisson's equation*,

$$\theta = T\left(\frac{p_0}{p}\right)^{R/c_p}$$

We typically substitute R_d (287.04 J kg^{-1} K^{-1}) for R in the above, such that R_d/c_p = 0.2854.

The *water vapor mixing ratio* defines the ratio between the mass of water vapor present in an air parcel and the total mass of dry air. Since density is defined as the mass per unit volume, the water vapor mixing ratio is typically expressed in terms of density, i.e.,

$$r_v = \frac{\rho_v}{\rho_d} = \frac{\varepsilon e}{p - e}$$

The second portion of this relation can be derived by substitution with the ideal gas law. Here, ε is given by the ratio of the gas constants for water vapor and dry air (approximately equal to 0.622). e is the pressure resulting only from water vapor. Water vapor mixing ratio is also conserved for dry adiabatic processes: no water vapor is exchanged between an air parcel and its environment by mixing, entrainment/detrainment, fallout, etc., and there is no phase change from vapor to liquid or solid.

The *ideal gas law*, or *equation of state*, relates pressure, density, and temperature. Generically, it is given by:

$$p = \rho RT$$

R is the gas constant and depends on the composition of the air, particularly its water vapor content. A more common expression of the ideal gas law is given by:

$$p = \rho R_d T_v$$

R_d is the gas constant for dry air and T_v is the *virtual temperature*. The virtual temperature is the temperature that a sample of dry air must have to have the same density as a sample of moist air at the same pressure. Mathematically, it is given by:

$$T_v = T\left(\frac{1+\frac{r_v}{\varepsilon}}{1+r_v}\right) \approx T(1+0.61r_v)$$

The virtual temperature is always larger than the air temperature, typically by ~1 K near the surface and approaching ~0 K at higher altitudes where water vapor mixing ratio approaches zero. Virtual temperature must be larger than air temperature because water vapor has less mass (and thus lower density) than dry air. We will revisit virtual temperature later when we discuss buoyancy.

The *dew point temperature* is the temperature at which the air becomes saturated if cooled while pressure and water vapor mixing ratio are held constant. Contrast this with *wet bulb temperature*, which is the temperature at which the air becomes saturated if cooled by evaporating water into it (i.e., water vapor mixing ratio is not held constant).

[Recall that evaporation (liquid to vapor), sublimation (solid to vapor), and melting (solid to liquid) all require heat input, such that the surrounding air cools under each process.]

Dew point temperature is approximated by:

$$T_d \approx \frac{243.5}{\left[\frac{17.67}{\ln\left(\frac{e}{6.112}\right)}-1\right]}$$

This expression provides T_d with units of °C and assumes e has units of hPa.

Wet bulb temperature can be determined iteratively or graphically using a skew T-ln p diagram.

Recall that for dry adiabatic processes, air parcels cool as they ascend but their water vapor content does not change. Consequently, if an air parcel ascends over a sufficiently large vertical distance, it will become saturated. Upon saturation, phase changes occur (vapor to liquid and/or solid) and condensate forms. For vapor to liquid (condensation), vapor to solid (deposition), or liquid to solid (freezing) phase changes, the water substance cools while releasing heat to its surroundings. This heat release violates the assumptions underlying dry adiabatic motions; there is now heating of the surrounding air. In general, this defines a moist or saturated adiabatic process.

There are two types of saturated adiabatic processes. The first, said to be reversible, assumes that total water (vapor plus condensate) is conserved following the motion; i.e., no condensate fallout. The second, said to be irreversible, assumes that condensate falls out as soon as it forms. We define the former as a *reversible moist adiabatic process* and the latter as a *pseudoadiabatic* process. In the former, latent heat released can be absorbed by the air and/or the condensate; in the latter, latent heat released can only be absorbed by the air because the condensate falls out. In both cases, the rate of cooling of an ascending saturated air parcel is less than for a dry adiabatic process; however, the rate of cooling is slightly larger for a pseudoadiabatic process because the condensate does not remain within the ascending air parcel.

Expressions for the moist adiabatic lapse rate can be derived from the first law of thermodynamics and Clausius-Clapeyron equation and are given by the following for reversible moist adiabatic and pseudoadiabatic processes, respectively:

$$\Gamma_m = \frac{g}{c_{pd} + r_t c_1} + \frac{1}{c_{pd} + r_t c_1} \frac{d l_v r_v}{dz}$$

$$\Gamma_m = \frac{g}{c_{pd} + r_v c_1} + \frac{1}{c_{pd} + r_v c_1} \frac{d l_v r_v}{dz}$$

In the first, r_t is the total water mixing ratio (vapor + condensate). It is larger than the water vapor mixing ratio, which decreases with further ascent beyond saturation as vapor is converted to liquid and/or solid. c_1 is the specific heat of liquid water at constant pressure and is approximately 4200 J kg^{-1} K^{-1}. l_v is the specific latent heat of vaporization; it is a function of temperature and equals 2.501 x 10^6 J kg^{-1} at 0°C. The distinction between the two lapse rates appears in the denominators of each term: r_t for a reversible moist adiabatic process and r_v for a pseudoadiabatic process. As $r_v \leq r_t$, the pseudoadiabatic moist adiabatic lapse rate is larger than the reversible moist adiabatic lapse rate.

In reality, condensate neither entirely falls out upon formation, nor is perpetually maintained. Thus, neither the pseudoadiabatic or reversible moist adiabatic lapse rates are exactly correct. In general, we assume pseudoadiabatic processes, and the moist adiabats on a skew *T*-ln *p* diagram are drawn accordingly. The moist adiabatic lapse rate is approximately 6 K km^{-1} in the lower troposphere and approaches the dry adiabatic lapse rate at higher altitudes.

The *equivalent potential temperature* is the moist analog to potential temperature. It is conserved for dry adiabatic and approximately conserved for pseudoadiabatic moist adiabatic processes. It can be expressed mathematically by:

$$\theta_e = T\left(\frac{p_0}{p}\right)^{\frac{R_d}{c_p}(1-0.28 r_v)} \exp\left[r_v\left(1+0.81 r_v\right)\left(\frac{3376}{T^*} - 2.54\right)\right]$$

Here, T^* is the temperature of the air parcel when it first becomes saturated due to ascent (i.e., the air temperature at the lifting condensation level).

Buoyancy

The vertical momentum equation, in the absence of friction and neglecting the Coriolis force, can be written in terms of a buoyancy force and a vertical pressure gradient force, i.e.,

$$\frac{dw}{dt} = -\frac{1}{\rho}\frac{\partial p'}{\partial z} + B$$

This describes the change in vertical velocity *w* following the motion (note that the total derivative is defined here by the lowercase script *d*). Buoyancy *B* is given by:

$$B = -\frac{\rho'}{\rho} g$$

It is a function of density variations within a column. In both equations above, prime terms denote deviations from a horizontally-homogeneous balanced (generally hydrostatic) base state. Prime terms are said to be small relative to the total field. We can use this to substitute for ρ with its base-state version $\overline{\rho}$ in the above; doing so and applying the logarithm-differentiated version of the ideal gas law, we obtain:

$$B = -\frac{\rho'}{\overline{\rho}} g \approx \left(\frac{T_v'}{\overline{T_v}} - \frac{p'}{\overline{p}}\right) g$$

This definition assumes no hydrometeors are present within the vertical column. In precipitating phenomena, there are hydrometeors falling out of the column. If we assume that this fall out occurs at the hydrometeor terminal velocity, the downward acceleration (i.e., buoyancy reduction) due to hydrometeor drag is equal to gr_h, where r_h is the hydrometeor mixing ratio. Incorporating this into the definition for buoyancy, we obtain:

$$B = \left(\frac{T_v'}{\overline{T_v}} - \frac{p'}{\overline{p}} - r_h\right) g$$

In general, the deviation in pressure from the base-state pressure is much less than the deviation in virtual temperature from the base-state virtual temperature. For a vertically-moving air parcel, if we assume that the base-state virtual temperature is that of the environment, virtual temperature is that of the vertically-moving air parcel, and perturbations are the difference of the two, we obtain:

$$B = \frac{T_v'}{\overline{T_v}} g = \frac{T_v^{parcel} - T_v^{env}}{T_v^{env}} g$$

In other words, *buoyancy is directly proportional to the magnitude of the difference between the virtual temperature of the ascending air parcel and that of its environment*! This underlies the definitions of *convective available potential energy* (CAPE) and *convective inhibition* (CIN). The former, or *positive area*, is defined by the integrated buoyancy over the vertical column where the air parcel's virtual temperature is greater than that of its environment. The latter, or *negative area*, is defined by the integrated buoyancy over the vertical column from the level from which the air parcel is lifted (here, taken to be the surface) to where the air parcel becomes positively buoyant.

$$CAPE = \int_{LFC}^{EL} B dz \approx g \int_{LFC}^{EL} \frac{T_v^{parcel} - T_v^{env}}{T_v^{env}} dz$$

$$CIN = -\int_{0}^{LFC} B dz \approx -g \int_{0}^{LFC} \frac{T_v^{parcel} - T_v^{env}}{T_v^{env}} dz$$

In the above, *LFC* is the level of free convection, or level at which the air parcel becomes positively buoyant (i.e., parcel virtual temperature becomes warmer than that of its surroundings), and *EL* is the equilibrium level, or level beyond the LFC at which the air parcel ceases to be positively buoyant (i.e., parcel virtual temperature is no longer warmer than that of its surroundings).

CAPE is proportional to the kinetic energy that a parcel can gain from its environment as a result of the contribution of buoyancy to the vertical acceleration. CIN is proportional to the work that must be done against the (stable) stratification to lift an air parcel to its LFC; in stability analysis, this work is often assumed to be accomplished by lifting, although heating (in isolation and/or with lifting) can also accomplish this work.

If we assume no vertical pressure gradient force in the vertical momentum equation, i.e.,

$$\frac{dw}{dt} = B$$

an expression for the maximum vertical velocity w_{max} due to realized buoyancy can be obtained:

$$w_{\max}^2 = 2\int_{LFC}^{EL} B dz = 2CAPE$$

In other words, the maximum-possible vertical velocity that an ascending parcel can achieve due to buoyancy is directly proportional to the CAPE. Larger values of CAPE imply stronger-possible updrafts. But, we have neglected several things in obtaining this expression, a downward-directed vertical pressure gradient force and hydrometeor drag chief among them. Thus, we define this w_{max} as the maximum-possible, rather than maximum-realized, vertical velocity. We will return to this idea later in this lecture.

Sounding Analysis: Thermodynamics

An *air parcel* is a self-contained, infinitesimally small volume of air. In effect, it has no dimension. When an air parcel moves vertically within the atmosphere, it is said to not disturb the surrounding air. The consideration of an infinitesimally small air parcel enables us to, as we did above, neglect the vertical pressure gradient force in the vertical momentum equation. This neglect, coupled with neglecting hydrometeor drag and the Coriolis force, underlies *parcel theory*, which forms the basis for conventional sounding and stability analysis.

We now wish to define, or expand on earlier definitions, of common sounding and stability-related parameters, particularly as it relates to determining each on a skew *T*-ln *p* diagram.

The *convection condensation level* (CCL) is the height to which an air parcel, if heated from below (i.e., by dry convection) will rise adiabatically until it becomes saturated. In summertime, the CCL typically defines the base of fair-weather cumulus clouds that dot the afternoon landscape on warm and sunny days. Identification of the CCL on a skew *T*-ln *p* diagram proceeds as follows:

- Identify the dew point temperature at a desired height (e.g., the surface).
- Follow a constant mixing ratio line upward until you intersect the observed temperature profile. The height at which this occurs defines the CCL.

The *convective temperature* (T_c) defines the temperature that must be reached in order for an air parcel to reach its CCL. Identification of T_c on a skew *T*-ln *p* diagram proceeds as follows:

- Identify the CCL as above.
- Following a dry adiabat, descend from the CCL to the original parcel level. T_c is defined by the isotherm intersecting the dry adiabat at the original parcel level.

Because the CCL and T_c imply dry convection, the originating parcel level is almost always taken to be the surface for both.

The *lifting condensation level* (LCL) is the height at which an air parcel, if mechanically forced to ascend, will become saturated. Identification of the LCL on a skew *T*-ln *p* diagram proceeds as follows:

- Identify the temperature and dew point temperature at a desired height.
- From the temperature, ascend parallel to a dry adiabat.
- From the dew point temperature, ascend parallel to a constant mixing ratio line.
- The intersection of the dry adiabat and constant mixing ratio lines above defines the LCL.

The *level of free convection* (LFC) is the height at which an ascending air parcel becomes less dense, or positively buoyant, than its surroundings or environment. This is commonly interpreted as the height at which an ascending air parcel becomes warmer (in terms of temperature) than its environment. From this interpretation, identification of the LFC on a skew *T*-ln *p* diagram proceeds as follows:

- Identify the LCL as above.
- Follow a moist adiabat (a pseduoadiabat) upward from the LCL until the air parcel becomes warmer than the environmental temperature. This height defines the LFC.

The LFC is always found at or above the LCL. It is found at the height of the LCL only when the air parcel is immediately positive buoyant upon reaching the LCL. If the air parcel never becomes positively buoyant, then the LFC is undefined.

The *equilibrium level* (EL) is the height at which an ascending air parcel is no longer positively buoyant relative to its surroundings or environment. This is commonly interpreted as the height at which an ascending air parcel ceases to be warmer (in terms of temperature) than its environment. From this interpretation, identification of the EL on a skew T-ln p diagram proceeds as follows:

- Identify the LFC as above.
- Follow a moist adiabat upward from the LFC until the air parcel ceases to be warmer than the environmental temperature. This height defines the EL.

It is possible, although not common, for a parcel to have multiple LFCs and ELs (e.g., when there exists an elevated temperature inversion). Once the LCL, LFC, and EL have been identified, CAPE and CIN can easily be identified from the area cut out between the environmental and air parcel temperature traces:

- CAPE is directly proportional to the area between the LFC and EL.
- CIN is directly proportional to the area between the original parcel level and the LFC.

In the above, we have been purposefully vague as to what defines the original parcel level when finding the LCL, LFC, EL, CAPE, and CIN. This is because the choice of original parcel level is not always immediately clear. There are four common choices for the original parcel level:

- *Surface-based*, or a parcel that originates at the surface.
- *Mixed layer*, or a near-surface air parcel with temperature and dew point representative of the entire planetary boundary layer (PBL). The PBL is typically viewed, at least in warm-season daytime environments, as a mixed layer given the vertical homogenization of θ and r_v that commonly results from dry convection-induced turbulent vertical mixing.
- *Most unstable*, or the parcel within the lower troposphere with the maximum amount of CAPE.
- *User-defined*, or a parcel that originates at some user-chosen level (e.g., perhaps the level where a synoptic analysis indicates that forcing for ascent is comparatively maximized).

Unless otherwise stated, we will assume a surface-based parcel in our consideration of atmospheric stability. The above methods for determining the LFC, EL, CAPE, and CIN are all approximations as they are based off of the parcel and environment temperature, rather than virtual temperature, profiles. However, these quantities are all formally functions of virtual temperature, and thus the virtual temperature should be used to identify these quantities. This results in some modifications to the above-stated processes.

First, compute the environmental virtual temperature. Next, we need to compute the parcel virtual temperature. Below the LCL, this is done by first identifying the original parcel level, computing its virtual temperature, and following a dry adiabat upward to the LCL. Above the LCL, this is done by identifying the moist adiabat along which an air parcel would ascend as before and then computing the virtual temperature for that parcel trace (where r_v is equal to its saturated, rather than environmental, value in doing so). Note in this process that the LCL is not computed different than before.

The LFC is the height where a parcel's virtual temperature becomes warmer than its surroundings. The EL is the height where a parcel's virtual temperature stops being warmer than its surroundings.

CAPE and CIN are the areas between the newly-defined LFC and EL and original parcel level and newly-defined LFC, respectively. Figure 2.9 from the course text provides a graphical depiction of the differences involved in identifying the LFC, EL, CAPE, and CIN between the conventional and more accurate methods, illustrating the potential importance to assessments of deep, moist convection potential in so doing. Most computed-based CAPE and CIN computation routines use virtual temperature, while most manual analyses use temperature, in their analyses.

Stability Analysis: Parcel Theory

Parcel theory can be used to infer atmospheric stability. In the discussion above, we stated that air parcels ascend along a dry adiabat until saturated, after which they ascend along a moist adiabat. If we consider *infinitesimally small* upward parcel displacements, we find that stability depends on (a) whether a parcel is saturated or unsaturated and (b) its lapse rate relative to the dry and/or moist adiabatic lapse rates, i.e.,

	Unsaturated	**Saturated**
Stable	$\Gamma < \Gamma_d$	$\Gamma < \Gamma_m$
Neutral	$\Gamma = \Gamma_d$	$\Gamma = \Gamma_m$
Unstable	$\Gamma > \Gamma_d$	$\Gamma > \Gamma_m$

Γ is the environmental lapse rate, Γ_d is the dry adiabatic lapse rate, and Γ_m is the moist adiabatic lapse rate. Under unsaturated conditions, if the environmental temperature decreases less rapidly with height than the dry adiabatic lapse rate, the environment is absolutely stable to upward displacements; an air parcel ascending along a dry adiabat would be cooler than the environment, and work must be done to overcome the inhibition. If the environmental temperature decreases at the dry adiabatic lapse rate, the environment is neutrally stable. Finally, if the environmental temperature decreases more rapidly with height than the dry adiabatic lapse rate, the environment is absolutely unstable to upward displacements; an air parcel ascending along a dry adiabat would be warmer than the environment and thus positively-buoyant. This latter scenario is uncommon; it typically only occurs near the surface under strong warm-season daytime heating.

If saturated, when the environmental temperature decreases less rapidly with height than the moist adiabatic lapse rate, the environment is absolutely stable to upward displacements; an air parcel ascending along a moist adiabat would be cooler than the environment. When the environmental temperature decreases at the moist adiabatic lapse rate, the environment is neutrally stable. Finally, if the environmental temperature decreases more rapidly with height than the moist adiabatic lapse rate, the environment is absolutely unstable to upward displacements; an air parcel ascending along a moist adiabat would be warmer than the environment and thus positively-buoyant.

The term *conditionally unstable* is often used to describe a situation where the environmental lapse rate is between the dry and moist adiabatic lapse rates. In that setting, the *condition* is saturation: it is unstable if saturated and stable if unsaturated.

Note that an environment can contain large CAPE but still be absolutely stable to upward parcel displacements. In this setting, there is abundant potential energy that may be converted to kinetic energy, but it can be released if and only if the air parcel can be lifted to its LFC. Thus, the presence of large CAPE does not imply an unstable situation. Stability depends upon the evaluation above, and in particular its vertical integration (as becomes manifest in both CAPE and CIN calculations).

Stability Analysis: Limitations of Parcel Theory

There have been many assumptions made throughout this lecture. Each of these contribute to errors in stability analysis through the lens of parcel theory, although it is generally not straightforward to quantify these errors. For instance, we assumed an infinitesimally small air parcel that does not disturb or modify the surrounding environmental air. But, ascending air requires compensating subsidence, which can influence the environment in terms of buoyancy and/or pressure. This can alter interpretations of stability.

Before, we assumed buoyancy as the only contribution to vertical motions (i.e., we simplified the vertical momentum equation such that only buoyancy remained). This resulted in the neglect of the vertical pressure gradient force, including both hydrostatic and non-hydrostatic components from buoyancy and dynamical effects. Here, let us consider only the buoyancy contribution to the vertical pressure gradient force. A positively-buoyant layer is warmer than its surroundings. In a hydrostatic atmosphere, the thickness of a layer is directly proportional to the layer-mean T_v. A positively-buoyant layer thus has greater thickness than its surroundings, implying anomalously high pressure above (as isobaric surfaces are displaced upward) and anomalously low pressure below (as isobaric surfaces are displaced downward) the positively-buoyant layer. The resulting anomalous pressure gradient force is directed downward, opposing the buoyancy force.

There is also a relationship between updraft width and the magnitude of this anomalous pressure gradient force. Wider updrafts result in the virtual temperature anomaly being much wider than it is tall (or deep). In this case, the atmosphere tends toward hydrostatic balance (describing a balance between gravity and the vertical pressure gradient force with no vertical parcel accelerations), such that the vertical pressure gradient force must entirely counterbalance the buoyancy force.

Parcel theory also assumes that an ascending parcel is entirely self-contained, isolated from (i.e., no interaction with) the surrounding air. In reality, this does not occur. Mixing of the updraft's air with the environment occurs for momentum, temperature, and moisture. The updraft can be viewed as a local maximum in upward momentum, temperature, and moisture; environmental values of each of these quantities are lower. Thus, mixing with the environment, or *entrainment*, weakens the updraft. Lower upward momentum directly slows the updraft. Entrainment reduces the relative warmth and CAPE of the ascending parcel by both cooling the ascending parcel and warming its environment. Entrainment of lower moisture air also reduces the relative warmth and CAPE of the ascending parcel due to cooling resulting from evaporation of updraft condensate. Entrainment is anecdotally most evident with updrafts in relatively dry, relatively low CAPE environments.

There is also a relationship between updraft width and entrainment. Wider updrafts shield the core updraft from the environmental air to a greater extent than is seen with narrower updrafts. This is particularly true for tilted updrafts. This effect opposes that of the vertical pressure gradient force: here, wider updrafts are favored. This implies that deep, moist convection favors updrafts of an intermediate width: not too large to increase the buoyancy-related vertical pressure gradient force, not too small to increase the deleterious effects of entrainment.

In general, parcel theory also neglects the presence of hydrometeors. In practice, hydrometeors act as a drag on upward parcel accelerations to a magnitude proportional to hydrometeor mass. When we assumed pseudoadiabatic ascent when saturated, we assumed that all hydrometeors fall out as they are created, such that CAPE derived from such principles is not affected by hydrometeors. In practice, hydrometeors neither fall out immediately nor are maintained perpetually in an updraft. The net effect of hydrometeors under the pseudoadiabatic assumption is to reduce buoyancy, and thus reduce CAPE, from that considered in their absence.

The pseudoadiabatic moist adiabatic lapse rate defined earlier in this lecture is a function of l_v, the latent heat of vaporization. It thus only considers vapor to liquid phase changes; the solid phase and accompanying latent heat release is altogether neglected. As the latent heat of fusion is almost an order of magnitude smaller than the latent heat of vaporization, this approximation generally works well. If latent heat release associated with the solid phase were included, an ascending air parcel would have a slightly warmer temperature than if it was neglected. Thus, the neglect of the solid phase slightly reduces the inferred buoyancy.

Stability Analysis: Potential Instability and the Layer Method

An alternative means of stability analysis is given by the *layer method*. In contrast to the parcel method, which considers only a parcel at a single originating altitude, the layer method considers an entire vertical layer. It quantifies how the lapse rate of a layer changes as it ascends or descends. Traditionally, the layer method considers two parcels, one each at the top and bottom of a layer. It follows the tenets of the parcel method in so doing: parcels ascend at the dry adiabatic lapse rate until saturated, after which they ascend at the moist adiabatic lapse rate. Both the top and bottom of the layer are lifted by an equivalent amount – usually on the order of 50-200 hPa – and then the new lapse rate is evaluated. Fig. 3.3 from the course text and Figs. 2-6 from the Stability Analysis lecture notes linked at the outset of this document provide examples of the layer method in action.

So long as the entire layer starts and remains unsaturated upon lifting, lifting an initially unstable layer makes it less unstable. Lifting an initially stable layer makes it less stable. Lifting an initially neutral stability layer does not change its stability. The situation is a bit more complex when all or part of the layer becomes saturated upon lifting.

General stability criteria, independent of saturation, for the layer method are given by:

Stabilizing: $\frac{\partial \theta_e}{\partial z} > 0$ **Destabilizing:** $\frac{\partial \theta_e}{\partial z} < 0$ **No Change:** $\frac{\partial \theta_e}{\partial z} = 0$

These represent *potential stability* criteria, indicating that a layer has the *potential* to become more or less stable if it ascends (or, technically, descends). Why they are defined in terms of equivalent potential temperature becomes clear if we consider how equivalent potential temperature can be identified on a skew *T*-ln *p* diagram:

- For a parcel at any level, find its LCL as described above.

- From the LCL, ascend following a moist adiabat to sufficient altitude such that moist and dry adiabats parallel each other.
- Descend along a dry adiabat from this altitude to the reference pressure (1000 hPa). The value of the isotherm intersecting the dry adiabat at 1000 hPa defines θ_e.

The first two steps describe the method by which air parcels ascend: first along a dry adiabat, then along a moist adiabat. The last step is only needed to identify the specific θ_e value of the air parcel. We can generalize this by considering only the first and part of the second steps. Once the LCL has been identified, determine the moist adiabats along which ascent would occur. If the moist adiabat to be followed from the layer's bottom lies to the right of the moist adiabat to be followed from the layer's top, θ_e decreases with height (destabilizing). If the opposite is true, θ_e increases with height (stabilizing). If the two moist adiabats are identical, then θ_e is constant with height.

The layer method can help describe how instability is released with mesoscale precipitation band formation in the precipitation shields of synoptic-scale cyclones. It can also help describe how subsidence inversions form in proximity to middle-to-upper tropospheric anticyclones. While it can help describe the weakening of an inversion due to ascent, it is unlikely that this is the primary means by which destabilization occurs preceding deep, moist convection formation.

Sounding Analysis: Hodographs and Horizontal Vorticity

A vertical wind profile consists of wind speed and direction observations at selected altitudes. The resulting wind vectors can be plotted on a diagram with u along the x-axis and v along the y-axis. If lines are drawn between the tips of these vectors, starting at lower altitudes and working upward, a *hodograph* can be obtained.

The lines drawn between the tips of the wind vectors represent the vector difference in wind with height. Within the context of geostrophic balance, this describes the *thermal wind*, which we will describe in more detail in our next lecture. Indeed, much of the vector wind change with height conveyed on a hodograph is driven by the thermal wind (and thus synoptic-scale horizontal layer-mean temperature gradient), although other effects such as friction and parcel accelerations not included within geostrophic balance can result in local departures from this relationship.

As the difference in altitude between two successive vectors becomes infinitesimally small, the lines between the tips of the wind vectors (i.e., tangent to the hodograph) also define the vertical wind shear vector over the layer. As a result, a hodograph describes *how the vertical wind shear vector changes with height*. The shape, or curvature, of a hodograph depends on the change in the vertical wind shear vector with height, not necessarily the change in the wind direction alone with height. The construction of a hodograph is illustrated in Fig. 2.11 in the course textbook; note that some artistic liberty was taken to curve the hodograph. We will use the convention that a straight hodograph describes unidirectional *vertical wind shear* and a curved hodograph describes *vertical wind shear* of varying direction.

The mean wind velocity over some vertical layer can be found by numerically averaging the winds within that layer. This can also be done approximately by visually inspecting the hodograph. The mean wind over some layer should lie between the wind observations at its top and bottom; this will lie along the hodograph itself for a straight hodograph and will lie on the concave side of the hodograph for a curved hodograph. If one assumes that a feature moves with the mean wind over a given layer, then this mean wind also provides an estimate of the feature's motion. Deep, moist convection tends to move with the mean wind over its depth early in its lifecycle. Later, due to a myriad of processes we will consider later this semester, it may move in a different direction and/or at a different speed from this mean wind. From an estimate of the feature's motion – whether equal to or deviant from the mean wind over a given layer – one can compute the feature-relative wind speed and direction: full wind **v** minus the feature motion **c**, i.e., $\mathbf{s} = \mathbf{v} - \mathbf{c}$.

The vorticity $\boldsymbol{\omega}$ is defined as the curl of the velocity vector, i.e.,

$$\vec{\omega} = \nabla \times \mathbf{v} = \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \dfrac{\partial}{\partial x} & \dfrac{\partial}{\partial y} & \dfrac{\partial}{\partial z} \\ u & v & w \end{vmatrix} = \left(\frac{\partial w}{\partial y} - \frac{\partial v}{\partial z} \right)\mathbf{i} + \left(\frac{\partial u}{\partial z} - \frac{\partial w}{\partial x} \right)\mathbf{j} + \left(\frac{\partial v}{\partial x} - \frac{\partial u}{\partial y} \right)\mathbf{k}$$

The last of these terms is our familiar vertical vorticity ζ. The first two of these terms define the *horizontal vorticity*, reflecting rotation in the meridional-vertical (**i**) and zonal-vertical directions (**j**), respectively. If we assume that horizontal gradients in vertical motion are negligible, such as in the absence of deep, moist convection, then the horizontal vorticity vector can be written as:

$$\vec{\omega}_H = \left(-\frac{\partial v}{\partial z}, \frac{\partial u}{\partial z} \right) = \mathbf{k} \times \mathbf{S}, \text{ where } \mathbf{S} \text{ is the vertical wind shear vector } \frac{\partial \mathbf{v}}{\partial z}$$

The lack of a multiplication factor means that the horizontal vorticity has the same magnitude as the vertical wind shear vector. The definition of the cross-product means the horizontal vorticity vector is directed 90° left of the local vertical wind shear vector (and thus hodograph). Because warm (cold) air lies to the right (left) of the thermal wind vector, the horizontal vorticity vector points down the layer-mean temperature gradient from relatively warm to relatively cool air.

Given the feature-relative wind defined above, one can partition the horizontal vorticity into along- and across-wind components, given respectively by the streamwise and crosswise vorticity. The streamwise vorticity is the component of horizontal vorticity along the feature-relative flow and is directly proportional to the change in *feature-relative wind direction* with height. The crosswise vorticity is the component of horizontal vorticity across the feature-relative flow and is directly proportional to the change in *feature-relative wind speed* with height.

We can illustrate these relationships through some thought experiments. Consider a unidirectional westerly vertical wind shear profile; the hodograph for such a vertical shear profile is straight. The horizontal vorticity vector – 90° to the left of the vertical wind shear vector – is directed to the left of the hodograph. Let the feature-relative motion lie along the hodograph so the feature-relative wind is also along the hodograph. Thus, the horizontal vorticity is entirely crosswise: none is along the feature-relative wind.

Now, consider a clockwise-curved hodograph defined by easterly wind at the bottom of the layer and westerly wind at the top of the layer, with uniform wind speed with height. The horizontal vorticity vector is directed to the left of the hodograph. The feature-relative motion in this instance is zero, such that the full wind is equal to the feature-relative wind. Thus, the horizontal vorticity is entirely streamwise: none is across the feature-relative wind.

In general, observed hodographs exhibit changes in feature-relative wind speed and direction with height: greater directional (speed) changes results in greater streamwise (crosswise) vorticity. The question can be asked; why do we care? As we will show later in the semester, the primary means by which vertical vorticity is generated in the vicinity of deep, moist convection is by the tilting of horizontal vorticity into the vertical by updrafts or downdrafts. The tilting of horizontal vorticity into the vertical results in both anticyclonic and cyclonic vertical vorticity. In the streamwise case, the feature-relative flow advects the cyclonic vorticity couplet under the updraft or downdraft, where vortex tube stretching can act to amplify it. In the crosswise case, the feature-relative flow advects neither the cyclonic or anticyclonic vorticity couplet under the updraft or downdraft.

Finally, the closely-related *storm-relative helicity* is a measure of the degree to which the direction of motion is aligned with the horizontal vorticity; thus, it is closely related to streamwise vorticity. While we will cover this in more detail later this semester, we wish to briefly describe how storm-relative helicity may be estimated from a hodograph:

- Obtain an estimate for the feature motion (direction and speed) and place a marker on the hodograph chart accordingly.
- Draw a line from this feature motion to the tip of the vector given by the wind at the bottom of a given layer (e.g., the surface).
- Draw another line from this feature motion to the tip of the vector given by the wind at the top of a given layer (e.g., commonly 1 km or 3 km).
- The storm-relative helicity is directly proportional to the area cut out between these lines and the hodograph between them. Positive values denote streamwise vorticity and negative values denote antistreamwise vorticity.