Module 5.1

Slide 1: This lecture series reviews some of the many applications of radar for meteorology and also includes a short module about lidar. Radar is generally now used as a word itself, but is technically an acronym for "radio detection and ranging". Primitive radar-like instruments were first developed as early as the late 1800s, but rapid development of radar occurred in the 1930s and 1940s prior to and during World War II, with the Naval Research Laboratory playing an important role in radar systems that would eventually evolve into usage for meteorological purposes. In this lecture series, we will review some of the fundamentals concerning how radar works, then we will move on to discuss some specific topics related to weather radar.

Slide 2: Radars used for meteorological purposes operate in the same general wavelength range as the passive microwave instruments we have discussed, from about 3 GHz for S-band weather radars to near 100 GHz for W-band cloud radars. They are active sensing instruments that transmit a signal then listen for return echo. This is the same principle used by scatterometers except that in the case of radar, a beam travelling along some path may return backscattered power from multiple targets along its path. In this picture, there might be two clouds—or other targets—along the same path, and some of the transmitted signal might continue to propagate beyond the first cloud and reflect back to the radar off the second.

Slide 3: The center frequencies for radars used for atmospheric science are denoted here. S-, C-, and X-band radar are used for detection of precipitating clouds. These wavelengths tend to be too long to be efficiently scattered by small cloud drops but are scattered by larger rain drops. However, of the three, S-band radar is least susceptible to attenuation—which just means reduction of the transmitted signal—by non-precipitating cloud. K_u-band is also used for precipitation radar in space since it requires a smaller antenna. Ka- and W-band radars are more susceptible to scattering by smaller hydrometeors and are thus known as cloud radars. Their signal is highly attenuated by rain drops but may be able to penetrate non-precipitating clouds well enough to reveal details about their structures. All of the bands are located outside of the narrow water vapor or oxygen absorption bands because we want the transmitted signal to be able to propagate as far as possible while also making the return trip to the radar receiver. However, absorption by water vapor increases as microwave frequency increases in this range, so cloud radar signals can be heavily attenuated in moist environments such as in the tropics.

Slide 4: A table of radar bands, including the frequency and corresponding wavelength range for each, is shown here. The K-band is generally not used because it is situated in a water vapor absorption band.

Slide 5: The next few slides introduce some of the terminology that will be useful when discussing radars throughout this lecture series. We will not spend much time discussing the details of signal processing and the electrical equipment on which radar depends. Instead, we will mostly focus on what happens to a transmitted signal after it leaves a radar and interpretation of the return signal, or echo. Collectively, the hardware involved in producing waves at the desired frequency is called a transmitter, which include klystrons, magnetrons, or solid state transmitters.

Slide 6: An example of a radar antenna is shown here. This picture was taken in the field of the National Center for Atmospheric Research S-PolKa radar system. The entire dish, or reflector, pictured is approximately 8.5 meters in diameter. This is a reflector for an S-band radar, which transmits radiation with a wavelength of 10 cm, meaning that the reflector is about 85 times wider than wavelength of transmitted signal. A larger reflector allows for a narrower beamwidth; this particular system permits a beamwidth of a little less than 1 degree. We will discuss the definition of beamwidth shortly. Waves that are generated by the transmitter travel up a waveguide depending on their polarization and are emitted through a feedhorn toward the reflector. Waves scatter off of the reflector and generally propagate into the direction that the radar is pointing.

Slide 7: Some other terms listed two slides previously were beam, ray, target, and echo. The beam refers to the conical area emanating from the radar that expands as it moves farther from the radar. The ray refers to the path along the center of the main beam. The target is any reflecting object that backscatters radiation back to the radar dish, and the backscattered radiation received back at the radar is called an echo. The range is the distance along the ray from where an echo is detected.

Slide 8: The range is one of three coordinates for describing the location of radar echoes in polar coordinates. The other two are the azimuth and elevation angles. The elevation angle is sometimes also called the tilt.

Slide 9: A weather radar pedestal contains gears that enable the antenna to rotate both horizontally and vertically. The elevation angle describes the angle above the horizontal at which the antenna is pointing and transmitted radiation is travelling. An overhead scan would have an elevation angle of 90°, while a perfectly horizontal scan has a tilt of 0°. The azimuthal angle is related to the horizontal direction in which the antenna is pointing. An azimuth of 0° generally is defined as a northward pointing antenna. Together, the azimuth angle, elevation angle, and range describe where in three-dimensional space an echo is located relative to the radar antenna.

Slide 10: A scan strategy refers to a usually pre-programmed set of azimuth and elevation angles at which the radar will transmit. For a weather radar, it consists of several sweeps that together are combined into a single volume of data. Scan strategies are discussed a little more in another module. A sweep typically refers to one full rotation at one elevation angle for all azimuth angles sampled.

Slide 11: As a radar rotates, its beam will sample an area that fills in a conical volume containing a hollow center. The red lines in this diagram represent the rays in a vertical cross-section drawn through the radar site.

Slide 12: Sweeps can be made at numerous elevation angles.

Slide 13: For a typical scan strategy used by a weather radar, the area in the center of the cone at the highest elevation angle is known colloquially as the “cone of silence”. It is an area directly above the radar that is not sampled.

Slide 14: Collectively, all of the sweeps make up a volume, which can denote any three-dimensional collection of radar data collected within a short time frame. There may be parts of the volume where beams at successive elevation angles do not overlap, leaving gaps in the radar echoes, especially at long ranges from the radar. Thus, the number and spacing of tilts used in a scan strategy determines the vertical resolution of radar data.

Slide 15: Next, let’s look at an example of gates and gate spacing, which effectively denote a radar data point, and the spacing between data points along a ray.

Slide 16: Visually, consider a beam shaded in yellow, with a ray represented by the red line. Each blue line represents the end of a gate. This drawing illustrates 6 gates. The gates become wider as the beam spreads with increasing range; however, the gate spacing is fixed, and represents the length along the ray covered by a single gate. Here, it is shown as the distance between the centers of two gates. The number of gates is dependent upon the signal processing hardware of the radar, and the gate spacing is determined by this and the maximum unambiguous range of the radar, which we will discuss shortly.

Slide 17: Like we saw in the module on scatterometers, the transmitted and received power refer to the average power transmitted per pulse, and the backscattered power received at the antenna that is reflected by targets. The weather radars we will talk about are pulse radars, meaning that they transmit successive short-duration pulses of radiation. The pulse period describes the period of time between successive pulses. It is typically about 1 millisecond. The pulse duration, also known as pulse length, describes the period of time during which the radar is actually transmitting the pulse. It is about 1 microsecond. For such a radar, in every millisecond, a signal is transmitted and reflected back to the radar antenna. If the pulse length is 1 microsecond, then the dwell time, or the time the radar is listening for return echo, is 999 microseconds.

Slide 18: As a visual, see this example of two successive pulses transmitted by a radar. The two pulses are shown in boxes as the red sinusoidal curve. The length of time during which the signal is transmitted is contained within one box. The signal consists of a wave packet, containing several wavelengths of radiation at the desired frequency in a single pulse. The pulse period is denoted as the time elapsed between the beginning of successive transmissions.

Slide 19: The sampling rate refers to the number of signals that can be transmitted before a ray moves the distance of one beam width. A typical sampling rate for weather radar is about 60. The power received by the radar for a set of gates is an average power over the many samples. This is required because backscattered radiation from all targets in a volume will interfere randomly. Taking a variety of samples within a short time window allows for targets to spatially redistribute within the volume making the random constructive and destructive interference of the backscattered radiation to approximately cancel so that the radar backscatter cross-section approximately equals the sum of the backscatter cross-sections of the individual targets in a contributing volume.

Slide 20: To visualize this, suppose that each red line represents the ray of a different sample. In this example, the sampling rate is 7. For a sampling rate of 60, imagine fitting 60 evenly spaced red lines within one beamwidth. This means that the average power returned at a gate centered on the middle ray, which is the black line, actually consists of some power returned to the antenna from targets outside the beamwidth shaded in yellow.

Slide 21: Finally, the pulse repetition frequency, or the PRF, is simply 1 divided by the pulse period. A period of 1 millisecond corresponds to a frequency of 1000 Hz. Radar meteorologists often use PRF instead of pulse period to describe how often pulses are transmitted. A high PRF means that pulses are transmitted frequently. This reduces the farthest distance that the radar can see unambiguously because the radar needs to receive a signal from a previous transmission before it transmits a new signal in order to unambiguously know from which transmitted signal the return signal originated. This unambiguous range is simply the speed of the radiation propagation, divided by 2 times the PRF, where the 2 is required because the signal must make both an outbound and inbound trip from and to the antenna.

Slide 22: We have already discussed wavelength and range, noting the difference between range as one axis in a polar coordinate system and the maximum unambiguous range, which is a specific value of range in that coordinate system. The dielectric constant will be important when we discuss stratiform brightbands in a different module. It is a property of liquid—generally water—or ice targets. The dielectric constant is a complex number, consisting of parts that describe both an object's or medium's permittivity or conductivity. The norm of this number will be squared in the radar equation as we will see shortly. Finally, all radar meteorologists must be familiar with conversion between linear units and decibel—or logarithmic—units. The potential mean power returned to a radar antenna often spans several orders of magnitude. Therefore, the processed raw signal, when converted to variables related to the size of targets, are often presented in logarithmic units. If Z represents radar reflectivity, then dBZ Is just 10 times the logarithm of Z. Likewise, we can convert values of dBZ to Z. Note that a change of 3 dB roughly represents a doubling or halving of the variable of interest. Also note that we cannot do a linear average of number in a logarithmic scale. Any linear operators applied to reflectivity, such as summation, require that decibel values are converted to a linear scale first.

Module 5.2

Slide 1: In this module, we will take a look at the weather radar equation, as well as a couple of other terms that are related to properties of the radar itself and are important in this equation.

Slide 2: The average power returned to an antenna after it transmits a signal is dependent upon several factors: The power transmitted, the antenna gain, the pulse duration, the beamwidth, the wavelength, the dielectric constants of the targets, the radar reflectivity factor, and the range. The radar reflectivity factor is denoted by Z and is perhaps the most important variable that radar meteorologists consider. It is often simply referred to as reflectivity by many people. It is related to the sum of the sixth power of the diameters of all the targets in a contributing volume. This equation describes the power returned to an antenna by a set of targets distributed throughout a contributing volume and can be derived by starting with the equations shown in blue at the top, which we first encountered when discussing scatterometers. Note that the contributing volume is that contained within a single gate.

Slide 3: The antenna gain is a property of the reflector itself. An isotropic radiator would transmit equally in all directions. The gain represents the ratio of the actual power transmitted as a function of direction relative to a lossless isotropic radiator. A perfect reflector would direct radiation into a narrow beam in the direction the radar was pointing. However, radar antennas are not perfect reflectors. A small amount of power is reflected by the dish in unintended directions, away from the main lobe. That radiation can be returned back to the antenna and picked up by the receiver; however, it is assumed that all backscattered radiation collected during a radar's dwell time comes from somewhere along the main lobe. We will briefly explore the ramifications of sidelobes in a different module that discusses various types of anomalous echoes.

Slide 4: Finally, the beamwidth is the angle over which the antenna gain function is one-half of its maximum value, which is located along the center of the main lobe. As stated before, the beamwidth decreases as the antenna size increases. The beamwidth is also proportional to the wavelength. This means that high-frequency radars can operate with much smaller antennas to achieve an acceptable beamwidth for remote sensing of clouds.

Slide 5: As a result of the sixth power dependence of radar reflectivity factor on size, Z is highly sensitive to the size of targets. Large targets can have outsized influence on Z in a volume. Consider the example shown here. The two blue dots are drawn to scale, with the blue dot on the right being 10 times the diameter of the small dot on the left. Because of the sixth power dependence of radar reflectivity factor on size, the large drop would yield a reflectivity 10 to the sixth, or 1 million, times larger than the single small droplet. Said another way, 1 million small droplets would be required in a single volume to yield the same reflectivity as the 1 large drop even though the 1 million small droplets would collectively have 10 thousand times the total volume. This means that when looking at radar reflectivity factor, you will be looking at information that mostly describes only the largest targets in a contributing volume. This drop-to-drop comparison is for visual purposes only. A single drop in a contributing volume will typically not return enough backscattered power to produce a signal that is able to be processed. We would say in such a case that the return signal is less than the sensitivity of the radar. However, antennas with extremely high gain or very large transmitted power can be sufficiently sensitive to detect individual raindrops. The Mid-Course Radar, or MCR, is operated by the Navy near Cape Canaveral and was previously used to track debris from space shuttles during launch. It has a high sensitivity primarily because it transmits a high-powered signal.

Slide 6: Our weather radar equation can then be rearranged to think about how properties of the radar itself and of the target impact the radar reflectivity factor. In this formulation, Z depends on the power received, which is related to the backscatter coefficient of the target, which is itself related to size.

Slide 7: This example shows two plots of reflectivity derived from data from two different radars near the same location. The left panel represents S-band reflectivity, while the right panel illustrates X-band reflectivity. Remember, X-band radiation is at a higher frequency. The approximate locations of the radars are denoted by gray arrows. They are looking northwestward toward the same squall line. The S-band radar detects a wide area of stratiform precipitation behind a leading squall line. The squall line is denoted by purple to white colors that represent reflectivity up to a very large value exceeding 70 dBZ, and lighter rain is denoted to the northwest of the apparently southeastward moving line. The X-band radar detects the same squall line, although its reflectivity is lower than the S-band value. Furthermore, the X-band radar sees very little behind the squall line, even though the S-band radar clearly denotes that rain drops are present. This happens because the X-band signal is attenuated by liquid water. Attenuation just refers to the loss of transmitted signal by absorption or scattering away from the ray.

Slide 8: To understand this process for active sensors, let's consider attenuation as a function of the familiar range of frequencies we discussed when learning about passive microwave sensors in a previous lecture series. The one-way attenuation is plotted here on the y-axis; double this number to yield the attenuation from and back to the radar. Larger attenuation means that less signal is available for backscatter to the radar. A few features are evident. First, attenuation increases as frequency increases. This is related to the increased ability of water vapor to absorb radiation as microwave frequency increases. This is the primary reason that we cannot use high frequency radar, such as K_a-band or W-band in the tropics to see at long ranges; too much signal is simply absorbed over long distances to capture a high enough signal-to-noise ratio to make a meaningful measurement except for within a few thousand meters of the antenna. The different solid lines represent various mixing ratios of water vapor, and we can see that for all frequencies, attenuation increases as water vapor concentration increases. Second, some particular frequencies seem to have enhanced attenuation. The first is around 22 GHz, in a water vapor absorption band. The second is in the oxygen absorption band near 60 GHz. Other absorption bands are seen at higher frequencies. Third, the attenuation at low frequencies such as S-band is almost negligible compared to higher wavelengths. In other words, S-band radiation is not strongly absorbed or heavily scattered away from the radar by an atmospheric target. Finally, other higher frequency bands are situated in parts of the spectrum where attenuation is minimized. W-band radar falls in this dip in attenuation.

Module 5.3

Slide 1: In this module, we will explore the usage of the Doppler effect in using pulse radar systems to determine radial velocity of targets, such as hydrometeors. Specifically, we will discuss the pulse-Doppler system.

Slide 2: In previous modules, we saw how a pulse radar system transmits packets of wave energy periodically, but a radar spends most of its time listening for a return signal. A Doppler radar is not only capable of measuring the intensity of backscattered power, but also the phase and frequency of the returned signal. The range, as defined before, is the distance between the antenna and a target. However, hydrometeors within a volume are usually moving; the movement causes a slight phase shift of the signal which may be used to determine the component of the velocity of targets parallel to the ray. Frequency shifts also occur; however, they are usually too small to determine radial velocity given the pulse length of about 1 microsecond used.

Slide 3: The goal of Doppler radar for weather is to measure the radial velocity. The radial velocity is the mean speed of a volume of targets either toward or away from the radar. Typically, signage is defined such as a negative radial velocity represents targets moving toward the radar. Suppose the blue circle represents a volume of raindrops and that it is moving toward the radar, and suppose that the raindrops move a distance *d* between successive pulses separated in time by about 1 millisecond. The radial velocity is then that distance divided by the pulse period. Of course, we cannot explicitly measure this distance with radar; however, we can detect a phase shift between successive pulses and use this to derive the frequency shift and radial velocity given the wavelength of transmitted radiation. The phase shift is characterized by the fraction of the distance traveled by radiation relative to its wavelength. Consider that transmitted radiation must travel to and from the target, thus the two appears in the numerator in the equation for phase shift. Simple algebra yields an equation for radial velocity that is a function of the phase shift given some fixed pulse period and wavelength.

Slide 4: A rudimentary visualization of the phase shift between successive pulses is shown here. In this example, the orange and blue lines represent different pulses of *returned* radiation. We will assume that the phase of the transmitted radiation leaving the radar is the same but that targets in the contributing volume will cause some shift in phase for the radiation once it returns to the antenna. If hydrometeors are moving, that phase will change between successive pulses. This is the phase shift that we insert into the equation on the previous slide to derive radial velocity. If the hydrometeors are moving neither toward nor away from the radar at all, then the phase shift will not differ between successive pulses. However, as we will soon see, hydrometeors with just the right magnitude of non-zero velocity components toward or away from the radar can cause the same phase shift, creating an ambiguity in determining Doppler velocity.

Slide 5: One way of visualizing this problem is by plotting the phase shift over a sequence of pulses. This plot shows the sine of the phase of returned signal at the radar for a series of pulses. Effectively, given a known pulse period, measuring a sequence of phase shifts allows us to recover the Doppler frequency. In this figure, the blue dots represent an individual phase shift measurement.

Slide 6: Next, let's zoom in on just a few pulses, where again the blue dots represent individual phase shift measurements. The red curve is one sinusoid that connects the blue dots. However, the green curve is a sinusoid with a different frequency that also satisfies the observed phase shifts. These different Doppler frequencies correspond with different Doppler velocities. In fact, infinitely many sinusoids, all with a different frequency, can satisfy the observed phase shifts.

Slide 7: The maximum unambiguous radial velocity that can be measured using this technique is called the Nyquist velocity. The Nyquist velocity corresponds with a phase shift of pi. This is because a phase shift of negative pi yields the same wave phase at the radar as a phase shift of +pi, and so the radar would not be able to distinguish between the two. The Nyquist velocity is a function of the PRF of the radar. As we increase the frequency of pulse transmission, or decrease the pulse period, the maximum resolvable velocity increases. The Nyquist velocity is also a function of wavelength, such that low frequency radar like S-band weather radar tends to have a higher Nyquist velocity than C- and X-band precipitation radars. If a radar operator needed to resolve high velocities unambiguously, they could increase the PRF. However, the trade-off is that the unambiguous range, seen in a previous module, decreases. This trade-off between maximum resolvable unambiguous range and velocity is known as the Doppler dilemma.

Slide 8: The ambiguity in Doppler velocity is illustrated here on a number line. Suppose, just for example, that the Nyquist velocity of a radar is 18 m/s. If so then a phase shift of ± pi yields a reported Doppler velocity of ± 18 m/s. Furthermore, a phase shift of ± 2 pi, which would correspond to a real radial velocity of ± 36 m/s, looks no different to a radar than a phase shift of zero. Therefore, the radial velocity reported at the ± 2 pi phase shift would be zero, even though the actual velocity is much stronger. This problem is called velocity folding. We would say that the 36 m/s observation is folded back into the ± 18 m/s Nyquist interval.

Slide 9: One challenge of working with Doppler radar is unfolding the radial velocities, especially in cases where strong winds occur. Often, real-time data will not be automatically unfolded. One example from a ship-based C-band radar looking at strong winds in a mesoscale convective system over ocean is shown here, with radar reflectivity factor shown on the left, and reported radial velocity illustrated on the right. The precipitation to the northeast of the ship consists of a leading line of convection with a large trailing stratiform region that extends off the image to the northeast because it is located beyond the maximum unambiguous range set by the PRF used. Although you cannot tell from a single image, a loop of imagery in real-time indicated that the rain was moving toward the ship. Within this precipitation, a jet of descending inflow was present. The flow was stronger than, in this case, 16.5 m/s, which was the Nyquist velocity for this radar when using this PRF. We can quickly tell that the velocity is folded given some background about the volume observed. There would be no reason to expect such strong horizontal shear to be present in a squall line as if the red in the right-hand plot were actually representative of strong wind away from the radar. Physically, it is more sensible to interpret this radial velocity as a strong, folded inbound velocity.

Slide 10: Another common example of velocity folding occurs in tornadoes. In strong tornadoes, folding can even occur with S-band radars. Again, reflectivity and radial velocity are respectively at left and right. This example is from May 3, 1999, the day a powerful tornado struck Moore, Oklahoma. In our radar lab, we will look at radar data from an EF-5 tornado that also occurred in Moore and followed a similar track on May 20, 2013. The cells in this image were moving generally eastward. Therefore, the radial velocities at bottom left of the image are folded. The radar indicates outbound, or positive, radial velocities, when the flow was actually strongly toward the radar. The fold occurs where the velocity reaches the Nyquist velocity. In this figure, purple colors, which indicate strong inbound motion are located directly next to reds, which indicate strong outbound motion. In the upper left quadrant, folding due to strong rotation occurs. Large positive radial velocities are reported in an area of generally negative radial velocities. If we unfold the velocities in the white circle, we would see that the radial velocity is actually strongly negative, which implies, in this particular case, the presence of cyclonic vorticity…

Slide 11: …as denoted by the white arrows that proportionally—based on their lengths—represent the true radial velocities.

Slide 12: In a contributing volume, millions of individual hydrometeors may be present. Not all of them move at the same speed. The Doppler radial velocity reported represents the mean radial velocity of targets in a contributing volume. Some radars can capture the Doppler spectrum in a volume and report a quantity called spectral width. The spectral width describes the width of a distribution of radial velocities in a volume. For the panel on the right, suppose that signal above the noise level is received with a range of phase shifts. The yellow bell curve represents the range of phase shifts detected. The spectral width of this curve is similar to twice the standard deviation of a Gaussian. A couple of examples from real data are shown at left. The top example shows the distribution of radial velocities for a volume of ice crystals that were moving toward the radar at an average of 1.5 m/s. However, ice crystals with velocities of between almost 0 and over 3 m/s toward the radar were observed. At bottom, is a Doppler spectrum for falling rain drops viewed by a vertically pointing radar. The mean of the spectrum was at about 6 m/s, but rain drops falling at between 0 and 9 m/s were observed. The spectral width is much smaller in the example of the ice crystals than of the rain drops.

Slide 13: Some biases can occur when measuring radial velocities that are seen when detecting the spectral width. Remember, the radial velocity reported is the average of all radial velocities of targets in a volume. The top example shows a Doppler spectrum that saddles the Nyquist velocity. In this case, the radar sees targets that are all moving away from the radar, but because some targets are moving away faster than the Nyquist velocity, the radar interprets the volume as containing some targets that are actually moving quickly toward the radar. Thus, while the true radial velocity is near the Nyquist interval, the reported radial velocity is skewed too low. The bottom example shows how ground clutter, which has zero radial velocity, can also contaminate a radial velocity estimate if actual hydrometeors are also present in the same contributing volume.

Module 5.4

Slide 1: This module will review various parameters that can be measured using radar data, including dual-polarimetric observations.

Slide 2: Recall that radiation can propagate with a variety of polarizations depending on the orientation of the electrical field, or at a quantum level, the spin of photons. Radiation that propagates along just one dimension is either horizontally or vertically linearly polarized. Dual-polarization radars detect, at minimum, the returned power in each polarization and the variation in the phase of each polarization. However, their transmissions vary depending on the complexity of the hardware used. Transmitting two polarizations of radiation allows a radar to determine, not just the size and radial motion of hydrometeors, but also their shape and composition.

Slide 3: We will look at examples of several radar-derived variables in this module for the same volume of rainfall. The variables listed in red here are several dual-polarimetric variables. We have already discussed the variables in blue in a previous module, but we will review them briefly.

Slide 4: The image shown is a map of radar reflectivity factor for an area of marine tropical precipitation. The parts of sweeps displayed will be from a tilt of 2.5°. Note again that the reflectivity is dependent upon the sixth power of the diameters of targets in a contributing volume, meaning that radar reflectivity factor is generally dominated by the largest drops in a volume. Note that you will sometimes see reflectivity referred to as “base” reflectivity and “composite” reflectivity. When used correctly, the term “base” reflectivity refers to data collected along the lowest sweep. Composite reflectivity is a derived field that illustrates the highest reflectivity in a vertical column and is never less than the base reflectivity.

Slide 5: The Doppler velocity is derived from the phase shift between two successive pulses. Negative values indicate that the wind in a gate has some component toward the radar, while positive values indicate wind with some component away from the radar. The true wind vector may also have some cross-ray component. For example, even a strong cross-ray wind has a radial velocity of zero. You may see the term “base” velocity used; this describes the radial velocity of the lowest tilt. Another field that you may encounter is called “storm-relative” velocity. It is the wind vector that remains after the background motion of the flow is subtracted from the total radial velocity vector. An example of the vector algebra is illustrated at the bottom right here. Suppose that the observed radial velocity is the blue arrow, but that the storm in which that observation was made was moving opposite the direction of the orange arrow. This storm motion would be governed by the mean flow. By subtracting the mean flow vector from the radial velocity vector, we back out the part of the total radial velocity vector that is caused by kinematics internal to the cell.

Slide 6: In the previous module, we also discussed spectral width. An example of a plan view of spectral width is shown here. Recall that this value represents the spread in the distribution of radial velocities of targets in a volume. The green colors here denote spectral widths of 2 m/s or greater.

Slide 7: The first dual-polarimetric variable we will discuss is differential reflectivity, also known as Z_{DR}. It is the ratio of the horizontal to vertical radar reflectivity factors, expressed here on a logarithmic scale. Two subscripted H and V letters in the fraction denote the polarization of the transmitted and returned signal. So, for example, Z_{HH} is derived from the return power in the horizontal polarization for a transmission that was also horizontally polarized. Z_{DR} in the logarithmic scale is measured in decibels, or dB. A typical value of Z_{DR} is often between about -2 dB and 5 dB. Z_{DR} of zero means that the targets are perfectly spherical. A non-zero Z_{DR} denotes an oblong hydrometeor. Positive Z_{DR} indicates that the targets are larger along the horizontal plane of the transmitted radiation than along the vertical plane. Large rain drops, for example, flatten as they fall and might have larger values of Z_{DR} than small cloud droplets. Vertically oriented hydrometeors are associated with various types of—although not all—frozen hydrometeors. Because the Z_{DR} is a ratio of reflectivity, it, like the radar reflectivity factor, is largely sensitive to the size of targets in a volume and the value of Z_{DR} reported is dominated by the largest hydrometeors.

Slide 8: The linear depolarization ratio, or L_{DR}, is only observed by radars that are polarization-agile, meaning that they separately transmit horizontally and vertically polarized radiation. L_{DR} is defined as the ratio of cross-polar radar reflectivity factor divided by the vertically polarized radar reflectivity factor. L_{DR} can be used to identify hydrometeors that may cause depolarization of the transmitted signal, such as mixed ice/liquid water phase processes. It may also be used to easily identify second trip echo, a subject discussed in a different module. In this example, the L_{DR} is generally uniform and small—less than -20 dB—indicating that little depolarization occurred in this particular volume, which is the same volume as shown in previous slides. Many radars transmit radiation at a 45° slant angle, meaning that the power sent out by the radar in either the horizontal or vertical polarization is half the signal that originated at the transmitter. National Weather Service WSR-88D radars, for example, do not measure L_{DR}. A radar that transmits at a 45° slant angle then simultaneously receives horizontally and vertically polarized components of the return signal. Such radars are not capable of measuring L_{DR} because the transmitted waves are not orthogonal to one another.

Slide 9: The differential phase shift, or ϕ_{DP}, is related to the difference in the phase shift between successive pulses between the horizontal and vertical polarizations. The difference between phase shifts in a single polarization are dominated by the radial velocity of the hydrometeors in a volume; however, the difference between the phase shifts in different polarizations can tell us information about the absorption of the transmitted signal. Therefore, it may be useful for estimating the attenuation experienced by a radar beam, especially at X- and C-band. ϕ_{DP} at one range is dependent upon its values at ranges closer to the radar, so you will often see its value change smoothly as a function of range.

Slide 10: The specific differential phase, or K_{DP}, is just half the radial derivative of ϕ_{DP}. K_{DP} is sensitive to exactly how the radial derivative is computed since it must be done across discrete range gates. It is measured in °/km, and is typically positive—ranging as high as 5 or more °/km. Like Z_{DR}, if it is positive, then oblate hydrometeors with their long axis oriented along the plane of the horizonal beam are present. K_{DP} tends to be close to zero if most hydrometeors in a volume are spherical or their total number—and therefore mass—is small. Since K_{DP} is not sensitive to spherical hydrometeors, in rain-hail mixtures in which rain is oblate but hail is closer to spherical, K_{DP} can provide an additional useful constraint on rainfall estimation. The retrieval of K_{DP} from ϕ_{DP} is a complicated topic that we will not delve into during this introduction to radar meteorology.

Slide 11: The last dual-pol variable we will discuss is the co-polar correlation coefficient, often just called the correlation coefficient, or ρ_{HV}. It describes the linear correlation between the returned signals in the vertical and horizontal polarizations of radiation transmitted, respectively, in the same polarization. The asterisk in the equation displayed denotes the complex conjugate, which just means that the imaginary part of the S is multiplied by $\mathbf{i}^2$. Values of ρ_{HV} are typically close to 1, with values above 0.95 common in a variety of rain and frozen precipitation. Smaller values of ρ_{HV} can occur in large hail, and a low correlation coefficient below 0.95 is a sign of a radar brightband, a topic that will be covered in another module. Non-meteorological echo generally has low co-polar correlation. Thus, ρ_{HV} can be used to identify ground clutter. It is often also used to determine when a tornado is occurring because the debris lofted by a tornado reduces the ρ_{HV}.

Slide 12: Finally, the dual polarimetric variables can be used together to identify the most prevalent type of hydrometeor present at a range gate. Such hydrometeor or particle identification methods typically use Z, Z_{DR}, and K_{DP} to make estimates. The specifics of how such algorithms work is beyond the point of this discussion; however, one example of a "fuzzy logic" algorithm is shown here. As an example, high Z_{DR} corresponding with low reflectivity in the top right plot corresponds with the light blue box, which represents a category called "ice crystals". The bottom left shows that this category typically has K_{DP} of less than 1. This is just one example. As you can see, many different types of ice and liquid water hydrometeor are classified by such algorithms.

Slide 13: An example of such an algorithm in action is shown here. The top left panel shows an example of C-band reflectivity from a research radar located at Darwin, Australia. The bottom left panel is a vertical cross-section through the purple ray indicated to the south-southwest of the radar in the top left panel. This cross-section was taken through a deep convective core, with echo extending as high as 16 km in altitude. The top right shows K_{DP} corresponding with the reflectivity shown at top left. Some K_{DP} values over 1.5 °/km are seen through the cross-section. The hydrometeor identification for the cross-section is shown at bottom right. Below the 0°C, rain was present and represented by dark blue. Near the 0°C level, the light blue color represents a category called "wet snow", which is essentially sticky, melting ice hydrometeors. The most intense convection is co-located with graupel that extend well up into the cloud. The remainder of the upper part of the cloud consists of "dry snow" or ice crystals. Thus, the different dual-pol variables used together can provide important information about the microphysical structure of observed clouds.

Radar Display: Another set of examples is shown here from WSR-88D, which does not include L_{DR}. Again, you see that correlation coefficient is generally high everywhere. Z_{DR} and K_{DP} are typically small but are locally large in the most intense convection. As we move to higher tilts, we see a ring of occasionally higher Z_{DR} and reduced ρ_{HV} associated with the phase change of ice to liquid water as it precipitates out of clouds. In this example, we cannot see the feature in the radar reflectivity field because so much convection is present, but the Z_{DR} and ρ_{HV} signals are associated with the radar brightband. Indeed, if we consult the hydrometeor identification in this example, we see that this area surrounds liquid water but itself is composed of various types of ice or melting ice.

<u>Module 5.5</u>

Slide 1: This module reviews scan strategies used for collecting meteorological radar data. A scan strategy is a pre-programmed set of azimuth and elevation angles at which the antenna will point, as well as specification of the PRF and antenna rotation speed, which controls the sampling rate given some PRF. The type of scan strategy used determines the spatial and temporal resolution of the data collected.

Slide 2: Two basic types of scan strategies are used for weather radar. The most commonly used one is called the plan position indicator, or PPI. In this type of scan pattern a pre-set list of elevation angles is used and the antenna sweeps through all azimuthal angles at each tilt. When the antenna has made a complete sweep at one tilt, it moves up to the next elevation angle and executes another complete sweep. After the last pre-programmed elevation angle, the antenna moves back down to the lowest tilt and repeats the process. A rotation at each tilt is a sweep, and all sweeps combined are considered one complete volume. PPI scans allow a radar to sample the environment across a large area—out to its maximum unambiguous range over a 360° circle—but the vertical resolution is usually small as some gaps can remain between individual tilts, especially at large ranges. The range height indicator, or RHI, rotates through a pre-set list of azimuthal angles and range of elevation angles. It makes a vertical scan across all elevation angles at a single azimuth then moves to the next azimuth and scans either back up or down the range of elevation angles. RHI scans provide high vertical resolution of convection but would take too long to execute to scan an entire 360° sector in a desirable amount of time. Because a radar antenna can only point in a single direction at any given moment, a full cycle of a scan strategy often takes 5 to 15 minutes to complete depending on the requirements of the radar operator. We will look at some specific types of scan strategies used for cloud radar in a later module, which will include several examples of RHI scans.

Slide 3: The current image is that of radar reflectivity factor from a research radar that used a 15-minute scan strategy that combined PPIs at several elevation angles seen at the right side of the figure and RHIs in sectors to the northeast and southwest of the radar, where you can see hash marks on the constant range rings. In this example, intense convection is seen to the northeast of the radar, with some scattered convection elsewhere. Next, we will look at a vertical cross-section through the convection along the yellow line shown at 12° azimuth.

Slide 4: The two figures show side-by-side cross-sections of the same convection through that yellow line. On the left is a synthetic cross-section made using PPI data that is interpolated to fill in gaps between rays. Note also that this data is collected over several minutes, so it does not represent an instantaneous cross-section of the hydrometeors. On the right, is the same cloud but sampled using an RHI. A couple of things stand out. Although the displays look qualitatively similar at first glance, the RHI clearly has better vertical resolution. For example, the radar brightband is seen in a stratiform precipitation area enclosed by the white circle at the right. This brightband consists of a high reflectivity region that occurs near the 0°C level and is caused by melting ice as it precipitates. The same feature is seen in the panel at left, but it is not as well-resolved. Additionally, the RHI on the right has a higher maximum elevation angle than the collection of sweeps used to create the image at left. Therefore, the RHI scan efficiently see high up into cloud close to the radar. Very high scan angles are seldom used with PPI scans because their primary point is to see the horizontal distribution of echo.

Slide 5: Most operational radars use PPIs as their primary scan strategy. The National Weather Service WSR-88D radars call their strategies volume coverage patterns, or VCPs. The VCPs are simply pre-programmed elevation angles and corresponding PRFs. The next couple of slides will take you through a few VCPs used. The one shown here is known as the clear-air mode. It is used for long-range surveillance in quiet conditions near the radar. Because the goal is to maximize the unambiguous range, the PRF is set to something low—as low as 314 Hz. However, the long pulse period means that the antenna must move more slowly than normal to capture an acceptable sampling rate. The slow antenna rotation also increases the sensitivity by decreasing the background noise. Various combinations of elevations and PRFs may be used depending on the requirements. This type of VCP generally features a small number of elevation angles, however. VCP 35 is shown in this image and is a little faster than other clear-air scan strategies.

Slide 6: Various types of other scan strategies are employed when precipitation echo is detected near a radar. In severe weather, the VCP is designed to repeat a cycle as quickly as possible. This is because events like tornadoes can develop rapidly. In fact, there are cases of tornadoes occurring between complete scan cycles and going undetected by radar. The VCPs mainly differ in the PRFs used. A high PRF allows for resolution of large radial velocities, but a lower PRF may be used if severe weather is more distant. These scans are designed to repeat in less than 5 minutes. The middle scan strategy is similar but includes an extra elevation angle to provide better vertical resolution. It is the general precipitation VCP and completes a cycle every 6 minutes. The bottom VCP is a special case used when tropical cyclones are near the radar. It uses a smaller number of elevation angles but repeats some of them, using different PRFs at low tilts. Again, the higher PRF is used to identify large Doppler velocities that may be present in a tropical cyclone.

Slide 7: As we discussed in a previous module, one limitation of PPIs is that the highest elevation angle used generally does not sample an area directly above the radar. This area is sometimes colloquially known as the “cone of silence”. If looking at a plan view of radar data, an empty region at the center of display illustrates where no elevation angle was high enough to sample. This example shows how the no-echo region expands laterally from the radar as we move up in altitude.

Slide 8: When designing scan strategies, we consider the beamwidth of the transmitted signal as well as the spreading of the beam. Recall that the beamwidth is defined as the angle that contains the half power points of the antenna gain function at the center of the main lobe. However, because the beam spreads as it moves outward, the spatial resolution of data far from the radar is less than it is close to the radar. Therefore, a small echo may be captured by multiple azimuthal angles if it located close to the radar but might fit completely within a single beam if located far from the radar. Azimuthal angles in a scan strategy are often separated by 1°, which is a typical beamwidth used for operational and many research precipitation radars.

Slide 9: We also must consider refraction of the radar beam as it moves through space. The atmosphere, particularly water vapor, causes the beam to bend downward toward the surface. However, additionally, the surface of the ground bends down and away from a horizontal beam path. This must all be considered when determining the height of an echo sampled at each range gate. Super-refraction can cause the beam to bend downward more than normal, and in extreme cases, the beam can actually bounce off the surface, returning power associated with stationary echo on the ground.

Slide 10: Finally, consider the following situation at try to design your own scan strategy for the given requirements. There is no one single answer, but the requirements are ones that are often considered by radar operators, and the answers you get to the questions at bottom will be close to values used for actual radar remote sensing.

<u>Module 5.6</u>

Slide 1: This module will discuss a variety of different types of special echoes, including the stratiform brightband and a variety of non-meteorological echo that you may encounter when viewing radar data.

Slide 2: We will discuss various types of anomalous propagation listed at the bottom.

Slide 3: First, we will discuss the stratiform brightband. It is a layer of enhanced reflectivity that occurs as frozen precipitation falls through the 0°C level and melts. Consider the weather radar equation that we saw earlier in this lecture series. The power received at an antenna is a function of properties of the radar, which remain constant, and properties of the target. The power received is then converted to some reflectivity by algebraically maneuvering this equation. The radar processor must assume some dielectric constant K when converting from received power to reflectivity. However, the dielectric constant for water is almost 5 times that of ice. This means that between the same-sized ice and water hydrometeors, the ice would return less power than the water. However, melting ice hydrometeors are often aggregates that are much less dense than liquid water drops and therefore have larger backscattering cross-sections. Since reflectivity is very sensitive to the size, the ice can still return plenty of power despite its lower dielectric constant. The region of melting ice experiences large melting ice hydrometeors that are surrounded by a liquid water shell that maximizes K. The radar sees the increase in returned power as an increase in reflectivity.

Slide 4: In visual form, consider falling ice hydrometeors above the 0°C level that have large Z but small K. As they melt, they initially retain their size but are surrounded by liquid water. The dielectric constant important at this point is that of water, but the backscattering cross-section is still large. Eventually, the ice completely melts and we are left with a distribution of denser completely liquid drops with large K, but smaller backscattering cross-section than in the melting layer. Thus, locally between altitudes just above and below the 0°C level, the maximum reflectivity is found where the melting occurs. As we stated in the lecture on dual-pol, this can also manifest itself in the differential reflectivity and co-polar correlation coefficient.

Slide 5: We'll take another look at some examples of a brightband that occurs in regions of stratiform precipitation where vertical motions are not strong enough to loft large hydrometeors. We'll look at vertical cross-sections along the yellow line to the southeast of the radar. PPI cross-sections of reflectivity, Z_{DR}, and ρ_{HV} are shown.

Slide 6: The maximum in reflectivity and Z_{DR} are clearly seen around an altitude of 4-5 km and are accompanied by minima in the correlation coefficient.

Slide 7: In looking at plan views of reflectivity or other fields, the brightband can be seen as a ring of enhanced reflectivity forming a ring at an approximately equidistant range, assuming the horizontal gradient in temperature in the radar domain is small. This example shows a radar brightband for a 5° tilt located at about 50 km from the radar site.

Slide 8: As we move to a higher elevation angle, we see that ring get closer to the radar because the height of the beam increases more rapidly with range.

Slides 9 and 10: With even higher tilts, the brightband appears even closer.

Slide 11: Next, we will look through various other types of anomalous propagation. The first is beam blockage. This example is taken from imagery in southwestern Alaska. Some terrain is located to the southwest of this particular radar. Echo can be detected between the radar and the terrain, but in the black ellipse behind the terrain, a relatively weak echo region is detected. The actual hydrometeors probably are more reflective, but because so much power was dissipated by the terrain, little power from farther along the ray remains that can be backscattered to the radar. The processor interprets the reduced power as reduced reflectivity. Radar sites are generally selected such that beam blockage by structures or trees is minimized; however, in mountainous areas, blockage of the beam at low elevation angles is sometimes unavoidable. This can partially be remedied by placing the radar atop small topographical features, but then the radar shoots over any echo that is located only at low altitudes. This is the case with the radar in the San Francisco Bay area, which is located high up in the Santa Cruz Mountains.

Slide 12: A more extreme example of beam blockage is seen here. This radar was located a few hundred meters to the east of a grove of trees that blocked the beam below the 2.5° tilt. The plan view shown here illustrates reflectivity at the 0.5° tilt. Even though convection is probably present to the west of the radar, so much radiation is scattered or absorbed by the trees, that the radar indicates no echo. Generally speaking, and this will be a common theme through the rest of this module, if you encounter strange looking echo features that are oriented lengthwise along one or more rays, you should interpret that data with caution because its cause is likely somehow non-meteorological.

Slide 13: Ground clutter is often observed close to a radar. Ground clutter is typically characterized by its radial velocity of zero. Clutter is caused by the beam intersecting objects on the ground. This can happen as the beam is moving upward away from the radar or as the beam is super-refracted back down to the surface. This example shows a clear example of ground clutter to the northeast of the radar with some moderate reflectivity. In this location, the return will be persistent because it is a wind farm that is not moving. The echo is caused by backscatter off of wind turbines.

Slide 14: Flying birds and insects can also be seen on radar. This example shows rings of reflectivity that are caused by power returned from birds leaving lakes in the early morning. If you follow the link shown in this slide, you can find several examples of how weather radar can see emergence of mayflies. In particular with insects that move with the wind, the radial velocity usually depicts the radial component of the actual flow.

Slide 15: A hail spike is an anomalous echo that is observed along a ray. It is also known as three-body scattering and happens in intense thunderstorms in which large hail is present. In this example, suppose the radar is located off the top right of the image, and the beam is pointing along the white arrow. The beam will encounter the strong scatterers in the hail core, which is denoted by the high reflectivity. However, for large, irregularly shaped hail, some of the radiation is scattered not back to the radar, but down to the ground. The radiation then reflects off the ground, back to the hail, and then back to the radar. Therefore, although much of the radiation is backscattered off the hydrometeors directly to the radar, there is some delay for a small amount of the transmitted signal. Since the radar processor determines the range of an echo based on the time delay between transmission and receipt, the processor interprets the delayed power return as reflectivity farther along the ray than where the hail core is actually located. Hail spikes such as these are indicative of a severe thunderstorm with particularly large hail.

Slide 16: Similarly, a sunrise or sunset spike occurs along a single ray in the direction of the sun when it is near the horizon. In this case, the echo is not reflecting off the sun of course, but rather the sun is a source of microwave emissions that the radar interprets as returned power. Technically, this is a form of radio frequency interference, in which the radar is detecting radiation from a different source than the signal it transmitted.

Slide 17: Side lobes occur because the antenna gain function is not perfectly aligned along the main lobe. Some power transmitted by the radar leaks into directions that are azimuthally displaced from the direction that the antenna points. As a result, a small amount of received power may actually come from somewhere other than the main lobe. This mainly occurs when the main lobe is pointing toward a region that is immediately alongside a particularly intense echo. In this case, the main lobe experiences little backscatter. Even though the side lobe contains less power, the backscattering by the large hydrometeors it encounters efficiently returns power to the radar. Thus, the radar receives comparable power from both the main lobe and the side lobe, but the processor interprets all of the power as having come from the main lobe. As a result, some non-existent echo may be depicted in clear-air along the edges of particularly intense convection containing large targets.

Slide 18: Second trip, or range-folded, echo happens when power is returned from targets located beyond the maximum unambiguous range of the radar. Suppose the maximum unambiguous range, which is controlled by the PRF given a radar with some wavelength, is 100 km. That means that radiation can travel out 100 km and back 100 km to the radar between pulse transmissions. However, radiation that is not scattered will continue to travel beyond the 100 km. Suppose that a tall cloud is present at 150 km and that radiation transmitted by one pulse scatters off the cloud and back to the radar. The problem is that the radiation doesn't get back to the radar until after a second pulse has been transmitted. The radar processor then sees the returned echo from the cloud 150 km away as being the returned signal of the second pulse. Therefore, the radar processor would think that an echo was located only 50 km from the radar. This type of false echo is called second-trip echo. It can be identified easily using L_{DR} for a polarization-agile radar, but without L_{DR}, second-trip echo is usually obvious by its elongated non-physical looking echo that extends along one or more adjacent rays. Third-trip echo is even possible for particularly tall clouds located well beyond the unambiguous range of a radar.

Slide 19: Finally, radio frequency interference, or RFI, can often occur at radar frequencies shared by other systems. For example, many communication systems operate in C-band. This reflectivity data is an example of C-band data collected aboard a ship with a C-band communication system. The communication system acted as an additional source of C-band radiation that the radar processor perceived as being a return echo from transmitted radiation. In this example, however, there is some meteorological echo that is embedded within the widespread RFI.

Slide 20: We can attempt to remove the RFI such as in this imperfect example, which still shows some RFI in various locations. However, the convective echo is much easier to pick out after a simple correction is made to remove any signal that does not exceed some threshold set by the background noise.

Module 5.7

Slide 1: This module will introduce some aspects of rainfall estimation via radar, with a particular emphasis on converting radar variables, including radar reflectivity factor and other dual-polarimetric variables, to rain rate.

Slide 2: We may be interested in estimating rainfall for several reasons. First of all, rain gauges only provide information about point locations on the ground, but radar data provides continual coverage in space, although with larger uncertainty. Therefore, the potential for flash flooding can be warned using radar data. Rainfall estimates derived from radar can be used for a variety of other research and technical applications as well. For example, because precipitation is related to latent heat release in the atmosphere, the amount of rainfall observed via remote sensing can provide some insight on the structure of the general atmospheric circulation.

Slide 3: The most common way to estimate rainfall from radar is to use the radar reflectivity factor in a Marshall-Palmer, or Z-R, relationship. It is simply a power law relationship between reflectivity and rain rate. Some examples of Z-R relationships used for conversion of WSR-88D reflectivity to rain rate are shown at the bottom. Each type of rainfall regime possesses its own pair of coefficients *a* and *b* seen in the equation above. The relationships are obtained empirically, as we will see an example of shortly. The relatively recent dual-polarimetric capabilities of WSR-88D radars allow for usage of Z_{DR} and K_{DP} in addition to Z to estimate rainfall. We'll see examples of such relationships soon. However, for radars without dual-polarimetric capability, only Z-R relationships can be used.

Slide 4: We'll start by looking at the simple Z-R relationships in more detail first. The table shows how rain rate (in mm/hr) increases for increasing reflectivity. That increase differs dependent upon the type of convection present. A 40 dBZ echo in deep convection in a moist environment will have a larger rain rate than the same echo in a region of gentler ascent. However, all of these estimates are deterministic. As we know, a variety of drop size distributions and shapes within a range gate can give us the same Z. However, that variety may represent vastly different volumes of water.

Slide 5: This figure illustrates observed rain rate as a function of reflectivity. Both variables were derived from observations made by a disdrometer, a ground-based instrument that directly measures the drop size distribution. From the drop size distribution, both the liquid water volume and the equivalent radar reflectivity factor can be determined. The red and blue dots on this plot represent data from two tropical islands that generally agrees well. The black solid and orange dashed lines represent Z-R relationships. Note again that these relationships may vary significantly depending on the location and type of rainfall. For example, Z-R relationships over the continental U.S. during winter would likely be very different.

Slide 6: Let's take a closer look at the estimated rain rate for a reflectivity of 30 dB*Z*. The estimated rain rate is 3 mm/hr, which is denoted by the horizontal green line intersecting the black and orange lines at 30 dB*Z*. However, if we consider the spread in observed rain rate for the bulk of observations at 30 dBZ, we might actually get a rain rate anywhere between about 1.5 and 6.5 mm/hr. That's almost a halving or doubling of the deterministic estimate! So clearly, Z-R relationships—or any empirically derived dual-pol relationships for that matter—are limited in their ability to absolutely estimate rainfall rate at a single point in time and space. At 40 dB*Z*, the absolute uncertainty in rain rate is even larger, with a potential rain rate roughly spanning about 8 to 30 mm/hr!

Slide 7: As opposed to a deterministic estimate of rainfall, which assigns a single estimate of rain rate based on one value of reflectivity, we can express potential rainfall as a probabilistic range if we have enough data to create a well-populated distribution of rain rate across a wide range of reflectivity. The plots shown in the top row of this slide are actually just zoomed in on 1 dB wide sections of the plot on the previous slide. For example, in the middle plot on the top row, we can see more clearly how the rain rate ranges from about 1.5 to 6.5 mm/hr. Other spreads can be seen in other 1 dB wide reflectivity bins. Thus, instead of reporting rainfall as some median or mean value, it can be reported as a probability distribution function, or PDF, three of which are seen on the bottom row. Each PDF peaks along the y-axis at the most likely rain rate listed on the x-axis. Uncertainty increases as the PDF widens. For the middle bottom panel, note how the PDF has a bit of a secondary bump near a rain rate of 4–5 mm/hr. Let's look at this more closely.

Slide 8: Precipitation can occur in regions of strong vertical motion, which is often alluded to as convective rainfall, or in regions of weaker vertical motion, which is called stratiform rainfall. Moderate reflectivity, such as around 30 dB*Z* is frequently observed in both types of rainfall regimes, whereas low reflectivity—say less than 20 dB*Z*—is more often found in stratiform, while very strong reflectivity—say in excess of 40 dB*Z*—is generally found in strong convection. In a little bit, we'll briefly discuss how the two rain-types can be determined using radar data. However, they can also be roughly estimated at a disdrometer based on the drop size distribution. When we compute rain rate PDFs for the 30–31 dB*Z* bin separately for convective and stratiform precipitation, we get two very different distributions that only overlap in their tails. Stratiform rainfall at 30–31 dB*Z*, denoted by the blue line, typically rains within a range of about 1 to 4.5 mm/hr, while convective rainfall with the same reflectivity typically rains at about 2 to 8 mm/hr. Thus, the primary peak in the total PDF, or the black line is mainly caused by stratiform rain, and the long tail in the distribution with the small bump near 4–5 mm/hr is caused by convective rainfall having the same reflectivity. This means that Z-R relationships can be separately computed for convective and stratiform rainfall to better reduce the uncertainty in an estimate.

Slide 9: There a few reasons that such spread exists in the rainfall PDFs. Part of the reason has to do with variability in drop size distributions. The top left figure shows outlines of raindrops captured by an aircraft-based cloud drop probe. As you can see, a wide variety of drop sizes can occur depending on the kinematics found within a cloud. But even within the same part of a cloud, a variety of drop and droplet sizes are present. The bottom right panel shows examples of drop size distributions seen in various tropical cyclones, which each color representing data from a different storm. The distributions all occurred when reflectivity was near 40 dB*Z*. Drop size distributions often follow approximate gamma distributions, which is more the focus of a cloud physics course than remote sensing. The main point here though is that at the tail end of the upper part of the distribution, the number of the largest drops varied—sometimes significantly—between different storms. For example, the orange line represents about one order of magnitude more 3 .5 mm wide drops than the other lines. The reflectivity is most sensitive to the size of these drops, meaning that the reflectivity is similar for all of the distributions. However, some distributions actually contain more water than others and would therefore likely result in a larger rain rate.

Slide 10: To get a 40 dB*Z* echo, a volume could consist of a very large number of small drops, although this is unlikely. More likely, the volume would contain a relatively small number of large drops…

Slide 11: …or somewhere in between, with a few large drops and a bunch of small drops that make little impact on reflectivity used to estimate rain rate. The canting angles of the drops matters as well because it affects the backscattering cross-section to the radar, which is especially important for radars without dual-polarimetric capability.

Slide 12: A few ways exist to get around the problem of having such wide-ranging possibilities for rainfall based on a single observation of reflectivity. One commonly used technique called kriging involves blending the radar estimates with relatively isolated rain gauge and disdrometer data located within a radar domain. This works quite well in places where you have ground-based data. Shown here is an example of 1-day observed rainfall product from NOAA over the southwestern United States that is an example of a product derived using such a blended technique.

Slide 13: However, over the ocean, we don't have rain gauge or disdrometer data outside of the occasional point on an island. In these cases, we are far less able to constrain rainfall rates satisfactorily. This is where a probabilistic estimate of rainfall can become particularly useful. However, deterministic estimates like the Z-R relationships shown earlier are still frequently used.

Slide 14: To the right is a diagram outlining one example of a decision tree for dual-polarimetric rainfall estimation that uses not just Z but also Z_{DR} and K_{DP} to estimate rainfall. Four different types of relationships can be formed based on the magnitude of each variable. An example of a Z- Z_{DR} -R relationship is seen at bottom left. Z_{DR} and K_{DP} are typically only used if they are large and indicate that oblate rain drops are present, which would skew the rain rate estimated using Z only and assuming a spherical drop. In the event that both Z_{DR} and K_{DP} are small, a standard Z-R relationship can be used, using separate ones depending on the type of precipitation occurring.

Slide 15: Spread is still present in rain rate even when using dual-pol relationships to estimate rainfall. Separate dual-pol relationships can be computed for convective and stratiform echo as well, which helps to constrain the uncertainty. The 2D histogram of Z and Z_{DR} is shown at top left, with an example of its split into convective and stratiform seen at top center. The other panels show disdrometer-derived fields associated with the drop size distribution in the Z-Z_{DR} space.

Slide 16: A probabilistic method again can account for the spread in potential rain rates given some value of reflectivity. Consider also that the data we looked at on the last slide was derived from ground-based data; however, radar data observes backscattered radiation at some altitude above the ground. Therefore, uncertainty associated with estimating rainfall from radar comes not just from the drop size distribution in a volume, but also the discrepancy between the reflectivity observed at some height and the corresponding reflectivity directly beneath at the ground where the rainfall is actually observed by a disdrometer from which rain rate relationships are empirically derived. The plot shows a 2D histogram that describes rain rate for a 31 dBZ echo observed at an altitude of 1300 meters. The actual reflectivity at the ground in this subset of data for such an echo may range from anywhere between 26 and 38 dBZ, which, given the distributions of rain rate we explored earlier, ends up corresponding to a potential rain rate between 1 and 8 mm/hr. The maximum likelihood is denoted by the center of the PDF, which is around 3–4 mm/hr for a most likely surface reflectivity of 33 dBZ, a little larger than the 31 dBZ measured above the ground.

Slide 17: This type of approach leads to its own set of complications as well, however, and we will not dive deeply into them all. However, spatial and temporal autocorrelation must be considered. For example, if we know something about the drop size distribution in one part of the radar domain, this probably allows us to narrow the likely width of PDFs used to estimate rainfall elsewhere in the radar domain. Mathematically modeling this autocorrelation is an active area of research.

Slide 18: Finally, we have mentioned convective and stratiform rainfall at length in this module because the drop size distributions found within the two basic types of rainfall differ substantially. The rainfall types can be identified subjectively in radar data, allowing us to apply different rain rate relationships in space.

Slide 19: One common way to do this is via evaluation of the two-dimensional reflectivity field at some height generally 1.5–2.5 km above the surface. Local peaks in reflectivity or very high reflectivity generally indicate convective echo, while weaker reflectivity that varies more smoothly in space is more likely to be stratiform. In this image, an example of reflectivity data is shown at left, with the corresponding rain type classification for the same echoes illustrated at right. Some isolated echoes, which tend to be shorter than deep convection, are indicated by shades of blue.

Slide 20: The performance of such methods can be evaluated if we know the vertical profiles of vertical motion or latent heating within them. These are difficult to measure directly but can be simulated in a model. This figure simply shows that the profiles of each in two different algorithms are roughly consistent with what we would expect from atmospheric dynamics and thermodynamics. Convective rainfall generally is coincident with latent heat release through a deep layer of the atmosphere, while stratiform rainfall generally occurs in regions where diabatic heating is maximized above the 0°C level but cooling occurs beneath. The top and bottom rows show the results of mean vertical velocity and latent heating profiles using two different algorithms and shows that the separate classification of shallow convection is important in order to not incorrectly classify it as stratiform rainfall, which would thereby cause an incorrect rain rate relationship to be used.

Module 5.8:

Slide 1: So far in this lecture series, our discussion has primarily been focused on precipitation radar. In this module, we will explore higher-frequency radar operating in the Ka- and W-bands, which are used for detection of smaller hydrometeors.

Slide 2: Going back to the transmittance of an atmosphere containing an average amount of water vapor, the attenuation of a radar beam will generally increase with increasing frequency. In a previous module, we saw that attenuation at S-band was negligible for most applications, but that attenuation at X-band was noticeable, especially in heavy precipitation and at long ranges. Ka-, and certainly W-band, radar is much more susceptible to attenuation than X-band, so its use is limited to short ranges and constrained by radar properties such as the intensity of transmitted power and the antenna gain. The high frequency radars are called cloud radar because their shorter wavelengths are more susceptible to scattering by small non-precipitating cloud drops that only have minimal impact on longer wavelengths.

Slide 3: Many cloud radars are vertically pointing, meaning that the antenna is fixed, only looks upward, and observes clouds as they pass over the radar. This picture shows two different cloud radars next to each other. On the left is a W-band radar; on the right is a Ka-band radar. Because of its longer wavelength, the antenna for the Ka-band radar is larger.

Slide 4: Vertically pointing Ka-band radar data may look something like that shown here. At the top is the radar reflectivity factor, and at bottom is the mean Doppler velocity. As before, negative radial velocities indicate motion toward the radar, which in this case is downward. The time axis progresses from left to right over the course of a day. Focusing on the reflectivity first, we see the leading edge of a precipitation system reach the radar, and high reflectivity—in excess of 35 dB*Z*—is seen early on near the ground. In subsequent hours, the reflectivity is lower in the stratiform part of the precipitation system. Finally, by 1800 UTC, rain mainly stops, but the radar can observe non-precipitating anvil cloud with reflectivity less than 0 dB*Z*. Typical weather radars would not see the non-precipitating cloud with such detailed structure. The 0°C level is also very obvious in vertically pointing radar data. Above this level, the power returned to the radar is generally lower due to the lower dielectric constant of ice, and the radar processor interprets this as a reduction in reflectivity. The background reflectivity of -30 dB*Z* or less is noise and can be removed in post-processing. At the bottom, the radial velocity is generally downward toward the radar where precipitation is occurring. The red colors likely are folded velocities, such that they are actually strongly downward but are depicted as upward because they exceed the Nyquist velocity.

Slide 5: The data can also provide other fields such as spectral width, shown at top, and the signal to noise ratio, illustrated at bottom, is a useful tool for helping to quality-control the radar fields and eliminate background noise.

Slide 6: Precipitation radars that transmit extremely high power also have increased sensitivity and can detect features within precipitating and non-precipitating clouds that typical weather radars cannot. One example from the pulse Doppler Mid-Course Radar (MCR) near Cape Canaveral, FL, is shown here. This radar is operated by the U.S. Navy and was previously used for tracking debris from space shuttle launches. By transmitting large amounts of power, the sensitivity of the radar is increased, such as seen by the -20 dB*Z* echoes observed in the raining cloud pictured at top left. In fact, this radar can see individual raindrops, which is useful for remotely sensing the drop size distribution in clouds—a critical step toward improvement of microphysical parameterization. Most radars, however, do not have this capability.

Slide 7: Scanning cloud radars are similar to the vertically pointing radars that we saw earlier, but they are mounted on a pedestal with gears that allows the antenna to rotate just like a precipitation radar. These are sometimes mounted in dual-frequency configuration, such as the two examples shown here, which pair Ka- and W-band antennas on the left, and X- and Ka-band antennas on the right.

Slide 8: Compared to precipitation radar, the properties of scanning cloud radar are quite different. For example, scanning W- and Ka-band radars often operate with PRFs of nearly 5000 Hz, compared to roughly 1000 Hz for S- or C-band radar. This greatly reduces the maximum unambiguous range of these radars, but this is acceptable considering that W- and Ka-band radiation get attenuated and would not complete two-way propagation over 100–150 km anyway. Even though the PRF is much larger than precipitation radar, the Nyquist velocity is still smaller due to its wavelength dependence. In other words, the PRF increases by a factor of 5, but the wavelength of Ka- or W-band is a factor of 10 or more less compared to S-band. The beamwidth and gate spacing for cloud radar tend to be small, allowing for high-resolution observations over the shorter distance observed. And finally, the sensitivity for cloud radars, while dependent upon the properties of the transmitted radiation, is typically somewhere between -20 and -30 dB*Z*. Newer radars are more commonly equipped with polarization-agile transmission.

Slide 9: While precipitation radars are largely used in operations and research, cloud radars today are still generally used almost exclusively for research applications. Therefore, the scan strategies can be adapted to the particular research objective, and PPIs like those used by weather radar are often not used with cloud radars. Some different examples of scan strategies are shown here. The top left is a standard PPI, with the hollow cones representing different tilts, and the horizontal cross-sections representing flow that might be derived from the single-Doppler data. The top right is a standard RHI scan strategy, the bottom left is a vertically pointing strategy, and the bottom right is an example of how a reflector with known shape and size can be used for calibrating the radar.

Slide 10: Other types of scan strategies usually are some type of RHI that may be used depending on the applications. The top left is called a hemispheric RHI, in which the antenna samples an entire 180° sector—including overhead—before moving to the next pair of azimuths. The top right is a standard RHI scan strategy, but with elevation angles that only cover the boundary layer. The bottom scan strategies are sector RHIs that point vertically and at angles slightly off of zenith. These scans can be configured in a cross-wind or along-wind orientation.

Slide 11: An example of the vertically pointing sector RHI data is shown here as a cross-wind scan for a convective cloud passing over a Ka-band radar.

Slide 12: Unlike vertically pointing fixed antennas, a scanning radar can collect a three-dimensional volume of non-precipitating clouds, analogous to how a precipitation radar views a three-dimensional volume of raining targets. Such non-precipitating clouds might include shallow cumuli or stratocumuli, stratus clouds, mid-level cumuliform and stratiform clouds, or cirriform clouds.

Slide 13: High frequency radars used for cloud detection are also operated from space. A precipitation radar, which operated in Ku-band, was aboard the Tropical Rainfall Measurement Mission satellite, which operated from 1997–2015. Although technically not a cloud radar, its operation at Ku-band was much more susceptible to attenuation than typical weather radar wavelengths. TRMM was able to provide spatial maps of rainfall through the tropics, and has since been replaced by the Global Precipitation Measurement, or GPM, satellite, which carries a dual-frequency Ku-/Ka-band radar system. An example of the swaths of the radar data compared to the swath of a microwave imager on the same platform is seen as right. The Ka-band radar is capable of detecting non-precipitating clouds more readily than the Ku-band radar. The differential attenuation experienced between the two bands can also be leveraged to identify rain and snow.

Slide 14: To obtain a resolution in the troposphere of 5 km x 5 km at nadir, however, both antennas must be much larger than corresponding ground-based radars operating at the same frequency. The size of the radar system on GPM is seen at top left next to a person for scale, while the bottom right image shows the companion ground-based radar operating at the same frequencies.

Slide 15: GPM is part of a family of satellites that carry microwave imagers, such as those that we discussed in the previous lecture series. However, GPM is unique in that it carries downward pointing precipitation radar. It operates in low-Earth orbit, but it is in a prograde, non-sun-synchronous orbit so that it can sample locations at various times of day—allowing for characterization of the diurnal variability in rainfall. It also operates at a higher inclination than the TRMM satellite, which enables GPM to estimate precipitation over a wider band of latitudes.

Slide 16: Examples of Ku- and Ka-band radar data along a swath following the GPM ground track is shown here. The swath of data—converted from reflectivity to rain rate—is seen at left. The cross-sections at right follow the ground track denoted by the black line in the left panel. At right is radar reflectivity factor, although without numbers attached to the color scale. Warmer colors denote higher reflectivity. Generally, the Ku-band reflectivity is higher than the Ka-band reflectivity because the Ka-band beam is more heavily attenuated as it propagates downward through the atmosphere.

Slide 17: An example of a global rain rate map is shown here. It is updated every 30 minutes and is derived partially from the GPM radar data, but largely from a combination of passive microwave and infrared brightness temperatures.

Slide 18: Another downward pointing radar is called CloudSat. It is a W-band radar that is more sensitive than the radars aboard GPM and is more capable of identifying optically thin cirrus clouds and shallow cumuli. In the example shown here, CloudSat, which is now part of the C-train constellation with CALIPSO, passed over a tropical disturbance off the southeastern coast of the United States. The satellite was on its descending node, and encountered three regions of clouds. A cross-section of the nadir pointing cloud radar data is seen at bottom, with time shown along the x-axis. Non-precipitating cloud is seen at both high and low altitudes. Note that the in the deeper convection, the reflectivity is largest near the tops of the convection. This happens because the W-band beam is attenuated rapidly at it moves through the cloud. As a result, a minimal amount of transmitted power remains for backscattering in the lower troposphere. Thus, even though drop size distributions are probably large close to the surface, the space-based cloud radar cannot detect this. In contrast, in the same echo, a ground-based cloud radar would see high reflectivity close to the surface and would be attenuated rapidly with height such that it would not clearly see the top of the cloud.

Slide 19: As a result, space-borne cloud radar is most effective for detecting the tops of deep convection and optically thick cirrus cloud. Vertically pointing ground-based radars can also detect non-precipitating clouds at multiple levels. Ground-based or space-borne lidars, the subject of a later module, operate at visible or near-IR wavelengths and are especially useful for detecting optically thin cloud that radars are not sensitive enough to detect.

Slide 20: Precipitating cloud is best observed from space using lower frequency precipitation radars, such as those operating in Ku-band, or using ground-based or aircraft-based precipitation radar operating in X-, C-, or S-bands.

Slide 21: Shallow non-precipitating convection in the boundary layer is best observed using ground-based or aircraft-based cloud radars to minimize the impact of attenuation on reducing signal strength compared to space-based radar; however, in the absence of heavy precipitation, space-based radar or lidar might detect such clouds. Any of these instruments may detect any of these types of clouds depending on the properties of transmitted radiation and the range of the target. However, these are the most common applications for the various instruments discussed in this lecture series.

Module 5.9

Slide 1: This short module discusses phased array radar systems.

Slide 2: A phased array consists of several antennas that are controlled by a central computer. Unlike a traditional pulse radar, which has its antenna mounted on a rotating pedestal, a phased array is stationary, and the direction of the collective beam of the array is determined by the time lag between transmission between the antennas in the array. A delay in transmission at one end of the array compared to the other causes constructive and destructive interference that causes the beam to propagate into the direction of the last transmission. In this animation, the bottom antenna transmits first, followed immediately after the other antennas, proceeding upward on the figure. In this example, the combined waveform moves at an angle toward the direction of the last antenna to transmit. All of the antenna are linked to the same transmitter and receiver, but each antenna is controlled by a different phase shifter that controls the phase of the emitted wave such as to also contribute to the propagation direction of the combined transmitted waveform.

Slide 3: The major advantage of a phased array radar system is its ability to scan in many directions in rapid succession without the need of moving an antenna around. A typical volume is collected every 5 to 15 minutes. Phased array radar, using multiple arrays facing different directions on the same platform, can scan the same volume in less than 1 minute.

Slide 4: An example of a phased array radar in Oklahoma is seen here with the protective radome pulled up. While this radar has a rectangular shape, radars can also have different shapes with multiple surfaces containing small phased arrays. An example of Doppler velocities seen by this phased array radar compared to those detected using WSR-88D in the same storm are seen at far right. In the top panel, the same supercell is observed simultaneously. Over the next five minutes, the phased array radar captures a volume of the supercell four times before the WSR-88D again scans the cell at a low elevation angle. Within those few minutes, dual, counter-rotating tornado vortex signatures were seen by the phased array radar. By the time the WSR-88D scanned the area again, the signatures had weakened. Thus, the phased array radar enables more rapid and consistent detection of short-lived events.

Slide 5: An example of the phased array radar data collected during the 2013 El Reno tornado is seen here. The EF3 El Reno tornado contained some of the most powerful winds ever recorded in a tornado; these winds were obtained via mobile Doppler radar. Here is shown the radar reflectivity factor, updated roughly once a minute. Again, the temporal resolution provided by phased array is far superior to standard WSR-88D. In the future, the radar network in the U.S. may be replaced by phased arrays as they become more affordable.

Slide 6: The diagram shown here summarizes the many capabilities of phased array radar. It can rapidly perform surveillance PPI-like scans like a rotating weather radar but can also be committed to tracking certain features of interest in numerous narrow sectors over very short amounts of time, such as when observing storm cells or tracking aircraft.

<u>Module 5.10</u>

Slide 1: This module introduces lidars, which are active remote sensing instruments that transmit visible or near-infrared radiation. As we mentioned in a previous module, for meteorological purposes, lidars are used most often for detection of very small targets, such as small ice crystals, cloud droplets, or aerosols that act as cloud condensation nuclei. A couple of lidars are pictured here: One is a downward pointing lidar inside an aircraft, and the other is a vertically pointing lidar located inside a trailer.

Slide 2: Like radars, lidars transmit radiation and process the part of the transmitted signal that is backscattered to the sensor. The lidar equation is shown here, color coded by term, with the power received, the fundamental variable detected by the receiver, on the left-hand side of the equation. The most useful variable that we generally want to derive from the observed power is the backscatter scatter, denoted by the bolded beta in the equation; it is conceptually analogous to the radar reflectivity factor in the radar equation. This requires that we know values for the remainder of the terms in the lidar equation. The power received is, of course, directly proportional to the power transmitted. It also scales with the area of the telescope and inversely-squared with range. The yellow term is the lidar constant that is a function of properties of the lidar itself such as pulse length. Sometimes the telescope area and transmitted power are wrapped into the lidar constant. The final term depends on properties of the medium through which the lidar beam propagates and represents two-way attenuation of the lidar beam. This term should look similar to some that we encountered when discussing Schwarzchild's equation earlier in the course. The overlap function, shown in blue, represents the range-dependent fraction of the transmitted signal's cross-section that is contained within the field of view of the receiver. We won't delve into the lidar equation with much detail, but we will look at some examples of lidar backscatter shortly.

Slide 3: Recall from earlier in the course this figure that depicted the scattering susceptibility and type of scattering predominant for several combinations of radiation wavelength and scatterer size. Because lidar operates in visible light or near-infrared, it is susceptible to scattering by much smaller objects such as aerosols and cloud drops. We can see this by looking for at what size scatterers on the y-axis the size parameter is greater for visible and near IR wavelengths. Any larger objects, such as liquid water cumulonimbus clouds, or even the ground, will also cause a strong scattering interaction.

Slide 4: An example of backscatter coefficient, plotted in microns per steradian, is shown here for a lidar operating at 532 nm. Detailed structure is seen in the lowest 2.5 kilometers where the backscatter is primarily caused by aerosols, which could be sea salt, dust, smoke particles, or other small particles. Above this layer, cloud is present throughout the time series. Sometimes low cloud extinguishes the beam at low levels. At most times, the beam is able to partially penetrate to altitudes above 2.5 km. For much of the time, the lidar sees cirrus cloud located between 10 and about 16 km. However, between 04 and 10 UTC, the lidar beam encounters optically thick cloud. The beam becomes attenuated, so only power from the bottom of the cloud is returned to the sensor and any cloud above this region of backscatter is not detected.

Slide 5: Lidar can also be used to map topography and coastal bathymetry like this example derived from aerial lidar near the Bixby Bridge along the Big Sur coastline using the same principles as discussed before. Topographic lidar generally uses near-infrared radiation at 1064 nm, while bathymetric lidar uses green light at 532 nm that can partially penetrate water.

Slide 6: Doppler lidar is capable of measuring Doppler velocity to and away from the sensor along with backscatter. The following is an example of Doppler lidar data showing backscatter and vertical motion in clear-air aerosol, low cloud, and some upper level ice cloud. Like radar radial Doppler velocities, negative values indicate targets moving toward the lidar. For aerosols that are suspended in the atmosphere, the Doppler velocities are very close to the actual vertical component of air motions while within clouds, in-cloud dynamics superimposed onto the fall speed of hydrometeors determines the Doppler velocity, just as for cloud radar.

Slide 7: A water vapor differential absorption lidar, or DIAL, is a special type of lidar that is used to detect humidity at short range. The DIAL uses two narrow bands near each other. One of the bands is located within a water vapor absorption band near 727 or 815 nm, and the second band is located adjacent to the first but just outside of the water vapor absorption band. The top panel shows an example of this concept with one band located in one of four water vapor absorption bands depicted and another band outside. The two adjacent wavelengths experience similar propagation through the atmosphere except that one is impacted by water vapor and the other is not. Thus, the difference in backscatter between the two can be attributed to the extinction coefficient in the lidar equation, which is related to the volume extinction coefficient, which is determined by the water vapor concentration. Like a radar, the DIAL is able to operate as a function of range, with radiation in the band not impacted by water vapor having a slightly larger return, with the discrepancy between the power returned in the two bands increasing with range unless the beam encounters completely dry air. The gate spacing depends largely on the magnitude of power transmitted. An eye-safe DIAL may yield gate spacing on the order of 100 meters.

Slide 8: An example of water vapor DIAL data from a vertically pointing instrument is seen here in the top panel. The middle panel shows satellite-derived water vapor concentration of the same location at the same time, and the bottom panel shows the same but derived from a ground-based vertically pointing. microwave radiometer. While the three agree to first order, the DIAL captures much more structure in the humidity field, especially in terms of its improved range resolution.

Slide 9: A Raman lidar is a complicated, sensitive instrument that transmits ultraviolet radiation and detects its backscatter. Like the DIAL, it transmits in narrow bands near molecular nitrogen, oxygen, and water vapor absorption bands in order to collect remotely sensed measurements of temperature and humidity.

Slide 10: Well-calibrated Raman lidars are extremely accurate sensors. On the left is an example of a timeseries of water vapor mixing ratio profiles measured by a Raman lidar. Again, detailed structure in the humidity field is visible. On the right is a comparison of the mixing ratio derived from weather balloon data—the blue line—and the mixing ratio derived by the Raman lidar in red. The two match very closely, proving the utility of lidar for profiling the thermodynamic structure of the atmosphere.

Slide 11: Space-based lidar may be used to detect aerosols and extremely small hydrometeors that may be invisible to the naked eye—especially in sub-visible cirrus at high altitudes. CALIPSO is a companion satellite to CloudSat in the C-train and operates at both 532 and 1064 nm. It uses a very narrow pencil beam pointing at nadir and captures profiles of aerosols and clouds with very high resolution. Like cloud radars, the lidar is strongly attenuated in deep convection and will not see low in the atmosphere in such cases.

Slide 12: An example of CALIPSO data over a swath near the Date Line from nearly 50°N to 30°S shows several high and low clouds over the Pacific Ocean. Note how the cloud top is lower at higher latitudes than it is in the tropics, where cloud top extends to as high as 17 km. Many shallow clouds and/or aerosols are seen, especially in the center of the image. In the presence of high cloud that is optically thick to visible or near-IR light, the beam is attenuated enough such that it does not detect returned power from beneath the cloud in many places. The dark blue vertical stripes you see beneath clouds are caused by extinction of the beam and represent where the lidar cannot detect scatterers.

Slide 13: The image shown at left is an example of radar and lidar derived fields for a co-located vertically pointing radar and lidar observing a passing mesoscale convective system. The top panel shows the Ka-band reflectivity. It sees some high cirrus cloud as well as low-level precipitating cloud. However, the cloud radar is attenuated heavily when the heaviest rain passes over the instrument, and this denoted by the white area of missing data during the rain event. The second panel denotes lidar backscatter. It is almost completely attenuated where rain occurs, but it detects some high cirrus cloud that the radar was not sensitive enough to detect. Combined, the two instruments can be used to derive fields related to water content and create products that indicate where clouds are present. Another example of a cloud mask is shown at right. It is taken from a cross-section of CloudSat and CALIPSO data as they passed over a cloud system in rapid succession. Red and blue shading indicate parts of the cloud seen by radar, and blue and green colors indicate parts of the cloud seen by lidar. The CALIPSO lidar is effective at detecting high-altitude cirrus that cloud radar cannot detect; such sub-visible cloud is a common occurrence, especially in the tropics, and has important implications for radiative transfer in the context of Earth's climate system. Other than a few very shallow clouds at low altitudes with clear-air above, the lidar in space does not detect backscatter from scatterers in the lower atmosphere. Likewise, a lidar on the ground could not see through precipitating cloud to detect the optically thin cirrus near the tropopause.