Readings to accompany the PHYSICS of the ATMOSPHERE

Parts A-B and appendices

Robert A. Cohen
Physics Department
East Stroudsburg University

January 20, 2014
Version 6.0

Contents

1. Weather Forecasting

The purpose of these readings is to show how physics (and also chemistry and math) can be used to describe and make predictions about the Earth's atmosphere. Before we get down to the nitty-gritty details, let's first examine what it means to use physics to make predictions about the atmosphere.

When most people hear about making predictions about the atmosphere, they think of **weather forecasting**. Weather has to do with the day-to-day conditions, like temperature, precipitation and wind. However, we can also make predictions about **climate**, which refers to long-term temperature, precipitation and wind patterns. We could also make predictions about what happens to the air if we cool it, warm it, compress it, expand it, move it, or whatever.

All three types of predictions will be periodically mentioned in these readings. For now, though, let's consider weather forecasting and how physics, chemistry and math can be used to forecast the weather.

The reason for doing so is two-fold. First, it provides a context for the chapters that follow. Second, it serves to clarify exactly how weather forecasts are made (i.e., with physics, rather than with statistics).

1.1 Local forecast method

Before explaining how physics is used to make weather forecasts, I will first describe the process that many people mistakenly think is used to forecast the weather. Let's call this technique the *local forecast method*. In this method, we compare today's weather at our location with the entire weather record for our location. We then search for a match, where today's weather is similar to the weather experienced sometime in the past. The matched days could be based on lots of different parameters, like the temperature and precipitation, or whether the temperature was rising or falling.

The forecast for tomorrow, then, is based upon what the weather was like following the “matched” day. If there are several “matched” days, all with different weather the following day, our forecast would be couched in probabilities. For example, if we found that 50% of the days were followed by rain, we would forecast a 50% probability of rain tomorrow.

The *local forecast method* is similar to the forecast techniques reflected in weather sayings, like “red sky at night, sailor’s delight,” in that like the *local forecast method* the sayings use observations made at a particular location to make forecasts for that particular location. No physics is used to make the prediction. Rather, it is simply based on identifying patterns in the historical record and matching a particular day’s weather to that pattern.

> ✍ The fact that many forecasts today also use the probability language (e.g., a 50% probability of rain) reinforces the misconception that the *local forecast method* is used to forecast the weather.

1.2 Synoptic forecast method

Except for the past two hundred years or so, the *local forecast method* was the predominant method of weather forecasting. Nowadays, however, meteorologists do not use this method for day-to-day forecasting, as it is very unreliable compared to two other techniques that I will now describe.

The first of these other techniques is to look beyond our local area and identify the weather all around us. This allows us to identify organized precipitation patterns and their movement. For example, Benjamin Franklin had recognized that Boston appeared to get rain a day or two after it rained in Philadelphia. He then surmised that rain must be associated with a storm that travels from place to place.

I’ll refer to this method as the *synoptic forecast method* because the method requires us to observe what the weather is at all locations at the same time. The word **synoptic** comes from *syn*, meaning same (as in synonym), and *optic*, meaning view (as in optician).

1.2.1 Synoptic maps

When the observations are all plotted on a map, we call the map a **synoptic map**. An example of such a map is shown in figure 1.1, which displays surface observations at various locations over the United States for 14Z on 04 January 2014, using the convention illustrated in figure 1.2. Only a few of the observations are shown, as putting all the observations would overcrowd the map.

> ✍ This is a *surface* map, which means that every point on the map is at the surface. It is as though we are looking *down* on the surface.

Since all measurements are taken at the same *global* time, they are taken at different times *locally*. For example, 1 PM Eastern Standard Time is at the same global time as 10 AM Pacific Standard Time. Because of this, synoptic maps are, by convention, identified according to Greenwich Mean Time (GMT), which is the local time at Greenwich, England (zero degrees longitude). This time zone is also sometimes called Z (for Zulu) or UTC (for Universal Time Coordinated).

> ✍ Notice that the map itself indicates the time in Zulu (Z) whereas I use UTC in the caption. They refer to the same time.

The local time is roughly one hour shifted for every 15 degrees of longitude (divide 360 degrees of longitude by 24 hours). Thus, a place at 75 degrees west (e.g., Philadelphia), is likely to be five hours behind GMT.[i]

Check Point 1.1: The information in figure 1.1 is for 14 UTC on January 4, 2014.
(a) What local time would it be in East Stroudsburg, Pennsylvania?
(b) What local time would it be in Greenwich, England?
(c) What local time would it be in San Francisco, California?

---

[i] Indeed, East Standard Time (EST) is five hours after GMT. During Daylight Savings Time, it is four after GMT.

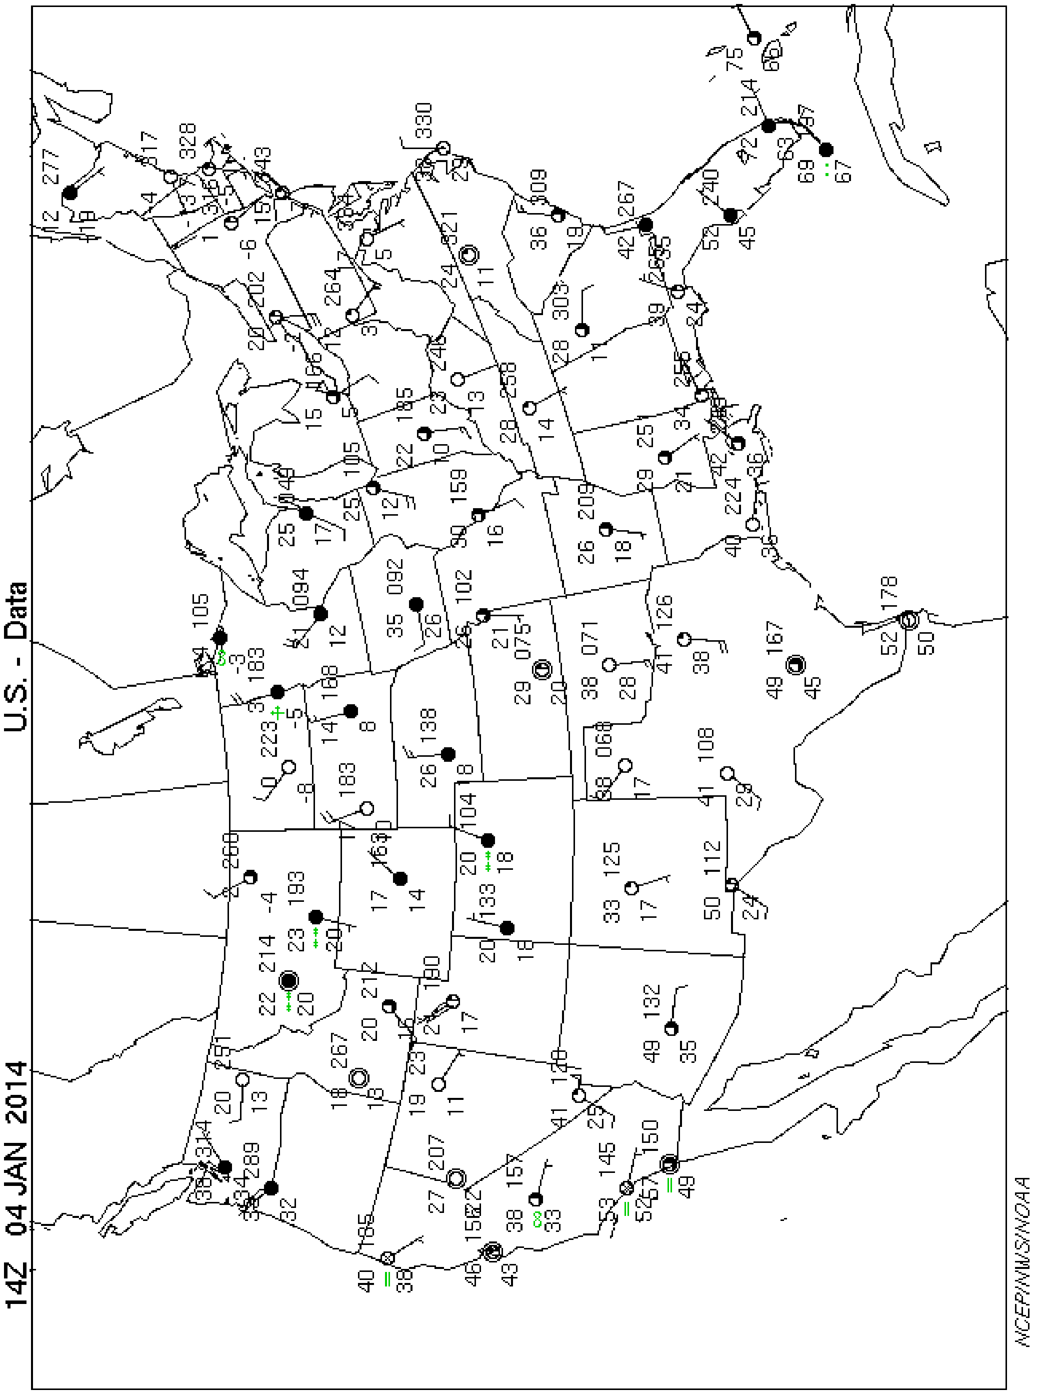

Figure 1.1: A map of surface observations for 14 UTC on 04 January 2014 (source: University of Illinois WW2010 ProjectDataStreme).

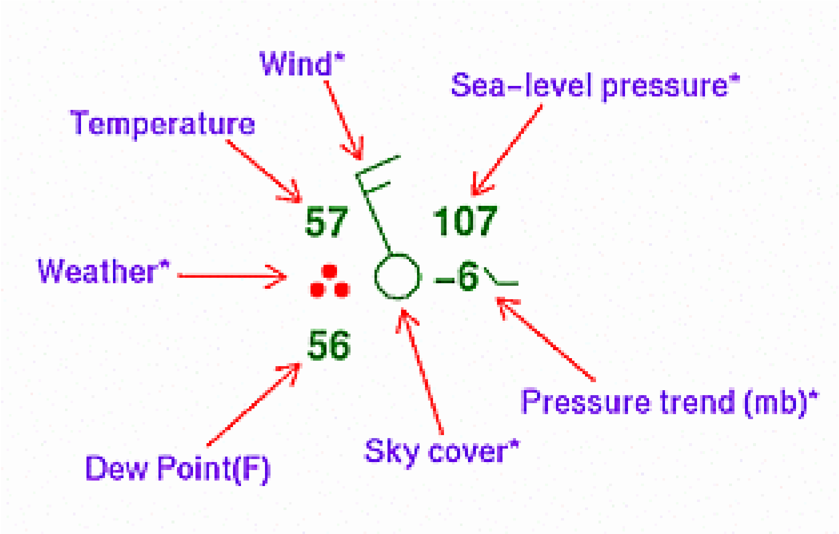

Figure 1.2: A sample station plot. For information about the asterisked items, go to http://www.hpc.ncep.noaa.gov/html/stationplot.shtml. (source: National Weather Service)

1.2.2 Polar front model

About one hundred years ago, the synoptic technique was greatly improved upon by the identification of a "typical" structure and evolution of a precipitation system, in which physics was utilized to explain how the system evolved and propagated. This model, which is known as the **Norwegian Cyclone Model** or the **Polar Front Model**, is still used today by many meteorologists to forecast the weather. The model essentially states that weather systems have a characteristic structure. By knowing that structure, we can make rough predictions of the weather.

For example, the air pressure at the surface is lower underneath the weather system than outside the weather system.[ii] Consequently, by tracking the trend of the local air pressure, one can anticipate the coming of a weather system.

Another part of the model states that the air moves around the center, counter-clockwise in the northern hemisphere, and that weather systems tend

---

[ii] For this reason, meteorologists tend to refer to weather systems as **lows**.

to move toward the east. That means that in the northern hemisphere the wind will be coming from the south prior to the arrival of the system and thus one can use the direction of the wind to determine whether a weather system is coming.

Finally, the model states that the wind ahead of the system (coming from the equator) will bring warm air with it while the wind behind the system (coming from the poles) will bring cold air with it. This means that it will be warmer ahead of the system and colder behind it. Thus, one can also use the trend of the temperatures to determine whether a weather system is coming.

As part of the temperature structure, the model states that there will be a relatively sharp boundary called a **front** between the warm and cold air. Ahead of the system, where the warm air is moving poleward, there is a warm front. Behind the system, where the cold air is moving equatorward, there is a cold front.

You don't need to understand much physics to use this model and make weather predictions. As such, I won't be discussing it much (beyond what I've already discussed), except when we encounter physics or math that can be used to understand it.

Check Point 1.2: Which of the following would expect after a weather system has passed?
(a) Increasing air pressure at the surface and decreasing temperatures
(b) Increasing air pressure at the surface and increasing temperatures
(c) Decreasing air pressure at the surface and decreasing temperatures
(d) Decreasing air pressure at the surface and decreasing temperatures

---

1.3 Numerical weather prediction

The final method I'll discuss, and the one that is most relevant to these readings, is called **numerical weather prediction**. The reason why it is called "numerical" is because it is very number intensive, so much so that it typically requires a computer to crunch the numbers.

Nowadays, almost all day-to-day weather forecasts (e.g., forecasts from one to five days) are based upon this technique.[iii] And, while the computer models used nowadays can be pretty complicated, the basic operation can be explained using the principles learned in your physics and chemistry courses.

To start, we need to know the temperature, humidity, wind and pressure throughout the entire atmosphere. I'll review each of these in part A of the textbook. You'll not only learn what they are but also what units are used and how they are measured.

Although we need to know the values of these things throughout the entire atmosphere, it is impossible to make a measurement everywhere. Instead, we make measurements at select locations and then assume that the atmosphere varies between our observation points in some uniform way.

At the surface, we either use a human being to make measurements or an *Automated Surface Observation Station* (**ASOS**). These are placed at various locations. Unfortunately, that only gives us information about the air near the surface. To get information above the surface, we release balloons (called **radiosondes**) that transmit information by radio waves indicating the temperature, pressure and humidity of the air it passes through. By tracking the balloon as it travels, we can also measure the winds (here we assume that the balloon moves with the air). In section 3.6.3, I introduce the method we typically use for displaying the information obtained by these radiosondes.

✍ We also use satellites and aircraft to obtain information about the atmosphere.

The problem with the observations is that they are not spaced equally apart. Some places have lots of data. Other places do not. What we want is a **grid** of data, where the state of the atmosphere is known at equally-spaced intervals both horizontally and vertically. Such a grid might consist, for example, of data points every 100 km in the horizontal and every 1 km in the vertical. At each data point, we'd like a particular value of the temperature, humidity, pressure and wind.

To "fix" this problem, we use an interpolation process. How this is done is beyond the scope of this course. However, it essentially consists of interpolation

---

[iii] The technique has only been widely used for the last fifty years or so because this technique utilizes computers.

(i.e., estimating the value at a location that lies between two measurement locations) while ensuring that the values are realistic. At the same time, the process must take into account the uncertainties and possible errors in the measurements. The whole analysis procedure is part of the **initialization** process.

The next step is to predict how the temperature, humidity, pressure and wind at each data point will change over a very short time period, called a **time step**. For a grid with data points every 100 km in the horizontal and every 1 km in the vertical, the time step might be only one or two minutes.

The reason why the prediction is only made for a short time period is because the prediction for each data point is made using only the values at neighboring data points. This is because for short time periods the temperature, humidity, pressure and wind at a particular point is only influenced by the air immediately surrounding the point. The assumption is that the atmosphere doesn't change all that much during the time step.

Once the predictions are made, the temperature, humidity, pressure and wind are updated. We then repeat the process, moving forward in time until we reach the forecast time.

Where is the physics in this process?

The physics is used to make the predictions during each time step.

For example, there are several mechanisms that can change the temperature. If it is sunny, radiation from the sun might be absorbed, warming the air. If there is moisture, any condensation or evaporation will cool or warm the air (due to latent cooling and heating). Based upon the wind, colder or warmer air might be brought into the region. Finally, if the pressure changes, the air might expand or compress, leading to cooling or warming.

These mechanisms, which I will refer to as radiation, latent cooling and heating, advection, and adiabatic expansion and compression, are all governed by relationships you've learned in physics and chemistry (e.g., the first law of thermodynamics). Part C addresses these mechanisms.

So, the future temperature depends a great on the wind. However, once the temperature changes, that changes the wind. The process for determining the new wind can be broken into two parts.

The first part uses a relationship between the temperature and pressure to obtain a new pressure value (from the new temperature value). The relationship is called the ideal gas law, which is discussed in part B of the readings.

The second part uses a relationship between the pressure and wind to obtain a new wind value (from the new pressure values). The relationship is based upon Newton's second law, which is discussed in part D of the readings.

So once the temperature has been updated the pressure must likewise be updated. And, once the pressure has been updated the wind must likewise be updated. And, since the revised wind pattern changes how the temperature changes, the temperature change during the next time step will be slightly different from what it was before. In this way, the evolution of the atmosphere is simulated, step-by-step, with each step governed by physics and chemistry.

> ✍ Notice that numerical weather prediction doesn't deal with probabilities. Since people like to have the forecasts in terms of probabilities, the forecaster might assign a probability based on his/her experience and/or a comparison of several computer models. Alternatively, computers can also be used to generate lots of forecasts (called **ensemble forecasts**) that vary only slightly in the initial conditions. Such forecasts can indicate the likelihood of a particular forecast.

How similar is this process to actual computer forecasts?

While actual computer forecasts are a lot more complicated than what I've presented here, the process is pretty much the same. Some additional complications, for example, are moisture (which is needed to predict precipitation), viscosity and friction (which is needed to predict the wind). They also have special routines to determine the effects of processes occurring on scales smaller than what can be resolved by the grid.

Does this process actually work?

Yes, it works (perhaps surprisingly so). That is why nowadays most weather forecasts are based upon numerical weather prediction models.

Are special "super" computers needed to carry out these simulations?

Yes and no. It depends on how fine our grid is.

As mentioned above, the forecast process treats the atmosphere as consisting of little "blocks" of air with a single value of temperature and wind. The more blocks, and the more complete the physics being used, the faster the computer has to be to make the forecast in a reasonable amount of time.

For example, let's suppose you want to make a forecast using a grid that is 100 km by 100 km in the horizontal (i.e., has data points every 100 km) and has levels every 1 km in the vertical up to 15 km. For a 30° by 50° domain (about the size of the United States), that would mean you'd need at least 22,500 points.

Then, for each time step, you'd have to calculate the new temperature, humidity, pressure and wind for each point. Assuming one equation for each variable[iv], that would be almost 100,000 calculations for each time step. Then, to make a 5-day forecast at 2-minute time steps, you'd need to make 3600 iterations over the course of the forecast.

For this scenario, then, you'd need to make 360 million calculations for the forecast. That sounds like a lot but computers are so good nowadays that a desktop computer can probably do this in a reasonable amount of time (i.e., a couple of hours at most).

However, add in the complications we've left out and/or increase the resolution and size of the grid, and then a really fast computer is necessary.

YOU'VE BEEN ASSUMING A 100-KM HORIZONTAL RESOLUTION, 1-KM VERTICAL RESOLUTION AND 2-MINUTE TIME STEPS. ARE THESE REPRESENTATIVE OF ACTUAL NUMERICAL MODELS?

It depends. The larger the domain, the poorer the resolution (so the number of points isn't too large). The values above are probably representative of what is used today for large domain models but thirty years ago were representative of what was used for small domain models.

Operational models (those used by government agencies and run on supercomputers), tend to have large domains and fine resolution. Examples of the most popular operational models[v] (i.e., the ones used by the weather services) are the **ETA**, **NGM** (Nested Grid Model), **MRF** (Medium Range Forecast) and **ECMWF** (European Centre for Medium-Range Weather Forecasts) models. The first three are run by the National Center for Environmental Prediction (**NCEP**), which is part of the U.S. National Weather Ser-

---

[iv] Actually, three equations would be needed for the wind, one for each dimension.

[v] The web page at http://www.rap.ucar.edu/weather/model/ lists various places you can obtain output from numerical models.

vice (**NWS**). Models like the MRF and ECMWF try to model the entire atmosphere (called **global models**) and/or have a longer forecast period. Consequently, they have a lower spatial resolution in order to keep the total number of grid points manageable.

Research models (i.e., those used by research organizations like universities), tend to have smaller domains. Examples of the most popular research models are the **MM5** (Mesoscale Model 5) and **RAMS** (Regional Atmospheric Modeling System). The smaller domains allow them have relatively high spatial resolution (i.e., the grid points are very close together), which provides an opportunity to explore events that are smaller is scale.

Each model has its own advantages and disadvantages.

> ✍ Everybody has access to the same model output. The difference in forecasts that you get from, say, two different television stations, is usually just due to a difference in one's interpretation of the model output.

Check Point 1.3: In general, computer models for the entire globe have a lower resolution than computer models that simulate the atmosphere over just a region, like North America. Why is this?

Problems

Problem 1.1: Suppose you have a synoptic map using data in Chicago, IL, from 10 PM on December 31, 2013. What would be the date of the map in UTC?

Problem 1.2: According to the polar front model, which way would the wind likely be blowing after a weather system passes?

Problem 1.3: Research one of the operational or research models mentioned above. Try to answer as many of the questions below as you can.
(a) What is the horizontal resolution of the model?
(b) How many horizontal grid points does the model have?
(c) How many vertical levels does the model have?
(d) To what height (or to what pressure) does the model domain go (above the surface)?

(e) What is the time step?
(f) For how long is the simulation made?

Problem 1.4: How many calculations would be needed to make a five-day forecast for the entire globe with a 100-km horizontal resolution and 1-km vertical resolution up to 40 km at 2-minute time steps?

Part A

Measurements

2. Wind

2.1 Introduction to part A

Before we get into the physics of the atmosphere, we need to observe the atmosphere. For the atmosphere, we focus on four variables: wind, temperature, pressure and moisture (including water vapor, clouds and precipitation). This is the focus of this part of the text

We'll start by examining the wind, followed by temperature (chapter 3), pressure (chapter 4), and moisture (chapters 5 and 6).

2.2 Direction

At first glance it would seem pretty simple to describe the direction of the wind. However, there are two conventions for doing so: the meteorological convention and the mathematical convention. As we study the physics of the atmosphere, we will encounter both conventions, so you need to be comfortable with both and be able to convert from one to the other.

The two conventions differ in several respects. The first aspect is in regard to whether the direction indicates where the air is *coming from* or where the air is *going to.*

- The **meteorological convention** is to specify the direction *from which* the wind is blowing. For example, a north wind means the wind is coming from the north (as in "cold north wind") and westerlies refer to winds coming from the west.[i]

---

[i] Note that when storms come up the eastern coast of the U.S., the winds along the coast will be from the northeast and thus are known as northeast winds. It is from this convention that such storms are called Nor'easters.

- The **mathematical convention** is to indicate the direction based on the direction the wind is blowing to. In much the same way as we would say a "force toward the north", we would likewise say a "wind toward the north".

As you can see, the meteorological convention is to emphasize the direction from which the wind is coming from. This is because that information tells us a bit about the weather, like whether it is likely to be colder or warmer than normal.

Mathematicians don't care about that stuff. They want something that is easy to work with mathematically. For that reason, they focus on where the wind is going.

The second difference between the two approaches has to do with the angle assigned to particular directions. You might think there is something special about setting zero degrees to be toward the right but that is only the mathematical convention. There is no special reason why we *have* to do it that way. Indeed, meteorologists do not do that.

- The meteorological convention is to use 0° to indicate wind from the north and increase the angle clockwise around the circle (so that 90° indicates a wind from the east). Thus, a 180° wind would be coming from the south.

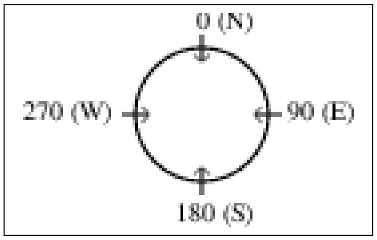

- The mathematical convention is to set zero degrees toward the right, increasing counter-clockwise. Since maps are oriented with north toward the top, that means that zero degrees is oriented toward the east with 90 degrees toward the north. Thus, a wind blowing toward the north would have a direction of 90° using the mathematical convention.

Check Point 2.1: Indicate the angle (in degrees) associated with wind coming from the west in the (a) meteorological convention and the (b) mathematical convention.

2.3 Wind speed

Like the direction, you would think describing the wind speed would be a rather straightforward task. Unfortunately, there are several different units used to measure the wind speed. You need to be comfortable with three in particular.

- In the United States, speed is typically measured in units of *miles per hour.*
- Standard scientific practice is to measure speed in metric units of *meters per second.*
- Traditionally, the meteorological community has followed the nautical practice of using **knots**, which are equivalent to **nautical miles** per hour.

In this text, we will typically follow the scientific practice (i.e., m/s). However, meteorological observations are typically in knots so you need to be able to convert between the different units.

To convert, you need to know the difference between **statute miles**, nautical miles and meters.

> ✍ Typically, when we refer to "mile" we mean statute mile. I will try to remember to include the "statute" adjective in order to distinguish it from the nautical mile but, if I forget, the word "mile" without any adjective is more than likely a statute mile.

First, let's look at the relationship between statute miles and nautical miles. The statute mile is a little smaller. This is because it is based upon furlongs (one statute mile equals eight furlongs) while the nautical mile is based upon degrees of latitude (one nautical mile equals one-sixtieth of a degree of latitude). To a good approximation, a degree of latitude is equal to 69 statute miles (actually $69.04676\bar{6}$) or 60 nautical miles.

Converting between statute miles per hour and nautical miles per hour (knots) is the same as converting between statute and nautical miles. Since a nautical mile is about 1.15 times bigger than a statute mile, one knot (nautical mile per hour) is about 1.15 times bigger than a mph (statute mile per hour).

> ✍ Since they are roughly the same, many times people will treat a knot as the same as a mph. However, technically, a knot is a little faster than a mph.

Let's now look at the relationship between statute miles and kilometers. A kilometer is smaller. There are 111.12 kilometers in a degree of latitude. And, as you should know, there are 1000 meters in one kilometer.

Based on the relationships, then, one can find that there are 1609.344 meters in one statute mile[ii] and, based on this, that 1 mph is a little less than 0.5 m/s.[iii]

One can find the relationship between knots and m/s in a similar way. I'll leave this as an exercise for the reader (see problem 2.1).

Check Point 2.2: Which is fastest: 10 m/s, 10 mph or 10 kts? Which is slowest?

2.4 Synoptic maps

There are two general ways that wind data is indicated on a map. One follows the meteorological convention and the other follows the mathematical convention.

2.4.1 Meteorological convention

The map in figure 1.1 on page 4 follows the meteorological convention for wind direction, which uses an arrow such that the "feathers" of the arrow indicate the direction from which the wind is blowing. For example, according to the map, the wind in Nebraska is coming from the north whereas the

---

[ii] You need to use the value that one statute mile is equal to 1/69.04677 of a degree of latitude. Then, since a degree of latitude is also equal to 111.12 kilometers or 111,120 meters, divide 111,120 meters by 69.04677 to get 1609.344 meters in one statute mile.

[iii] One mile is equal to 1609 meters whereas one hour is equal to 3600 seconds. Consequently, one mile per hour is equivalent to 1609 meters per 3600 seconds, or 1609/3600 m/s.

wind in Oklahoma is coming from the south. Following the meteorological convention, the direction would be from 0° over Nebraska and from 180° over Oklahoma.

Check Point 2.3: At 14 Z on 04 January 2014 (see figure 1.1), from which direction is the wind coming over Florida?

The wind speed is indicated via the use of little barbs, which come off the wind arrow like little flags. The flags, lines and half-lines are used to indicate 50 kts, 10 kts and 5 kts, respectively (figure from the National Weather Service). Note that the values are in knots.

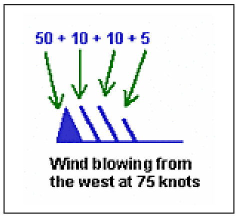

Check Point 2.4: What is the wind speed indicated by the station model in figure 1.2 on page 5? Include units.

2.4.2 Mathematical convention

The second way to display wind data is via **wind vectors**. This is typically used when displaying gridded data (as with that produced by a numerical model). An example is shown in figure 2.1.

In this case, the wind direction is given by an arrow, with the head of the arrow pointing in the direction of the wind, with the length of the arrow being proportional to the wind speed (a “reference” arrow is usually provided to indicate what speed a particular length corresponds to).

I refer to this as **gridded data** because the data is plotted a regular intervals along a grid. If you focus on the tails of each arrow, you’ll see that they line up along a north-south and an east-west grid.

Check Point 2.5: In figure 2.1, toward which direction is the wind blowing over Kansas?

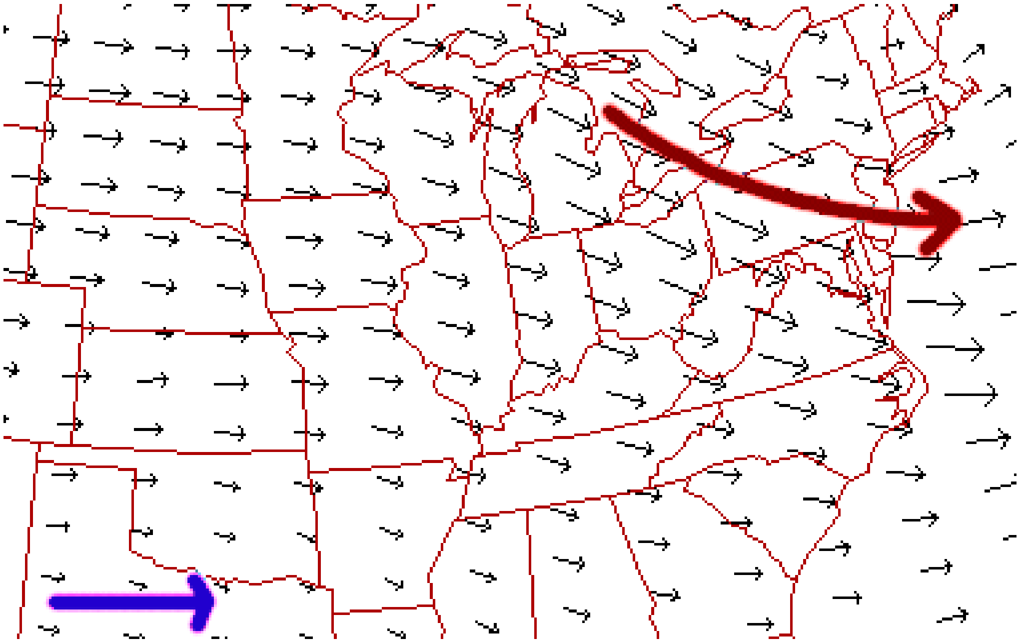

Figure 2.1: A vector wind plot. (source: University of Illinois WW2010 ProjectDataStreme)

2.5 Wind components

In physics, you may remember that there are two ways that two-dimensional vectors can be described. One way is in terms of magnitude and direction. The other is in terms of components.

Since the wind velocity is a vector, we can express it in terms of components instead of wind speed and direction. Indeed, this is how numerical models treat the wind.

To describe via components, we must first explicitly identify our coordinate system. A **coordinate system** is used as a basis for describing location and displacements. In a three-dimensional world, the coordinate system has three coordinates.

Since there are two conventions for describing the direction of the wind, there are technically two different coordinate systems we could use. However, it turns out that the mathematical system is used almost exclusively when components are used. Consequently, we'll focus on that.

In the mathematical system, one axis points toward the east (i.e., the zero degree direction) and another toward the north (i.e., the 90 degree direction). The third component (since we live in three dimensional space) either points upward (toward higher heights) or downward (toward higher pressures), depending on the purpose. For our purposes, we'll have the third component pointing upward. All three components are at right angles to each other[iv].

Mathematically, the three directions will be indicated as $\hat{\imath}$, $\hat{j}$, and $\hat{k}$ (or $\hat{x}$, $\hat{y}$ and $\hat{z}$).

In three dimensions, an object's position will be given in terms of three components: an x component, a y component and a z component. The three components represent how far the object is from the origin along the three component directions.

> ✍ Note that a positive value of x corresponds to a position towards the east. A negative value corresponds to a position towards the west.

Just as position in three dimensions is given with respect to three coordinates (x, y and z), velocity in three dimensions is also described via three components. For velocity, we will label the three components u, v and w. These components correspond to dx/dt, dy/dt and dz/dt.

> ✍ As with position, a positive value of u corresponds to a velocity towards the east and a positive value of v corresponds to a velocity toward the north. Consequently, for a north wind (i.e., a wind coming from the north) of 10 m/s, v would be *negative* 10 m/s and u would be zero.

Using $\hat{\imath}$ and $\hat{\jmath}$ to indicate the x and y directions, respectively, the **horizontal wind** (in vector notation) can be indicated as $u\hat{\imath} + v\hat{\jmath}$.

[iv] Actually, since we are on a sphere, this coordinate system differs slightly from a **Cartesian** (or rectangular) coordinate system. In a Cartesian coordinate system, the direction of each coordinate is the same no matter where one is. This is not true for our coordinate system. The atmosphere surrounds the earth, which is a sphere. Thus, the x-direction in one location may differ from the x-direction in another location. For example, one side of the earth is experiencing sunset when the other side is experiencing sunrise. For the former, east is directed away from the sun whereas for the latter east is directed toward the sun.

> ✍ Typically, we will deal with the horizontal wind components separately from the vertical wind component. Following convention, we will use $\vec{V}$ to represent the horizontal velocity vector, which includes both the magnitude V and direction θ (as discussed previously). In other words, $\vec{V}$ is not the *total* three-dimensional velocity but rather the *horizontal* velocity.

Check Point 2.6: Suppose the horizontal wind vector is given as $A\hat{\imath} - B\hat{\jmath}$*, where* A *and* B *are both positive numbers. In what direction is the wind blowing? Is* v *positive or negative?*

Given the magnitude and direction of the wind, how do we determine the u and v components?

One can use sines and cosines to convert from magnitude/direction to components. The cosine of the direction gives the component in the direction of zero degrees. The sine of the direction gives the component in the direction of 90 degrees.

Taking zero degrees to be toward the east (i.e., a wind coming from the west) and increasing counter-clockwise, a wind velocity of magnitude V and direction θ, $V\cos\theta$ would be equivalent to the component directed toward the east and $V\sin\theta$ would be equivalent to the component directed toward the north:

$$\begin{aligned} u &= V\cos\theta \\ v &= V\sin\theta \end{aligned}$$

Suppose we were given the u and v components. How would we find the magnitude and direction?

We can calculate V and θ given u and v:

$$V = [u^2 + v^2]^{1/2}$$

and

$$\theta = \tan^{-1}\left(\frac{v}{u}\right).$$

Check Point 2.7: For an angle of 60°, and without a calculator, which is *greater: the cosine of 60° or the sine?*

Problems

Problem 2.1: A typical wind speed near the surface might be 7 knots.
(a) Convert this to units of mph (miles per hour).
(b) Convert this to units of m/s.
Show your work in both cases, using the fact that there are 3600 seconds in an hour and that one degree of latitude is equal to 69.045677 mi, 60 nm and 111.12 km.

Problem 2.2: Examine the map in figure 1.1 (page 4). Determine the wind speed in km/h at Chicago, Illinois, on January 4, 2014.

Problem 2.3: (a) Suppose you have a map of wind vectors, with the key stating that the maximum wind vector has a length of 0.5 cm and a magnitude of 30 m/s. If the wind at a particular location was from the SE at 8 mph, what would be the length of the wind vector at that point? Remember to convert to m/s first (see previous problem).
(b) Indicate the orientation of the wind vector in degrees according to the (i) meteorological system and the (ii) mathematical system.
(c) Determine the u and v components of the wind vector.
(d) Double-check your answer to (c) by using those components to calculate the speed and direction.

3. Temperature

3.1 Definition of temperature

What is temperature?

The **temperature** reflects the average kinetic energy of the molecules. Even if a volume of air is just sitting there, the individual molecules are moving around and bouncing into each other. The faster they move, the higher the temperature of the air.

Next to the wind, day-to-day variations in temperature are probably what most people are concerned about. This is because changes in temperature can be easily perceived. Other variables like pressure, density and composition may also change from day to day but most people are relatively unaware of how they change from day to day.

3.2 Measuring temperature

If temperature reflects the average kinetic energy of the molecules, how do we measure it?

To measure temperature, we don't measure the kinetic energy of each molecule and then take the average. Rather, another property is found that depends on temperature and then that is measured instead. For example, the volume taken up by mercury varies with its temperature. By measuring the volume of the mercury, we can then infer the temperature. This, of course, is just how a mercury thermometer works.

> ✍ Some thermometers use a **thermistor**, which is a piece of material with a resistance that depends upon temperature. Even with a thermistor, we do not measure the temperature directly. Rather, we measure the resistance which, in turn, depends upon temperature.

3.3 Typical values

The air temperature near the earth's surface typically varies between 0°C (the freezing point of water) and 37°C (normal body temperature). Room temperature is typically around 20°C.

3.4 Units of temperature

I AM USED TO DEALING WITH TEMPERATURE IN °F. WHAT DOES 0°C AND 37°C MEAN?

Three common temperature scales are the degree **Celsius**[i] (°C) scale, the **kelvin**[ii] (K) scale and the degree **Fahrenheit** (°F) scale. You will most likely be using all three so it helps to investigate them a little now.

The following table shows the temperature in all three scales for some typical situations.

Freezing point of water	32°F	0°C	273 K
Cool autumn day	50°F	10°C	283 K
Typical room temperature	68°F	20°C	293 K
Normal body temperature	98.6°F	37°C	310 K
Boiling point of water	212°F	100°C	373 K

Each scale has its own offset and increment. The **offset** refers to the temperature at which the scale is set to zero. This can be set to the temperature where water freezes (as with °C), the temperature where no thermal energy exists (as with K) or the lowest temperature one could get with a mixture of salt and ice (as with °F).

The **increment** refers to the difference in temperature represented by one increment on the scale. For example, the kelvin and degree Celsius scales use the same increment.[iii] This means that a temperature *difference* of 1°C

---

[i] Before 1948 this scale was known as the degree centigrade scale.

[ii] Unlike the other common temperature scales, the kelvin scale (and unit) is not preceding by the "degree" term or symbol. The name is also not capitalized.

[iii] At the 13th General Conference on Weights and Measures (CPGM), 1967, this was set to 1/273.16 of the thermodynamic temperature of the triple point of water.

is equivalent to a temperature *difference* of 1 K. However, since they use different offsets, a temperature of 1°C is not equivalent to a temperature of 1 K.

3.4.1 Unit conversion

HOW DO WE CONVERT FROM ONE SCALE TO ANOTHER?

There is a linear relationship between the different scales. So, one way of mathematically representing the relationship between two scales is as follows:

$$T_1 = AT_2 + B \tag{3.1}$$

where T_1 is the temperature in one scale, T_2 is the temperature in a second scale, A is the conversion from one increment to the other, and B is the difference in offset. In a way, this is just like the equation of a straight line, where A is the slope and B is the y intercept.

For example, if T_1 was kelvin and T_2 was degree Celsius, A would be (1 K/°C), since an increase of 1 K is equal to an increase of 1°C, and B would be 273.15 K, because 0°C is equal to 273.15 K.

WHAT ARE A AND B IF WE WANT TO CONVERT BETWEEN DEGREE CELSIUS AND DEGREE FAHRENHEIT?

The degree Fahrenheit scale has both a different increment and a different offset than the degree Celsius scale. The different increment means that a temperature difference of 1°C is not equivalent to a temperature difference of 1°F and the different offset means that a temperature of 0°C is not equivalent to a temperature of 0°F.

To find the Celsius equivalent to a temperature in Fahrenheit, one uses the same equation but with A equal to 5/9 °C/°F, since an increase of 5°C is equal to an increase of 9°F, and B equal to 32°F, because 0°C is equal to 32°F.

Check Point 3.1: Which corresponds to the offset, A or B? Why?

WHY DO THE CELSIUS AND FAHRENHEIT SCALES HAVE DIFFERENT OFFSETS?

They have different offsets because they use zero degrees to correspond to a different state. One sets zero to be the freezing point of water while the other sets zero to be the lowest temperature one could get with a mixture of salt and ice.

WHY DOES THE DEGREE FAHRENHEIT SCALE HAVE A DIFFERENT INCREMENT?

They have different increments because the Fahrenheit scale has 180 degrees between freezing and boiling whereas the Celsius scale has only 100 degrees between freezing and boiling.

To understand why the Fahrenheit scale has 180 degrees, which seems kind of strange, we first we need to recognize why the degree Fahrenheit scale sets 0°F as it does. Historically, Fahrenheit wanted all reasonable temperatures to fit between 0 and 100. Consequently, he set 0°F as the lowest temperature one could get with a mixture of salt and ice. For his second point, he either used 30 (or 32 degrees, depending on your source) for the freezing point of water or 96 (or 100 degrees, depending on your source) for human body temperature.

Celsius, on the other hand, wanted to use reference temperatures that were more objective. Consequently, he used the freezing point of water and the boiling of water. In Celsius' scale, the difference between freezing and boiling was set[iv] at 100 °C. A difference of 100 °C corresponded to a difference of 180°F.

Check Point 3.2: Why is it said that the Fahrenheit scale has 180 degrees between freezing and boiling when boiling is 212 degrees, not 180 degrees?

3.4.2 Derivation

HOW DO YOU KNOW THAT A WOULD BE (1 K/°C) AND B WOULD BE 273.15 K FOR THE CELSIUS TO KELVIN CONVERSION?

[iv] Apparently, Celsius decided to use 100 increments in order to illustrate the more precise thermometer that he had developed (or perhaps because he expected it to be used over a wider range of temperatures).

To determine the values of A and B, someone needs to give you the values for each scale at two separate temperatures. Mathematically, you need two independent equations in order to solve for two unknowns.

For example, the freezing point of water is at 0°C or 273.15 K. The two values, 0°C and 273.15 K, represent the same temperature. Plugging these two values into the mathematical expression, we get

$$273.15 \text{ K} = A(0°\text{C}) + B.$$

This information alone is not sufficient to determine A and B because there are two unknowns and only a single equation. To determine the values of A and B, we need to know something else. For example, the boiling point of water is at 100°C or 373.15 K. Plugging these two values into the mathematical expression, we get

$$373.15 \text{ K} = A(100°\text{C}) + B.$$

We now have two expressions,

$$\begin{aligned} 273.15 \text{ K} &= A(0°\text{C}) + B. \\ 373.15 \text{ K} &= A(100°\text{C}) + B. \end{aligned}$$

and two unknowns, A and B. With an equal number of equations as unknowns, we can solve for the unknowns. From the first expression, we get that $B = 273.15$ K. Plugging this into the second expression, we get that $A = 1$ K/°C.

WHAT ABOUT CONVERTING BETWEEN DEGREE CELSIUS AND DEGREE FAHRENHEIT?

The degree Fahrenheit scale has both a different increment and a different offset than the degree Celsius scale. The different increment means that a temperature difference of 1°C is not equivalent to a temperature difference of 1°F and the different offset means that a temperature of 0°C is not equivalent to a temperature of 0°F.

As before, we need two data values for which both the Fahrenheit and Celsius values are known. For example, the freezing point of water is 32°F in the Fahrenheit scale and 0°C in the Celsius scale. In addition, normal body temperature is at 98.6°F in the Fahrenheit scale and 37°C in the Celsius scale. Using these two pairs of values into the mathematical relationship gives us the following two expressions:

$$\begin{aligned} 32^\circ\text{F} &= A(0^\circ\text{C}) + B \\ 98.6^\circ\text{F} &= A(37^\circ\text{C}) + B \end{aligned}$$

From the first expression, we can solve to get $B = 32^\circ\text{F}$. Using that value for the second and solving for A, we have

$$A = \frac{98.6^\circ\text{F} - 32^\circ\text{F}}{37^\circ\text{C}} = \frac{66.6^\circ\text{F}}{37^\circ\text{C}} = 1.8^\circ\text{F}/^\circ\text{C}.$$

Now that we have the values of A and B, that means the relationship between a Celsius temperatures and the corresponding Fahrenheit temperature is

$$T_F = (1.8^\circ\text{F}/^\circ\text{C})T_C + (32^\circ\text{F}) \tag{3.2}$$

Check Point 3.3: The boiling point of water is $212^\circ F$ in the Fahrenheit scale and $100^\circ C$ in the Celsius scale. Suppose we had used

$$212^\circ F = A(100^\circ C) + B$$

instead of

$$98.6^\circ F = A(37^\circ C) + B$$

in the derivation above. What would the values of A and B be?

3.5 Horizontal variation

To examine how the temperature varies in the horizontal, we will examine a synoptic map (see previous chapter). An example of such a map is shown in figure 3.1, which displays observed values of surface temperature over the United States for 20 UTC on 19 May 2004.

On this particular map there are also lines that connect places that have similar values of temperature. The generic name for such lines are **isolines** because each place along the line has the same (*iso*) value. Depending on

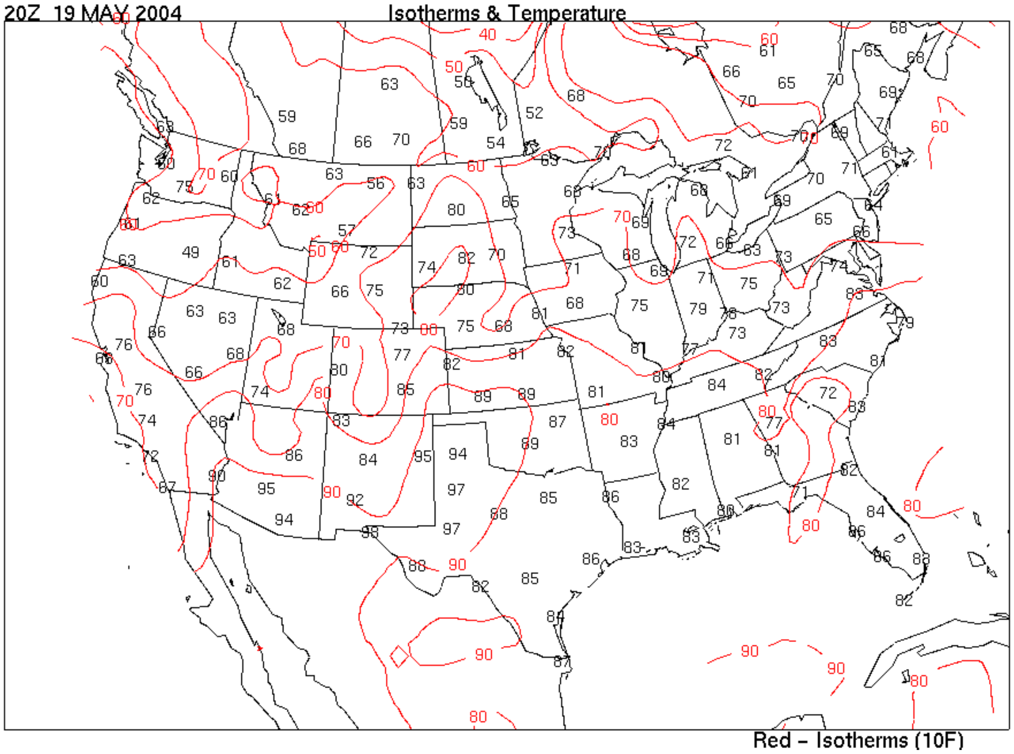

Figure 3.1: A map of surface temperature for 20 UTC on 19 May 2004 (source: DataStreme).

what the parameter is, they can also be called something specific to the parameter, such as **isobars** (“bar” = pressure) or **isotherms** (“therm” = temperature).

✍ The only time we stray from this naming convention is when we look at lines of equal height (as in a topographic map). While we could call these lines **isoheights**, in general we call them **contours**.

3.5.1 Gradients

The value of isolines is that, once we know the units and intervals of the lines, we can then use them to quickly get a general sense of where the air is warm vs. cold. They can also be used to determine the value at any location and the **gradient**.

A gradient is like a slope (i.e., how quickly the "vertical" component changes over a given "horizontal" distance) except that it represents how quickly a *variable* changes over a given distance. If we plotted a topographic map with contours (lines of constant height), the height gradient would be greatest where the contours are close together.

Similarly, when we plot a temperature map with isotherms (lines of constant temperature), the temperature gradient is greatest where the isotherms are close together.

HOW DO WE MEASURE THE TEMPERATURE GRADIENT?

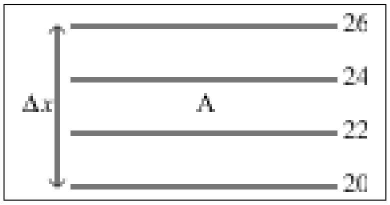

The temperature gradient is obtained by identifying how close the isotherms are. For example, consider the isotherms drawn to the right. Four isotherms are drawn, from 20°C to 26°C.

Let's suppose that the distance from the 26°C isotherm to the 20°C isotherm is 100 km (the distance is indicated as Δx in the figure). That means that the temperature changes by 6°C over a distance of 100 km. That would correspond to a temperature gradient of (6 °C)/(100 km) or 0.06 °C/km. Basically, we divide the temperature difference by the position difference.

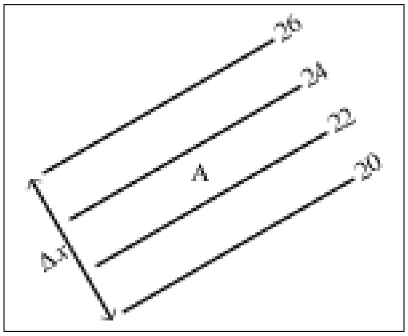

It doesn't matter how the isotherms are oriented. The drawing above has the temperature increasing toward the top of the page with the isotherms drawn left to right. Consider the figure to the right, where I've simply rotated the drawing from before. The temperature gradient in this figure is the same as the temperature gradient in the drawing before.

Check Point 3.4: Suppose you are examining a region of a synoptic map where the isotherms are close together. Is this a region of large temperature gradient or small temperature gradient?

HOW DO WE MEASURE Δx?

Since distance is not usually indicated on synoptic maps, you'll need to use something of known distance. For example, if you know the size of Nebraska (333 km north to south)[v], then you can use that as a reference. Measure the distance from one isotherm to the next and then compare that distance to the reference distance.

Check Point 3.5: Suppose isotherms are drawn every 4° C. At a particular location two isotherms are separated by a distance measured to be twice the distance from the northern to southern boundary of Nebraska. What is the temperature gradient at that location?

3.5.2 Direction

As mentioned before, the temperature gradient represents how quickly the temperature changes over a given distance. Mathematically, it is a ratio: the temperature difference divided by the position difference.

Mathematically, we represent a temperature difference as ΔT and a position difference as Δx. So, the temperature gradient would be indicated as $\Delta T/\Delta x$.

I should point out that, technically, $\Delta T/\Delta x$ represents the gradient in the $\hat{x}$ direction. I've been treating the $\hat{x}$ direction as being perpendicular to the isotherms. That is perfectly fine, but we could instead have stated that the $\hat{x}$ direction is toward the east. In that case, $\Delta T/\Delta x$ would *not* be the same when we rotate the isotherm pattern.

As we know from the previous chapter, the standard convention does use $\hat{x}$ to mean the eastward direction. In a similar way, the standard convention uses $\hat{y}$ to mean the northward direction and $\hat{z}$ to mean upward. Thus, $\Delta T/\Delta y$ would be the temperature gradient in the north/south direction and $\Delta T/\Delta z$ would be the temperature gradient in the vertical direction.

---

[v] The southern border of Nebraska is at 40°N while the northern border is at 43°N. Consequently the height of Nebraska is 3° of latitude or, since each degree of latitude is 111.12 km, about 333 km. One can also use Wyoming, with southern and northern boundaries of 41°N and 45°N, respectively (for a height of about 444 km).

Meteorologists tend to use special names for vertical gradients. For example, the “vertical gradient of wind” is called the **wind shear** and the (negative) vertical gradient of temperature is called the **lapse rate**.

3.5.3 Gradient as a derivative

Speaking of conventions, the “Δ” notation means that a *finite* difference was used to calculate the gradient. While this is usually the case in practice, in theory we do not need to use a finite difference. We could, instead, use an **infinitesimal**[vi] difference. This is particularly useful if the temperature gradient is different from one place to the next.

For example, the temperature gradient in the figure to the right is greater at A than at B. That is because the isotherms are closer together at A than at B. In this case, what gradient is represented by $\Delta T/\Delta x$, which uses the total difference in temperature and distance?

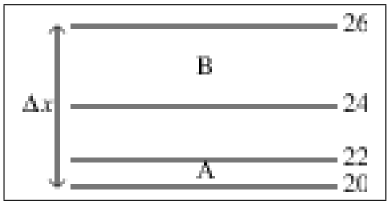

The answer is that it gives the *average* temperature gradient over the entire distance Δx. The average would be less than the gradient at A and greater than the gradient at B.

The gradient at a particular point, then, would use the *infinitesimal* differences at that point. To indicate that, we use the symbol “∂” instead of “Δ”.

The ratio $\partial T/\partial x$ is known as the derivative of T with respect to x. It is pretty much the same derivative you learned in calculus class as dT/dx.

The difference is mostly irrelevant to much of what we will cover. However, you’ll notice that sometimes I’ll write the derivative one way and sometimes I’ll write it the other way. So, it might help to know the difference.

First of all, $\partial T/\partial x$ is called the *partial* derivative while dT/dx is called the *total* derivative. The partial derivative is used when a variable (like T) depends upon several things (like x, y and z) and I only want to refer to how

---

[vi] The word “infinitesimal” means extremely tiny – so tiny, in fact, that we can treat it as zero in comparison to any **finite** value (where “finite” means neither infinite nor infinitesimal.

it depends on *one* of those things. In those cases, a derivative like $\partial T/\partial x$ means that things like y and z are held fixed and only x is changed.

The total derivative is used when none are held fixed or if a variable (like T) only depends on one thing (so we don't have to worry about which are held fixed). More will be mentioned about the total derivative when we get to a situation where we have to make that distinction.

> ✍ When we use x, y and z to indicate directions, then $\partial T/\partial x$ would represent the component of the gradient in the eastward direction whereas $\partial T/\partial z$ would represent the *vertical* temperature gradient.

Check Point 3.6: Suppose you are in a room where the temperature is higher on the south side of the room (and cooler on the north side of the room). If the $\hat{y}$ direction is toward the north, is $\partial T/\partial y$ positive, negative or zero?

3.5.4 Notation

As mentioned earlier, the convention is to use $\hat{x}$ as representing the eastward direction. So, $\partial T/\partial x$ would represent just the component of the gradient in the eastward direction.

Most often, we just want to know what the gradient is in a direction perpendicular to the isotherms (i.e., across the isotherms), regardless of how they are oriented. This is called the **total horizontal gradient**.

The mathematical convention used to indicate the total horizontal gradient is ∇T. Notice that this is not the symbol "Δ", which is used to represent a difference. It actually uses an inverted "Δ".

From now on, I will use ∇T to represent a horizontal gradient of temperature and $\partial T/\partial z$ to represent a vertical gradient of temperature. The same notation will be used for the gradients of other variables (e.g., ∇u for the horizontal gradient of u and $\partial u/\partial z$ for the vertical gradient of u).

Check Point 3.7: Suppose you are examining a region of a synoptic map where the isotherms are close together. Is ∇T large or small?

3.6 Vertical variation

3.6.1 U.S. standard atmosphere

Although the temperature profile varies from day to day, an average profile has been identified by the United States Committee on Extension to the Standard Atmosphere (COESA).

On 15 October 1976, COESA, representing 29 U.S. scientific and engineering organizations, selected values to represent the **U.S. Standard Atmosphere**. The equations and parameters used are documented in a book entitled *U.S. Standard Atmosphere,* 1976 published by the U.S. Government Printing Office, Washington, D.C. The U.S. Standard Atmosphere is sort of like an ideal average. The actual values will vary from day to day and from height to height.

In the U.S. Standard Atmosphere, temperature values are given for every kilometer, from 5 km below sea level to one thousand kilometers above sea level. For example, the standard sea-level temperature is 15°C. At 10 km, the standard temperature is −50°C.

> ✍ The heights given in the standard atmosphere file are given in both geometric and geopotential heights. The former are actual heights (i.e., radar heights). For our purposes, we'll treat them as the same thing.

The U.S. Standard Atmosphere separates the atmosphere into two parts, a lower part up to 86 km and an upper part from 86 km to 1000 km altitude.

If one plots the standard temperature with height (see project), you'll notice that the U.S. Standard atmosphere can be broken down into finer layers by examining how quickly the temperature changes with height. Each layer has a vertical temperature gradient ($\partial T/\partial z$) that is the same throughout the layer (or changes in some characteristic way throughout the layer).

Check Point 3.8: On any particular day, will the temperature of the real atmosphere at various levels equal the temperature given in the U.S. Standard Atmosphere?

3.6.2 Lapse rate

Since the temperature typically decreases with height, the vertical temperature gradient is typically negative. Rather than deal with negative numbers all the time when discussing the vertical temperature gradient, meteorologists tend to define the **lapse rate** as the negative of the vertical temperature gradient:

$$\Gamma = -\partial T/\partial z \tag{3.3}$$

where the Greek letter Gamma (Γ) is used to represent the lapse rate.

> ✍ Many people interpret the derivative to be the slope of a line. This is true only if one creates a graph and if the variable in the numerator is on the y axis and the variable in the denominator is on the x axis. In the project, you are asked to create a graph with T on the x axis, so $\partial T/\partial z$ would not correspond to the slope of the line on that graph.

Check Point 3.9: In the morning, it is common to have cold air at the surface with warm air on top. Would that be a situation of positive lapse rate or negative lapse rate.

3.6.3 Actual temperature profile

The U.S. Standard Atmosphere represents the average temperature profile. What does an actual temperature profile look like?

To see what an actual temperature profile looks like, we need to make some measurements of the atmosphere.

A common way to measure the temperature at various levels is to release a balloon with sensing instruments. If the data is received via radio waves then such a balloon is called a **radiosonde** (or a rawinsonde if wind values are also obtained). The data that is received is called a **sounding**.

There are three ways I could show you the sounding data. One way would be via a list of temperatures (as with the U.S. Standard Atmosphere data used in the project). Another is by showing you several synoptic maps, each at a different level. A third way is via a single graph (as in project 3.1).

The simplest way is via a single graph. Unfortunately, the ground-level temperatures are usually much warmer than the temperatures above the surface, particularly if we want to focus our attention on the lowest 15 km of the atmosphere, where most of the weather takes place (see chapter 11 to learn why). Plotting the temperature on a normal graph would require a chart that was wide enough to contain both the warm temperatures at the ground and the cold temperatures above.

To deal with this, specialized graphs are used. The specific one we'll use here is the **skew-T/log-P** chart. An example is shown in figure 3.2.

There are three characteristics that make this type of graph different than a regular temperature vs. height graph:

1. First of all, the vertical coordinate is not height but pressure. Whereas height increases as one goes up in the atmosphere, the pressure decreases. Consequently, the numbers on the left decrease as one goes up the vertical axis.
2. Second, the pressure lines are not equally spaced. This is because the vertical coordinate is actually the logarithm of pressure (i.e., the lines corresponding to 10^5 Pa, 10^4 Pa, 10^3 Pa, etc., would be equally spaced).
3. The third strange thing about the graph is that the temperature lines are skewed. Thus, colder temperatures are to the upper left. This is done so that the colder upper-level temperatures can still fit on the figure.

Why is the vertical coordinate the logarithm of pressure, not simply the pressure?

We will examine pressure in the next chapter. For now it is sufficient to just recognize that plotting the pressure in that way makes the graph look roughly the same as it would be if we had made height the vertical axis.

What are all of the extra curves and such on the graph?

Besides the temperature and pressure lines, there are several more sets of lines, corresponding to other variables one might wish to know about. These are introduced later in the text, as needed. Because of all the lines, they are frequently color-coded to distinguish one set from another.

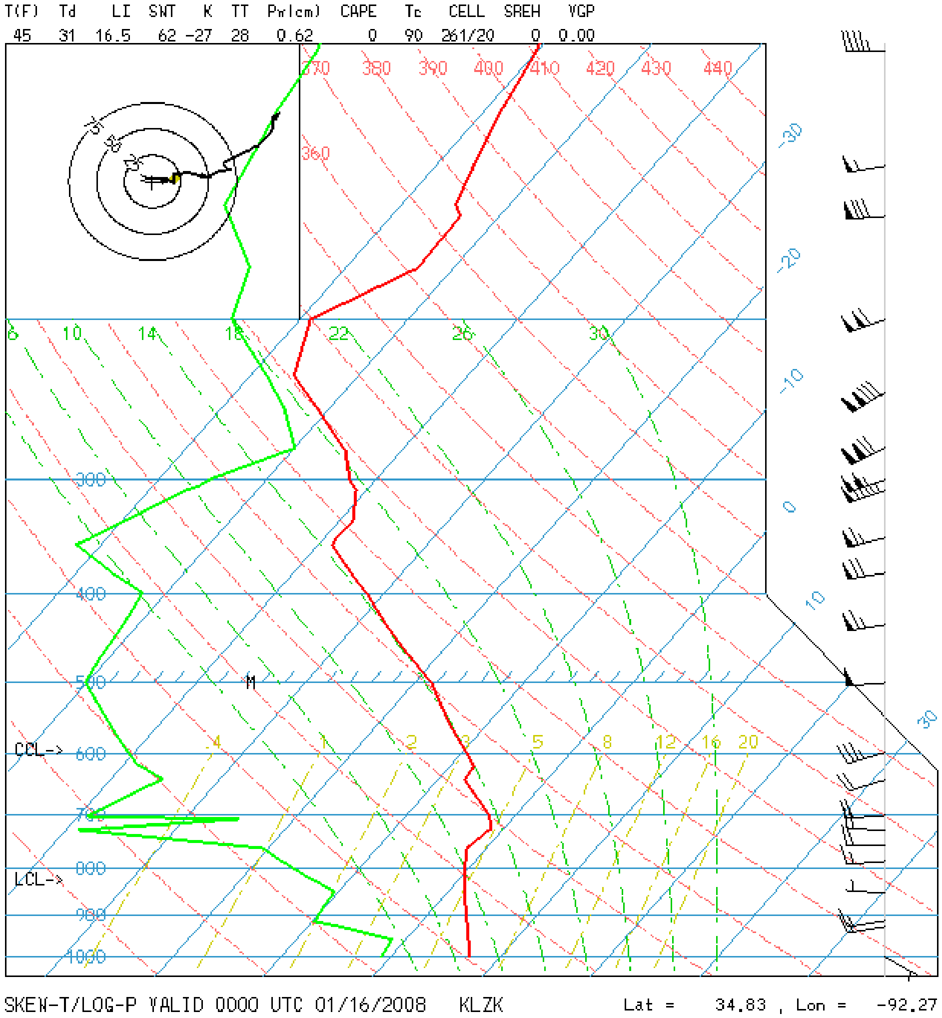

Figure 3.2: A skew-T log-P graph for Little Rock, Arkansas. Data is from 0000 UTC, 16 Jan, 2008 (local time would be 6 PM the previous day). Horizontal solid lines are pressure (in mb). Skewed, solid lines are temperature (in °C).

On a skew-T log-P, there are two thick squiggly lines. The one on the right (black in the figure) is the measured temperature of the atmosphere. The one on the left (gray in the figure) represents the amount of moisture measured in the atmosphere (this is discussed in chapter 5).

Check Point 3.10: (a) Why is a skew-T log-P graph called a skew-T log-P graph?
(b) According to the skew-T log-P graph shown in figure 3.2, what is the temperature at a height where the pressure is 700 mb?

Project

Project 3.1: Create a graph of the average temperature profile of the atmosphere from sea-level up to 120 km. Do this by downloading the temperatures of the **U.S. Standard Atmosphere** *from the course home page and importing it into a spreadsheet. Arrange the graph so that height is on the vertical axis.*

Project 3.2: You should notice that the U.S. Standard atmosphere can be broken down into several layers, each with its own lapse rate. For example, for the first 11 km or so, the lapse rate from one kilometer to the next is roughly the same at about 6.5 °C/km. It then changes abruptly to a lapse rate of 0 °C/km. How many different layers of constant lapse rate regions are there? Identify the lapse rate of each of those layers.

Project 3.3: The names **troposphere** *and* **stratosphere** *are used to describe the lowest two layers of the atmosphere. The troposphere is characterized by a positive lapse rate (cooling with height) and the stratosphere is characterized by zero and/or negative lapse rates (warming with height). To what height does each layer extend in the U.S. Standard atmosphere? The boundary between the two layers is called the* **tropopause**.

Problems

Problem 3.1: Figure 3.1 on page 31 displays observed values of surface (ground-level) temperature over the United States for 20 UTC on 19 May 2004.

(a) What interval is used for the isotherms on the map? Include units.
(b) Identify the temperature at East Stroudsburg. Include units.

Problem 3.2: From Figure 3.1, what is the range of temperatures that are observed? Include units.

Problem 3.3: (a) According to Figure 3.2, what is the temperature at a height where the pressure is 850 mb?
(b) According to the figure, does the temperature tend to decrease or increase with height?

Problem 3.4: (a) Calculate the maximum horizontal temperature gradient that is present in figure 3.1 (see the 90°F observation along the border between Arizona and California and the 67°F observation along the Pacific coast near the border between California and Mexico).
(b) Which is greater (absolute value): the maximum horizontal temperature gradient observed in figure 3.1 or a typical vertical temperature gradient in the troposphere (see U.S. Standard troposphere).

Derivations

Show-me 3.1: Using equation 3.2, one can convert the temperature in °C to °F. Derive an equation that can be used to convert temperature in °F to °C (i.e., the reverse) and use your equation to show that a typical **room temperature** *of 68°F is equivalent to 20°C.*

Show-me 3.2: Originally, Celsius' scale had used 100°C as the freezing point and 0°C as the boiling point, the reverse of what we use today[vii].
(a) Derive the relationship between the Fahrenheit scale (where the freezing point of water is $32°F$ *and the boiling point of water is* $212°F$*) and the* original *Celsius scale. In other words, in the equation* $T_F = AT_C + B$*, where* T_F *is the temperature in °F and* T_C *is the temperature in the original Celsius scale, what would be the value of the increment ratio* A *and the offset* B*?*
(b) Use your relationship to determine what room temperature was in the original Celsius scale.

[vii] Using zero for the freezing point of water would mean that some observed temperatures would be negative. His "reverse" system would mean that all observed temperatures (in his region, at least) would result in positive values. This was reversed soon after (once people got comfortable using negative temperatures, I guess).

4. Pressure

4.1 Introduction

Pressure is one of the variables typically plotted on a weather map (see, for example, the sample station plot illustrated in figure 1.2 on page 5). Although we are not as sensitive[i] to day-to-day variations in air pressure as we are to the wind and temperature, we measure the pressure because day-to-day variations in pressure are related to the weather (e.g., if the pressure is decreasing, that typically indicates that a weather system is approaching). We'll leave for later chapters a discussion of how the pressure is related to the weather. For now, we'll focus on what air pressure is, the typical values of air pressure, and how we measure it.

4.2 Pressure vs. force

The air exerts a force on us. This force is due to the random motions of the molecules hitting against us. Each molecule exerts a different force depending on its speed and mass. What we feel is the total force exerted by the molecules.

While this force is usually very large, we normally don't notice this force because the force is the same all around us (see next section).

Why, then, do we measure the pressure instead of force?

Certainly force *could* be used. However, the force exerted by the air depends upon the size of the surface upon which the force is being exerted. The larger the surface, the greater the amount of air that is interacting with surface and the larger the pressure.

To get rid of that dependence on the size, we use the pressure instead of the force. Pressure is the force *per area.* Using F, A and P to indicate force,

---

[i] Some people can feel pain in their joints when the air pressure changes.

area and pressure, respectively, we can write this definition as follows:

$$P = \frac{F}{A} \tag{4.1}$$

where F represents the total force exerted by the molecules on the surface of area A.

By dividing by area, that means that pressure is independent of the area and thus can be assigned to a single point in space. This is similar to what we did with temperature. As mentioned in chapter 3, temperature is the average kinetic energy of the molecules (i.e., the kinetic energy *per molecule*). By defining it as the kinetic energy per molecule, we get rid of the dependence on how many molecules we measure.

The pressure at a single point, then, is the force exerted on an **infinitesimal** area. Although the force on an infinitesimal area would be infinitesimal also, the *ratio* of the two is not.

> ✍ We will use the same reasoning for defining density as the mass per volume (see next chapter).

It turns out that using pressure instead of force is a common approach when dealing with fluids like air.

Isn't air a gas?

The word "**fluid**" is used for both **liquids** and **gases** because both flow. In fact, the word "fluid" comes from the Latin word *fluere*, which means "to flow." This ability to flow is the main feature that distinguishes fluids (liquids and gases) from solids.

Because fluids, unlike solid objects, flow and don't hold their shape, we generally do not use the same terms to describe fluids as we do to describe solids. Whereas for solids we might use mass and force, for fluids we use density and pressure.

Check Point 4.1: Pressure is defined as the force per area. Does each point in space have a pressure value, even though a point technically has no area? If so, how? If not, why not?

4.3 Isotropy

As mentioned in the previous section, the pressure exerted by the air is due to the random motions of the molecules hitting against the surface. Since the molecules are moving in all directions, the force exerted by the molecules is the same in all directions. In a similar way, the pressure exerted by the air is the same in all directions. This is called **isotropy**.

This doesn't mean the pressure is the same all over the place. Certainly, the pressure is different on top of a mountain than down in a valley. Rather, it means that at a particular location, the air pushes up as much as it pushes down or horizontally. In other words, if could enclose a tiny piece of air in a cube, the pressure on each face of the cube would be the same.

Check Point 4.2: Suppose we orient a piece of paper horizontally, parallel to the ground. If the pressure is isotropic, is the pressure exerted on one side be equal to the pressure on the other side? Explain.

4.4 Typical values and units

Since pressure is force per area, pressure has units of N/m^2. The standard sea-level pressure (i.e., from the U.S. Standard Atmosphere) is 101,325 N/m^2 (or 1.01325×10^5 N/m^2).

Since pressure is used so often, the units (N/m^2) are frequently replaced with a single unit called the pascal (abbreviated as "Pa" and named after French scientist Blaise **Pascal**). So, we can say that a typical sea-level pressure is around one hundred thousand pascals.

Since typical sea-level pressures are so many pascals, atmospheric pressure is usually written in terms of other units. For example, instead of pascals, the pressure might be given in **kilopascals**. So, the standard sea-level pressure would be 101.325 kPa.

Another common practice, however, is to use units of bars to measure atmospheric pressure. One **bar** is set equal to 10^5 Pa such that one bar is approximately the standard sea-level pressure at the earth's surface. So, the standard sea-level pressure would be 1.01325 bars.

In meteorology, the convention is to use **millibars** instead of bars. This is because pressure decreases as one goes up in the atmosphere[ii] and so, by using **millibars**, meteorologists don't have to deal with fractions or decimals. The unit abbreviation for millibars is "mb," so the standard sea-level pressure is 1013.25 mb.

> ✍ Because millibars are not strictly part of the metric system, many people, particularly in Europe, use **hectopascals** (hPa) instead of millibars (mb). They are equivalent. You can use either one.

The actual pressure at sea-level, of course, is not exactly 1013.25 mb. Indeed, it varies from day to day. However, it stays pretty much within the range from 985 mb to 1035 mb, and is rarely below 950 mb or above 1050 mb. For that reason, only the last two digits, plus the tenth, are plotted on synoptic weather maps, like the one shown on page 4. This is illustrated in the station model shown on page 5, where the digits "107" indicate a pressure of 1010.7 mb.

Check Point 4.3: What is the standard sea-level pressure? Provide the answer in pascals, kilopascals and millibars.

4.5 Measuring pressure

With all of the units one can use with pressure (see previous section), you might be surprised to learn that until recently the most popular way of representing pressure by TV meteorologists (and still used by pilots) is in inches or millimeters. For example, they would say that a typical sea-level pressure is 760 mm or about 30 inches.

Technically, pressure cannot be measured in inches or millimeters. After all, those are units of distance, not pressure.

The reason for doing so can be traced to the traditional method for measuring pressure. In that method, the air pressure is related to the height of a column of mercury. Thus, when the pressure was given in terms of inches or

---

[ii] For example, at the cruising altitude of commercial jet aircraft, the air pressure is about one-quarter the pressure at sea-level.

millimeters, they were not really telling you the pressure. Rather, they were telling you the height of the mercury column that the pressure corresponded to.

To explain how the height of a mercury column is related to air pressure, let's first explore how one might measure air pressure.

One way to measure the air pressure would be to take a piece of paper and physically remove all of the air on one side. The force needed to keep the paper stationary would equal the force exerted by the air on the other side. The pressure is just the force divided by the area of the paper.

As most people can guess, it would take a great deal of force to keep the paper stationary if one side of the paper was evacuated. After all, "nature abhors a vacuum." It is only because air exerts the same force on both sides of the paper that we don't feel the air pressure.

A more practical way to measure the air pressure is via a **mercury barometer** (see figure to right). A vacuum is created by inverting a long tube full of mercury and letting it fall out until the weight of the column is balanced by the force of the air pushing it back in (we assume there is a vacuum at the top of the column where the mercury used to be).

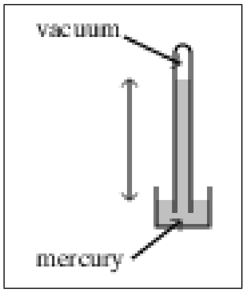

With a mercury barometer, we find that, at sea-level, the column of mercury is typically about 760 mm high (i.e., about 30 inches). That means that the 760-mm high column of mercury has the same "weight" as the atmosphere. The mercury barometer essentially is a two-pan scale, balancing the weight of the mercury on one side and the weight of the air on the other.

> ✍ A **barometer** is an instrument that measures pressure ("bar" means pressure). Thus, a mercury barometer is an instrument that measures pressure by using mercury. Such barometers need to be held strictly vertical so the mercury doesn't pour out. Nowadays, most barometers do not use mercury.

Check Point 4.4: What would weigh more: a container of mercury 30 inches deep or a container of air 30 inches deep?

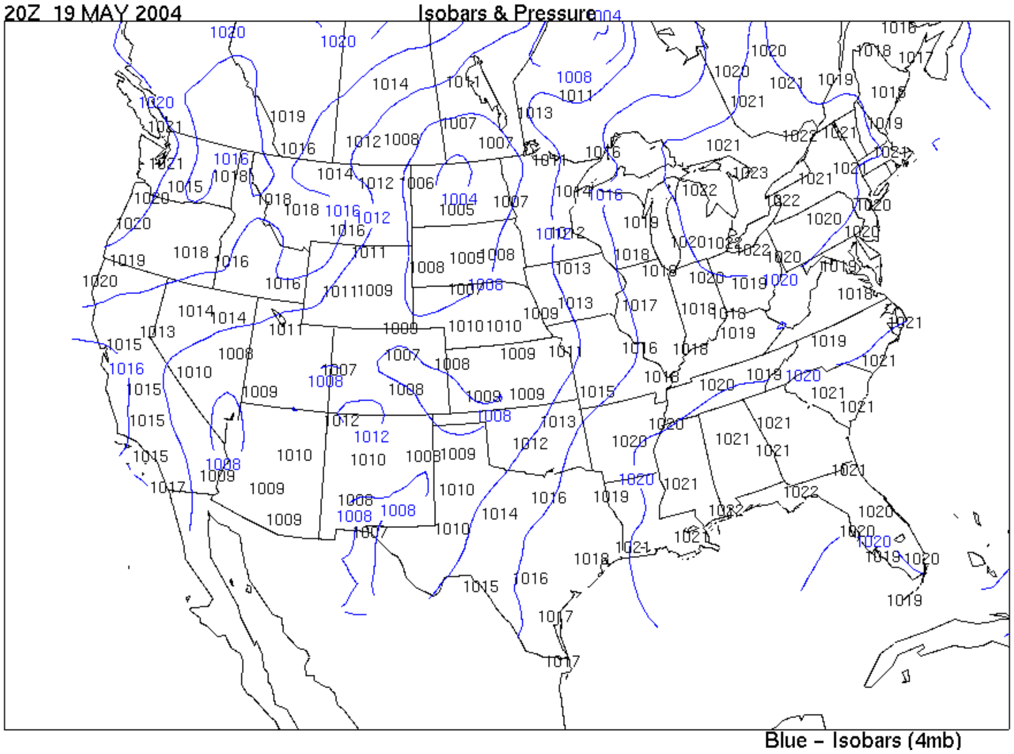

Figure 4.1: A map of sea-level pressure for 20 UTC on 19 May 2004 (source: DataStreme).

4.6 Horizontal variation

The value given above for the standard sea-level pressure is 1013.25 mb. The actual sea-level pressure, of course, varies from place to place and from time to time. Typical values can vary from, say, 950 mb to about 1050 mb.

As with temperature, we will examine the horizontal variation of pressure via a **synoptic maps**. An example of such a map is shown in figure 4.1. The **isolines** in this case are called **isobars** because they connect locations that have the same pressure.

Important: The values depicted in figure 4.1 are not those of the air pressure at the surface but rather the estimated pressure *at sea level* (more on that in section 15.3.3).

WHY IS SEA-LEVEL PRESSURE SHOWN ON THE MAP INSTEAD OF GROUND-

LEVEL PRESSURE?

As mentioned before, the air pressure decreases as one goes higher in the atmosphere. Consequently, the pressure will be less where the ground is very high (like in the Rocky Mountains) than where the ground is very low (like along the coast). A map of ground-level pressure wouldn't show us much more than where the ground is high and where it is low (i.e., how the pressure changes in the vertical). A map of sea-level pressure, on the other hand, tells us something about how the pressure changes in the horizontal.

Check Point 4.5: On a map of sea-level pressure for the United States, does the pressure shown on the map equal the pressure measured at the surface of the Earth? If not, which is greater: the sea-level pressure or the pressure at the surface?

In section 3.5.1, it was mentioned that synoptic maps provide information about the temperature **gradient** as well as the temperature. Similarly, synoptic maps provide information about the pressure gradient as well as pressure. The pressure gradient is measured in the same way as the temperature gradient. Mathematically, the pressure gradient in direction $\hat{x}$ is indicated as $\partial P/\partial x$. Thus, the vertical pressure gradient is written as $\partial P/\partial z$.

✍ As discussed with temperature, the proper way to represent the total horizontal gradient, not just the gradient in the $\hat{x}$ direction, is ∇P.

Check Point 4.6: Suppose you are examining a region of a synoptic map where the isobars are close together. Is this a region of large pressure gradient or small pressure gradient?

4.7 Vertical variation

As mentioned in section 3.6, the vertical coordinate on a skew-T log-P graph is not pressure but rather the logarithm of pressure.

The reason we use the logarithm of pressure is because the pressure tends to decrease exponentially with height.

Why does the pressure decrease exponentially with height?

This will be explained in chapter 15. For now, we simply need to accept that it does.

What does it mean to decrease exponentially with height?

If the pressure decreases exponentially with height, that its value can be modeled by the following equation:

$$P(z) = P(0)e^{-z/H} \tag{4.2}$$

where H represents the height at which the pressure falls to $1/e$ of its initial value (H is called the **scale height**).

The letter e represents a number (about 2.718282). This number has a special property. When this number is raised to a power, say e^x, we find that the value changes at a rate equal to its value. In other words, $d(e^x)/dx$ is equal to e^x. So, as it gets larger and larger, so does its rate of change.

> ✍ Population growth, for example, tends to be exponential since the growth rate is greater when the population is greater.

With atmosphere pressure, the pressure *decreases* exponentially. That is why there is a negative sign in equation 4.2. As the pressure gets smaller, the rate at which it continues to decrease slows. So, near the surface, where the air pressure is high, the pressure changes quickly with height. Very high in the atmosphere, where the air pressure is low, the pressure doesn't change much at all with height.

In equation 4.2, $P(0)$ represents the value of the pressure when z is equal to zero. When z is equal to zero, $e^{-z/H}$ equals one (since e^0 is equal to one). By multiplying by $P(0)$ we ensure that the pressure equals the surface pressure when z equals zero.

As mentioned above, the letter H in equation 4.2 represents the scale height. When the height z equals the scale height H, $e^{-z/H}$ is equal to e^{-1}, which is equivalent[iii] to approximately 1/2.718282 or a little more than one-third. Thus, the scale height represents the height at which the pressure falls to $1/e$ (a little more the one-third) of what the pressure is at the surface.

[iii] Remember that x^{-1} is equivalent to $1/x$.

Check Point 4.7: Why is it said that the scale height represents the height at which the pressure is about one-third the pressure at the surface?

SO WHAT DOES THIS HAVE TO DO WITH PLOTTING THE LOGARITHM OF PRESSURE ON A SKEW-T LOG-P GRAPH?

The logarithm of a number essentially represents the **order of magnitude** of a number. Every increase of one in the logarithm represents an increase by a factor of ten in the number (e.g., if the logarithm increases from 3 to 4, the number increases by a factor of ten).

For example, the logarithm of 100 is 2 and the logarithm of 1000 is 3. A ten-fold increase from 100 to 1000 is reflected in an increase in the logarithm from 2 to 3. Similarly, the logarithm of 10,000 is 4 and the logarithm of 100,000 is 5. A ten-fold increase from 10,000 to 100,000 is reflected in an increase in the logarithm from 4 to 5.

Notice that in both cases, the logarithm increased in a uniform manner (an increase of one in both cases) while the number itself changed a great deal more in one case (10,000 to 100,000) than in the other (100 to 1000).

In a similar way, the pressure decreases much more quickly near the surface of the Earth than higher up in the atmosphere, while the logarithm of the pressure changes in a more uniform manner. Consequently, the vertical scale on the skew-T log-P chart can be interpreted somewhat like height.

It is for this reason that the logarithm of the pressure is plotted instead of the pressure itself.

WHAT DOES IT MEAN TO INTERPRET THE VERTICAL SCALE LIKE HEIGHT?

To see what I mean by interpreting the vertical scale as though it represented height, examine the skew-T log-P chart on page 39. Notice that the horizontal lines indicating 1000 mb and 900 mb are rather close, especially when compared to the horizontal lines indicating 400 mb and 300 mb. This is because the heights where 1000 mb and 900 mb exist in the atmosphere are much closer together. In other words, if you were in a very tall elevator and moving upward at a constant speed, you'd move from 1000 mb to 900 mb in less time than it would take you to move from 400 mb to 300 mb.

Mathematically, the 1000 mb and 900 mb lines are closer together than the 400 mb and 300 mb lines because 1000 mb to 900 mb represents only a

10% decrease in pressure (a decrease of 100 mb out of 1000 mb) whereas 400 mb to 300 mb represents a 25% decrease in pressure (a decrease of 100 mb out of 400 mb). The height difference is roughly proportional to the fractional decrease in pressure. For example, you should be able to see that the spacing between 1000 mb and 750 mb (you'll have to estimate where that is) is roughly the same as the spacing between 400 mb and 300 mb. In our ideal elevator moving upward at a constant speed, you'd move from 1000 mb to 750 mb in roughly the same amount of time it would take you to move from 400 mb to 300 mb.

✍ A more mathematical treatment is left for the reader to complete in Show-me 4.2.

Check Point 4.8: Suppose you took a ruler and measured the distance from 1000 mb to 800 mb on a skew-T log-P diagram. How would that compare to the distance from 800 mb to 600 mb on the diagram? What about the distance from 800 mb to 640 mb (80% less than 800 mb)?

Project

Project 4.1: Create a graph[iv] *of the average pressure profile of the atmosphere from sea-level up to 86 km, assuming the pressure decreases exponentially according to equation 4.2 on page 50):*

$$P(z) = P(0)e^{-z/H}$$

where H *is 8000 m and* $P(0) = 1013.25$ mb. *If you do this correctly, you'll find that the pressure at the U.S. Standard Atmosphere tropopause (where the temperature no longer decreases with height) is approximately 250 mb. Note: In Excel, one can use "EXP(A1)" to determine the value of e raised to the value stored in cell A1.*

Project 4.2: From your U.S. standard atmosphere, how high does one have to go to reach a pressure one-half that at the surface?

[iv] It is recommended that you use the same spreadsheet that you used for the project in the previous chapter. You will use the same data for later projects.

Project 4.3: Suppose the pressure decreases exponentially with height according to equation 4.2. Answer the following questions assuming the scale height is 8000 m and the surface pressure is 1013.25 mb.
(a) Use the equation and solve for the height at which the pressure is equal to one-half the value at the surface. Compare your answer to that observed in the project.
(b) Use the equation to solve for the height at which the pressure is equal to 1/*e of the value at the surface, where e is equal to 2.718282.*
(c) Compare your answers in (a) and (b). Are they the same? Should they be? Explain.

Problems

Problem 4.1: (a) Express the typical sea-level pressure (see page 45) in Pascals and millibars.
(b) Estimate the air pressure (in mb) at the tropopause (use the project information).

Problem 4.2: Figure 4.1 displays observed values of sea-level pressure over the United States for 20 UTC on 19 May 2004.
(a) What interval is used for the isobars on the map? Include units.
(b) Identify the estimated sea-level pressure at East Stroudsburg. Include units.
(c) Should the surface pressure at East Stroudsburg be larger than that identified in (b), equal to, or less? Why?
(d) What is the range of sea-level pressures that are observed? Include units.
(e) Where is the largest horizontal pressure gradient located on the map? Calculate the value of the horizontal pressure gradient at that location.

Problem 4.3: (a) Determine the weight[v] *of a mercury column of height 760 mm and area 1 cm*^{2}*. Assume a density of 13.5951 g/cm*^{3}*. Be careful about your units.*
(b) Using this value, determine the pressure the air must exert on the bottom of the mercury column in order to keep it from falling. Compare it to the value given on page 45 for the typical sea-level pressure. Explain the difference, if

---

[v] The weight of an object on Earth is equal to its mass times g, which is 9.80665 N/kg at sea-level.

any.
(c) One square centimeter was used for the area in part (a). Try using a different cross-sectional area. Does your answer in (b) depend on the cross-sectional area specified in part (a)? Should it? Why or why not?

Problem 4.4: If water was used in the barometer instead of mercury, how high of a column of water would be needed to balance the typical sea-level air pressure given on page 45?

Derivations

Show-me 4.1: (a) Using the fact that 1 bar equals 10^5 Pa and there are 1000 millibars in one bar, show the conversion process needed to convert 101,325 Pa into 1013.25 mb.
(b) Show that hectoPascals and millibars are equivalent.

Show-me 4.2: Mathematically, we can show that the logarithm of a product is equal to the sum of the logarithms

$$\log(AB) = \log A + \log B$$

and the logarithm of a number raised to power equals the power times the logarithm of the number

$$\log(A^B) = B \log A.$$

Using these two properties, show that one can apply the logarithm to equation 4.2 to get

$$\log P(z) = [-1/(H \ln 10)]z + \log P(0)$$

which shows that the logarithm of the pressure $\log P(z)$ *varies linearly with height* z*, with slope equal to* $-1/(H \ln 10)$ *and* y*-intercept equal to* $\log P(0)$*.*

5. Humidity

5.1 Introduction

So far we have looked at wind, temperature and pressure. With those three, we can do a lot of physics. However, we have skipped over one aspect of the atmosphere that is possibly the biggest concern of weather watchers – water.

There are three aspects of water that we are concerned about: humidity, clouds and precipitation. All three are indicated in the weather station plot illustrated in figure 1.2 on page 5. The humidity is indicated by the "Dew Point", the clouds are indicated by the "Sky Cover" and the precipitation is indicated by the "Weather" (although the weather symbol can also indicate other things like haze and fog).

I will focus on the humidity in this chapter and will look at clouds and precipitation in the next chapter.

5.2 Water and its states

Whereas clouds and precipitation consist of solid or liquid water, humidity is a measure of the water **vapor** in the air.

WHAT IS THE DIFFERENCE BETWEEN WATER VAPOR AND LIQUID WATER?

Water vapor is water in the gaseous state. Like the other gases in the air, we cannot see gaseous water. For example, there is water vapor in the breath you are exhaling right now. Only when the water vapor condenses will we then be able to see the liquid water (as on a cold day).

WHAT IS THE DIFFERENCE BETWEEN VAPOR AND GAS, THEN?

Typically the term vapor is used to describe the gaseous phase of a material that can exist as liquid at typical temperatures. So, for example, since water can exist both as a liquid and a gas at typical temperatures, we will use the term **water vapor** for the gaseous phase of water. In comparison, since nitrogen does not normally exist as a liquid, we will use the term nitrogen **gas** for the gaseous phase of nitrogen.

Despite the difference in terms, however, the gaseous phase of water and the gaseous phase of nitrogen behave pretty much the same. In other words, both look and behave like, well, like air, which is not too surprising as air is a combination of gases.

Since water vapor is invisible, when it condenses into liquid water the newly formed water droplets appear to form from "thin air". Thus, clouds (which are made up of many water droplets) appear to form from "nothing". In a similar way, liquid water appears to just disappear when it evaporates.

Can water vapor exist below the boiling point of water?

Yes.

Since one cannot see the gaseous phase, one may mistakenly believe that the gaseous phase cannot exist below a liquid's boiling point or that the vapor, if it exists, must have a temperature above the liquid's boiling point.

On the contrary, we know that there is water vapor in the air we exhale (as a result of the respiration process) and our breath is certainly[i] below 100°C (the boiling point of water). Similarly, there is water vapor in the room you are in right now – you just can't see it[ii].

Couldn't the water vapor have a higher temperature than the air temperature?

No. For our purposes, we will assume that if the water vapor is coexisting with the air, the water vapor must have the same temperature as the other gases in the air. If that was not the case, the air would warm up and the water vapor would cool.

What about steam? Is steam the same as water vapor?

Both **steam** and mist consist of tiny drops of liquid water drops. While some people might use the term *water vapor* to describe them, such usage is not

---

[i] The temperature of our breath is closer to the temperature of our body, which is about 37°C.

[ii] As another example, consider that liquids can't burn yet we seem to "burn" **gasoline**. The gasoline evaporates (even if you don't boil it) and it is the gasoline vapor that burns. That is why gasoline is such a hazard.

consistent with the definition of water vapor as a gas. Gases are *not* made up of very tiny liquid drops. To be a liquid, even a very tiny drop of liquid, the molecules need to be bound in some way. In a vapor, like a gas, the molecules are not bound to each other.

Check Point 5.1: Can one "see" water vapor?

5.3 Humidity

When there is a lot of water vapor in the air, the air feels sticky. We'll examine later why that is. For now, let's look at how the water vapor content is measured.

The water vapor content is best reflected in the pressure exerted by the water vapor. However, this is hard to measure directly, so we typically report something called the dew point instead. The higher the dew point, the greater the vapor pressure and the stickier it feels.

5.3.1 Vapor pressure

DOES THE WATER VAPOR EXERT THE SAME PRESSURE AS THE AIR?

No. The pressure exerted by the water vapor is a small fraction of the air pressure. This is because the water vapor pressure (which we typically just call the **vapor pressure**) is the pressure of the vapor molecules alone. It would be the air pressure only if we could remove all of the other molecules, in which case the air pressure would be a small fraction of what it normally is.

> ✍ Whereas the air pressure near the surface of the Earth is about 1000 mb, the pressure of the water vapor is only 25 mb or so at most.

Since the water vapor pressure is different than the air pressure, we'll use different variable abbreviations for the pressure of each.[iii] In particular, whereas

---

[iii] In comparison, the water vapor has the same temperature as the rest of the air, so we can use T for the temperature of each.

we used P_d for the pressure of the dry air, we'll use e to represent the vapor pressure.[iv]

> ✍ In order to avoid confusion with the e used for exponential relationships, many people use exp instead of e for the exponential one.

Check Point 5.2: What letter do we use for the variable abbreviation of vapor pressure?

5.3.2 Dew point

If you cool the air, you eventually reach a point where some of the water vapor will **condense** and form liquid water (which you can then see).

The temperature when that occurs is known as the **dew point**, because dew refers to the liquid water droplets that form on surfaces when the temperature decreases to the dew point.[v]

> ✍ In chapter 12, I provide an explanation for why condensation occurs when the temperature cools.

In practice, the dew point is determined by taking a metal plate and cooling it. The temperature at which droplets appear on the plate is then the dew point.

As mentioned earlier, for a particular vapor pressure, there is a particular temperature at which the water will start to condense. I use the word "start" because not all of the water vapor condenses when the dew point is reached. Rather, only some water vapor condenses, which decreases the vapor pressure. Lowering the temperature will get more water vapor to condense.

This is why the dew point is almost always at or below the air temperature, as can be seen from the skew-T log-P chart on page 39 (figure 3.2); the dew point is indicated by the vertical squiggly line to the left of the temperature

---

[iv] The reason for using a different letter is because we will eventually need to distinguish between the actual vapor pressure and the vapor pressure that would be needed at saturation. Using e instead of P_v allows us to avoid needing multiple subscripts.

[v] This typically happens during the night as the air cools due to radiation (see chapter 13) and the dew is seen in the morning on plants and other surfaces.

line). If the dew point were above the temperature then water would condense and very quickly lower the vapor pressure and, subsequently, lower the dew point down to the temperature.

Check Point 5.3: According to the skew-T log-P diagram on page 39, what is the dew point at 900 mb?

As the air is cooled, we assume that the vapor pressure remains the same. This means that no water vapor condenses until the dew point is reached.[vi]

In other words, the dew point does not depend upon the temperature. It only depends upon the vapor pressure. Since there is a one-to-one correspondence between the dew point and the vapor pressure, we can obtain the vapor pressure by measuring the dew point.

GIVEN THE DEW POINT, HOW DO WE GET THE VAPOR PRESSURE?

The simplest way is to just look up the values in a table. An example of such a table is available from the course web page.

The table simply lists the vapor pressures for various dew points. The advantage of this is that one can determine the vapor pressure very quickly if the dew point is provided. The disadvantage is that only dew points every 1°C are provided. Consequently, one needs to estimate the value for dew points in between those values.

Check Point 5.4: According to a vapor pressure table, what vapor pressure corresponds to a dew point of 20°C?

There are also equations that can give you the vapor pressure. The simplest is the **Tetens formula**. This formula is not based on any physics. It is **empirical**, which means that it just happens to provide a fit to the saturation curve, rather than being derived from basic principles. Tetens' formula is as

---

[vi] This also means that the air pressure itself doesn't change. This is an adequate assumption as long as the air is free to contract and/or expand in order to maintain the same pressure.

follows.[vii]

$$e_s(T) = e_0 \exp\left[\frac{b(T-T_1)}{T-T_2}\right] \tag{5.1}$$

where

$$\begin{aligned} e_0 &= 0.611 \text{ kPa (6.11mb)},\\ b &= 17.2694,\\ T_1 &= 273.16 \text{ K, and}\\ T_2 &= 35.86 \text{ K.} \end{aligned}$$

As you should be able to see from the project, the Tetens' formula reproduces the actual values pretty well. However, it won't necessarily reproduce correct values for temperatures lower than 0°C or higher than 100°C.

Check Point 5.5: According to the Tetens formula, what is the vapor pressure for a dew point of 20° C?

5.3.3 Relative Humidity

While dew point is plotted on the skew-T log-P chart and the weather map (as part of the station plot), non-meteorologists tend to use the **relative humidity**, rather than the dew point, to indicate the water vapor content.

The relative humidity indicates what the water vapor pressure is relative to the maximum it could be at that temperature. For example, if the dew point equals the temperature then the relative humidity would be 100%. In that case, the vapor pressure equals the maximum that could exist at that temperature.

The maximum is called the **saturation vapor pressure**. That is the vapor pressure that would exist if the dew point equalled the temperature. Thus,

---

[vii] The original reference is Tetens, O., 1930: Uber einige meteorologische Begriffe. *Zeitschrift fur Geophysik*, Vol. 6:297. The constants used here are from Stull, *Meteorology for Scientists and Engineers*, 2000, Brooks/Cole. Different sources have slightly different values of the constants.

to find the saturation vapor pressure, simply use the same table (or equation) as before but use the actual air temperature rather than the dew point.

Once you have the vapor pressure and the maximum vapor pressure values, then, the relative humidity is the ratio of the two:

$$\mathrm{RH} = \frac{e}{e_{\mathrm{s}}} \tag{5.2}$$

where e_{s} is the saturation vapor pressure.

> ✍ The relative humidity does not indicate the amount of moisture directly but rather expresses the amount as a percentage of what would be present if the air was saturated.

Is relative humidity also equal to T_{d}/T?

No. This would be true only if the saturation vapor pressure varied linearly with temperature. It doesn't.

Check Point 5.6: (a) If the air is saturated and the temperature is 20°C, what is the vapor pressure, saturation vapor pressure and relative humidity?
(b) Suppose the relative humidity is 50% and the air temperature is 20°C. What is the vapor pressure and saturation vapor pressure?

5.3.4 Specific humidity

Although dew point and relative humidity are the most popular ways of describing how much water vapor is present, a third way will prove to be more useful when dealing with how much water changes state and the impact on heating and cooling the air.

In this approach, we examine the *mass* of water vapor and compare that to the total mass of the air. This ratio is called the **specific humidity** and is indicated by q:

$$q = \frac{m_{\mathrm{v}}}{m_{\mathrm{T}}}. \tag{5.3}$$

Like density, specific humidity is independent of size and we can treat the air as though it has a specific humidity at each location in space. Another advantage of specific humidity is that, unlike the vapor pressure, the specific humidity remains the same as a parcel expands or contracts (i.e., the ratio of water vapor to the total air remains the same). Such expansion and contract can occur, for example, when the parcel rises (and encountering lower air pressure) or falls (and encountering higher air pressure).

For the earth's atmosphere, the specific humidity is usually a small number since the water vapor makes up such a small portion of the air. For this reason, it is usually expressed in terms of grams of water vapor per kilograms of air. Thus, if 0.6% of the air mass is water vapor, we'd write that the specific humidity is 6 g/kg instead of 0.006.

> ✍ The specific humidity value is typically obtained from the dew point value or the vapor pressure value (see, for example, Show-me 7.2).

Check Point 5.7: Why is the specific humidity typically written with units of g/kg?

5.3.5 Mixing ratio

While it is sufficient to specify the amount of water vapor in the air via the specific humidity, there is another very similar parameter that is frequently used instead: the **mixing ratio**.

The mixing ratio, like the specific humidity, is a ratio of masses. However, rather than the ratio of water vapor to total air, the mixing ratio is the ratio of water vapor to the *dry air only*:

$$r = \frac{m_{\mathrm{v}}}{m_{\mathrm{d}}} \tag{5.4}$$

where r is used to represents the mixing ratio[viii].

Since water vapor typically makes up a small portion of the air, m_{d} is usually very close to m_{T} and so the mixing ratio is typically close to the specific humidity. The exact relationship between the two is as follows:

$$q = \frac{r}{1+r} \tag{5.5}$$

[viii] Some people use w to represent the mixing ratio.

Since r is typically much less than 1, the denominator is approximately 1 and $q \sim r$.

Where did equation 5.5 come from?

Equation 5.5 can be easily derived as follows. Since

$$q = \frac{m_{\mathrm{v}}}{m_{\mathrm{T}}} = \frac{m_{\mathrm{v}}}{m_{\mathrm{v}} + m_{\mathrm{d}}},$$

we can write

$$\begin{aligned} \frac{1}{q} &= \frac{m_{\mathrm{v}} + m_{\mathrm{d}}}{m_{\mathrm{v}}} \\ &= 1 + \frac{1}{r} \\ &= \frac{r+1}{r} \end{aligned} \tag{5.6}$$

Simply invert to get equation 5.5.

Why would one use mixing ratio over the specific humidity?

Since their values will be very similar, there really isn't any advantage of mixing ratio over specific humidity except that, as will be shown later, there are some equations that are a little simpler when expressed in terms of mixing ratio instead of specific humidity.

Check Point 5.8: It was mentioned that the specific humidity value is roughly the same as the mixing ratio value. Which is larger?

For every dew point and pressure, there is a specific mixing ratio. Because of this, the skew-T log-P chart also provides lines of mixing ratio. These are the dashed lines that are skewed like the temperature lines but not as great and they only go up to 600 mb (see figure 3.2 on page 39). They are labeled in units of g/kg and in increments of 0.4, 1, 2, 3, 5, 8, 12, 16 and 20 g/kg.

Check Point 5.9: Answer the following questions based on the skew-T log-P

chart on page 39,
(a) What is the mixing ratio when the air pressure is 900 mb and the temperature is 0° C?
(b) What is the actual *observed mixing ratio at 900 mb (use observed dew point)?*
(b) What is the mixing ratio that would be present at 900 mb if the air was saturated (use observed temperature)?

5.4 Boiling

If you examine the vapor pressure table, you may have noticed that the values at the end of the table look a little familiar. The temperature is 100°C and the pressure value is 101.325 kPa (i.e., 1013.25 mb). The temperature corresponds to the boiling point of water and the pressure corresponds to the standard sea-level pressure.

This isn't just a coincidence. To understand what is going, we have to take a small detour and examine **boiling**. Many students, even in college, do not understand what boiling is, and mistakenly think that the gaseous phase of a substance (e.g., water vapor) can't exist below the liquid's boiling point.

During boiling, little bubbles form in the water when the water reaches a certain temperature (called the **boiling point**). These bubbles are made of water vapor.

How do we know they are made of water vapor?

The bubbles form from water molecules that were separated from the other water molecules. At any temperature, there are some water molecules that are energetic enough to break from the other molecules and enter the vapor state. Those bubbles, therefore, only contain water vapor. As such, they must have a pressure equal to the saturation vapor pressure at that temperature. Anything less and additional liquid water would evaporate into the bubble. Anything more and some water vapor in the bubble would condense.

What is the pressure in the liquid water?

The pressure in the water is roughly equal to the atmospheric air pressure.

It is for this reason why we usually don't see bubbles form in water unless the temperature is at the boiling point. Below the boiling point, the vapor pressure inside the bubbles is less than the atmospheric pressure and the bubbles are "squashed." At the **boiling point**, the vapor pressure inside the bubble equals the atmospheric air pressure and any bubbles that form will remain and the liquid water will continue to evaporate into the ever-increasing bubbles.

> ✍ During boiling, the water stops warming, as all the energy goes toward evaporating the liquid into the bubbles. The bubbles, meanwhile, being less dense than the liquid, are buoyant and rise to the surface.

It turns out that the saturation vapor pressure equals that the standard sea-level pressure (1013 mb) when the temperature is 100°C. That is why water doesn't boil at standard sea-level pressure until the temperature reaches 100°C.

Check Point 5.10: (a) Based upon the above discussion, what is the saturation vapor pressure at 212°F?
(b) Suppose the air pressure was less than 1013 mb. What would happen to the boiling point, increase or decrease?

Project

For this project, you need to download the saturation vapor pressures (in kPa) from the course web page and import the values into a spreadsheet program like Excel.

Project 5.1: Create a graph of the saturation vapor pressure over water with respect to temperature (for the temperatures for which values are available). The curve that results is called the **saturation curve***.*

Project 5.2: Add a column to the spreadsheet. In the column, calculate the saturated vapor pressures at each temperature according to the Tetens' formula (5.1):

$$e_{\mathrm{s}}(T) = e_0 \exp\left[\frac{b(T - T_1)}{T - T_2}\right]$$

where

$$\begin{aligned} e_0 &= 0.611 \text{ kPa } (6.11 \text{ mb}), \\ b &= 17.2694, \\ T_1 &= 273.16 \text{ K, and} \\ T_2 &= 35.86 \text{ K.} \end{aligned}$$

Add a curve for those values to your graph.

Problems

Problem 5.1: For each of the following, specify which state the water is in: vapor, liquid or solid.
(a) Steam (the stuff you see coming out of a boiling teapot).
(b) Clouds (the white puffy things you see in the sky).
(c) Snow.

Problem 5.2: Which is more likely, an air pressure of 1000 mb or a vapor pressure of 1000 mb, or are they equally likely?

Problem 5.3: Suppose a bubble forms inside a liquid via evaporation. Using equation 5.1 or your graphs from problem 5.1, for what temperature will the pressure inside the bubbles (i.e., the saturation vapor pressure) be equal to standard sea-level pressure (see appendix B.2)?

Problem 5.4: (a) Based on the discussion, what is the boiling point in a place like New York City, where the elevation is near sea-level?
(b) What should be the boiling point in a place like Denver, where the surface pressure is about 850 mb? Explain.
(c) At which location would it take longer to "cook" an egg by placing it in boiling water (hard-boiling)? Explain.

Problem 5.5: Determine the saturation vapor pressure when the temperature is 100° C via (a) the information downloaded in project 5.1 and (b) the Tetens' formula (5.1). Do they agree? Why or why not?

Problem 5.6: (a) Give an estimate for the temperature of the water vapor is this room.
(b) Suppose an endless source of liquid water was present in the room. What

would be the pressure of the water vapor in the room (i.e., vapor pressure) once equilibrium is reached (i.e., as much vapor evaporating from the liquid as condensing into the liquid)?

Problem 5.7: Using the information downloaded in project 5.1, estimate the dew point when the vapor pressure is 5 mb and the air temperature is 20° C. Explain your reasoning.

Problem 5.8: On the skew-T log-P charts, the observed dew point at each level is indicated by an additional squiggly line to the left of the temperature line. In order to get a cloud to form, the temperature must be less than the dew point. However, the dew point curve on the skew-T never gets to the right of the temperature curve, even when there are clouds. Explain.

Problem 5.9: If the dew point equals the temperature, what is the relative humidity?

Problem 5.10: (a) Identify the observed dew point and air temperature at the surface for Little Rock, Arkansas, according to skew-T in figure 3.2 (on page 39).
(b) Based on your values in (a), what is the vapor pressure and saturation vapor pressure at the surface? Feel free to use the saturation table or the Tetens' formula.
(c) Based on your values in (b), what is the relative humidity at the surface?
(d) Calculate the ratio of the dew point and the air temperature at the surface. Try it once with units of ° C and once with K. Does the ratio equal the relative humidity? Do you expect it to? Explain.

Problem 5.11: Suppose a given parcel of air has a total mass of 1 kg and, of that, the water vapor has a mass of 3 g.
(a) What is specific humidity in g/kg?
(b) What is the mixing ratio in g/kg? Pay attention to the units.

6. Clouds and Precipitation

6.1 Clouds

As we know from the previous chapter, water vapor will condense and form liquid water when the temperature drops to the dew point. At the Earth's surface, the liquid water forms on the leaves of plants and other objects, forming dew. What happens above the Earth's surface, where there are no such objects to condense onto?

The answer is that the water condenses onto particles in the air. These are called **Cloud Condensation Nuclei** (or **CCN**, for short). The number and size of these CCN will vary depending on location but there is usually no shortage of such CCN.

For example, in a typical liter of air, there are about 1 million CCN over oceans, 10 million CCN over land and 100 million CCN over polluted air over cities. The concentration goes down as one goes higher in the atmosphere but is still quite high (over 10,000/liter) up into the stratosphere.

A cloud, then, consists of millions of tiny little water droplets per liter. Since these are *liquid* droplets and not vapor, we can see them. However, because they are so small, they scatter light and the cloud appears white (rather than clear, like liquid water).

> ✍ Some CCN particles (like smoke) attract water vapor molecules even before the dew point is reached. Consequently, they can grow big enough to impact visibility even if the relative humidity is less than 100%. This is the mechanism behind **haze**

Check Point 6.1: (a) What are CCN and how many are typically in a liter of air?
(b) Based on this, how many cloud droplets per liter would you expect in a cloud?

Why do the clouds float in the air?

Liquid water, being more dense than air, falls. Consequently, so do the cloud droplets. However, there is **air resistance** (or air **drag**) and this prevents them from falling too quickly.

According to Newton's second law, the gravitational force on a droplet acts to accelerate it downward. However, as the droplet picks up speed, the air resistance increases and quickly reaches a point where it counters the gravitational force. At that point the droplet no longer speeds up. We call the maximum speed the **terminal velocity** or the **fall velocity**.

Because of the low terminal velocity of cloud droplets, it doesn't take much upward motion of the air to keep them up. Even if there was no upward motion the cloud droplets would fall so slowly that they'd still appear to "float" in the sky.

Check Point 6.2: When something is falling at its fall velocity, what is the net force exerted on it: upward, downward or zero?

What is the size of a typical cloud droplet?

Typical cloud droplets have diameters between 1 and 100 μm in diameter (1 micrometer or **micron**, abbreviated as 1 μm, is equal to 1×10^{-6} m).

This is consistent with the assumption that the moisture that condenses is distributed equally among all of the cloud condensation nuclei. In fact, as you can show in Problem 7.4 and Show-me 7.1, if the mixing ratio decreases by 1 g/kg (due to condensation) and the moisture is distributed equally among CCN with number density 10^6 per liter, then the radius of the resulting drops is equal to about 7 microns.

Check Point 6.3: Assuming the moisture is distributed equally among all of the CCN, what happens to the size of each droplet if there are a greater number density of CCN (for the same decrease in mixing ratio)?

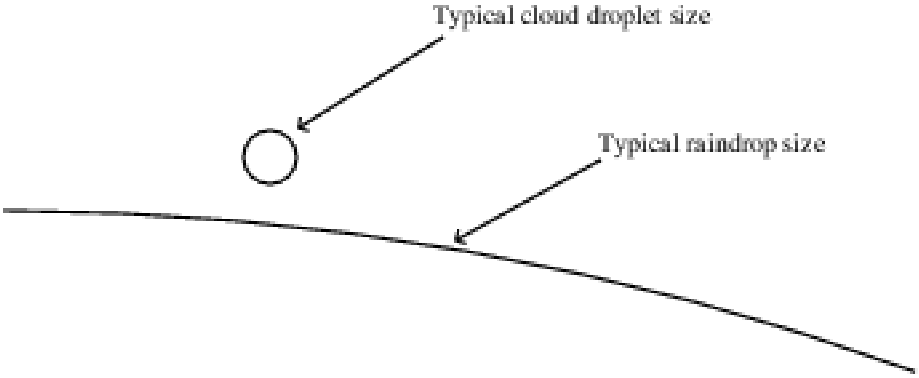

Figure 6.1: The relative sizes of a cloud droplet and a raindrop.

6.2 Precipitation

The difference between raindrops and cloud droplets is that raindrops are much bigger (see figure 6.1). Raindrops range in size between 1 and 5 mm in diameter (compared to 0.001 and 0.1 mm for cloud droplets).

Because a raindrop is larger, there is a larger gravitational force on the raindrop. This allows the raindrop to attain a larger terminal velocity.[i] For example, the terminal velocity of raindrops is about 5 to 10 m/s whereas the terminal velocity of cloud droplets is less than 0.3 m/s.

As a result, the raindrops fall out of the sky whereas the cloud droplets do not.

How do the raindrops form?

You might think that raindrops are just large cloud droplets. However, that isn't the case. As you can show in Show-me 7.5, there are just too many CCN's in the atmosphere to cause cloud droplets to get large enough to fall. Distributing the moisture across all those CCN's means the droplets are too small.

So, how do we get raindrops?

---

[i] The gravitational force is proportional to the drop's volume whereas the air resistance is only proportional to the radius or cross-sectional area of the drop. This means that as the drop grows in size, the gravitational force goes up faster (i.e., as the cube of the radius vs. the square or linear).

While condensation nuclei provide the mechanism for starting the cloud droplets, we definitely need another mechanism to get the cloud droplets to grow to raindrop size. Actually, there are two mechanisms. One is called **diffusion** and the other is called **collision and coalescence**.

Diffusion refers to the process by which water evaporates off the smaller cloud droplets (making them shrink in size) and condenses onto the larger droplets (making them grow). It turns out that the surface of small droplets are curved more than the surface of large droplets. The greater curvature means that there is more opportunities for a water molecule to evaporate (i.e., the saturation vapor pressure is greater for smaller droplets).

Compare, for example, the curvature of the raindrop to that of the cloud droplet in figure 6.1. From that perspective, the raindrop hardly curves at all (although this is a consequence of the magnification used).

This means that, given a particular vapor pressure in the atmosphere, it is possible for it to be more than enough to condense onto the raindrop (i.e., greater than saturation with respect to the raindrop) but *less* than what is needed to condense onto the cloud droplet (i.e., less than saturation with respect to the cloud droplet). For that vapor pressure, then, the small cloud droplet will evaporate whereas the larger raindrop will grow.

In other words, given a wide distribution of droplet sizes, the little droplets will evaporate and the bigger droplets will grow, eventually leading to a smaller number of bigger-sized droplets. Consequently, the big droplets will grow at the expense of the little droplets.

Collision and coalescence refers to the process by which droplets collide and merge (coalesce) to form bigger droplets.

✍ As we will see in section 6.3.3, when the temperature is less than 0°C there is another method of creating large raindrops: from ice crystals. This other method is actually more efficient than the collision and coalescence method.

Check Point 6.4: (a) How much bigger is the radius of a raindrop than the radius of a cloud droplet? (b) Name the two processes responsible for the formation of larger droplets.

6.3 Precipitation types

Now that you understand the basic process by which raindrops are formed, let's examine the processes by which other forms of precipitation are formed.

6.3.1 Freezing rain

When rain falls onto ground that is less than 0°C, it can freeze upon contact, forming a layer of ice. This is called **freezing rain**.

If the temperature is less than 0°C, why doesn't the rain freeze before hitting the ground?

Freezing rain occurs when the *ground* is less than 0°C. That doesn't necessarily mean that the *air* is less than 0°C. The air may actually be above freezing, even though the ground is below freezing.

Even if the air just above the ground is below freezing, the air above that can be above freezing. This is called an **inversion**. And, if the below-freezing layer is shallow, the raindrops may not have enough time to freeze before hitting the ground.

Further complicating matters is the fact that raindrops and cloud droplets don't necessarily freeze even if they cool to temperatures below freezing.

How can this be?

Contrary to popular belief, liquid water can exist below 0°C. Indeed, pure water will not freeze on its own until the temperature reaches −40°C or so. Such unfrozen water is called **supercooled** water.

✍| Keep in mind that supercooled is not the same as frozen!

More typically, some material is available, like an object on the ground, that initiates freezing. This is the case with freezing rain, where the rain freezes upon contact with the ground.[ii]

So rain doesn't freeze automatically when the air is less than 0°C?

[ii] Such freezing can exist when the drops encounter an airplane. This effect is appropriately called **icing**.

Not if the raindrop is made up of pure water. The creation of an ice crystal in a large water droplet is just as unstable as the creation of a water droplet in a large area of water vapor. As with the creation of water droplets, the creation of ice is enhanced via the use of impurities known as **freezing nuclei**.

Some impurities can also freeze water droplets on contact. Such impurities are called **contact nuclei** and initiate freezing much like contact with the ground.

ISN'T 0°C THE FREEZING POINT?

Yes, but that only means that freezing is possible. It is not guaranteed. Above 0°C, it isn't even possible.

In a similar way, condensation is possible below the boiling point. It isn't possible above the boiling point (since the air pressure isn't large enough; see section 5.4).

SO FREEZING CAN OCCUR BELOW 0°C. CAN MELTING OCCUR BELOW 0°C ALSO?

No. Even if a block of ice is introduced into an area that is very warm, the ice will first need to warm to 0°C before the ice melts.

This is different from evaporation, which can occur at any temperature.

SO CAN SOLID ICE EXIST ABOVE 0°C?

No. The ice itself will not warm to a temperature above 0°C before it melts entirely.

In a similar way, liquid water cannot exist above the boiling point.

Check Point 6.5: A popular demonstration is to place a test tube with distilled water in a beaker full of crushed ice, water and salt. Even when the temperature of the crushed ice mixture falls below 0°C, the water in the test tube remains unfrozen. After several minutes of being supercooled, a small piece of ice is dropped into the test tube and the distilled water suddenly freezes. Why is distilled used instead of regular tap water?

6.3.2 Sleet

Another precipitation type is called **sleet** (or **ice pellets**).

As mentioned above, there are impurities in the air called **contact nuclei** and **freezing nuclei** that, upon contact with the water droplet, initiate freezing much like contact with the ground.

This causes the droplets to freeze into little spheres of ice before they hit the ground. The longer the raindrops fall within the sub-freezing layer, the more likely this will occur. Consequently, when the cold layer is thick enough, sleet is produced instead of freezing rain.

Check Point 6.6: What is the difference between freezing rain and sleet?

6.3.3 Snow

WHAT IS THE DIFFERENCE BETWEEN SNOW AND SLEET?

Snow looks very different from sleet. Whereas sleet are frozen spheres of water, snow are ice crystals, hexagonal in shape.

The difference has to do with how they are formed. Whereas sleet forms when liquid raindrops freeze, **snow** is produced when the ice crystals form directly from the vapor.

HOW IS THAT POSSIBLE?

Just as there is a back-and-forth process between the liquid and vapor states (called condensation and evaporation), there is also a similar back-and-forth process between the vapor and solid states. The conversion from vapor to solid is called **deposition** and the conversion from solid to vapor is called **sublimation**[iii].

When water is frozen, it takes the form it had before it was frozen. Consequently, a frozen water droplet is round, not hexagonal. The deposition process, on the other hand, produces a crystal. The specific form of the crystal depends on the temperature but the basic crystal structure is hexagonal.

Most **snowflakes** are made up of aggregates of individual crystals.

CAN ICE CRYSTALS FORM IF THE TEMPERATURE IS ABOVE FREEZING?

No. Deposition only occurs for temperatures below freezing.

[iii] Some references use the term *sublimation* to refer to both processes.

It was mentioned before that cloud droplets require the presence of CCN. Do ice crystals also require such nuclei?

Yes. Getting water vapor to form an ice crystal directly is difficult for the same reason that it is difficult to form a water droplet directly from water vapor – due to the size of the initial crystal, large supersaturations are needed. Consequently, ice crystals are routinely not created without the help of impurities known as **deposition nuclei**.

The generic name for contact, freezing and deposition nuclei is **ice nuclei**.

> ✍ Once a droplet is frozen, additional deposition can now occur (i.e., the frozen droplet acts as a deposition nucleus).

Do the CCN also act as ice nuclei?

No. Most CCN will not act as ice nuclei. Whereas there may be a million CCN per liter, there may only be one ice nucleus per liter.

So, if the temperature is below freezing, wouldn't the overwhelming number of CCN favor the formation of droplets instead of ice crystals?

Yes.

Then how do we get snow?

All we need are a couple of ice crystals. Then, the ice crystals will grow at the expense of the water droplets.

Why?

Because, for a given temperature, evaporation from liquid water will be greater than the sublimation from ice. Thus, at the same temperature, a water droplet will evaporate faster than an ice crystal will sublimate. If the vapor pressure is sufficient to counter the sublimation rate, the ice crystal won't grow or shrink but the water droplet will evaporate.

That leads to an increase in vapor pressure which, in turn, leads to a growing ice crystal.

This is known as the **Bergeron-Findeisen** process.

So the ice crystals continue to grow. With only one ice nucleus per liter, they can get pretty large. Is there anything stopping their growth?

Just their weight. After a certain size, their fall velocity becomes large enough that they fall out of the sky.

So ice crystals will form instead of water droplets if the temperature is 0°C?

It depends. Usually, if the temperature is warmer than −15°C, the cloud consists of supercooled water droplets. On the other hand, if the temperature is colder than −20°C, the cloud usually consists of ice crystals.

Aren't most clouds at heights where the temperature is below −20°C, even in the summer?

Yes, in the extratropics.

In fact, this process (the **ice crystal process**) is much more efficient than the **collision and coalescence** process described in section 6.2 because, unlike the water droplets, the process doesn't slow down as the crystals get larger. Consequently, this is the more likely way of producing precipitation.[iv]

So why don't we get snow all the time?

Because the snow melts when it encounters air that is warmer than 0°C.

Check Point 6.7: The temperature profile in the skew-T graph shown on page 39 shows that, in that case, the temperature in the upper troposphere was less than 0° C whereas the surface temperature was greater than 0° C. In such situations, which process is likely responsible for rain: the collision and coalescence of cloud droplets or the melting of ice crystals (possibly created by the Bergeron-Findeisen process) as they fall?

6.3.4 Hail

Hail is produced when the rain gets pushed back up into the cold areas due to strong vertical motions. Repeating the process adds additional layers to the hail stone. Thus, it only occurs if the lower troposphere is above freezing and the upper troposphere is below freezing. In addition, it requires a large instability so that updrafts are strong enough to push the rain/hail back up to the upper troposphere.

[iv] This is more true for the extratropics than the tropics where the initial CCN are large and sparse enough with lots of moisture available.

6.4 Observing Precipitation

Measurements of precipitation are not used for numerical models since what is measured is the water that is no longer "in" the atmosphere. However, precipitation observations are needed for hydrological forecasts, like the potential for flooding, as well as for short-term forecasts (e.g., to see whether a storm is approaching).

There are two methods for observing precipitation.

One method has to do with the collection of the precipitation on the ground. The amount of precipitation can be determined by placing a bucket or graduated cylinder outside and just seeing how much water is collected. With the Automated Surface Observation Stations (ASOS), this process is carried out via a tipping bucket mechanism, where a bucket tips over when a certain amount of water is collected. The station can then determine the total amount of precipitation by counting the number of times the bucket tips. For frozen precipitation, a heater is applied to melt the precipitation first.

The other method has to do with "seeing" the precipitation while it is the process of falling. This is done via **radar**, which stands for Radio Detection And Ranging. An example of a radar map is shown in figure 6.2.

To understand how this process works, we need to first learn a bit about radiation. If you've already taken the second semester of an introductory physics sequence, this will just be review.

6.4.1 What is radiation?

Radiation is essentially a wave in the **electromagnetic field**.

Huh?

Consider an object that has charge, like an electron. Any other object that has a like charge will be repelled away from the electron.[v] If the electron moves, the force on the other object changes. Moving the electron back and forth causes an oscillatory force on the other object.

---

[v] And, conversely, if an object has an opposite charge, it will be attracted toward the electron.

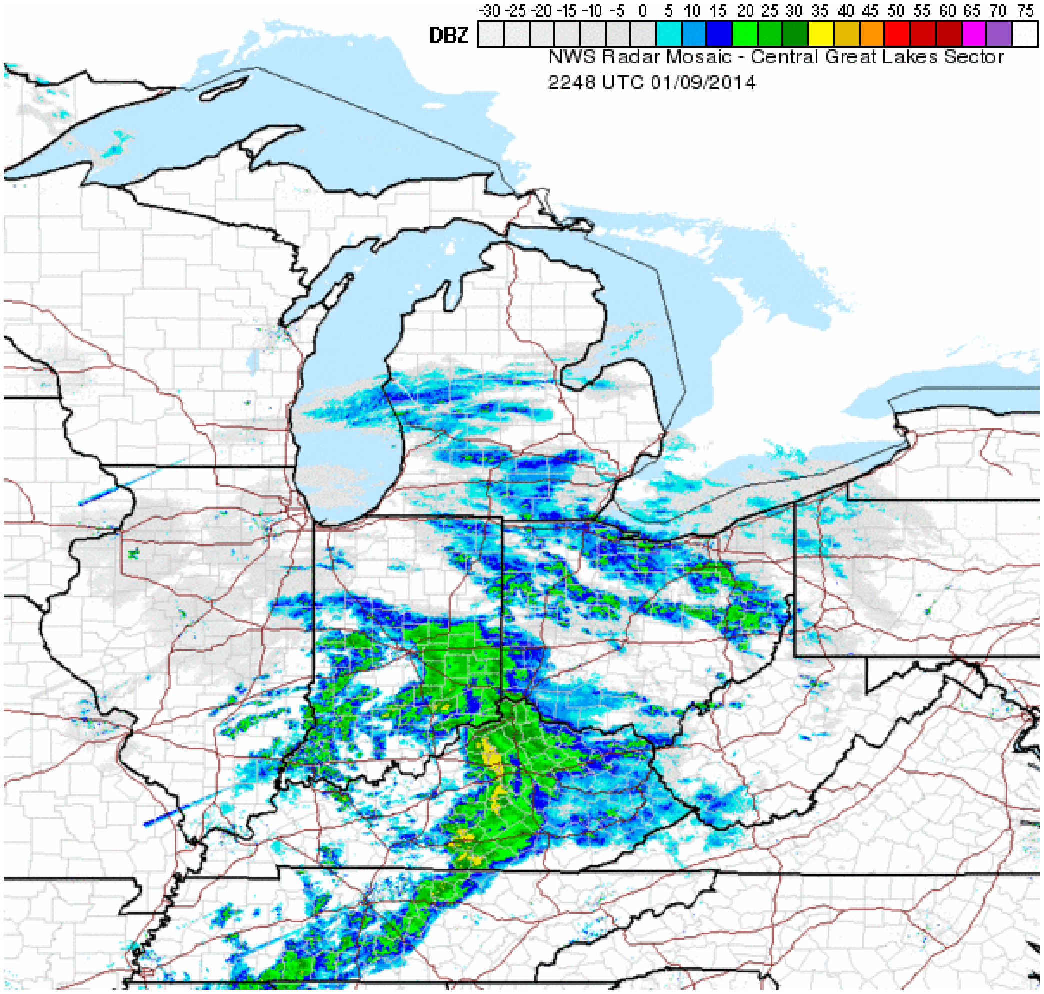

Figure 6.2: Reflectivity for the Midwest United States on January 09, 2014, at 2248 UTC.

The “information” from the electron does not reach the other object instantaneously. Rather, it travels at a speed known as the **speed of light**. Consequently, there is a slight time delay between the movement of the electron and the forcing on the other object. In this way, it is similar to making waves in one part of a pool and watching the movement of a toy boat in another part of a pool.

In both cases (the water wave and the electromagnetic wave), there is a “status quo” or condition where the environment wants to be and the wave comes about when one attempts to disturb that equilibrium position. For example, in a pool, the water prefers to be at one height. One makes waves by changing the height of the water. This “desire for the status quo” gives rise to a water wave that travels away at a particular speed, v.

A similar thing happens with the electric force. By changing the position of the electron, one is essentially changing the “force field” (or **electromagnetic field**) associated with the electron (i.e., its influence on other charged objects). Just like nature prefers the “status quo” for the height of the water in a pool, nature prefers the “status quo” for electromagnetic fields. Consequently, changing the electric field gives rise to a wave in the electromagnetic field that travels away at a particular speed v (i.e., the speed of light).

The general term for **electromagnetic waves** is called **radiation**. As we will see in section 6.4.3, visible light is one type of electromagnetic wave. Consequently, I will tend to use the words “electromagnetic waves,” ”radiation” and ”light” interchangeably.

Check Point 6.8: What is the difference between light and electromagnetic radiation?

6.4.2 Speed of light

For our purposes, it isn’t crucial to know how fast electromagnetic waves travel. It suffices to know that electromagnetic waves travel very, very fast! Of course, we already know that, as the time it takes between turning on a flashlight and the light beam hitting the wall is very small indeed!

With careful measurements, scientists have found that in the vacuum of space this speed is close to 300,000,000 meters per second and is typically indicated as c:

$$c = 2.99792458 \times 10^8 \text{ m/s}$$

as used with the equation $E = mc^2$ (see equation 9.1). The qualifier "in the vacuum of space" is used because this speed is only really valid in a region in which there are no molecules, atoms or other particles that may interact with the radiation.

What about if we don't have a vacuum?

Of course, we know that light can also travel in other materials, like air, glass and water. However, if something is present, the speed of the light will be slower. For example, it has been found that it travels 25% slower in water than it does in vacuum, and about 33% slower in glass than it does in a vacuum.

What about the speed in air? After all, we're talking about the atmosphere here, aren't we?

Yes, and fortunately (for us) the difference between the speed in a vacuum and the speed in air is very small. In fact, they differ by less than 0.03%. Thus, we typically use the speed in a vacuum for radiation in the atmosphere.

Check Point 6.9: Does light travel in a vacuum?

6.4.3 Frequency and wavelength

If light is an example of electromagnetic radiation, does that mean it is harmful?

When the term **radiation** is used, many people think of something harmful. Actually, we (as human beings) are normally oblivious to electromagnetic waves. We cannot see them, hear them, taste them or touch them. Usually, we need specialized instruments to detect them, with each instrument designed to detect a specific range of electromagnetic frequencies.

Here we have been using the word "radiation" as a general term for **electromagnetic waves**. Some confusion arises because the public tends to use the word "radiation" for only the very, very high frequency electromagnetic waves (such as X-rays)[vi] that at high doses can have an negative effect on living cells.

What distinguishes visible light from other types of electromagnetic radiation is its **frequency**.

WHAT IS FREQUENCY?

Frequency is the number of wave cycles that occur in a given amount of time. For example, in a pool, one can make water waves by tapping up and down on the surface of the water with a stick. The frequency of the wave would be given in terms of how many times the water goes up and down at a particular point in a given amount of time. For example, if you're a duck floating on the water and, as the wave passes, you find yourself going up and down twice every second, the frequency of the water wave would be 2 cycles per second or 2 Hertz (where **Hertz** is short for "cycles per second").

The frequency is typically indicated as f in equations (although some sources use ν). The frequency is the inverse of the **period** T (which is the time per cycle):

$$f = \frac{1}{T}$$

Just as water waves can be made with different frequencies, electromagnetic waves can be different frequencies, also. Since waves of different frequencies can have different properties, certain frequency ranges are "named" depending on our use (see figure 6.3) and the entire range of frequencies is called the **electromagnetic spectrum**.

For example, we tend to use the frequencies from about 1 to 10^{12} Hz for communication. Consequently, this range is called **radio waves** (although only a small part of that range is actually used for the **AM** and **FM** radio stations).

Unlike at higher frequencies, radiation at these frequencies do not appear to cause any ill effect on the human body. There are radio waves bombarding your body right now yet we are blissfully unaware of them (until someone turns on a radio or some other detector).

The human body is not built to detect electromagnetic waves except for a narrow range from about 5×10^{14} to 10×10^{14} Hz. We detect these waves

---

[vi] X-rays have a frequency greater than 10^{16} Hz.

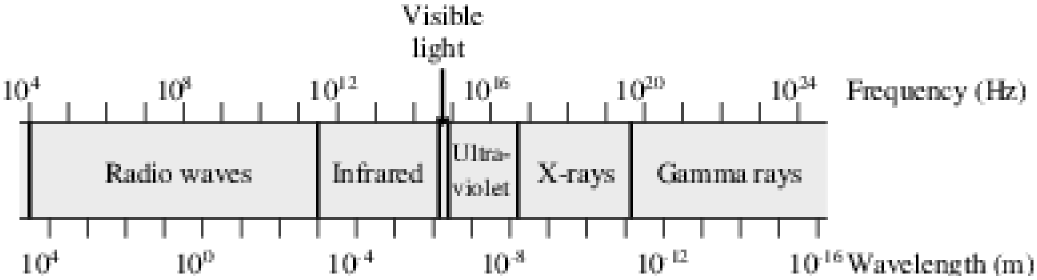

Figure 6.3: The portion of the electromagnetic spectrum from 10^4 Hz to 10^{24} Hz (wavelength range from 10^{-16} m to 10^4 m). Higher frequencies (and smaller wavelengths) are on the right.

through our eyes and so we call this range **visible light**. The higher frequencies are interpreted as blue and violet colors and the lower frequencies are interpreted as red and orange colors.

What frequency is white light?

White light is really the combination of all of the colors.[vii]

Check Point 6.10: Which color is associated with a higher frequency of electromagnetic waves: blue or red?

Each frequency has a particular **wavelength**.

To visualize what is meant by wavelength, imagine what we'd see if we took a snapshot of a water wave. At an instant in time, we'd observe that the surface of the water is high at certain points and low at other points. The distance from one peak to another is called the wavelength, λ. The wavelength is related to the frequency and speed as follows:

$$v = f\lambda \tag{6.1}$$

This relationship is known as the **wave equation** and holds for all waves. As such, it holds for all electromagnetic waves, regardless of the frequency.

[vii] Our eyes consist of three different types of cones, each able to sense a different frequency range. The frequency ranges roughly correspond to the red, green and blue regions. When all three types of cones are excited, we see white. Consequently, we can "trick" our eyes into see white by shining red, green and blue light together. Indeed, your computer or television screen takes advantage of this and only uses red, blue and green light.

Given the speed of the electromagnetic wave, each part of the electromagnetic spectrum can be identified by its wavelength instead of its frequency.

Check Point 6.11: As the frequency of a wave increases, what happens to its wavelength?

Is it better to use frequency or wavelength?

In general, it is better to use the frequency. The frequency of the wave remains the same as it travels from one medium to another. Since the speed changes and the frequency stays the same, the wavelength must change. Thus, any reference to the wavelength depends on the medium.

> ✍ Since we are limiting ourselves to observations in air (or in the vacuum of space) and since the speed in air is very close to the speed in a vacuum, the wavelength values we'll use will be the wavelength values in a vacuum.

On the other hand, many people are more familiar with the wavelength values than the frequency values. In addition, how the radiation interacts with objects depends on the wavelength of the radiation relative to the size of the object.

If the wavelength is much, much larger than the size of the object, the object will have no impact on the radiation. The object is essentially invisible. For example, FM radio waves have wavelengths about 3 meters in length. Those waves travel easily through the air, since the molecules that make up the air are so much smaller than 3 meters.

On the other hand, if the wavelength is smaller than the size of the object, the object will reflect the radiation. For example, visible light has a wavelength around half a micron (i.e., about 500 nm). This is much smaller than the size of, say, us. Consequently, visible light reflects off us.

At wavelengths in between, we get a phenomenon called **scattering**, which means that the radiation reflects off the object in all directions, both forwards and backwards.

Weather radars use radiation of wavelength around 10.7 cm. This is about ten to 100 times bigger than the size of raindrops (0.1 to 0.5 cm in diameter) and about one thousand to one hundred thousand times bigger than the size of cloud drops (0.0001 to 0.01 cm in diameter). Consequently, the radiation scatters off of raindrops but is unaffected by cloud drops.

Check Point 6.12: What is the purpose of using a wavelength of radiation that is so much larger than cloud drops?

6.4.4 Radar

A weather radar emits a short pulse of radiation, which travels away from the radar in a straight line until it encounters something that is big enough to either **scatter** the radiation or **reflect** it. As discussed above, the weather radars use a particular wavelength such that it scatters off raindrops and other types of precipitation and not the cloud drops.

The **reflectivity** refers to how much of the radiation is scattered back to the detector (colocated with the radar). This not only depends on the size of the particles but also on the concentration. Thus, more intense precipitation will result in a higher reflectivity.

Based on how long it takes for the scattered signal to return to the radar, one can determine how far away the precipitation is. Since the speed of light is equal to the total distance travelled by the radiation divided by the time, the precipitation must be at a distance equal to the speed of light multiplied by half the elapsed time (since the radiation must travel to the precipitation and back). The radar map is built up from this information.

> ✍ The radar map shown on page 6.2 is created by combining the signals of many radars together. The radars, part of the National Weather Service's Nexrad system, are called WSR-88D's (Weather Surveillance Radar). The "D" stands for Doppler, which means the radars can also measure the velocity of the precipitation toward or away from the radar. While the radars are able to measure the velocity, the velocity information is hardly ever shown on TV. Rather, only the reflectivity information is shown.

Check Point 6.13: On the radar map shown in figure 6.2, the reflectivity is *given in units of DBZ. What is the reflectivity of the most intense precipitation measured on the map?*

6.5 Observing clouds

There is no physics involved in identifying the *type* of cloud that is present, so I won't comment on that. Instead, I want to focus on how we determine the fraction of the sky that is covered by clouds and the height of the cloud tops (which is related to the severity of the storm). These measurements do involve some physics.

6.5.1 Sky condition (cloud fraction)

One can estimate the fraction of the sky that is covered by clouds simply by looking at the sky. However, an automated system like ASOS can't do that. Instead, it looks upward and determines the fraction of the time that a cloud is directly overhead.

In other words, instead of doing a space average (i.e., fraction of space that is occupied by a cloud at a given time), ASOS does a time average (i.e., fraction of time that a cloud is seen at a given position).[viii] It then assigns one of five descriptors based on the fraction of time that a cloud was seen: CLR (if $\leq 5\%$), FEW (if $\leq 25\%$), SCT (if $\leq 50\%$), BKN (if $\leq 87\%$) or OVC (if $> 87\%$).

Still, how does the ASOS instrument know if there is a cloud overhead?

It turns out it uses a process much like the WSR-88D radar – it sends out a pulse of radiation and waits for a return signal. However, it must use a smaller wavelength (0.9 microns vs. 10.7 cm), since the cloud droplets are so much smaller than raindrops. That wavelength is within the infrared, this is called a LIDAR system (from Light and radar) instead of a RADAR system. In fact, the system essentially consists of a laser, which produces a short pulse of light.

[viii] More information is available at http://www.nws.noaa.gov/asos/pdfs/aum-toc.pdf.

If a return pulse isn't sensed within a certain amount of time, the instrument concludes that there are no clouds.

Of course, this only means that there are no clouds within that distance. The "cutoff" distance happens to be 12,000 ft. Consequently, CLR simply means there are no clouds below 12,000 ft.

> ✍ ASOS actually reports on each level of clouds that it can observe. If it can't see beyond 2000 ft then it reports that the sky is obscured and the vertical visibility (VV) in hundreds of feet is reported.

Check Point 6.14: Why does ASOS use a much smaller wavelength to determine the cloud cover than the WSR-88D radar uses to determine the precipitation?

6.5.2 Cloud top heights

Stronger storms tend to have clouds that reach higher into the troposphere. Thus, it is useful to have a system that can identify how high the clouds go.

Whereas ground-based observers can determine the base of the clouds, cloud tops are determined by satellite systems. For example, figure 6.4 displays a picture taken by a satellite high above the Earth's surface. It shows the height of the cloud tops, with brighter areas indicating higher cloud tops and gray areas indicating lower cloud tops. This particular picture is enhanced, which means that it uses color to enhance the difference between the highest and lowest cloud tops (as opposed to just using a gray scale).

Before I explain how the satellite gets this picture, first notice that the time of the picture is 0215 UTC. This means it is night in East Stroudsburg, PA, local time being 9:15 PM. So, the satellite isn't using visible light from the sun that has scattered off the cloud tops.

The satellite also isn't sending out a pulse, like the LIDAR described in the previous section, to sense where the cloud tops are.

Instead, the satellite is simply observing the radiation that is being emitted by the cloud tops.

Clouds emit radiation?

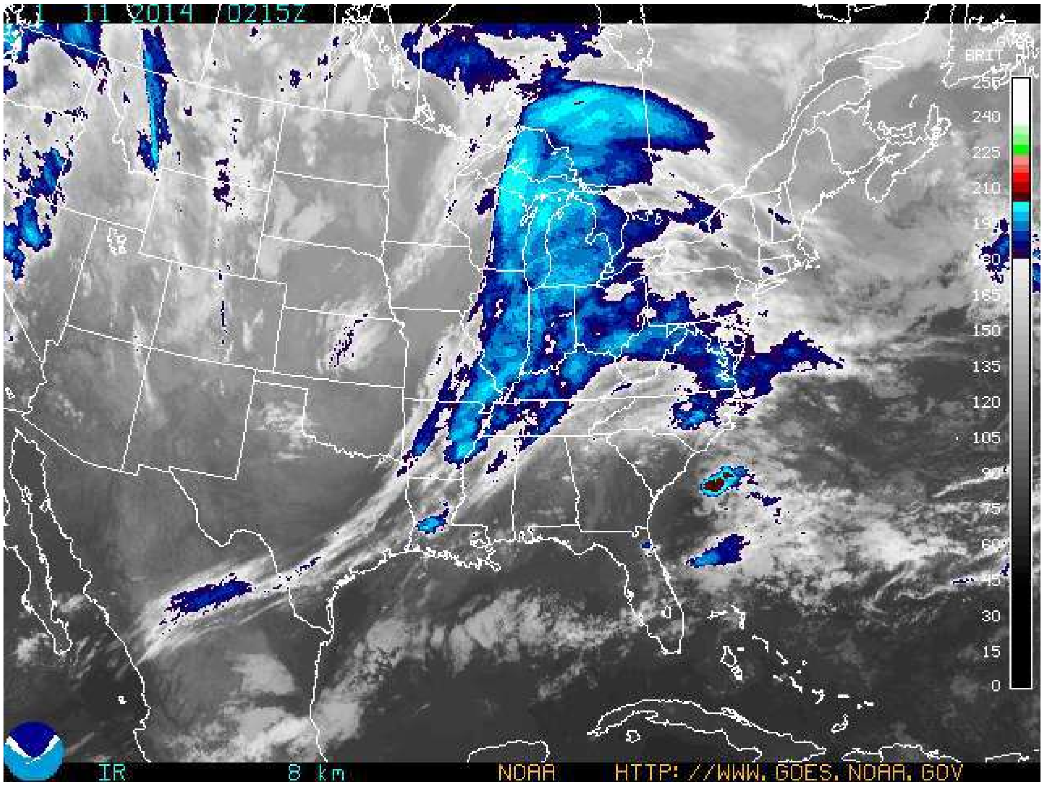

Figure 6.4: Infrared satellite picture for the Eastern Conus Sector on January 11, 2014, at 0215 UTC.

Yes. All objects emit radiation. We aren't aware of clouds emitting radiation because the radiation is at a frequency we can't see. In other words, clouds, like most objects, don't emit *visible* light but they still emit radiation. Remember that visible light is just one type of radiation.

To understand how the satellite image is obtained, there are three aspects of radiation that we need to recognize:[ix]

1. The warmer the object, the more radiation overall that is emitted by the object.
2. The radiation that is emitted is spread among a range of frequencies.

[ix] We'll examine these relationships in more detail in chapters 13 and 14.

3. The warmer the object, the more radiation that is emitted at higher frequencies (i.e., shorter wavelengths).

The first aspect means that a very hot object, like the sun (whose temperature is about 5800 K), emits a lot more radiation than a cooler object, like a cloud. The second and third aspects mean that both the sun and a cloud emit radiation over a range of frequencies but the sun emits most of its radiation at a higher frequency than a cloud does. In particular, the sun emits much of its radiation in the visible range, whereas the cloud does not emit any radiation in the visible range. The cloud still emits radiation – it is just at a lower frequency (longer wavelength) and our eyes can't detect those frequencies.

> ✍ Cool objects like clouds, chairs and people, do not emit radiation at a frequency perceived by our eyes. Consequently, we cannot see them unless we shine visible light on them, in which case the objects reflect the light back to us, allowing us to see them.

So, what the satellite does is detect the radiation emitted by clouds. This radiation corresponds to infrared radiation, with wavelengths longer than visible light.

The warmer the cloud, the more radiation that is emitted (see aspect #1 above). Those areas correspond to darker areas on the satellite picture. Conversely, the colder the cloud, the less radiation that is emitted. Those areas correspond to lighter areas on the satellite picture.

So the satellite detects the cloud temperature, not the cloud height?

Technically, yes. However, we know that the temperature of the atmosphere decreases with height (within the troposphere, which is where the clouds are). So, a colder cloud is most likely at a higher altitude.

> ✍ Because the satellite detects the temperature, cold ground temperatures (like those nearer the poles or at night) will appear brighter and thus can be confused with clouds.

Check Point 6.15: What is the difference between a low cloud and a high cloud that allows an infrared satellite to distinguish between them?

Project

Project 6.1: In project 5.1, you downloaded the saturation vapor pressures as a function of temperature. One of the columns displays the saturation vapor pressure relative to ice (rather than water).
(a) Add a column showing the difference *between the two saturation vapor pressures for the temperature range they have in common (for the temperatures for which* <u>*both*</u>[x] *are available).*
(b) Create a graph showing the difference between the saturation vapor pressure over water and that over ice.
(c) Are there any temperatures where the saturation vapor pressure over water is greater than the saturation vapor pressure over ice? If so, at what temperature is the difference the greatest?
(d) At the temperature indicated in (c), is it possible to have a vapor pressure that is below the saturation vapor pressure for water but above the saturation vapor pressure for ice? If so, give an example of a vapor pressure value for which that is true.
(e) At the temperature indicated in (c) and the vapor pressure indicated in (d), would a water droplet grow (due to condensation) or diminish (due to evaporation)?
(f) At the temperature indicated in (c) and the vapor pressure indicated in (d), would an ice crystal grow (due to deposition) or diminish (due to sublimation)?

Problems

Problem 6.1: Why are CCN necessary to produce clouds?

Problem 6.2: How many typical cloud droplets are required to make a typical rain droplet? Note: compare volumes, not diameters.

Problem 6.3: Estimate a height of a typical cloud. How long does it take for (a) a typical cloud drop and (b) a typical rain drop to fall out of the atmosphere assuming no vertical motion of the atmosphere?

Problem 6.4: Cirrus clouds are very high in the troposphere. What particles

---

[x] Between −15°C and 0°C.

are cirrus clouds made of and how did these particles form? Consider the environment of cirrus clouds for your answer.

Problem 6.5: Describe the process whereupon rain is produced in (a) a warm cloud (> $0^\circ C$ throughout) and (b) a cold cloud (< $0^\circ C$ in the upper-regions where the precipitation forms).

Problem 6.6: Describe the process whereupon sleet, snow, and hail are produced.

Problem 6.7: Knowing that the speed (of light) is 0.03% smaller in air than in a vacuum, calculate the speed of light in air.

Problem 6.8: Suppose a water wave travels at a speed of 10 m/s. A toy boat on the water is seen to oscillate up and down with a period of 2 s. What is the wavelength of the water wave?

Problem 6.9: A typical Helium-Neon laser produces light with a wavelength of about 632 nm. What color is this light?

Problem 6.10: Are the clouds shown on the radar map shown in figure 6.2? If so, where? If not, why not?

Problem 6.11: The infrared satellite picture in figure 6.4 (on page 88) has the Gulf of St. Lawrence being brighter than the Gulf of Mexico. Why is that, given that both are at sea-level?

Derivations

Show-me 6.1: Show that the frequency range from 10^{12} Hz to about 10^{16} Hz can be written as a wavelength range from 30 nm to 300,000 nm (or 0.03 μm to 300 μm).

Part B

Ideal Gas Law

7. Density

7.1 Introduction to part B

The physics of the atmosphere can be summarized by three relationships: the ideal gas law, Newton's second law of motion, and the first law of thermodynamics.

The ideal gas law is the focus of this part. The **ideal gas law** describes how the temperature, pressure and density of the air are related. Knowing the values of two of these variables, one can use the ideal gas law to determine what the other one must be.

The ideal gas law is perfectly suited as the first relationship to examine because it involves four of the five main ways we describe the state of the atmosphere: temperature, pressure, density and composition. The fifth way is the wind (see chapter 2).

Since we have already examined temperature and pressure, we only have two more variables to look at. In this chapter, we'll examine density and in chapter 8 we'll look at composition.

> ✍ The reason it is called the ideal gas law, and not just "the" gas law, is because technically it only works for "ideal" gases. In the real world, there are no ideal gases. Fortunately for us, however, air is pretty close to an ideal gas.

7.2 Definition

Density is the term typically used to describe the ratio of mass to volume:

$$\rho = \frac{m}{V} \tag{7.1}$$

where the lower-case Greek letter ρ (rho) is used to represent the **density** (m and V are used to represent mass and volume, respectively).

Density, like pressure and temperature, can be measured at a single point because, like the others, it is a *ratio*. Although the mass and volume of a point are both infinitesimal, the ratio of the two is not.

Check Point 7.1: The density is a ratio of mass to volume. When dealing with a point of infinitesimal volume, the density is still finite. The same is true of pressure and temperature. Name one other common ratio that has a finite value even though the numerator and denominator are both infinitesimal.

7.3 Typical values

A typical density of the air (near sea-level) is about 1.2-1.3 kg/m^3. What this means is that a cubic meter of air typically contains 1.2 kg or so of air molecules. A cubic meter is the volume of a cube one meter on each side.

> ✍ Liquid water is much more dense that air. The density of water is about 1000 kg/m^3.[i]

The atmosphere does not have the same density throughout. In general (but not always), the density is lower the higher one goes.[ii]

Just as the density of the air is not the same throughout the atmosphere, neither does it remain the same with time. In other words, the density is neither constant in space nor constant in time.

> ✍ In comparison, the density of the ocean is practically the same throughout in comparison to the atmosphere. The density of the water in the ocean varies mostly with salinity rather than height.

Check Point 7.2: (a) What is the density of air near sea level?
(b) What is the density of water?
(c) Which is larger?

---

[i] Be careful. The density of water is not 1 kg/m^3. Rather, it is 1 g/cm^3. To compare densities, the units must be the same.

[ii] Exceptions to this generality tend to lead to unusual optical effects called mirages (as do situations where the density changes are very large; see problem 15.6).

7.4 Ideal gas law

As mentioned before, all three properties (density, temperature and pressure) can be measured for a given point in the atmosphere. By “point” I mean an arbitrarily small region of air (see section 7.2).

To measure **density**, for example, we could take a small volume and measure the mass of air enclosed by the volume. The ratio would then represent the average density of the air in that volume.

In practice, however, we rarely measure density in this way. Rather, we first measure the temperature and pressure at that point and then determine the density based upon the temperature and pressure.

To do this, we use the **ideal gas law**.

Check Point 7.3: Which of the three properties (density, temperature and pressure) are typically measured and which is usually calculated via the ideal gas law?

The ideal gas law relates the four properties of a gas: temperature, pressure, density and composition. We’ll examine composition in the next chapter.

The ideal gas law can be written as follows:

$$P = \rho R T \tag{7.2}$$

where P, ρ and T are the pressure, density and temperature, respectively.

✍| The ideal gas law is also known as the **equation of state**.

What is R?

That is the **gas constant**. Different gases have different values of R because the value depends on the composition of the gas. For dry air (i.e., air without any water vapor in it), R has a value of 287.06 $\mathrm{J \cdot kg^{-1} K^{-1}}$. More will be said on this in the next chapter.

Why the strange units for R?

The units of R are such that the units on the left of the ideal gas law equal the units on the right. Notice, for example, that R is multiplied by the temperature. Temperature is measured in kelvin. Since kelvin is not part of the units for pressure, R must have units of inverse kelvin (K^{-1}) in order to cancel out the units of kelvin (K) in the temperature.

Do we have to use kelvin for temperature?

Yes.

According to the ideal gas law (equation 7.2), the pressure equals zero when the temperature equals zero. Since air pressure is certainly not zero at a temperature of 0°C (32°F), we obviously cannot use °C in equation 7.2. A temperature scale must be used that equals zero only when the pressure exerted by the gas goes to zero.[iii] The kelvin scale is designed such that the temperature, in kelvin, is zero when the pressure is zero.

It turns out that the point of zero pressure is at about −273°C. Thus, zero K is equal to −273°C. Or, conversely, 0°C is equal to 273K.

> ✍ Just as most of our constants will use units such as meters, kilograms and seconds, so will most of our constants use kelvin. Consequently, whenever temperature is used (such as in equation 7.2), we should use kelvin. However, since the temperature *increment* in °C is the same *increment* in kelvin, we can use either °C or kelvin when a problem involves the *change* in temperature.

Check Point 7.4: Suppose we are using the ideal gas law to determine the pressure of a gas when its temperature is 10° C. What value do we use for T *in equation 7.2?*

The ideal gas law tells us how the temperature, pressure and density are related. So, knowing two of them, one can determine what the remaining one must be (assuming we know the composition of the air). Similarly, if one property changes, at least one of the others must change as well.

For example, for a particular density, if one increases the temperature, the pressure will increase (as the molecules hit the walls with a greater speed). Likewise, for a particular temperature, if one increases the density, the pressure will also increase (as more molecules are hitting the wall).

[iii] It is at that point that the average molecular motion is the lowest it can possibly be (technically, in quantum physics terms, this is the lowest energy state).

Check Point 7.5: According to the ideal gas law, is it possible for the pressure to increase when the temperature decreases? If so, how? If not, why not?

7.5 Properties of an ideal gas

WHY IS THE RELATIONSHIP CALLED THE IDEAL GAS LAW?

The reason it is called the ideal gas law, and not just "the" gas law is because technically it only works for ideal gases. In the real world, there are no ideal gases. Fortunately for us, however, air is pretty close.

An ideal gas is one that has the following properties:

- The average kinetic energy of the molecules is proportional to the absolute temperature of the sample.
- The molecules have no volume themselves.
- There are no intermolecular attractions.
- The collisions between molecules are completely elastic.

These postulates may seem very unrealistic but it turns out that if a gas is composed of a huge number of molecules (but not so large that they are squeezed too close together) that are in constant, random motion, most of these postulates hold true.

In retrospect, one can see why. The first postulate means that the average molecule has a speed that is much greater than the wind speed. In other words, the molecules are moving quickly but not all in the same direction (i.e., they keep bumping into each other). This is the case for air.

The second postulate means that the molecules are so small that their size takes up very little of the volume. In other words, if we could stop the molecules from moving around and compress them together, they'd form just a minute fraction of the total volume. Again, this is reasonable for air.

The third postulate means that as the molecules are bouncing around, they don't really feel the effects of the other molecules until they happen to bump into one. Again, considering the size of each individual molecule and assuming the molecules are neutral (i.e., have no excess positive or negative charge), this seems like a reasonable approximation.

The fourth postulate means that the average kinetic energy of the molecules does not increase or decrease as the molecules are allowed to collide with each other. If the average kinetic energy was impacted by collisions, the temperature (which is a measure of the average kinetic energy) would start increasing or decreasing without doing anything else. Again, this seems to be a reasonable assumption for air.

Such postulates define what is called an **ideal gas**. Even though the atmosphere is not made up of ideal gases, the gases are very close to ideal gases. Consequently, we will make this assumption for the rest of the text.

How close to an ideal gas is the air?

To illustrate how close it is, consider the following: for an ideal gas, Avogadro's number ($6.02214199 \times 10^{23}$) represents the number of molecules that would exist in a volume of 22.4 liters at standard temperature and pressure (0°C and 1 atm). This is also true for air, within about 0.01%.

There are two remarkable things about this fact. The first is that the volume can be predicted to be 22.4 liters. It turns out that the ideal gas law can be used for this. However, one needs a version that is written in terms of volume. We'll do this in the next chapter.

The second is that molecules must be incredibly tiny in order to get so many in such a space (22.4 liters is less than one cubic foot)!

Check Point 7.6: Can we consider air to be an ideal gas?

7.6 Gas constant

As mentioned earlier, the gas constant depends on the gases that are present. For air, the only gas that varies a lot, depending on where you are, is water vapor. Thus, we typically focus on either the gas constant for dry air[iv] (i.e.,

---

[iv] See appendix B.2.

the air without the water vapor)

$$R_{\mathrm{d}} = 287.06\ \mathrm{J} \cdot \mathrm{kg}^{-1}\mathrm{K}^{-1}$$

or the gas constant for water vapor[v] alone

$$R_{\mathrm{v}} = 461.52\ \mathrm{J} \cdot \mathrm{kg}^{-1}\mathrm{K}^{-1}.$$

Then, if we want to use the ideal gas law for the actual air content (i.e., dry air plus water vapor), we'll use some combination of the two.

Fortunately, water vapor makes up a small part of the air (typically less than 4%; see chapter 5). Consequently, the gas constant for air will be close to the gas constant for just the dry portion of the air (since the relative amounts of the other gases are pretty much the same from day to day).

✍ If the air has some water vapor content, then we will refer to the air as **moist air**. This does not mean that there is necessarily some liquid water present. It only means that water exists in the gaseous state.

Check Point 7.7: Will the gas constant for moist air be closer to the gas constant for dry air or closer to the gas constant for water vapor?

Although the actual gas constant for air will be close to the value for dry air alone, if we want to be accurate (and, of course, in many cases we do), we'll need a method of handling the variation in water vapor content. In other words, we need a way of first determining the value of R in order to use the ideal gas for the air:

$$P = \rho R T \ [\text{ total }]$$

where I've added the descriptor "total" to indicate that the variables correspond to the total air content, not just the dry air or the water vapor.

The value of R will depend on how much water vapor there is (i.e., the vapor pressure e). What we need is to derive an expression that gives us the value of R as a function of the vapor pressure e.

To do this, we apply the ideal gas law to the dry air and the water vapor, separately. We can do this because each portion of the air also acts like an

---

[v] See appendix B.3.

ideal gas.

$$\begin{aligned} P_{\mathrm{d}} &= \rho_{\mathrm{d}} R_{\mathrm{d}} T \text{ [dry]} \\ e &= \rho_{\mathrm{v}} R_{\mathrm{v}} T \text{ [water vapor]} \end{aligned}$$

We'll assume we can measure the total pressure P, the vapor pressure e and the temperature T (all three temperatures are the same). We also know the values of R_{d} and R_{v}. We are thus left with five unknowns: P_{d}, R and the three densities. These are too many unknowns (to solve for n unknowns, you need at least n independent equations). We need two more relationships, involving the same variables.

The first relationship is known as **Dalton's law of partial pressures**, which states that the total pressure exerted by a mixture of gases is equal to the sum of the pressures exerted by each gas individually (assuming the gases do not interact chemically). When applied to the dry air and water vapor, this means that the total air pressure is equal to the pressure of the dry air plus the vapor pressure:

$$P = P_{\mathrm{d}} + e.$$

The second relationship is between the densities. Since all of the gases occupy the same volume, the total density ρ must equal the sum of the two individual densities:

$$\rho = \rho_{\mathrm{d}} + \rho_{\mathrm{v}}.$$

We now have five equations and we can solve for any of the five unknowns in terms of the known variables.

For example, suppose we want to know what ρ_{v} is, using only the variables that are measurable or known (P, T, e, R_{d} and R_{v}). We can use the ideal gas law for water vapor for that. Just rearrange as follows:

$$\rho_{\mathrm{v}} = e/(R_{\mathrm{v}} T)$$

It is likewise very easy to determine what P_{d} is, using only the variables that are measurable or known (P, T, e, R_{d} and R_{v}). For that, use the law of partial pressures and just rearrange as follows:

$$P_{\mathrm{d}} = P - e$$

To determine the other three variables (ρ_{d}, ρ and R), we need to combine two or more relationships. To solve for ρ_{d}, we need to use the ideal gas law for dry air and the law of partial pressures:

$$\begin{aligned}\rho_{\mathrm{d}} &= P_{\mathrm{d}}/(R_{\mathrm{d}}T) \\ &= (P-e)/(R_{\mathrm{d}}T)\end{aligned}$$

To solve for ρ, we need to use the relationship between the densities, the law of partial pressures and multiple versions of the ideal gas law:

$$\begin{aligned}\rho &= \rho_{\mathrm{d}} + \rho_{\mathrm{v}} \\ &= (P-e)/(R_{\mathrm{d}}T) + e/(R_{\mathrm{v}}T) \\ &= \Big((P-e)/R_{\mathrm{d}} + e/R_{\mathrm{v}}\Big)/T \\ &= \Big(P - e + e(R_{\mathrm{d}}/R_{\mathrm{v}})\Big)/(R_{\mathrm{d}}T)\end{aligned}$$

To solve for R, we need the ideal gas law for moist air (total air), and the relationship for ρ obtained above:

$$\begin{aligned}R &= P/(\rho T) \\ &= R_{\mathrm{d}}P/\Big(P - e + e(R_{\mathrm{d}}/R_{\mathrm{v}})\Big)\end{aligned}$$

By the way, the ratio $R_{\mathrm{d}}/R_{\mathrm{v}}$, which is in two of the relationships, will turn up a lot. Since the values of R_{d} and R_{v} are known, it is common to replace the ratio with a single variable indicated by ϵ:

$$\epsilon = \frac{R_{\mathrm{d}}}{R_{\mathrm{v}}} \approx 0.622 \tag{7.3}$$

Both R_{d} and R_{v} are known. The ratio happens to be about 0.622.

If you want, you can replace every occurrence of $R_{\mathrm{d}}/R_{\mathrm{v}}$ with ϵ. Thus, the equation for the density ρ can be written as

$$\rho = \frac{P - e + \epsilon e}{R_{\mathrm{d}}T} \tag{7.4}$$

and the equation for the gas constant R can be written as

$$R = R_{\mathrm{d}}\frac{P}{P - e + \epsilon e} \tag{7.5}$$

> ✍ You can always solve for R by first solving for ρ and then using the ideal gas law (or, visa-versa, you can solve for ρ by first solving for R and then using the ideal gas law). The equations above simply give a way of getting the unknown in one step.

Check Point 7.8: Of all of the variable discussed in this section (P, e, P_d, ρ, ρ_d, ρ_v, R, R_d, R_v and T), which ones are actually measured?

Problems

Problem 7.1: (a) From problems 3.1 (see page 41) and 4.2 (see page 53), and assuming they correspond to the same approximate height, calculate the density that would be observed at East Stroudsburg using the ideal gas law (equation 7.2) and the value of R for dry air (see appendix B.2). Compare your value of density to the range mentioned in section 7.3 and comment on any differences.
(b) From problems 4.2 and 3.2, and assuming they correspond to the same approximate height, calculate the range of densities that would be observed near the surface. Remember that the largest density will correspond to the location that has a large pressure and small temperature. Compare your range of density to that mentioned in section 7.3.

Problem 7.2: (a) Using the fact that the density of water is 1 g/cm^3, calculate the density of water in units of kg/m^3.
(b) How much mass would be in one cubic meter of water?
(c) How much mass would be in one liter of water? Note: 1 ml = 1 cm^3.
(d) Estimate the volume of a single drop of water. What is the mass of a single drop of water?
(e) Estimate the surface area and average depth of the Earth's oceans. From those estimates, obtain an estimated volume and, from that, determine an estimated mass of the Earth's oceans, assuming the oceans were pure water.

Problem 7.3: Using the fact that there is 1 mole of atoms in 22.4 liters[vi]*, calculate the average distance between each atom.*[vii] *Look up the average*

[vi] A liter is equal to 1000 ml and a milliliter is equal to 1 cm^3.

[vii] First find the space taken up by each atom and then assume each space is a cube.

diameter of an atom.[viii] *How many atomic diameters are there between each atom?*

Problem 7.4: In section 6.1, it was stated that if the mixing ratio decreases by 1 g/kg (due to condensation) and the moisture is distributed equally among CCN with number density 10^6 *per liter, then the radius of the resulting drops is equal to about 7 microns. The relevant equation for the radius* R *of the resulting drops is as follows:*

$$R = \left[\frac{3\Delta r \rho_{\text{air}}}{4\pi n_{\text{CCN}} \rho_{\text{water}}}\right]^{1/3}$$

where Δr *is the decrease in mixing ratio,* n_{CCN} *is the CCN number density (per volume) and* ρ_{air} *and* ρ_{water} *are the densities of air and water, respectively. Verify that if the mixing ratio decreases by 1 g/kg and the moisture is distributed equally among CCN with number density* 10^6 *per liter, then the radius of the resulting drops is equal to about 7 microns. Make sure your units work out.*

Problem 7.5: Suppose the air temperature is 10° C and the relative humidity is 100%. From a skew-T diagram, one finds that the mixing ratio at 1000 mb is about 8 g/kg.
(a) If somehow we could condense all of the moisture out, so that 8 g/kg goes completely into water droplets, how big (diameter or radius) would each droplet be? Use the relationship provided in Problem 7.4.
(b) How does your size compare to the size of a typical raindrop?

Problem 7.6: As one can show (see Show-me 7.2), the mixing ratio r *and vapor pressure* e *are related as follows:*

$$r = \epsilon \frac{e}{P - e}$$

where P *is the total pressure of the air. Suppose a given parcel of air has a total mass of 1 kg and, of that, the water vapor has a mass of 3 g. If the air pressure is 1000 mb, what is the vapor pressure?*

Problem 7.7: Suppose the vapor pressure was 6 mb on a day when the total air pressure was 1013 mb and the air temperature was 15° C.

Knowing the volume of the cube, you can determine the length of each side of the cube, which must represent the distance from one atom to the next.

[viii] You should find it is around one angstrom, which is 10^{-10} m or 10^{-8} cm.

(a) Use the ideal gas law to calculate the density of the water vapor.[ix]
(b) Use the ideal gas law to calculate the density of the dry air.[x]
(c) Based upon your answers to parts (a) and (b), should the gas constant for the total air be greater, or smaller, than the gas constant for dry air alone? Should it differ by a lot, or a little? Explain.
(d) Calculate the value of the gas constant for the total air and compare your answer to your prediction in (c).
(e) What is the air density? (remember to use the correct units)

Derivations

Show-me 7.1: Derive the relationship in Problem 7.4. Hint: the volume of each droplet is $V = (4/3)\pi R^3$*. So, to get the total* volume *of water that condenses per volume of air, multiply this by the number density of CCN,* n_{CCN}*. One can then use the densities of water and air to get the total mass of water per mass of air (which should equal to the drop in mixing ratio).*

Show-me 7.2: Show that the relationship between mixing ratio r *and vapor pressure* e *is*

$$r = \epsilon \frac{e}{P - e}$$

where ϵ *is about 0.622. Hint: use the ideal gas law twice, once for the dry air and once for the water vapor, then take the ratio of the two.*

[ix] Since you are applying the gas law to water vapor only, remember that the value of R that you use must be appropriate for water vapor, not dry air (see appendix B.3)].

[x] See previous footnote. In this case, you are applying the gas law to the dry air only.

8. What is in the Air?

8.1 Introduction

As mentioned in chapter 7, the ideal gas law can be expressed as follows:

$$P = \rho RT$$

and relates the four atmospheric variables of temperature, pressure, density and composition. We've already discussed the first three. In this chapter, we examine composition.

8.2 The gas constant

The composition is reflected in the variable R, which is called the **gas constant**. For dry air (i.e., air without any water vapor in it), R has a value of $287.06\ \mathrm{J \cdot kg^{-1}K^{-1}}$.

In this section, I'll try to give you a mental picture of what R represents. I'll do this by deriving the ideal gas law as written above

$$P = \rho RT$$

from the version you probably learned in chemistry class:

$$PV = nR^{*}T$$

where R^{*} is not the same thing as the gas constant R (more on this later).

The reason for the two different versions is because in chemistry we typically deal with a confined volume of gas that has a particular volume and temperature (thus a version that has the volume V) whereas the atmosphere is not a confined volume of gas with a particular pressure and temperature (thus we use a version that has the density ρ).

It is straightforward to derive one from the other. All you have to do is rearrange the variables and introduce some new variables. I am going to go through the derivation now.

Why?

I am doing the derivation because doing so provides some meaning to the gas constant R. Consequently, do not just skip over the derivation, thinking that only the final relationships are important. There is physics within the derivation that will provide insight into the meaning of the different variables.[i]

First, we need to recognize what n represents. That is the total number of atoms or molecules that are present within the volume V. To make the units work out, the number is given in terms of moles, where one **mole** represents 6.0221415×10^{23} atoms or molecules. The unit abbreviation is "mol".

> ✍ This number is called **Avogadro's number** ($N_\text{A} = 6.0221415 \times 10^{23}$ atoms per mole).

If we divide both sides by V, we get

$$P = \frac{n}{V} R^* T$$

The quantity n/V represents the number of atoms or molecules per volume. At standard pressure (1013.25 hPa) and temperature (273.15K), the quantity n/V is equal to one mole per 22.4 liters.

This quantity is like the mass density discussed in the previous chapter but in terms of number, not mass. To get it in the form of a mass density, we can multiply the numerator and denominator by m, the mass. By multiplying both top and bottom by the same thing, we aren't changing the value of the quantity. Doing so, we get:

$$\frac{n}{V}\frac{m}{m} = \frac{m}{V}\frac{n}{m} = \rho\frac{n}{m}$$

where I used the definition of density ($\rho = m/V$) to make the replacement to ρ (I also took advantage of the fact that the order of multiplication does not matter: $n \times m$ is the same thing as $m \times n$).

---

[i] It is also a fun academic exercise for your brain.

Making this replacement, we get the following:

$$P = \rho \frac{n}{m} R^* T$$

which looks very similar to equation 7.2 except that it has an extra term (n/m) and it uses R^* instead of R.

Let's interpret each.

First, let's look at n/m. This quantity is easier to interpret if we invert it. The inverse, m/n, is the total mass that is present divided by the number of atoms or molecules that are present. Since n is given in units of moles, the quantity m/n represents the average mass of a mole of atoms or molecules.

The average mass of a mole of particles is called the **molar mass** and is represented by the Greek letter "mu":[ii]

$$\mu = \frac{m}{n}$$

For example, the molar mass of oxygen is 15.9994 g/mol. If you look on a periodic table (see inside back cover), the number listed below the letter abbreviation for an atom is the molar mass of that atom.

> ✍ The molar mass is sometimes referred to as the **atomic weight**, which is a bit misleading since the quantity has units of mass, not weight.

In any event, making this replacement, we get the following:

$$P = \rho \frac{R^*}{\mu} T$$

which looks even more similar to equation 7.2 except that it has R^*/μ instead of R.

So now let's look at R^*. This quantity is known as the **universal gas constant**. It is called the "universal" gas constant because, unlike R, the *universal* gas constant has the same value regardless of what gas is present. In other words, the universal gas constant, R^*, does *not* depend on the make-up of the gas. Its value is given as follows:

$$R^* = 8.3145 \text{ J} \cdot \text{mol}^{-1}\text{K}^{-1}$$

[ii] There is no standard abbreviation for the molar mass. In fact, many people use M.

Notice that it has units of inverse moles and inverse kelvin. When multiplied by n and T, those units will cancel and leave units of joules. That is a unit of energy. It turns out (and we will explore this in chapter 11) that the product of pressure and volume has the same unit as energy.

Although R^* does not depend upon the gas, when we divide by the molar mass μ, we get a quantity that does. That quantity is known as the **gas constant** and is indicated by the letter R:

$$R = \frac{R^*}{\mu} \tag{8.1}$$

Thus, if we know the molar mass of a gas, we can determine the value of the gas constant for that gas.

Check Point 8.1: The constants R and R^* have different units. Is that consistent with the relationship between the two given in equation 8.1?

8.3 Nitrogen and Oxygen

Since we are told that the gas constant for dry air is 287.06 $\mathrm{J \cdot kg^{-1}K^{-1}}$ (see previous chapter) and we know the relationship between the gas constant and the molar mass (equation 8.1),

$$R = \frac{R^*}{\mu}$$

it is rather straightforward to solve for μ and plug in the values of R and R^* to determine the molar mass of dry air:

$$\mu = \frac{R^*}{R} = \frac{8.3145\ \mathrm{J \cdot mol^{-1}K^{-1}}}{287.06\ \mathrm{J \cdot kg^{-1}K^{-1}}}$$

which gives a molar mass of 0.028964 kg/mol.

What does this mean?

First of all, it may have more meaning if we write the result in units of grams per mole, rather than kilograms per mole. In grams per mole, the molar mass is 28.964 g/mol. This is very close (within rounding) to the official value of 28.96443 g/mol given by the CRC Handbook (1983).

Let's compare this to some elements, as listed in the periodic table (see inside back cover). The element with a molar mass closest to that of air is silicon, which has a molar mass of 28.0855 g/mol. That doesn't make sense, though. After all, the atmosphere isn't made up of silicon. Silicon isn't even a gas.

> ✍ The gases that make up the atmosphere are invisible to the naked eye. Each gas is made up of lots and lots of tiny molecules, which we could see only if we were to examine it on a very small scale.

The atmosphere is mostly nitrogen and oxygen. Yet, the molar masses of nitrogen and oxygen are 14.0067 g/mol and 15.9994 g/mol, respectively (see periodic table). Why is the molar mass of air so much more than the molar masses of nitrogen and oxygen?

The reason for the discrepancy is that the nitrogen and oxygen in the air are in the form of **diatomic** nitrogen molecules (N_2) and diatomic oxygen molecules (O_2). The word "diatomic" means that there are two atoms in the molecule (thus the subscript "2").

Each diatomic molecule, then, has a molar mass that is twice the individual atoms. That means that diatomic nitrogen has a molar mass of 28.0134 g/mol and diatomic oxygen has a molar mass of 31.9988 g/mol. As you can see, these values are closer to the molar mass of air (28.964 g/mol).

Check Point 8.2: What does it mean for a molecule to be diatomic?

8.4 Molar mass of dry air

Not only can we see that the molar mass of dry air is roughly the same as that of diatomic nitrogen and diatomic oxygen, but we can also be more specific.

Notice, for example, that the molar mass of air is *closer* to the molar mass of diatomic nitrogen than it is to the molar mass of diatomic oxygen. The reason for this is that dry air is mostly nitrogen. In fact, about 78% of the dry air is nitrogen while only 21% is oxygen.

With a little mathematics, we can show *how much* closer it should be to diatomic nitrogen.

Basically, we are doing an average. We want to find the average molar mass, given that 78% of the molecules have a molar mass of 28.0134 g/mol and 21% of the molecules have a molar mass of 31.9988 g/mol.

To find the average molar mass, we weight each molar mass by its relative abundance. Mathematically, this is done by first multiplying the molar masses of oxygen and nitrogen by their relative abundance and adding them together. Using μ_{d} to represent the molar mass of dry air:

$$\begin{aligned}\mu_{\mathrm{d}} &= (28.0134\ \mathrm{g/mol}) \times (0.78) + (31.9988\ \mathrm{g/mol}) \times (0.21) \\ &= 28.85045 + 6.71975\end{aligned}$$

which gives a value of 28.5702 g/mol.

This is very close to the actual value (28.96443 g/mol) but it is a little low. There are two reasons for this.

One reason is that we are only considering 99% of the molecules, not 100%. Thus, the 78% and 21% values are not the abundances relative to each other. Mathematically, it would be more appropriate to use 78/99 and 21/99. This would give a molar mass of 28.859 g/mol, which is still a little low but much closer to the actual value.

The other reason it is off is because of the missing 1% that we've ignored. It turns out that the remaining 1% is made up of gases that tend to have a higher molar mass than nitrogen and oxygen.

In particular, Argon makes up most of that missing 1% and the molar mass of Argon is 39.948 g/mol. If we include Argon in our calculation, we get:

$$\begin{aligned}\mu_{\mathrm{d}} &= (28.0134\ \mathrm{g/mol}) \times (0.78) + (31.9988\ \mathrm{g/mol}) \times (0.21) \\ &\quad + (39.948\ \mathrm{g/mol}) \times (0.01) \\ &= 28.85045 + 6.71975 + 0.39948\end{aligned}$$

which gives a value of 28.9697 g/mol.

As expected, this is even closer to the actual value. It is still a little off, though, because I've used rounded values for the relative abundances. There are also additional gases in the atmosphere that I have not included.

Table 8.1: Molar masses (from physics.nist.gov) and proportion (by number; from CRC 1983) of molecule type in dry air (at sea level). Proportions are given in units of percent (or parts per hundred), ‰ (parts per thousand) and ‱ (parts per ten thousand).

Molecule	**molar mass (g/mol)**	%	‰	‱
N_2	28.0134	78.084	780.84	7,808.4
O_2	31.9988	20.9476	209.476	2,094.76
Ar	39.948	0.934	9.34	93.4
CO_2	44.0095	0.0314	0.314	3.14
Ne	20.1797	0.001818	0.01818	0.1818
He	4.002602	0.000524	0.00524	0.0524
CH_4	16.04246	0.0002	0.002	0.02
Kr	83.798	0.000114	0.00114	0.0114
H_2	2.01588	0.00005	0.0005	0.005
Xe	131.293	0.0000087	0.000087	0.00087

The actual fractions of each gas in dry air are listed in table 8.1 along with the molar masses.

Check Point 8.3: If the average molecular mass of air was calculated using only the nitrogen and oxygen, with values given in table 8.1, one gets a value of 28.8564 u. That is less than the accepted value of 28.96443 u. Why?

8.5 Distribution

Are the ratios of the gases the same no matter where the air resides? Wouldn't the relatively heavy gases (like nitrogen and oxygen) tend to settle to the bottom of the atmosphere while light gases (like helium and hydrogen) rise to the top?

Certainly, relatively heavy gases (like nitrogen and oxygen) have a tendency to settle to the bottom of the atmosphere while light gases (like helium and hydrogen) tend to rise to the top. And, consequently, we might expect the ratio to change from an abundance of heavy elements near the bottom of the atmosphere to an abundance of light elements near the top.

However, the ratios of the gases given above are relatively the same no matter where the air resides. In other words, the air composition in New York City is essentially the same as that at the top of Mt. Everest. It is not until one gets very high in the atmosphere, about 100 km, that the ratio of the gases changes. One hundred kilometers may not seem very large when compared to the radius of the earth (6,370 km) but it is significant when compared to the typical height of the tallest thunderstorm cloud (about 12 km or 35,000 ft; about as high as commercial jet aircraft might fly).

To see why the ratio of the gases is essentially independent of height below 100 km, one needs to consider that there are two competing mechanisms at work. These mechanisms are called **diffusion** (from the Latin *dis* + *fringere*: to break apart) and **mixing**.

Diffusion describes the process by which particles (those of a particular molecular species in our case) move from a region of higher concentration to one of lower concentration. Diffusion is high when there is a lot of space between the particles because less gets in the way of the migration. Another way of saying this is that the **mean free path** is large. A large mean free path means that each particle tends to travel a relatively long distance before hitting another particle. Diffusion is also high when the temperature is high. The higher the temperature, the faster the particles are moving (see chapter 3.1) and so the more quickly they can migrate into the low concentration area.

Countering the diffusion mechanism is the mechanism of mixing. Just as mixing helps spread milk throughout your coffee, mixing in the atmosphere helps homogenize the air.

Above 100 km, diffusion dominates and atmospheric composition varies with height. Above 100 km, the atmosphere is called the **heterosphere**, since settling produces a make-up that varies with height. As one goes even higher, there is a greater presence of ions as well.

Below 100 km, mixing dominates. Indeed, we can usually neglect diffusion below 50 km or so because the time needed for air to diffuse in the lower atmosphere is usually much smaller than the time needed for the winds to move the air to a different location. Below 100 km, the atmosphere is called the **homosphere**, since atmospheric composition is homogeneous (i.e., ratios are the same everywhere).

Check Point 8.4: At the top of Mount Everest, which mechanism is more dominant: mixing or diffusion? Explain.

HOW CAN THE AVERAGE MOLAR MASS BE CONSTANT IN THE LOWER ATMOSPHERE WHEN THERE ARE OBVIOUS DAY-TO-DAY CHANGES IN THE AMOUNT OF POLLUTANTS AND MOISTURE?

Usually, the pollutants have concentrations in parts per million and so the numbers given above don't really change all that much. Even **ozone**, for which concentrations are closely monitored both on the ground (as a pollutant) and in the stratosphere (as an absorber of ultraviolet radiation), makes up a relatively small portion of the air.

In addition, as pollutants and moisture are added to the air, the main gases (nitrogen, oxygen and argon) may take up a smaller percentage of the total air but their ratios (about 78:21:1 for nitrogen, oxygen and argon) remain the same.

The only other gas that is likely to reach appreciable percentages of the total air is water vapor. Because of the importance of water vapor, I discuss it in a separate chapter (see section 5.2).

WHY ISN'T WATER VAPOR INCLUDED IN THE LIST OF CONSTITUENTS LISTED IN TABLE 8.1?

Because the amount of water vapor, unlike the others, is highly variable.

WOULDN'T THE PRESENCE OF WATER VAPOR CHANGE THE AVERAGE MOLECULAR MASS OF AIR?

Yes. That is why everything we've done to this point is accurate only for **dry air**. The subscript "d" is used to indicate that it is represents only the *dry* air.

Check Point 8.5: What does it mean for the air to be dry?

8.6 Evolution

Why does nitrogen make up the bulk of the atmosphere? Why are gases like Argon, Neon, Krypton and Xenon even present?

To answer this, we need to examine how the atmopshere came to be. Whether the earth initially had an atmosphere or not, it seems likely that volcanic emissions would greatly affect the composition of the atmosphere. Surprisingly, though, the molecular make-up of volcanic emissions looks nothing like the molecular make-up of the current atmosphere. Gaseous emissions from volcanoes is roughly 85% water vapor, 10% carbon dioxide and up to a few percent nitrogen and sulfur (or sulfur compounds)[iii].

Part of the problem is that not everything that gets put in the atmosphere stays there. Some gases are so light they escape from the atmopshere altogether. Other gases get return to the earth via a process called deposition.

8.6.1 Escape

To escape from the atmosphere altogether, a molecule has to be going fast enough such that the earth's gravity is not strong enough to keep it down. This speed is known as the **escape velocity** and can be determined according to the following equation.

$$v_{\mathrm{e}} = \left[\frac{2Gm_{\mathrm{earth}}}{r}\right]^{1/2} \tag{8.2}$$

Here v_{e} is the escape velocity, G is the gravitational force constant (6.67×10^{-11} N·m^2·kg^{-2}), m_{earth} is the mass of the earth (5.98×10^{24} kg) and r is the distance of the molecule from the center of the earth.

Where does this equation come from?

As an object moves away from earth, its velocity decreases (if no other force besides gravity is present) because the gravitational force is directed opposite to its motion. Another way to say this is that its kinetic energy decreases because the work done by gravity is negative ($W = \Delta E_k$).

[iii] Wallace, John M., and Peter V. Hobbs, 1977: Atmospheric Science, An Introductory Survey, Academic Press, 467 pp.

As the object continues to move away from earth, the kinetic energy decreases until it eventually reaches zero. At that point, since the gravitational force extends out to infinity, there is always a slight gravitational pull and the object will always be pulled back toward the earth. The only way it will not get pulled back is if there is no point where the kinetic energy goes to zero. This is possible only if the maximum amount of (negative) work that can be done by the earth's gravitational force is less than the object's initial kinetic energy.

Since kinetic energy is given as $\frac{1}{2}mv^2$ and work is given as $F_{\text{avg}}\Delta x$ (where Δx is the distance moved and F_{avg} is the average force exerted) we have the following relationship.

$$\Delta\left(\frac{1}{2}mv^2\right) = F_{\text{avg}}\Delta x \tag{8.3}$$

The force is the gravitational force

$$F_{\text{g}} = G\frac{m_{\text{earth}}m}{r^2} \tag{8.4}$$

and this varies as the object moves away from the earth.

If the gravitational force was constant, we could just plug the expression for the gravitational force into equation 8.3. However, the gravitational force on the object is not constant, in the sense that the force decreases as the object moves away from the earth.

To find the average force, one must use integration. As with finding the average molar mass, to find the average force we look at the value of the force at each location along the path (as the object moves away from the earth). At each location, the force is multiplied by the relative portion of the path length for which that force applies (i.e., $F(x) \times (dx/\Delta x)$ where dx indicates an infinitesimal piece of the path length). We add these all up to get the average. This is indicated as follows.

$$F_{\text{avg}} = \int \frac{F(x)dx}{\Delta x} \tag{8.5}$$

Since $F(x)$ is given by (8.4) above (with r instead of x) we actually have an expression that looks like it is the derivative of something:

$$F_{\text{avg}} \quad = \quad \int \frac{F(x)dx}{\Delta x}$$

$$\begin{aligned}
&= \int G\frac{m_{\text{earth}}m}{\Delta x} \times \left(x^{-2}dx\right) \\
&= \int G\frac{m_{\text{earth}}m}{\Delta x} \times d\left(-x^{-1}\right) \\
&= G\frac{m_{\text{earth}}m}{\Delta x} \times \int d\left(-x^{-1}\right) \\
&= G\frac{m_{\text{earth}}m}{\Delta x} \times \Delta(-x^{-1})
\end{aligned}$$

Along the way, I used the fact that the derivative of $(-x^{-1})$ is $x^{-2}dx$ (see third line) and moved $G\frac{m_{\text{earth}}m}{\Delta x}$ outside the intregral because every "piece" is multiplied by the same thing (i.e., $G\frac{m_{\text{earth}}m}{\Delta x}$ doesn't vary along the path). I also utilized the fact that summing up all the little pieces of $-x^{-1}$ just gives the total change in $-x^{-1}$.

Plugging everything back into (8.3) we get

$$\frac{1}{2}mv^2 = Gm_{\text{earth}}m \times \frac{1}{r} \tag{8.6}$$

Note that the Δx's cancel. Also note that $\Delta(-x^{-1})$ gets transformed into $(-\infty^{-1}) - (-r^{-1})$ and that $1/\infty$ is zero.

Check Point 8.6: (a) Why isn't the force of gravity just the mass times 9.8 N/kg?
(b) Why can't we assume the gravitational force is constant when determining the escape velocity?
(c) Why is $1/\infty$ *equal to zero?*

Since the atmosphere is very thin compared to the size of the earth, the escape velocity at the top of the atmosphere is not too much different from the escape velocity at the bottom of the atmosphere.

It is not the escape velocity that distinguishes the homosphere from the heterosphere but rather the density. Near the earth's surface, there are so many other molecules in the way that even if a molecule there was moving at the escape velocity, it probably wouldn't get very far. It is not until about 500 km or so that molecules are far enough apart that one with the escape velocity can actually escape.

How likely is it for a molecule to reach a speed equal to the escape velocity?

To really answer this, we need to know more about temperature and pressure than will be discussed here. However, it turns out that the probability of reaching such speeds depends on the temperature and the molecular mass of the molecule. To obtain such speeds, the molecule has to be pretty hot or pretty light (since the lower the mass, the higher its speed for a given temperature). As such, it really only happens in the upper regions of the atmosphere (where the temperature are on the order of 600 K) and for light gases such as hydrogen. For atomic oxygen, the escape rate is negligible. That is one reason why oxygen and nitrogen are more prevalent than hydrogen and helium.

Check Point 8.7: Is the atmosphere losing hydrogen to space? What about nitrogen?

8.6.2 Deposition

Deposition is the process of removing constituents from the atmopshere via settling. There are two types of **deposition**: **dry deposition** and **wet deposition**.

Dry deposition refers to the settling of gases and particulates onto surfaces, such as the ground or buildings.

Wet deposition refers to the absorption of gases into water droplets (or the condensation of water vapor onto particulates) which then fall as precipitation. For example, sulfur and nitrogen can interact with oxygen and water to produce sulfur and nitrogen-type acids. This then returns to earth as **acid rain**. Gases like neon, argon, krypton and xenon are noble gases. These gases are unlikely to react with any other gases and thus will be more long-lived.

What about the increase of carbon dioxide that I hear about? Why is that?

The increase in carbon dioxide comes from the burning of **fossil fuels**. Whenever one burns a hydrocarbon (whether it is in the form of a fossil fuel such as the gasoline in a car or in the form of a carbohydrate that you have eaten), oxygen is utilized with the hydrocarbon to produce water vapor and carbon dioxide. This is why carbon dioxide is exhaled when animals respire (as do plants).

It is interesting to note that fossil fuels are essentially from material that had been alive at some time in the past (hence the name fossils). The hydrocarbons in the fuel are essentially derived at some point from sugars in plants. **Photosynthesis** is the process by which the carbon in carbon dioxide (from the air) is converted to the hydrocarbons (or carbohydrates). Thus, the carbon being released by burning (or respiration) is the same carbon that was removed from the air via photosynthesis at some time in the past.

This removal by photosynthesis is responsible in large part for the current concentration of CO_2 in the atmosphere (less than 0.04%) being so much less than that put out in volcanic emissions (about 10%). By burning the fuel, we are putting back the CO_2 that was taken out many, many years ago (via photosynthesis).

Check Point 8.8: Does burning of fossil fuels create carbon or does it just move it from one form to another?

Project

Project 8.1: Create a spreadsheet (e.g., in Excel) with two columns. In the first column type in the molecule abbreviations for the gases that make up the atmosphere (excluding water), as listed in table 8.1. In the second column type in the relative proportion of each gas, in percent.

Project 8.2: Add two additional columns to your spreadsheet, one containing the concentrations in ppm (parts per million) and the other containing the concentrations in ppb (parts per billion). To determine the concentrations in ppm and ppb, use a spreadsheet equation. Do not simply type in the numbers.[iv]

[iv] As you can see, whereas a unit like percent (parts per hundred) is okay for nitrogen and oxygen, the proportion for other gases like helium and methane (CH_4) are so small that listing their proportion in percent means the number will have lots of zeroes. For those, it is better to use units of ppm (parts per million). Indeed, when examining the

Project 8.3: Add a new column to your spreadsheet. In it, copy the average molecular mass for each gas from table 8.1. In a separate cell, use a spreadsheet equation to calculate the average of all of the molecular masses, remembering to weight the values by their relative proportions.[v] *See previous chapter for the technique of calculating the weighted average.*

Problems

Problem 8.1: (a) Given the molar masses of Hydrogen and Oxygen (as given in the period table on the inside back cover), what is the molar mass of water vapor (H_2O)*?*
(b) If water vapor was present in the air, would that lower, or raise, the average molar mass relative to the accepted value of 28.96443 g/mol for dry air? Explain.

Problem 8.2: (a) Determine the escape velocity for an object at the earth's surface.
(b) Determine the escape velocity for an object at a point 500 km above the earth's surface.
(c) Compare the two values in (a) and (b). Is it significant difference? Should it be? Explain.

Problem 8.3: Make a prediction (estimation) for what the molar mass of the atmosphere might be one thousand years from now. Support your answer by providing estimates for how the relative abundances of N_2, O_2, *Ar*, CO_2, H_2 *and* He_2 *might change.*

Problem 8.4: (a) What is the mass of a mole of N_2 *molecules? (b) What is the molar mass of* N_2*?*

Problem 8.5: Suppose you land on an unknown planet and find that the pressure at the surface is 7.5 hPa with a temperature of $-20°C$ *and a density of 0.0155 kg/*m^3*. Based on these measurements, what is your guess for the gas that makes up most of this atmosphere?*

---

concentration of pollutants, it is not unusual to have the concentration reported in ppb (parts per billion).

[v] You can check your answer by comparing it with the accepted value of 28.96443 u (source: CRC Handbook, 1983).

Problem 8.6: Most of the atmosphere is made up of diatomic nitrogen and oxygen. How does the molar mass of water vapor (H_2O) compare to the molar masses of diatomic nitrogen and oxygen?

Derivations

Show-me 8.1: In most physics texts, the ideal gas law is given as $PV = nR^*T$, where n is the number of moles of the gas and R^* is the universal gas constant ($R^* = 8.3145$ J mol^{-1} K^{-1}). Show that equation 7.2 can be derived from this expression if the R in equation 7.2 is equal to the universal gas constant R^* divided by the average molar mass density.

Show-me 8.2: (a) Show that the value of the gas constant for dry air is equal to the universal gas constant divided by the molar mass density of dry air. In other words, plug in the values for the universal gas constant and molar mass density of dry air (see appendix B.2) and compare the result with the gas constant for dry air (see appendix B.2).
(b) Determine the gas constant for water vapor, R_{v}, by dividing the universal gas constant by the molar mass density of water vapor. Compare the result with the value of the gas constant for water vapor provided in appendix B.3.

Show-me 8.3: Show that equation 8.2

$$v_{\mathrm{e}} = \left[\frac{2Gm_{\mathrm{earth}}}{r}\right]^{1/2}$$

can be obtained from equations 8.3, 8.4 and 8.5. In other words, repeat the derivation produced in section 8.6.1 (stating the mathematics taken in each step) and add the algebra needed to go from equation 8.6

$$\frac{1}{2}mv^2 = Gm_{\mathrm{earth}}m \times \frac{1}{r}$$

to equation 8.2.

Show-me 8.4: Although about 78% of air molecules are nitrogen molecules, less than 78% of the air's mass is nitrogen. Show that the atmosphere is roughly 76% nitrogen by mass.

Show-me 8.5: Use the information in table 9.1 to show that the atomic mass of carbon dioxide is 44.0095 u (see table 8.1).

9. Molar masses

9.1 Introduction

In the previous chapter, I showed how the molar mass of dry air can be determined from the molar masses of each atom that makes up the dry air. In this chapter, I show how the molar mass of each atom can be determined from the masses of the sub-atomic particles that make up the atom. This information is not needed for the chapters that follow and thus can be reasonably skipped. However, the physics involved is very interesting and provides some insight into why the molar masses are what they are for each element on the periodic table.

9.2 Protons, neutrons and electrons

Each atom is made up of a positive-charged nucleus surrounded by a cloud of negative-charged **electrons**. The nucleus, in turn, is made up of positive-charged **protons** and neutral **neutrons**. The general term for particles in the nucleus are **nucleons**. Thus, both protons and neutrons are considered to be nucleons.

On the periodic table (see inside back cover), each element is assigned an **atomic number**. This is the number on the upper-left portion of each element box on the table. It corresponds to the number of *protons* that are in the nucleus.

> ✍ Elements are identified by the number of protons in the nucleus because that determines the electric charge of the nucleus and thus is the biggest influence on the chemical properties of the atom.

The charge of each electron is equal and opposite to the charge of each proton. We assume each atom is neutral, which means it has an equal number of electrons as protons.

> ✍ If the number of electrons does not equal the number of protons, the atom is not neutral and is called an **ion**.

Like charges repel and opposite charges attract. Consequently, the electron cloud is attracted to the protons[i] and the protons repel one another.

If the protons repel one another, how do they stay together in the nucleus?

In addition to the repulsive electric force, there is an attractive nuclear force. The nuclear force, by itself, is not sufficient to keep the protons together. That is what the neutrons are for. They provide some additional "glue" to keep the nucleus together.

In general, there are at least as many neutrons as there are protons. More will be said about this in the next section.

Check Point 9.1: (a) What is the difference between a nucleon and a neutron? (b) If the nitrogen nucleus has seven protons and seven neutrons, how many electrons does a neutral nitrogen atom have?

9.3 Isotopes

As mentioned in the previous section, elements are defined according to how many protons are in the nucleus, not how many total nucleons (protons and neutrons) are in the nucleus. For example, a nitrogen nucleus has seven protons. The number of neutrons is irrelevant.

Although a nitrogen nucleus can have any number of neutrons and still be considered a nitrogen nucleus (as long as there remains seven protons), it turns out that having seven neutrons is more stable than having any other number. So that version is more popular, so to speak.

How do we distinguish between the various "versions" of nitrogen nuclei that are possible?

---

[i] The electrons would "fall" into the nucleus if not for their kinetic energy.

In order to distinguish between the various "versions" of nitrogen nuclei that are possible, we indicate the version by the total number of nucleons in the nucleus.

So, for example, if a particular nitrogen nucleus has seven protons and eight (instead of seven) neutrons, for a total of 15 nucleons instead of 14, we would call it nitrogen-15. Nitrogen-14 is more stable but nitrogen-15 is possible. Similarly, oxygen-16 is more stable but oxygen-17 is possible, as is oxygen-18.

The various versions of an atom are called **isotopes**.

> ✍ The most stable configuration is not always the one with equal numbers of protons and neutrons. As the nucleus gets larger, extra nucleons are needed in order to maintain a stable nucleus.

Check Point 9.2: The most stable configuration of a nitrogen nucleus has seven protons and seven neutrons.
(a) If the number of protons in the nucleus changes, is it still considered a nitrogen atom?
(b) If the number of neutrons in the nucleus changes, is it still considered a nitrogen atom?
(c) If the number of electrons surrounding the nucleus changes, is it still considered a nitrogen atom?

9.4 Mass of protons, neutrons and electrons

To determine the molar mass of the atom, we need to know the masses of the particles that make up the atom. The masses of a proton (m_p), neutron (m_n) and **electron** (m_e) are[ii]:

$$\begin{aligned} m_p &= 1.67262171 \times 10^{-27}\ \text{kg} \\ m_n &= 1.67492728 \times 10^{-27}\ \text{kg} \\ m_e &= 9.1093826 \times 10^{-31}\ \text{kg} \end{aligned}$$

Notice that the electrons are about one thousand times less massive than the protons and neutrons. For that reason, we can neglect the mass of the

---

[ii] Source: physics.nist.gov

electrons if we do not need to be accurate within 0.1% (i.e., one out of a thousand).

Check Point 9.3: What fraction of oxygen-16's mass do the electrons provide?

9.5 Molar masses of sub-atomic particles

As you can see, sub-atomic particles have very little mass. In order to deal with more reasonable numbers, we can consider the mass of a **mole** of particles, rather than just one.

As mentioned on page 108, a mole represents 6.0221415×10^{23} objects. This number is known as **Avogadro's number** (N_{A}).

To determine the mass of a mole of electrons, then, we simply multiply the mass of a single electron by Avogadro's number. Similarly, we can do the same for the proton and neutron. Doing so, we get the following values:

$$\begin{aligned} \mu_p &= 1.00727646 \text{ g/mol} \\ \mu_n &= 1.00866591 \text{ g/mol} \\ \mu_e &= 0.00054858 \text{ g/mol} \end{aligned}$$

Notice that I've written the molar masses in units of grams per mole, rather than kilogram per mole, to be consistent with the way molar masses are given in the periodic table.

WHY IS A MOLE SET TO THIS PARTICULAR NUMBER?

This number was chosen such that a carbon-12 atom, with 12 nucleons in its nucleus, would have a molar mass of 12 g/mol. We could have instead defined it such that a hydrogen-1 atom would have a molar mass of 1 g/mol or an oxygen-16 atom would have a molar mass of 16 g/mol. Either way the we'd get that a mole of protons (or a mole of neutrons) would be close to 1 gram. It is just that 12 g/mol has been defined as the molar mass of a carbon-12 atom and that is where that conversion comes from.

> ✍ Actually, in the past, oxygen was indeed used (although many people preferred using hydrogen). However, that was before people realized that isotopes existed. Carbon-12 was chosen (in 1961) in part because that resulted in just a slight change in the definition (i.e., the average molar mass of oxygen was shifted from exactly 16 g/mol to 15.9994 g/mol).[iii]

Check Point 9.4: Is the determination of a mole simply a matter of convenience or does it represent something that is set by the physical universe?

9.6 Mass-energy equivalence

At this point, you should have noticed a problem.

I stated that carbon-12 has a molar mass of 12 g/mol. However, carbon-12 has 12 nucleons, each of which has a molar mass greater than one gram per mole. Shouldn't the molar mass of carbon-12 be equal to the sum of the molar masses of 6 protons, 6 neutrons and 6 electrons:

$$\begin{aligned} 6\mu_p + 6\mu_n + 6\mu_e &= 6(1.00727646 \text{ g/mol}) \\ &\quad +6(1.00866591 \text{ g/mol}) \\ &\quad +6(0.00054858 \text{ g/mol}) \\ &= 12.0989397 \text{ g/mol} \end{aligned}$$

instead of 12 g/mol?

The reason why the actual molar mass of carbon-12 is *less* than 12.0989397 g/mol is because of something called the mass-energy equivalence. In particular, the missing mass is associated with the energy that was "released" when the individual nucleons were brought together to form the nucleus.

> ✍ The mass difference is called the **mass defect**.[iv]

[iii] Source: Holden, N., *Atomic Weights and the International Committee: A Brief Historical Review*, Chemistry International, **26**(1), available via www.iupac.org (click on "Publications").

[iv] By the way, in a sense there is also a mass defect associated with the neutron. A

Most of the time, we don't need to worry about this, since the energy transfer involved in physical and chemical reactions is very small compared to the masses involved.

However, in the case of nuclear forces and energy, the energy transferred can be significant in comparison to the masses involved and we can no longer ignore the fact that mass and energy are equivalent.

Is the energy transferred during nuclear reactions really that huge?

It is all relative. For a single nucleus it is very tiny. All we are doing is taking a bunch of tiny nucleons and putting them together in the form of a nucleus. There isn't much energy released during that.

However, the particles involved are also very tiny. They are so tiny that the energy becomes significant, when compared with the masses. In comparison, the energy involved when two atoms come together (as in a chemical reaction) is much smaller. That is why don't need to consider the mass-energy equivalence during chemical reactions.

Why is energy released when nucleons are combined together?

The nucleons are attracted to each other via the nuclear force. Thus, energy is released when they come together. It is like dropping a book or having two magnets of opposite poles come together. When two objects that are attracted to each other come together, energy is released (perhaps in the form of heat or sound).

Conversely, it takes energy to break the nucleus apart, just like it takes energy to break apart two magnets of opposite poles or lift a book off the floor. It takes energy to break the nucleus apart because the nucleons in a nucleus are stuck together due to the nuclear force.

What about the electric force of repulsion due to all of the positive-charged protons in the nucleus? Wouldn't this make it easier to take the nucleus apart?

Yes, but the reason the nucleus stays together despite the electric repulsion is that the nuclear force of attraction is greater than the electric force of

---

single, isolated neutron is somewhat unstable – it will decay into a proton and an electron (and a neutrino). The energy released by this decay is associated with a mass defect and so the neutron is slightly more massive than the sum of the proton and electron masses.

repulsion. Technically, the energy we are talking about is the net effect of the two together.

Check Point 9.5: During a chemical reaction, energy can be released or absorbed. The following three situations are examples of chemical reactions. For each, identify whether any mass has been converted to energy or visa-versa.

(a) During combustion, like the burning of a candle, energy is released in the form of heat and light. Does the mass of the reactants (hydrocarbon and oxygen) equal the mass of the products (carbon dioxide and water vapor) or has some mass been converted to energy?

(b) During respiration, like what happens during exercise, energy is released in the form of heat and work. Does the mass of the reactants (glucose and oxygen) equal the mass of the products (carbon dioxide and water vapor) or has some mass been converted to energy?

(c) During photosynthesis, energy is absorbed in the form of light. Does the mass of the reactants (carbon dioxide and water vapor) equal the mass of the products (glucose and oxygen) or has some energy been converted to mass?

9.7 Binding energy

In the case of carbon-12, 0.0989397 g/mol is converted to energy when the nucleons come together.

This seems like a small amount of mass but it is significant when compared to the masses involved (i.e., 12 g/mol for carbon-12). In addition, it is equivalent to a large amount of energy. The equivalency is given by the following equation:

$$E = mc^2 \tag{9.1}$$

where c is the speed of light:

$$c = 2.99792458 \times 10^8 \text{ m/s}$$

The speed of light is very large, which means that a small amount of mass is equivalent to a significant amount of energy.

In this case of carbon-12, the 0.0989397 g/mol is equivalent to 8.89226×10^{12} J per mole of carbon-12 nuclei (convert the grams to kilograms and then multiply by the square of the speed of light).

This is the energy released when the nucleons stick together (or, conversely, how much is absorbed in order to break them apart). The general term for this energy is the **binding energy**.

The amount of binding energy associated with a nucleus depends on the number and type of nucleons present. In general, the larger the number of nucleons, the larger the binding energy per nucleon and the greater the difference between the nucleus mass and the sum of its components.

A list of binding energies (per nucleon) for various atomic nuclei can be found at the Lawrence Berkeley National Laboratory website (ie.lbl.gov).

Check Point 9.6: According to NIST, the molar mass of oxygen-16 (eight protons and eight neutrons) is 15.9949146221 g/mol. How is it possible for the molar mass of oxygen-16 to be less than 16 g/mol when each nucleon in its nucleus has a mass greater than 1 g/mol?

9.8 Estimating the molar mass of isotopes

Based on what has been discussed so far, we can make a rough estimate of the molar mass of any isotope as long as we know how many nucleons are in the nucleus.

First of all, we'll ignore the mass of the electrons. Next, we'll estimate the molar mass of each nucleon as being 1 g/mol. We know the actual molar mass is greater than this. However, when they come together in the nucleus some mass will be converted to energy. So, treating the molar mass as 1 g/mol is a good, though rough, estimate.

Based on this, we can expect oxygen-16 to have a molar mass of about 16 g/mol (it is actually 15.995 g/mol) and nitrogen-14 to have a molar mass of 14 g/mol (it is actually 14.003 g/mol).

Check Point 9.7: According to the periodic table on the inside back cover, *the molar mass of Argon is 39.948 g/mol and has an atomic number of 18, meaning there are 18 protons in its nucleus. Based on this information, how many neutrons are probably in the Argon nucleus?*

9.9 Average molar mass

In the last section, I mentioned that the molar masses of oxygen-16 and nitrogen-14 are 15.995 g/mol and 14.003 g/mol, respectively. However, according to the periodic table on the inside back cover, the molar masses of oxygen and nitrogen are 15.9994 g/mol and 14.0067 g/mol, respectively. You might ask, then, what is the difference?

The difference has to do with the fact that oxygen-16 is a particular isotope of oxygen, and the value given in the periodic table is the *average* oxygen mass, regardless of isotope.

It turns out that almost all of the oxygen in the air is oxygen-16. Consequently, the average molar mass of oxygen is very close to what it is for oxygen-16. However, there are some other isotopes, namely oxygen-17 and oxygen-18, and their presence means that the average molar mass of oxygen is slightly greater than the molar mass of oxygen-16.

To determine the average molar mass, we need to know the **relative abundance** of the various isotopes.

Table 9.1 lists all of the natural isotopes of atmospheric gases and their relative abundance. As you can see, some isotopes are more prevalent than others. For example, oxygen-16 is much more prevalent than oxygen-18 (this is because oxygen-16 is a much more stable configuration than oxygen-18).

To find the average molar mass of oxygen, then, we take the molar masses of the three isotopes and take the average, with each isotope mass weighted by its relative abundance. Since oxygen-16 is most abundant, we will weight that molar mass more heavily in the average than the molar masses of the others. The following example illustrates the technique.

Example 9.1: What is the average molar mass of oxygen?

Table 9.1: Atomic masses and relative abundances of natural atmospheric isotopes. Source: physics.nist.gov (look under "Physical Reference Data" and then "Nuclear Physics Data").

Element	**protons**	**nucleons**	**molar mass (g/mol)**	**relative abundance (%)**
H	1	1	1.0078250321	99.9885
		2	2.0141017780	0.0115
He	2	3	3.0160293097	0.000137
		4	4.0026032497	99.999863
C	6	12	12.0000000	98.93
		13	13.0033548378	1.07
N	7	14	14.0030740052	99.632
		15	15.0001088984	0.368
O	8	16	15.9949146221	99.757
		17	16.99913150	0.038
		18	17.9991604	0.205
Ne	10	20	19.9924401759	90.48
		21	20.99384674	0.27
		22	21.99138551	9.25
Ar	18	36	35.96754628	0.3365
		38	37.9627322	0.0632
		40	39.962383123	99.6003
Kr	36	78	77.920386	0.35
		80	79.916378	2.28
		82	81.9134846	11.58
		83	82.914136	11.49
		84	83.911507	57.00
		86	85.9106103	17.30
Xe	54	124	123.9058958	0.09
		126	125.904269	0.09
		128	127.9035304	1.92
		129	128.9047795	26.44
		130	129.9035079	4.08
		131	130.9050819	21.18
		132	131.9041545	26.89
		134	133.9053945	10.44
		136	135.907220	8.87

Answer 9.1: First multiply the molar mass of each isotope by its relative abundance and add them together. Using $\bar{\mu}(\mathrm{O})$ to represent the average[v] molar mass of oxygen:

$$\begin{aligned}\bar{\mu}(\mathrm{O}) &= 0.99757(15.9949146221\ \mathrm{g/mol})\\ &\quad +0.00038(16.99913150\ \mathrm{g/mol})\\ &\quad +0.00205(17.99916040\ \mathrm{g/mol})\end{aligned}$$

which equals 15.9994 g/mol.

✍ In this case, the relative abundances add up to 1.0000, as it should. If it did not, it would be more proper to divide by the total abundance.

Check Point 9.8: In the note I state that the relative abundances should add up to one. Why should that be so?

Problems

Problem 9.1: (a) Knowing the masses of a proton, neutron and electron, and assuming there are equal numbers of each in your body, how many protons are in your body?
(b) From your estimates in the previous problem, how many drops of water are in the oceans?
(c) Which is greater, the number of protons in your body or the number of drops in the ocean (see problem 7.2)?

Problem 9.2: From the molar masses of the isotopes of nitrogen provided in table 9.1, calculate the average molar mass of diatomic nitrogen. Compare your answer to that given in table 8.1. Should it be the same? If so, why? If not, which should be larger and why?

Problem 9.3: Suppose a friend calculates the mass of a carbon-12 atom by simply adding together the mass of 6 neutrons, 6 protons and 6 electrons. The

---

[v] By convention, a horizontal line on top of a variable abbreviation represents an average of some kind.

friend then comes to you for help because the total mass, when multiplied by Avogadro's number does not exactly equal 12 grams. What is wrong with your friend's calculation? Would the mistake make the mass of one mole of carbon-12 greater than 12 grams or less than 12 grams? Why?

Problem 9.4: According to ie.lbl.gov, the binding energy per nucleon associated with nitrogen-14 (7 protons and 7 neutrons) is 7475.614 keV per nucleon. Using equation 9.1, determine the mass defect for the nitrogen-14 nucleus. Note that one eV is equivalent to $1.602176534 \times 10^{-19}$ *J.*

Problem 9.5: (a) Given the molar masses of an individual proton, neutron and electron, and the mass defect determined in problem 9.4, calculate the molar mass of the nitrogen-14 atom.
(b) Compare your answer to that given in table 9.1. Should it be the same? If so, why? If not, which should be larger and why?

Part C

Impacts on Temperature

10. Advection

10.1 Introduction to part C

In this part of the text, we look at how a numerical determines how the temperature changes over time. There are several processes that might lead to a change in temperature. In chapter 11, we'll examine how the temperature changes due to the expansion and compression that occurs when air rises and descends. Finally, in chapter 12 we will look at how evaporation and condensation impacts the temperature.

In this chapter, we'll look at the process by which the wind brings in air that is colder or warmer. Basically, when the wind is blowing from a colder area to a warmer area, the warmer area will get colder. Conversely, if the wind is blowing from a warmer area to a colder area, the colder area will get warmer. This process is called **temperature advection**.

> ✍ Many properties can be advected, including moisture, momentum and temperature.

10.2 Warm vs. cold advection

Temperature advection refers to what happens when the wind is blowing from a colder or warmer location and consequently leads to the temperature cooling or warming.

If the wind is bringing warmer air, we call that **warm advection**. Conversely, if the wind is bringing air that is cooler, we call that **cold advection**.

Mathematically, we'll use the convention that warm advection is positive. Consequently, cold advection will be negative.

Check Point 10.1: (a) If warm advection is occurring at a particular location, *what would be your short-term forecast for temperature at that location? (b) For warm advection, what is the sign of the temperature advection?*

For temperature advection to occur, three things have to happen:

- The temperature must be different somewhere else – this means that there must be a temperature gradient present (i.e., $\nabla T \neq 0$),[i]
- The wind must be blowing (i.e., $\vec{v} \neq 0$), and
- The wind must be directed, at least partially, along the temperature gradient (i.e., across the isotherms).

For example, if the air temperature is colder to our north and the wind is out of the north, there is cold advection. On the other hand, with the same temperature gradient, if the wind is out of the east of if the wind is calm (i.e., the air is not moving), there is no temperature advection.

To illustrate, consider the plot shown in figure 10.1, which provides the observations of wind and temperature at the height where the pressure is 925 mb. The solid isolines indicate the height where the pressure is 925 mb. The dashed lines indicate the temperature.

Notice that the winds in the central United States are blowing from the south, from warmer temperatures to colder temperatures. This is a region of temperature advection, since the wind is blowing across the isotherms. It happens to be warm advection, since the wind is blowing from warmer to colder temperatures.

Based on this, we would expect the temperature to decrease in the central United States. This is indeed what happened, as can be seen in the same plot for the following day (see figure 10.2).

WHAT IF THE WIND HAD BEEN BLOWING FROM THE EAST INSTEAD?

If the wind was blowing from the east, there would be no temperature advection because the wind would not be blowing across the isotherms.

WHAT IF THE WIND HAD BEEN BLOWING FROM SOME OTHER DIRECTION?

---

[i] Remember that $\vec{\nabla} T$ is short-hand for $(\partial T/\partial x)\hat{\imath} + (\partial T/\partial y)\hat{\jmath}$.

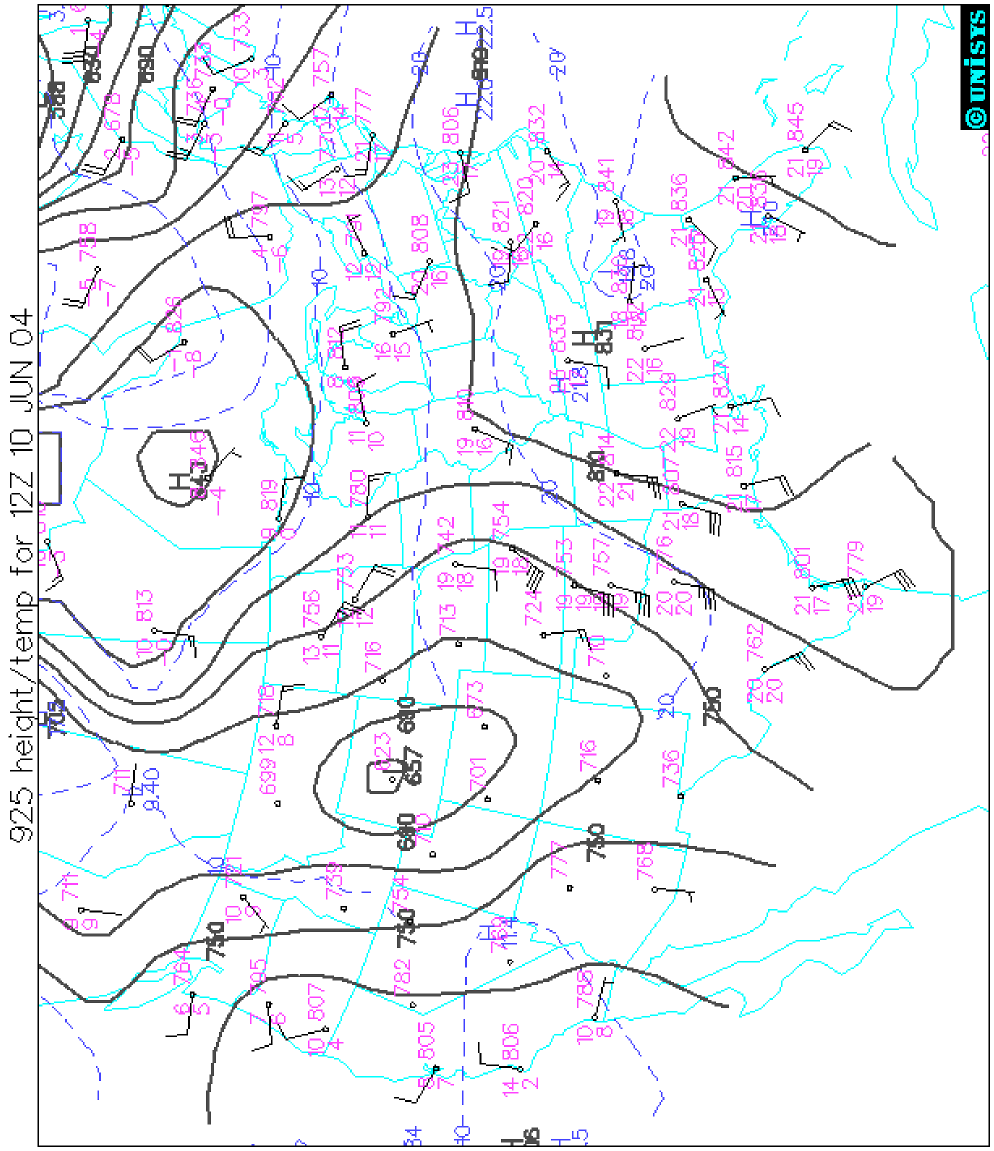

Figure 10.1: Heights (solid lines; every 30 meters) and temperatures (dashed lines; every 5°C) at 925 mb for 12 UTC on 10 June 2004. Heights, temperatures and wind bars at sounding sites are also provided.

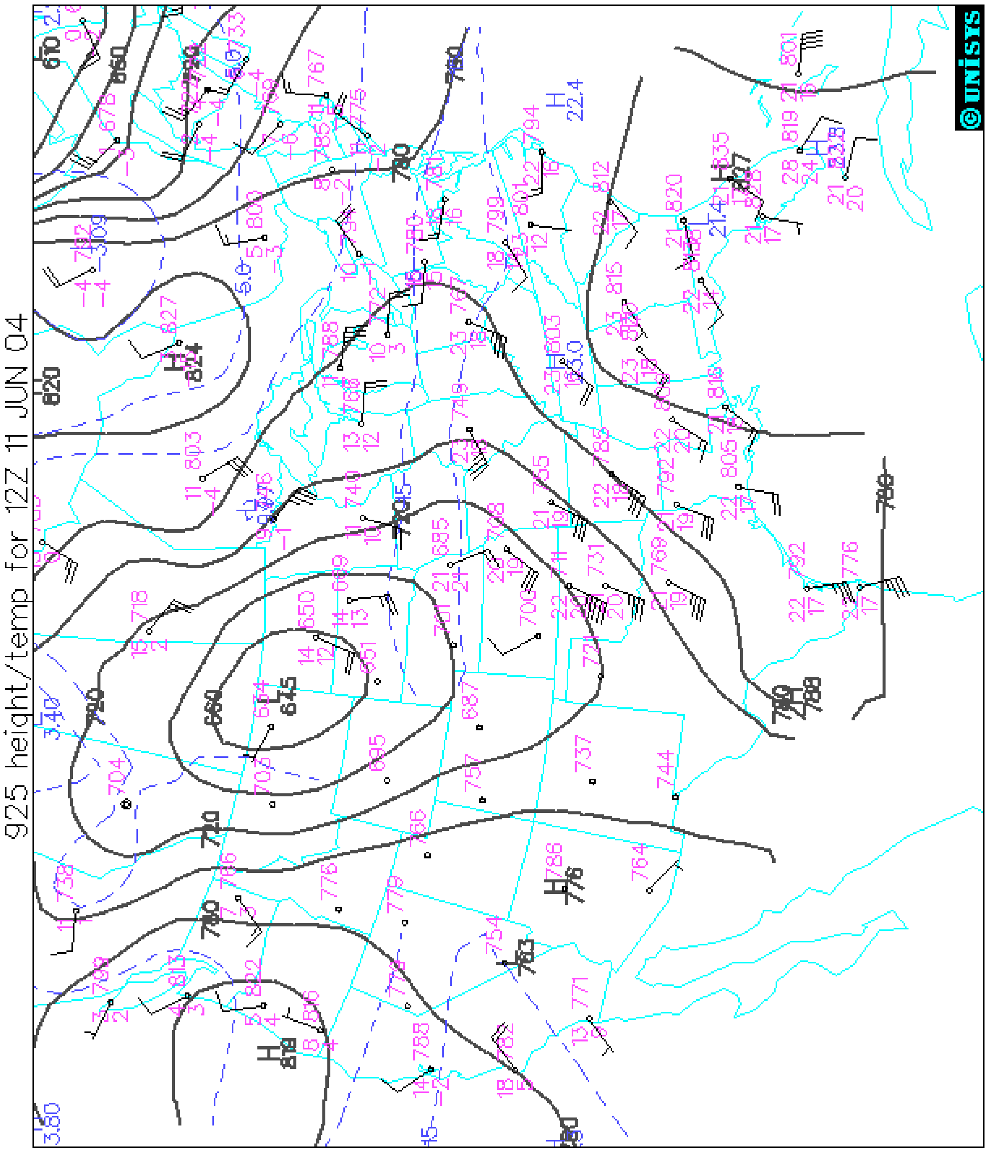

Figure 10.2: Same as in figure 10.1 except for 12 UTC on 11 June 2004 (24 hours later).

If the wind was blowing at an angle from north-south, it would be warm advection if some component of the wind was from the warmer region and it would be cold advection if some component of the wind was from the colder region.

WHAT IF THERE WAS AN EAST-WEST TEMPERATURE GRADIENT?

Then temperature advection would depend on whether the wind was blowing toward the east or toward the west.

Check Point 10.2: Suppose the temperature increased toward the west, with a wind blowing toward the east. Is this a situation of warm advection, cold advection or neither?

10.3 Time tendency equation

Before we examine how the temperature advection impacts the temperature, we first recognize that it will depend on time. The longer warm advection occurs, for example, the warmer it will get.

Because of this, we will typically focus on the **rate** at which the temperature changes rather than the temperature change itself. Knowing the rate, of course, we can then multiply by time to get the temperature change.

Because we will be predicting the rate at which the temperature changes, meteorologists have a special name for it. In meteorology, it is called the **temperature tendency**.

✍| In meteorology, any rate is called a **tendency**.

Check Point 10.3: What do meteorologists call the rate at which the temperature changes?

Mathematically, the temperature tendency is indicated as $\partial T/\partial t$. A positive temperature tendency means the temperature is warming whereas a negative temperature tendency means the temperature is cooling.

> ✍ I am using the "partial" notation, $\partial T/\partial t$, instead of the "total" notation, dT/dt, because I want to indicate that I am only considering the temperature change at a particular location. In other words, I am not following the parcel as it flows from place to place, so only the time will change, not the position. The partial derivative $\partial T/\partial t$ is used to indicate the change in temperature that occurs at a particular point in space whereas the total derivative dT/dt is used to indicate the change in temperature that occurs within the parcel, as it flows from place to place.[ii] We have used the partial notation before to represent gradients, like $\partial P/\partial x$ because we needed to know how a variable might change in a particular direction.

Check Point 10.4: Suppose the wind is blowing from an area that is warm to an area that is cold. Assuming there is no heating or cooling due to evaporation, condensation, mixing, radiation, expansion or compression, what is $\partial T/\partial t$*: positive, negative or zero?*

To predict *how much* warmer or colder the temperature will be, we need a **temperature tendency** equation.

There are two types of changes that can cause the temperature at a particular location to change. One is the temperature advection, where air of a different temperature is brought into the area. The other type of change is one that occurs as the air travels. An example of this would be solar radiation that is absorbed by the air.

The temperature advection can either be horizontal (due to horizontal winds) or vertical (due to upward or downward motion). To simplify things, we'll focus on horizontal advection.

Mathematically, we represent the horizontal temperature advection as $-\vec{v} \cdot \vec{\nabla} T$.

Let's interpret this in terms of what is happening in the central United States in figures 10.1 and 10.2. In particular, let's look at the expression in three parts.

- First we'll look at $\vec{\nabla} T$. Recall that $\vec{\nabla}$ represents a gradient. Consequently, $\vec{\nabla} T$ represents the temperature gradient. In the central

---

[ii] The partial derivative reflects a **Eulerian** view whereas the total derivative reflects a **Lagrangian** view.

United States there is a negative north-south temperature gradient ($\partial T/\partial y < 0$, where y is directed toward the north) because the temperature T is lower at greater values of y. This is typically what happens in the northern hemisphere.

- Next we will look at $\vec{v}$. This is the wind velocity. The wind in the central United States is directed toward the north (i.e., it is a southerly wind), which means that v is positive and u is zero.
- Finally, let's look at the presence of the negative sign. Notice that the product of v (positive in this case) and $\partial T/\partial y$ (negative in this case) is negative. Yet, this is a region of warm advection, which should be positive. The negative sign is added in order to be consistent with how we defined temperature advection.

> ✍ The dot between two vectors is used to indicate that the same components are used in the multiplication. This is known as the **dot product**.[iii] Mathematically, $\vec{v} \cdot \vec{\nabla} T$ is equivalent to $u(\partial T/\partial x) + v(\partial T/\partial y)$.

Check Point 10.5: If the temperature is warmer to the east and the wind is out of the east, what is the sign of the temperature tendency? Is this a case of warm advection, cold advection or neither?

What if the wind is not aligned along north-south or east-west?

In those cases, the temperature advection can be calculated in one of three ways:

- Use the total speed of the wind but only the component of the temperature gradient in the direction of the wind,
- Use the total temperature gradient but only the component of the wind in the direction of the temperature gradient,
- Calculate the product of the components of each and them together:[iv] $-u(\partial T/\partial x) - v(\partial T/\partial y)$.

[iii] Not too imaginative, is it?

[iv] For completeness, you could add a vertical term equal to $-w(\partial T/\partial z)$.

WHICH METHOD SHOULD I USE TO CALCULATE THE TEMPERATURE ADVECTION?

That depends on what you're given. On a synoptic map, you are given the total wind speed and direction. Since you have to calculate the temperature gradient in some direction regardless of which method you choose, the easiest method is to calculate the temperature gradient along the direction of the wind. Then use that with the total wind speed (method 1).

WHAT UNITS DOES THE TEMPERATURE ADVECTION HAVE?

When you multiply wind speed (in units of m/s) by the thermal gradient (in K/m), you get the same units as the temperature tendency (K/s).

WHAT DO WE DO IF THERE ARE OTHER FACTORS CONTRIBUTING TO THE TEMPERATURE TENDENCY?

If there are other factors contributing to the temperature tendency, like compression, expansion, radiation, etc., then we need to add in the temperature tendencies due to those factors. Since those factors contribute to the Lagrangian temperature tendency (i.e., dT/dt following the air parcel), we can modify our expression to include those factors as follows:

$$\frac{\partial T}{\partial t} = \frac{dT}{dt} - \vec{v} \cdot \vec{\nabla} T \tag{10.1}$$

For example, in figures 10.1 and 10.2 there is cold advection over eastern Canada. However, temperatures in that region didn't cool. This suggests that other processes were occurring in that region. We'll examine those other processes in the chapters that follow.

Check Point 10.6: If temperature advection is the only process responsible for the local change in air temperature, what is the value of dT/dt?

10.4 Wind chill

An important thing to notice about temperature advection is that the air doesn't necessarily get colder just because the wind blows. It depends on where the wind is blowing from. If the wind is blowing from a region that is warmer, the temperature advection would be positive and the temperature would rise.

Doesn't the temperature necessarily decrease when the wind blows?

No. It depends on which way the wind is blowing relative to the temperature gradient.

But when I go outside on a cold day, it certainly feels colder when the wind blows. Why is that?

This is known as the **wind chill** effect. The wind chill effect is due to advection but on a smaller scale. Remember that temperature advection will lead to cooling if the air being replaced is warmer than the air replacing it.

The air right next to your skin is usually warmer than the surrounding air because, well, you are warmer than the surrounding air and the right next to your skin warms up as a result.

The reason why we tend to wear clothes (like a sweater) when we go out in cold weather is because the sweater keeps the cold air out (away from our skin) and the warm air in (next to your skin). The air next to your skin is warm because it has warmed by contact with your warm skin. When the wind blows, it can advect the colder air toward you and replace the warmer air next to your skin. The faster the wind, the greater the temperature advection and the quicker the air next to your skin cools.

Does the actual temperature change?

Yes and no. Yes, the air right next to your skin tends to get colder when the wind blows. This is called **convection**. However, the air far away from you does not get colder or warmer.[v]

> ✍ Keep in mind that the temperature is a measure of the average kinetic energy of the molecules. If anything, increasing the wind speed increases their kinetic energy (and so temperature increases) but the effect is negligible because the kinetic energy associated with the wind speed is small in comparison with the kinetic energy of the molecular motion.

---

[v] Convection can change the temperature when the wind is mixing up the air with air from another level (see section 16.6).

What would happen if the air temperature was the same temperature as my skin? Would it still feel colder when the wind blows?

Yes, but not because of temperature advection (convection). In that case, the air may be drier than the air next to your skin (since you sweat and the water in your sweat evaporates into the air next to your skin). This allows more evaporation to occur and, as we discuss in chapter 12, evaporation is a cooling process.

Check Point 10.7: When the wind blows, does the temperature necessarily decrease? What about how it feels?

Project

Project 10.1: For this project, you will need a surface map that has temperature, pressure and wind. In particular, you need to obtain two surface maps, 24 hours apart. Print them out and identify a region that, on both maps, the wind is clearly blowing from an area with warmer temperatures to an area with colder temperatures or, visa versa, the wind is blowing from an area with colder temperatures to an area with warmer temperatures. The wind can be blowing either way but it must be the same (cold to warm or warm to cold) on both maps at the same location. Avoid areas near the water (as that may make it difficult to observe the impact of temperature advection) and time periods during which a front or other feature has passed through the region (as that will overwhelm the effect we are investigating).

Project 10.2: (a) At the location identified in project 10.1, which of the two surface maps show warmer temperatures – the earlier one or the later one?
(b) Is the wind blowing from cold to warm (cold advection), or is it blowing from warm to cold (warm advection), or is it neither?
(c) Is the situation identified in part (b) what you would expect to find for the situation identified in (a)? If so, why? If not, why not?

Problems

Problem 10.1: (a) Examine the map in figure 10.1. Determine the wind speed in km/h at Albany, New York, on June 10, 2004.
(b) For the same location, determine the component of the temperature gradient in the direction of the wind.
(c) From (a) and (b), determine the temperature advection at Albany in ° C/h.
(d) Calculate the change in temperature that would result in 24 hours.
(e) Based upon your answer to (d), what should the temperature be at Albany (at 925 mb) on June 11, 2004, assuming no heating/cooling due to compression/expansion, condensation/evaporation or radiation?
(f) Does the result in part (e) agree with the observed temperature 24 hours later in figure 10.2? If not, what could cause a difference?[vi]

Problem 10.2: In which of the following cases will a cup of water freeze quickest?
(a) 35° F and winds of 35 mph.
(b) 30° F and calm winds.
Explain your choice.

[vi] Note that if the air rises or falls, its pressure can change and that will lead to compression and/or expansion. Vertical motion is not provided on the map.

11. Compression and Expansion

11.1 Air parcels

In the previous chapter, we examined how the wind can bring in air that is cooler or warmer. This is known as temperature advection. However, there are other processes that can cause the temperature to change. In this chapter, we look at how air can warm due to **compression** and **expansion**.

Is compression and expansion a common occurrence?

Yes. In fact, it occurs whenever we have air moving up or down in the atmosphere.

Why?

To explain why, let's first distinguish between the air that is moving up or down and the air that represents the environment into which it is moving.

For the air that is moving up or down, imagine that we envelop a small portion of it with an imaginary plastic wrap-like surface. It would be like the air is inside a balloon, except our imaginary surface will move with the air, changing its shape as necessary in order to contain the air inside it. Such a region of air is considered an **air parcel**.

What would happen to the air parcel as we follow it on its journey through the atmosphere?

We'll assume the air parcel has the following properties:

- The air parcel will *not mix* with any air it encounters. It is as though the imaginary surface surrounding the air parcel is impermeable.
- The air parcel will expand or contract as necessary to always have exactly the *same pressure as the environmental air* surrounding it. Thus, if the air parcel flows into a region of lower pressure (by rising, for example), it will expand and in so doing push the air in the new environment out of the way.

How can we say that the air will not mix with its surroundings? Isn't that a pretty iffy assumption?

It is actually a pretty good assumption if the air parcel is big enough that any **diffusion** across its boundaries is too small to affect the average properties of the parcel as a whole. For example, in a typical numerical model (see page 10) each grid point can represent an air parcel of area 10,000 km^2 and a depth of about 1000 m. Even in operational models of higher resolution, the grid points are far enough apart to make this a pretty good assumption.[i]

✍| We will explore how mixing affects things in section 16.6.

A consequence of this assumption is that the air parcel can be considered to be thermally isolated (i.e., any temperature changes inside the parcel are not due to heat transfer across the boundaries). This means that the air parcel will not warm or cool due to mixing with the new environment *even if it encounters air that is colder or warmer.*

If temperature changes are not due to heat transfer across the boundaries, and we are assuming no condensation, evaporation or radiation, how can the temperature of the parcel change?

If the pressure doesn't change then the temperature can't either. However, when there are vertical motions, the air parcel enters an environment that has a different pressure. Thus, by the second assumption, the air parcel must either push against the environment (in an effort to expand) or get pushed on by the environment (as it contracts).

In the first case, the air parcel does work on the environment and, in so doing, the air parcel loses energy and cools down. In the second case, the environment does work on the air parcel and, as a consequence, the air parcel gains energy and warms up.

Both cases result in a change in air parcel temperature even though (according to the first assumption) the air parcel is thermally isolated from its surroundings. This effect is **adiabatic** because there is no heat transferred (*a-diabatic* literally means *not-passable*).

---

[i] Diffusion is not totally ignored in operational models of the atmosphere. The point being made here is that it is a small effect.

Check Point 11.1: When an air parcel moves into an environment where the surrounding air has the same pressure as the air parcel but is colder, what does the adiabatic assumption imply about the air parcel's temperature change?

11.2 Adiabatic changes in temperature

For the temperature of the air parcel to change adiabatically, it must enter an environment where the pressure is different from its own. This is most significant when an air parcel rises or falls (since the pressure necessarily decreases with height; this is discussed in section 15.1).

So, by how much does the temperature rise or fall when the air parcel compresses or expands?

Assuming an adiabatic process (i.e., no heat transfer), the small change in temperature dT depends upon the small change in pressure dP as follows:[ii]

$$dT = \frac{1}{c_\mathrm{p}\rho} dP \tag{11.1}$$

where ρ is the density of the air and c_p is the **specific heat** of dry air at constant pressure.

What is "the specific heat of dry air at constant pressure"?

The reason why the phrase "dry air" is included is because the value of the specific heat depends on the molecular makeup of the air. For dry air (no water vapor) it is $1004.67\ \mathrm{J \cdot kg^{-1} K^{-1}}$.

Exactly what "specific heat" means is described later. The meaning of "at constant pressure" will be discussed later also.

Where does this expression come from?

This expression comes from the **first law of thermodynamics**. The first law of thermodynamics is one of the basic relationships we will use. Because

---

[ii] Note that we are using Lagrangian (total) derivatives (e.g., dT instead of Eulerian (local) derivations (e.g., ∂T) because we are examining the temperature change as we follow the air parcel. In other words, the position is not being held fixed.

the first law of thermodynamics is so important, I'll examine it in more detail in the next section.

Check Point 11.2: According to equation 11.1, if the pressure decreases, what happens to the temperature?

11.3 First law of thermodynamics

In this section, I'll discuss the first law of thermodynamics.

The **first law of thermodynamics** provides a relationship between three things, which I will abbreviate as dU, dQ and dW.

dU is the change in internal thermal energy of the system (reflected in its temperature),[iii]

dQ is the energy added to the system due to *microscopic* processes (like conduction and radiation), sometimes called **heating** or **heat**, and

dW is the energy subtracted from the system due to *macroscopic*[iv] processes (like expanding against the environment), called the **work** done by the system on the environment.[v]

The specific relationship between the three is as follows:[vi]

$$dU = dQ - dW \tag{11.2}$$

[iii] We'll assume the parcel is moving slowly enough that its translational kinetic energy is a negligible fraction if its total internal energy (so its internal kinetic energy accurately reflects its temperature).

[iv] There is a fine and somewhat arbitrary line between microscopic and macroscopic processes.

[v] Typically, there is work done by gravity as the air parcel rises and falls but we are assuming that will only add to the translational kinetic energy, which will still be a small part of the total internal energy.

[vi] In some cases (e.g., in chemistry), the first law is written as $dU = dQ + dW$. In that case, dW represents the work done *on* the system (to increase U) rather than the work done *by* the system (which will decrease U).

Let's interpret this relationship and make sure it makes sense. Basically, it says that for an object's temperature to change (dU), we either have to add/extract heat to/from it (dQ) or have it do work or have work done on it (dW).

For example, for our "insulated" parcel of air (since we are assuming the parcel does not mix with the environment), no heat is added to the system so $dQ = 0$. For a rising parcel of air, which expands and pushes against the surroundings as it encounters lower pressure, it (the system in this case) does work on the environment and so $dW > 0$. With $dQ = 0$ and $dW > 0$, the first law of thermodynamics says that $dU < 0$. This means that the temperature of the system decreases. This is indeed what happens when a parcel of air rises in the atmosphere.

Check Point 11.3: According to the first law of thermodynamics (equation 11.2), what happens to the temperature of the parcel when the air parcel expands adiabatically?

We don't actually use the first law of thermodynamics as expressed in equation 11.2. Instead, we use the "idea" embodied in the first law to derive expressions that are appropriate for the situations we encounter.

To make it easier to derive relationships from the first law, however, we need to write it in terms of temperature and pressure (i.e., dT and dP). When we do so, it looks like:

$$c_{\mathrm{p}} dT = \frac{dQ}{m} + \frac{dP}{\rho} \tag{11.3}$$

where $c_{\mathrm{p}} = 1004.67\ \mathrm{J\ kg^{-1}\ K^{-1}}$ for dry air (no moisture).

The first term, $c_{\mathrm{p}} dT$, represents the change in internal thermal energy (which depends on temperature). The second term, dQ/m, represents the heating. The third term, dP/ρ, represents the work that the parcel does on the environment (due to the change in pressure).

This form allows us to derive equation 11.1 quite easily. If there is no heat added ($dQ = 0$) then a change in pressure dP will be related to a change in temperature dT such that $c_{\mathrm{p}} dT = dP/\rho$. Solve for dT to get equation 11.1.

The derivation of this expression is provided in appendix C.2. It is recommended that you go through the derivation yourself, as it shows what is represented by c_p (the **specific heat** of dry air at constant pressure). It also clarifies what it means to be "at constant pressure."

As you go through the derivation, you'll find that we need to utilize a couple of definitions.

$dW = Fdx$	This is the definition of work. Work is done by an object when it exerts a force F over a distance dx.
$P = F/A$	This is the definition of pressure. By replacing F with PA in the definition of work, we get $dW = PAdx = PdV$.
$C = dQ/dT$	This is the definition of heat capacity. It represents how much energy an object needs to absorb in order to raise its temperature one degree.
$c = C/m$	This is the definition of specific heat. It is the heat capacity per mass.

Check Point 11.4: In the derivation of equation 11.3, four definitions are used to get it in a form involving dP and dT. What are they?

One by-product of the derivation is the discovery that there are two specific heat values: one at constant pressure (c_p) and one at constant volume (c_v). As shown in the appendix, there is a rather simple relationship between them:

$$c_\mathrm{p} = c_\mathrm{v} + R. \tag{11.4}$$

Although only c_p is used in the version of the first law we are using, this relationship between c_p and c_v will come in handy later.

Check Point 11.5: According to equation 11.4, which is bigger: c_p or c_v?

11.4 Potential temperature

In this section, I introduce a adiabatic version of the first law of thermodynamics called Poisson's equation. The reason why we need another version is because the expression we currently have (equation 11.1)

$$dT = \frac{1}{c_\mathrm{p}\rho} dP$$

involves the density ρ. As we've noted before, the density is a problem because it isn't easily measured.

To get an expression we can work with, we need to do two things.

First, we need to get rid of the density, This is actually pretty easy to do. After all, we can use the ideal gas law to replace ρ with $P/(RT)$. Doing so, we get

$$dT = \frac{RT}{c_\mathrm{p}P} dP$$

or, rearranging,

$$c_\mathrm{p}\frac{dT}{T} = R\frac{dP}{P}.$$

This still has a problem in that it only works for small changes in temperature and pressure (i.e., dT and dP). What about large changes?

So, the second thing we need to do is convert this expression, involving infinitesimal changes, to one involving finite changes. This involves some calculus (see appendix C.3). Doing so, we can get:

$$T = T_0 \left(\frac{P}{P_0}\right)^{R/c_\mathrm{p}} \tag{11.5}$$

which gives the temperature T that the parcel will have at pressure P if it initially had a temperature T_0 at pressure P_0.

Check Point 11.6: According to equation 11.5, what happens to the temperature if the pressure decreases?

Example 11.1: Suppose 10°C air at 800 mb is compressed adiabatically to 1000 mb. What would its temperature be then?

Answer 11.1: Using equation 11.5, we have

$$\begin{aligned} T &= T_0(P/P_0)^{R/c_\mathrm{p}} \\ &= (283\ \mathrm{K})\Big((1000\ \mathrm{mb})/(800\ \mathrm{mb})\Big)^{2/7} \end{aligned}$$

which gives a temperature of 302 K or 29°C.

> ✍ Notice that I used 283 K for the temperature, not 10°C. This is because the ideal gas law, which was used in the derivation of equation 11.5, uses the absolute temperature (not the temperature relative to freezing).

In the example, the temperature T represents the temperature the air parcel would *potentially* have if the pressure was changed to 1000 mb such that no heat is added to or taken away from the air parcel. We could've changed the pressure to anything, but the resulting temperature depends on what pressure we choose.

I used 1000 mb, though, because it is common to use that as a reference point. It is so common, in fact, that the 1000-mb equivalent temperature is called the parcel's **potential temperature**.

To determine a parcel's potential temperature, then, we simply need to determine what the temperature would be if compressed (or expanded) to 1000 mb:

$$\theta = T \left(\frac{1000 \text{ mb}}{P} \right)^{R/c_\mathrm{p}} \tag{11.6}$$

where the temperature at 1000 mb is indicated by the Greek letter θ.

This is known as **Poisson's equation**. The variables P and T represent the pressure and temperature, respectively, of the parcel before the air is compressed/expanded to 1000 mb.

What advantage does the potential temperature have over regular temperature?

The potential temperature is useful for making comparisons between parcels. For example, we might want to know is which of two parcels, if brought to the same pressure, is warmer. The potential temperature allows us to do this.

For air parcels at altitudes above 1000 mb (i.e., pressure less than 1000 mb), the potential temperature is greater than the air temperature (since the air parcel would have to be compressed to get to 1000 mb). Conversely, for air parcels at altitudes below 1000 mb (i.e., pressure greater than 1000 mb), the potential temperature is less than the air temperature.

Check Point 11.7: If air at 800 mb has a potential temperature 300K, what is its potential temperature if brought up to 500 mb?

Problems

Problem 11.1: The pressure and temperature at the levels at which jet aircraft normally cruise are typically 200 mb and -60° C. Determine the temperature of this air if it were adiabatically compressed to pressure inside the cabin (1000 mb). If an aircraft were to use the outside air to replenish the air inside the cabin, would the aircraft need a heater or an air conditioner?

Problem 11.2: Upwelling, like that along the coast of California, is responsible for the cold surface water there. This is because the water below the surface is cooler than that above. How do we know this is not because the rising water cools adiabatically, as with air parcels?

Derivations

Show-me 11.1: For dry air (no moisture; and assuming the molecular make-up discussed in chapter 1), $c_\mathrm{v} = 717.62\ \mathrm{J \cdot kg^{-1}K^{-1}}$ *and* $c_\mathrm{p} = 1004.67\ \mathrm{J \cdot kg^{-1}K^{-1}}$. *Verify that* $c_\mathrm{p} = c_\mathrm{v} + R$ *(in keeping with equation 11.4) and* $c_\mathrm{p} : c_\mathrm{v} : R = 7 : 5 : 2$ *(in keeping with the footnote on page 286).*

Show-me 11.2: Derive equation 11.5,

$$T = T_0 \left(\frac{P}{P_0}\right)^{R/c_\mathrm{p}}$$

from equation C.4 (see appendix C.3),

$$c_\mathrm{p} d \ln T = R d \ln P.$$

12. Latent Heating and Cooling

12.1 Introduction

In the previous chapter, changes in an air parcel's temperature (in the Lagrangian sense) were due only to compression and expansion. There are other processes, however, that can change the temperature of the air. In this chapter, we examine the process of **latent heating**, associated with the evaporation and condensation of water. In the next chapter we will examine the process of radiation. In both of these cases, dQ will not be zero.

Every time water changes state (e.g., liquid to/from vapor, solid to/from liquid, or solid to/from vapor), energy is either released or absorbed. For example, evaporation is a cooling process and condensation is a warming process. In other words, when a vapor condenses, it releases energy to the surroundings, warming it up. When a liquid evaporates, it absorbs energy from the surroundings, cooling it off.

As long as the water (in whatever state) remains in the parcel of interest, the process is reversible (i.e., when the water changes back into its original state, ΔQ is reversed so the net ΔQ is zero) and so the process is still considered to be adiabatic (in that situation). Meteorologists consequently refer to this process as **moist adiabatic**.

In real life, once the water condenses into liquid from vapor, the water droplets can become large enough that they fall out of the parcel (as precipitation; see chapter 6). In that case, the process would no longer be adiabatic and would more appropriately referred to as non-adiabatic or diabatic.

Check Point 12.1: When the water vapor in the air condenses, does that warm up the air or does that cool down the air?

12.2 Latent heat

As stated earlier, the air warms up when the water vapor in the air condenses. And, conversely, the air cools down when water evaporates.

What we want to do is predict the change in temperature. Before doing that, however, let's examine the process of evaporation and condensation to get a qualitative sense for why there should be warming or cooling.

The difference between a liquid and a gas is that in a liquid there is an intermolecular force of attraction between the molecules. For liquid water, that is the attraction between the hydrogen of one water molecule and the oxygen of another water molecule. We call this a **hydrogen bond**.

In any event, it takes energy to separate the molecules against this intermolecular force. That energy comes from the surrounding molecules (i.e., the thermal energy associated with the molecules in the water and/or the molecules in the air). Since temperature is a measure of the thermal energy of the molecules, the temperature drops when water evaporates.

For the reverse situation of condensation, energy is released and the temperature warms up. This is analogous to what happens when you let two things that are attracted to each other come together. For example, if you drop a rock to the Earth, there is not only a release of energy via sound, but the ground warms up at the point where the rock fell.

Check Point 12.2: When a bond is broken, as during the transformation from liquid to vapor, it energy released or is it absorbed?

12.2.1 Cooling examples

Continuing with the analogy, we recognize that at any given time some students are leaving while others are entering. Whether a cooling or warming is experienced depends upon the *net* transfer. If more students are entering the classroom than leaving (i.e., net condensation), there will be a general warming. If more students are leaving the classroom than entering (i.e., net evaporation), there will be a general cooling.

Sweating

As long as the air is unsaturated, any liquid water will evaporate leading to cooling. This is the case with sweating. The body releases water to the skin surface and the water evaporates, helping to cool off the body. A similar effect occurs when one gets out of the water after swimming.

The air near the skin would quickly reach saturation (and the cooling effect would stop) if the air around the skin weren't continually replaced by drier air. This is why sweating is more efficient when the wind blows (or a fan is blowing).

Check Point 12.3: What process is principally responsible for the cooling associated with sweating: the process of creating the sweat or the process of evaporating the sweat?

Wet-bulb temperature

Another example of latent cooling is when you use a blow drier to send hot air through a moist towel or piece of fabric. The air on the other side will be relatively cool because the moisture in the fabric evaporates into the hot air, cooling it. The drier the hot air, the cooler the air on the other side because more moisture will evaporate into it. In fact, one way of determining the moisture content of the air is to send the air through a moist fabric and then measure the cooling that results.

The temperature of the moist air is called the **wet bulb temperature**. It is called the wet bulb because it is typically measured by wrapping the bulb of a regular thermometer with a wet cloth and then swinging the bulb in the air to allow the air to pass through the wet cloth to the bulb. At the same time, a regular thermometer measures the air temperature before it passes through the cloth. The difference between "wet" bulb reading and the "dry" bulb reading is indicative of how much moisture is present in the air (usually through a table, although the skew-T log-P diagram can also be used). Such an instrument is called a **sling psychrometer**.

Is the wet-bulb temperature the same as the dew point?

No, they are not the same. The wet-bulb temperature is higher. As the air passes through the wet cloth, moisture is added to the air. Consequently, the dew point temperature of the air must increase. After passing through the cloth, we assume the air is saturated. At that point, the wet-bulb temperature must equal the now higher dew point (i.e., a dew point that is higher than the original dew point).

Check Point 12.4: Which is higher: the dew point or the wet-bulb temperature?

12.2.2 Heating examples

In the atmosphere, cooling situations are mainly limited to areas close to a moist surface (such as above the ocean) or when there is rain falling through a dry layer. Most of the time, we are concerned about heating (due to condensation) rather than cooling (due to evaporation).

Condensation occurs when the air cools and, as we know from chapter 11, rising motion causes the air to cool. Such cooling can lead to condensation (and clouds) and this introduces latent heating that we cannot ignore. More is said about this in chapter 17.

12.3 Quantifying latent heating

Now that you hopefully have a good qualitative feel for whether the air will warm or cool, let's quantify the amount of heating or cooling that occurs. To do so, we will utilize the first law of thermodynamics, which relates the warming or cooling with the amount of energy that is absorbed or released.

The amount of energy absorbed or released depends on two things. One is the amount of water that changes state. The greater the mass of water that changes state, the more energy that is released or absorbed. We indicate the amount of mass by Δm.

Since the energy is represented by ΔQ, we can indicate the dependence on the mass as a proportionality:

$$\Delta Q_{\text{latent}} = L\Delta m$$

where L is the proportionality constant (and thus has units of Joules per kg). According to this relationship, the more mass that changes state, the more energy that is released/absorbed.

> ✍ Technically, L can depend on the temperature. Since the temperature changes when the water changes state, we technically should express the relationship in terms of infinitesimals as follows: $dQ_{\text{latent}} = Ldm$.

The second thing that influences the amount of energy that is released or absorbed is the type of change. There are three types of phase changes that water can undergo:

1. Vaporization/condensation (liquid to vapor or vapor to liquid)
2. Melting/freezing (solid to liquid or liquid to solid)
3. Sublimation/deposition (solid to vapor or vapor to solid)

The value of L will depend on which change of state is taking place. We will examine each of these separately.

12.3.1 The latent heat of vaporization

When water changes state from liquid to vapor (vaporization) or vapor to liquid (condensation), we represent L as L_{v} and refer to it as the **latent heat of vaporization**. For every kilogram of water that evaporates or condenses, about 2.5×10^6 J of energy is absorbed or released. In other words, L_{v} is equal to about 2.5×10^6 J/kg.

$$L_{\text{v}} \sim 2.5 \times 10^6 \text{ J/kg}.$$

> ✍ We could also call L_{v} the **latent heat of condensation** – it turns out that the energy per mass is the same for both evaporation and condensation.

Check Point 12.5: Why doesn't the latent heat of vaporization have units of heat?

12.3.2 The latent heat of sublimation

When water changes state from solid to vapor (**sublimation**) or vapor to solid (**deposition**), we represent L as L_{s} and refer to it as the **latent heat of sublimation**.

The solid form of water is called ice. It is the sublimation process that is responsible for the crystalline structure of snow as opposed to frozen water.

For water, it takes a little more energy (per mass) for the solid/vapor conversion compared to the liquid/vapor conversion (for water, the latent heat of sublimation, L_{s}, is about 2.83×10^6 J/kg).

$$L_{\mathrm{s}} \sim 2.83 \times 10^6 \text{ J/kg}.$$

Thus, for a given temperature, it is slightly easier for a molecule to enter the vapor phase from the liquid phase than from the solid phase. It follows, then, that less vapor is needed for saturation with respect to ice than with respect to liquid water. We will use that property to explain snow formation in section 6.3.3.

Check Point 12.6: Which is larger, the amount of energy required to evaporate 1 gram of water or the amount of energy required to sublimate 1 gram of ice?

12.3.3 The latent heat of fusion

When water changes state from solid to liquid (**melting**) or liquid to solid (**freezing**), we represent L as L_{f} and refer to it as the **latent heat of fusion**.[i]

For water, it is easier to go from liquid to solid (or back) than to make any of the other phase transitions. For water the latent heat of fusion is about 3.35×10^5 J/kg.

$$L_{\mathrm{f}} \sim 3.35 \times 10^5 \text{ J/kg}.$$

---

[i] The word *fusion* is related to *melting* and derives from the Latin word *fusion*, past participle of *fundere*, which means *to melt*. Consequently, fusion, like vaporization, describes the process from the more ordered to less ordered state.

It turns out that the energy required to go directly from solid to vapor (the latent heat of sublimation) equals the sum of both the latent heat of fusion (solid to liquid) and the latent heat of vaporization (liquid to vapor):

$$L_{\mathrm{s}} = L_{\mathrm{v}} + L_{\mathrm{f}}.$$

If freezing is exothermic (i.e., releases latent heat), how is freezing possible? Wouldn't it warm up too much to freeze?

All other things being equal, it is possible for the freezing process (like the condensation process) to warm up the water so much that further freezing is not possible. Conversely, melting (like evaporation) might cool the water so much that further melting is not possible.

In both cases, you must continue to add or extract heat simply to make promote the change of state and maintain the same temperature.

For example, consider a raindrop falling through a layer that is below freezing (see discussion on sleet in section 6.3). As the raindrop falls through the cold air, some of it freezes. The freezing releases latent heat that, in turn, warms up the environment (made up by the drop and the surrounding air). If it weren't for the transfer of this heat away from the falling rain droplet, the droplet might warm so much that it is no longer below freezing.

> ✍ Problem 12.2 shows that if kept insulated, only a fraction of the drop will freeze and we wouldn't get sleet (see section 6.3.2).

Check Point 12.7: Which is larger, the amount of energy required to evaporate 1 gram of water or the amount of energy required to melt 1 gram of ice?

12.3.4 The first law of thermodynamics

Now that we can quantify the amount of energy associated with the change of state, we can use that with the **first law of thermodynamics** to predict the change in temperature.

The first law of thermodynamics (equation 11.3)

$$\frac{dQ}{m_{\mathrm{T}}} = c_{\mathrm{p}} dT - \frac{dP}{\rho}$$

relates the heat transfer dQ with the change in internal energy (related to dT) and the work done (related to dP). A positive heat transfer ($dQ > 0$) leads to an increase in temperature ($dT > 0$). Likewise, work done on the parcel via compression ($dP > 0$) leads to an increase in temperature.

> ✍ Typically, we are heating/cooling the air. Consequently, we use c_p for air, which is 1004.67 J/(kg K), and the mass m corresponds to the mass of air that is heated or cooled. The energy can also be used to heat/cool the water. If we want to know the change in temperature of the water, we'd have to use c for water, which is 4160 J/(kg K) (this is sometimes given in units of calories and grams, which gives 1 cal/(g°C)).

If some water evaporates or condenses, there is heat absorbed or released proportional to the mass of the water involved:

$$dQ = L\ dm_\mathrm{v}$$

where L could be L_v, L_f, or L_s, depending on whether the conversion is between liquid/vapor, solid/liquid or solid/vapor. Notice that I use m_v here. In the first law, m refers to the material that is being heated/cooled. That would be the air. I use m_v to represent the mass of the water that is changing state. That may be different than the mass of air that is being heated/cooled. For that reason, I'll use m_T for the latter.

Plugging this into the first law of thermodynamics, we find that

$$\frac{L\ dm_\mathrm{v}}{m_\mathrm{T}} = c_\mathrm{p} dT - \frac{dP}{\rho}$$

or

$$dT = \frac{1}{c_\mathrm{p}} \frac{L\ dm_\mathrm{v}}{m_\mathrm{T}}. \tag{12.1}$$

What happened to the dP term?

We'll assume that the change in temperature is not accompanied by a corresponding change in pressure (i.e., the parcel simply expands or contracts in order to maintain the same pressure with the surroundings). Consequently, $dP = 0$.

I thought we were going to avoid the use of mass in an equation. Shouldn't equation 12.1 be written in terms of density?

Right you are. And we *could* write this in terms of density or pressure. However, it turns out to be simpler to write this in terms of q, the specific humidity, which is the ratio of vapor mass to the total air mass (see section 5.3.4).[ii]

Consequently, equation 12.1 becomes

$$dT = \frac{L\ dq}{c_\mathrm{p}} \tag{12.2}$$

In other words, given a change in the specific humidity (due to condensation, evaporation, sublimation or deposition), the temperature change that results is given by this expression.

Check Point 12.8: In equation 12.2, dT refers to the change in temperature experienced by the air as the moisture condenses or evaporates. What does L, dq and c_p refer to?

Problems

Problem 12.1: (a) Which is greater: the energy "absorbed" when water evaporates, the energy "released" when water vapor condenses, or are they equal?
(b) Which is greater: the energy "absorbed" when water evaporates (from liquid to vapor) at 0° C, the energy "absorbed" when ice sublimates (from ice to vapor) at 0° C, or are they equal?

Problem 12.2: (a) If the air is already saturated, should there be a difference between the wet bulb reading and the dry bulb reading? Explain.
(b) If the air is not saturated, which should be higher, the wet bulb reading or the dry bulb reading? Explain.

Problem 12.3: Suppose the air in this room was saturated. If all of the moisture were removed by condensing it out (at constant pressure), what would be the temperature increase in the room? Hint: Assume a reasonable temperature to get an estimate of the specific humidity (see problem 5.6).

[ii] Again, keep in mind that we are dealing with the air, since the air is warming/cooling. So, c_p is for air.

Derivations

Show-me 12.1: From the appendix, identify how much energy (per mass) is required to melt ice at 0° C, evaporate water at 0° C and sublimate ice at 0° C. Verify that $L_{\mathrm{s}} = L_{\mathrm{v}} + L_{\mathrm{f}}$.

Show-me 12.2: Suppose a blob of water has a temperature of -20° Cwhen freezing is initiated within the blob. Show that 23% of the water will freeze before the temperature rises to 0° C, assuming the water is insulated from its surroundings. Hint: assume X *represents the fraction that freezes and derive an expression that represents the amount of latent heat given off when a fraction* X *of the total mass freezes. That heat must go into warming both the frozen part and the unfrozen part to zero degrees (and each part has a different specific heat).*

13. Radiational Heating and Cooling

13.1 Introduction

So far, we have examined two processes that can change the air's temperature.

One process is compression and expansion. Even with no mixing, the air temperature will change if the air parcel expands against the environmental air or, conversely, the environmental air compresses the air parcel. This is called adiabatic cooling or warming and is the reason why rising air cools and descending air warms.

The other process we examined was latent heating and cooling, due to the condensation or evaporation of water. This is why we feel colder when water evaporates from our skin.

There is one more process that we will examine. That process is called **radiation**.

13.2 Irradiance

We know from section 6.5.2 that all objects emit radiation. If an object absorbs more radiation than it emits, it warms up. If it emits more than it absorbs, it cools down.

To predict how much an object warms up or cools down, then, we first need to quantify the "amount" of radiation it emits or absorbs. This we will do via a quantity called **irradiance**.[i]

---

[i] Some people call this the **radiant flux density** while others call it the **radiative flux**, mostly likely because it is similar to the **latent heat flux** and the **sensible heat flux**. I'm avoiding the term "flux" because there does not seem to be a consistent definition for it.

> ✍ Notice that "irradiance" does not mean "opposite of radiance." The difference between irradiance and radiance is more like the difference between irradiating an object (with radiation) and radiating confidence. The latter is more ambiguous whereas the former is more specific.

In a sense, radiation "carries" energy. The irradiance represents the amount of energy that "flows" (via radiation) through a region in a given amount of time.

If the irradiance is the energy, why can't we just call it the energy?

Because the irradiance does not just represent the energy. It represents the *rate* at which the energy is "delivered" by the wave and how "concentrated" the energy is.

After all, when you sit out in the sun, the energy you receive from the sun depends on how long you experience the radiation and also how large an area that receives the radiation. We need a term that is independent of time and area. That is why we are using the *irradiance*, which is the rate that energy passes through a certain cross-sectional area. Consequently, irradiance will have units of joules per second per square meter.

Isn't a joule per second the same thing as a watt?

Yes. So, irradiance will have units of watts per square meter ($\mathrm{W/m^2}$).

What symbol will we use for irradiance?

We will use the symbol F for the irradiance except for when we are dealing with the irradiance emitted by the sun. In that case, I will use S.

> ✍ I am using the letter F because that is the abbreviation used in PHYS304 (Modern Physical Astronomy), likely because of how the irradiance is sometimes called the radiative flux. Most references, however, seem to use E for irradiance. Unfortunately, that is likely to cause confusion with energy. An alternative symbol I've seen is I. However, that looks like the number one.

What are some typical irradiance values?

The **solar irradiance** (the irradiance from the sun) is about 1367.6 $\mathrm{W/m^2}$ just outside the atmosphere.[ii] For example, consider a one-square-meter

---

[ii] Source: http://nssdc.gsfc.nasa.gov/planetary/factsheet/Earthfact.html (NASA Fact Sheet).

surface oriented perpendicular to the sun's rays just outside the atmosphere. We can send up a satellite to measure the radiation energy incident upon that one-square-meter surface each second. It turns out that just outside the atmosphere, about 1367.6 J of energy is incident upon that one-square-meter surface each second. This value is known as the **solar constant**.[iii]

As mentioned above, I will use S to indicate the irradiance of the sun. Since we'll be using the solar constant value a lot, we'll give it a special symbol: S_0.

WHY IS IT IMPORTANT TO NOTE THAT THIS IS THE IRRADIANCE "JUST OUTSIDE THE ATMOSPHERE"? IS IT ANY DIFFERENT CLOSER TO THE EARTH?

As the radiation passes through the air, some of it might reflect back out into space (which we'll discuss later in this chapter) or get absorbed by the air (which we'll discuss in chapter 14). Consequently, the solar irradiance is less by the time the radiation gets to the surface of the Earth. We'll explore later how the irradiance changes within the atmosphere.

IS THE IRRADIANCE ANY DIFFERENT CLOSER TO THE SUN?

Yes.

WHY?

Because energy is conserved and the energy "spreads out" as it flows away from the sun. After all, the sun emits radiation in all directions not just on the one-square meter area outside the atmosphere.

Consider, for example, the following scenario. Suppose we enclose the sun with a big invisible balloon with a radius equal to the orbit of Earth. The radiation passing through the balloon would have an irradiance equal to 1367.6 W/m^2. Since irradiance is the power per area, the total power passing through the balloon surface would be the product of the solar irradiance at that location S_0 and the surface area[iv] of the balloon $4\pi R^2_{\text{Earth orbit}}$:

$$P_{\text{sun}} = S_0 4\pi R^2_{\text{Earth orbit}}$$

[iii] This is really a misnomer because it does not stay constant but varies slightly due to variations in the sun's output and distance (see http://www.ngdc.noaa.gov/stp/SOLAR/IRRADIANCE/irrad.html). Not only does it vary from day to day but it also varies over the solar cycle (greater during periods of higher sun spot activity). For our purposes, we will assume it is constant.

[iv] The surface area of a sphere is $4\pi R^2$.

> ✍ P_{sun} represents the **luminosity** of the sun, which is the rate at which energy is emitted by the sun. If you plug in the solar constant for S_0 and the radius of Earth's orbit for $R_{\text{Earth orbit}}$ (see appendix B.4), you'll find that the sun emits 3.846×10^{26} joules every second (see problem 13.1).

Now suppose we enclose the sun with a much smaller, but still pretty big, invisible balloon that has the same radius as the sun. Again, the total power passing through the balloon surface would be the product of the solar irradiance at that location $S_{\text{at sun's surface}}$ and the surface area of the balloon $4\pi R_{\text{sun}}^2$:

$$P_{\text{sun}} = S_{\text{at sun's surface}} 4\pi R_{\text{sun}}^2$$

Since energy is conserved, the power passing through the smaller balloon must be the same as that passing through the larger balloon.

$$S_0 4\pi R_{\text{Earth orbit}}^2 = S_{\text{at sun's surface}} 4\pi R_{\text{sun}}^2$$

Solving for the ratio of the two irradiances, we get

$$\frac{S_0}{S_{\text{at sun's surface}}} = \left(\frac{R_{\text{sun}}}{R_{\text{Earth orbit}}}\right)^2$$

Although the total energy (per time) is the same, the irradiances (the power per area) are different (since the radii are different). As the energy spreads out further into space, the irradiance goes down.

> ✍ Plugging in the solar constant for S_0 and the radii of Earth's orbit and the sun (see appendix B.4), you'll find that the solar irradiance at the sun's surface is 6.32×10^7 W/m^2 (see problem 13.1).

Check Point 13.1: How does the irradiance of the sun change as one moves away from the sun?

13.3 Emission

Whether an object warms due to radiation depends on whether it is absorbing more than it emits. We know that objects emit radiation. We just need to quantify how much.

For example, the air warms up during the day because it is absorbing radiation from the sun. However, the air is also emitting radiation. It warms up because it absorbs more than it emits.

On the other hand, at night the air cools down. This is because it is no longer absorbing radiation from the sun yet continues to emit radiation.

As mentioned in section 6.5.2, the amount of radiation emitted depends on the object's temperature. Warmer objects emit more radiation. The actual relationship is provided by the **Stefan-Boltzmann law**:

$$F = \sigma T^4 \tag{13.1}$$

where $\sigma = 5.67 \times 10^{-8}$ W·m^{-2}·K^{-4}.

DOES IT MATTER IF THE TEMPERATURE IS IN KELVIN OR DEGREES CELSIUS?

Yes. It must be in kelvin for the units to work out to be W/m^2.

> ✍ As pointed out before, the irradiance depends on how far you are from the object. Equation 13.1 gives the value of the object's irradiance at the object's surface. In fact, since we know the solar irradiance at the sun's surface (see previous section), we can use this expression to determine the sun's temperature (see problem 13.2).

Example 13.1: The surface of the Earth has an average temperature of 288 K. What is the irradiance emitted at the Earth's surface?

Answer 13.1: Use equation 13.1 with a temperature of 288 K.

$$\begin{aligned} F &= \sigma T^4 \\ &= (5.67 \times 10^{-8}\ \mathrm{W \cdot m^{-2} K^{-4}})(288\ \mathrm{K})^4 \end{aligned}$$

to get an irradiance of 390 W/m^2.

Check Point 13.2: According to the Stefan-Boltzmann law, which should be greater: the Earth's irradiance at its surface or the Sun's irradiance at its surface? Note: the Sun is a lot hotter than the Earth.

13.4 Equilibrium

So far, we've seen that the atmosphere is receiving radiation from the sun but, at the same time, since the atmosphere has a temperature, it is also emitting radiation. Depending on the temperature of the atmosphere, the emitted radiation might be more or less than the incident insolation and so the atmosphere might warm up or cool down.

In the previous section, the Earth was found to emit an irradiance equal to 390 W/m^2. Is that the same for the atmosphere?

No. On average, the atmosphere is cooler than the Earth's surface.

So what happens? Does the atmosphere emit more or less than it absorbs?

It varies, of course, but on the whole it emits the same as it absorbs.

How do you know?

Because if it emitted less, it would warm up and, as the temperature increases, the emitted radiation would likewise increase (see section 13.3) until the emitted radiation equaled the absorbed radiation. Conversely, if the atmosphere emitted more, it would cool down and, as the temperature decreases, the emitted radiation would likewise decrease until, again, the emitted radiation equaled the absorbed radiation.

So why keep track of all this then?

For two reasons. First of all, while the globe as a whole is in balance, at any given time one region might be gaining or losing energy.

For example, during the night, there is no solar irradiance and that part of the atmosphere cools down because, since it has a temperature, it emits radiation. When the sun comes up, there is absorbed radiation but it may still not be enough to counter the emitted radiation and so that part of the atmosphere may still cool for a short time. At some point, the absorbed radiation is greater than the emitted and that part of the atmosphere starts to warm up. Since the solar irradiance increases during the day (until a maximum at local noon), that region will continue to heat up.

At some point, typically around 3 o'clock in the afternoon, a balance is once again restored and that region stops warming. After that point, the region cools as the solar irradiance decreases.

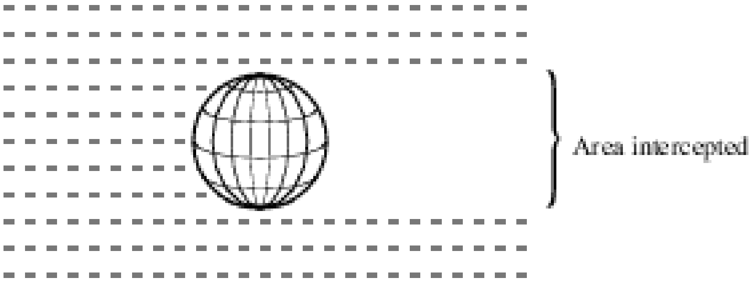

Figure 13.1: The portion of the incoming solar radiation that is intercepted by a sphere is equal to the cross-sectional area of the sphere (i.e., the area of a circle).

This daily cycle is explored in section 13.6 (and the project).

A second reason for keeping track of each is that we can use an assumption of global balance to determine an average global temperature. In other words, we assume that there is a balance between the absorbed energy due to the solar irradiance and the emitted energy due to its temperature. We'll explore this in this section.

Basically, all we need to do is determine the absorbed energy per time. From that, we get the emitted energy per time (since they must be equal). We can then divide the emitted energy per time by the total surface area of the Earth to get the emitted irradiance and then use the Stefan-Boltzmann law (equation 13.1) to find out what the temperature must be.

The tricky part is finding the absorbed energy per time.

It isn't the solar constant?

No.

The solar constant value of 1367.6 W/m^2 is the energy per time *per area.* We need to multiply the solar constant by the area.

What is the area?

The area corresponds to what is being intercepted. In the case of the Earth, that would be the cross-sectional area of the Earth. This is illustrated in figure 13.1.

For a sphere, the cross-sectional area is just the area of a circle. To find the total solar energy per time intercepted by the entire Earth, then, we need to multiply the solar constant S_0 by πR_e^2, where R_e is the radius of the Earth.

Check Point 13.3: (a) As seen from the sun, what is the cross-section of the Earth? (b) How does that compare to the total surface area of the Earth?

IS THE QUANTITY $S_0(\pi R_e^2)$ THE ABSORBED ENERGY PER TIME?

No. The quantity $S_0(\pi R_e^2)$ is just the *intercepted* energy per time. Some of that energy is reflected and not absorbed. The fact that we can see the Earth from space means that not all of the light from the sun was absorbed by the Earth.

In reality, about 30.6% (0.306) of the incident energy is reflected. This fraction is known as the **albedo**, a.

Check Point 13.4: What fraction of the energy incident upon the atmosphere is actually absorbed by the atmosphere and Earth on average?

Thus, only $(1 - a)$ or 69.4% of the incident energy is absorbed. Mathematically, that means we must multiply the intercepted energy per time by $(1 - a)$.

For equilibrium, this quantity must equal the total energy per time that is being emitted. From the Stefan-Boltzmann law, the emitted irradiance is σT^4. To get the energy emitted per time, multiply the emitted irradiance by the surface area of the Earth $4\pi R_e^2$.

Setting these two quantities together, we get:

$$(1 - a)S_0(\pi R_e^2) = \sigma T^4(4\pi R_e^2)$$

Cancelling the πR_e^2 on both sides, this simplifies to

$$T = \left(\frac{(1 - a)S_0}{4\sigma}\right)^{1/4}$$

Plugging in the value of the solar constant, 1367.6 $\mathrm{W/m^2}$, and σ we get an equilibrium temperature of 254 K (i.e., −19°C).

Check Point 13.5: To obtain the equilibrium temperature of 254 K, which was assumed to be larger (or were they assumed to be equal): the energy emitted by the Earth or the energy absorbed from the sun?

13.5 Blackbody temperature

How can we say that the equilibrium temperature is 254 K? Isn't that a bit too cold?

The temperature in the Stefan-Boltzmann relationship (equation 13.1) provides the temperature that would produce the given irradiance.

The reason it seems cold is because it doesn't correspond to the surface temperature. As discussed in section 6.5.2, much of the radiation emitted from the surface does not get through the atmosphere. Instead, the radiation that gets emitted to space (and thus sensed by a satellite) comes from different layers – some from the surface and some from the clouds. In a sense, then, the temperature of 254 K represents an average throughout the atmosphere.

> ✍ If we know how the atmosphere absorbs radiation, we can use that to determine the surface temperature. We explore this in chapter 14.

The temperature obtained via the Stefan-Boltzmann law is called the **blackbody temperature**. It is the "actual" temperature of the object only if (a) there is nothing between the object and the observer to absorb any radiation and (b) F does not include any radiation that was simply reflected toward the observer and not radiated by the object. In our case, it corresponds to a "representative" temperature of the Earth/atmosphere system.

> ✍ The reason why it is called the blackbody temperature is because it assumes that all of the radiation is emitted at a single source (i.e., temperature) and no radiation is from other sources. That means that it must absorb all the radiation incident upon it from other sources, just like a black absorbs all visible light incident upon it. Such objects are called **blackbody** objects.

Check Point 13.6: Why is the Earth's blackbody temperature less than its *surface temperature?*

Can we do the same with the Sun?

Yes (see problem 13.2). However, as we found with the Earth, the temperature found in this way is an **effective temperature** rather than the temperature of any particular surface. It is called the "effective" temperature because the Sun's atmosphere can actually be hotter than its surface and so some of the radiation comes from that.

NASA says is the effective temperature of the Sun is 5778 K (see problem 13.2). This is pretty close to the surface temperature of the sun because the solar atmosphere, though much hotter, is also relatively thin. Thus, much of the radiation we received is emitted from the surface.

> ✍ The inner sun is much hotter than the surface but its radiation is absorbed before reaching the surface.

Check Point 13.7: Which is warmer: the Earth's blackbody temperature or its surface temperature?

13.5.1 At the poles

The 254 K temperature is the equilibrium temperature over the entire globe. We can use the same idea, though, to get an estimate for the equilibrium temperature at different locations on the globe if we assume there is no transfer of energy horizontally. This is not a good assumption, since energy tends to flow from the hotter areas, like along the equator, to colder areas, like near the poles. Still, it can be insightful to look at what the equilibrium temperature would be if there was no horizontal transfer.

We'll first look at the poles. The Earth is tilted 23.5° relative to its orbit around the sun. Consequently, at some point during the Earth's orbit, the north pole is directed away from the sun at an angle of 23.5°. This is the northern hemisphere's winter solstice. As the Earth spins upon its axis, the north pole remains pointed away from the sun. On the other hand, at some other point during the Earth's orbit, the north pole is directed toward the sun at an angle of 23.5°. This is the northern hemisphere's summer solstice. As the Earth spins upon its axis, the north pole remains pointed toward the sun.

Consequently, one advantage of considering the poles is that the sun's **elevation angle** (i.e., the height of the sun above the horizon) doesn't change significantly during the day. At the equinox, the sun is right at the horizon and stays that way the entire day. At the solstice, the sun is either below the horizon all day (winter solstice) or 23.5° above the horizon all day.

Let's look at what the equilibrium temperature should be in these situation. We'll consider two situations: the winter solstice and the summer solstice.

During the winter solstice there is no incoming solar radiation at the poles. With no incoming radiation there should be no outgoing radiation either (under equilibrium). Thus, the equilibrium temperature should be zero kelvin (-273°C). Of course, this is not the observed temperature because it doesn't stay this way the entire year. In addition, it ignores energy that would be transferred from warmer areas on the globe.

During the summer solstice the sun's elevation angle is 23.5° above the horizon the whole day. Since the surface is not oriented perpendicular to the solar radiation, we cannot approach this exactly the same way as we did the global equilibrium temperature. However, the general idea is the same. We want to have a balance between the intercepted solar radiation and the emitted radiation.

In this case, the intercepted energy per time is the solar constant, minus the albedo, multiplied by the area of the surface. Let's suppose the surface area is A. Since that surface is not oriented perpendicular to the incoming radiation, we can't use the entire area A. This is illustrated in figure 13.2.

If we define Z to be the **zenith angle**, the angle between the solar radiation and the normal (perpendicular) to the surface, then the amount intercepted is proportional to the cosine of Z (see figure). Consequently, the *intercepted* solar irradiance is equal to $S_0(1-a)$ times $A\cos Z$.

The emitted energy per time is the emitted irradiance σT^4 times the surface area A. Setting these two quantities together, we get:

$$(1-a)S_0 A\cos Z = \sigma T^4 A$$

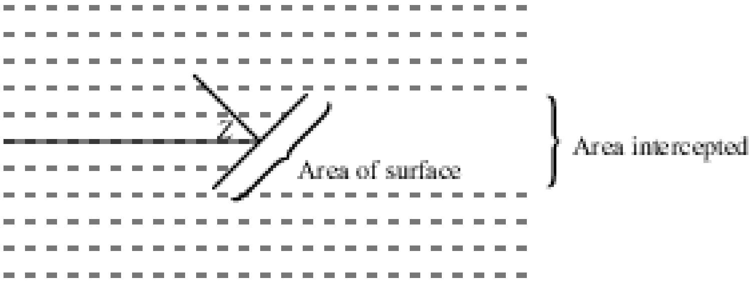

Figure 13.2: The portion of the incoming solar radiation that is intercepted depends upon the orientation of the surface relative to the incoming radiation (indicated by the zenith angle Z).

Cancelling the A's on both sides, this simplifies to

$$T = \left(\frac{(1-a)S_0 \cos Z}{\sigma} \right)^{1/4}$$

Plugging in the value of the solar constant, 1367.6 $\mathrm{W/m^2}$, and σ we get an equilibrium temperature of 286 K (i.e., 13°C).

This is pretty warm, considering it is at the pole and represents the effective temperature, not the surface temperature. Keep in mind that this represents the equilibrium temperature assuming it is summer solstice all the time. This is certainly not the case.

Check Point 13.8: During what part of the year (from when to when) is the incoming solar radiation zero at the north pole?

13.5.2 At the equator

At the poles, the average solar radiation over 24 hours is easy to calculate because the sun's elevation angle doesn't change during the course of the day. At most locations, however, the solar radiation is zero at night and then rises during the morning to a maximum at local noon and decreases back to zero at sundown. This makes determining an average a little more difficult.

To simplify the calculation, let's examine the situation for the equator during the equinox. At that time, the Earth is oriented as illustrated in figure 13.1 (on page 175) with the axis perpendicular to the direction of the solar radiation. We'll consider a strip of line of width W that runs along the equator. One can picture it as being a "belt" around the Earth, with a length equal to the circumference of the Earth $2\pi R_e$ and a width equal to W.

To find the radiation that is intercepted by this strip of land, we need to consider its cross-sectional area as "seen" by the incoming solar radiation. That area is just a rectangle of width W and length equal to the diameter of the Earth $2R_e$.[v] Consequently, the *intercepted* solar irradiance is equal to $S_0(1-a)$ times $W(2R_e)$.

The emitted energy per time is the emitted irradiance σT^4 times the surface area $W(2\pi R_e)$. Setting these two quantities together, we get:

$$(1-a)S_0 W(2R_e) = \sigma T^4 W(2\pi R_e)$$

Cancelling the $2R_e$ on both sides, this simplifies to

$$T = \left(\frac{(1-a)S_0}{\pi\sigma}\right)^{1/4}$$

Plugging in the value of the solar constant, 1367.6 W/m^2, and σ we get an equilibrium temperature of 270 K (i.e., -3°C).

Notice that this temperature is actually *less* than the equilibrium temperature for the pole during the summer solstice. Although the sun never gets very high at the pole, having it shine over 24 hours rather than 12 hours as at the equator makes a big difference.

Check Point 13.9: Does the 270 K determined for the equator correspond to the equilibrium temperature at the surface?

---

[v] This is equivalent to the shadow formed by that swath of land if we could cut away the rest of the Earth. It can also be derived by using calculus.

13.6 Diurnal temperature cycle

WHAT IF WE DO NOT HAVE RADIATION EQUILIBRIUM?

If there is more radiation being absorbed than emitted, or visa-versa, then the temperature will change. How much it changes depends on the radiation imbalance and what is absorbing the radiation.

Our interest, of course, is on the atmosphere. So, it makes sense to ask how the atmosphere will warm or cool based on a radiation imbalance.

It turns out that an expression for the temperature tendency due to a radiation imbalance can be derived from the first law of thermodynamics. The actual derivation is provided in appendix (see appendix C.4). For a irradiance imbalance F_{net} and a depth of atmosphere with pressure difference ΔP (between top to bottom), the temperature tendency is as follows:

$$\frac{dT}{dt} = \frac{gF}{c_{\text{p}}\Delta P} \tag{13.2}$$

The larger the layer, the more mass that needs heat up and the smaller the overall increase in temperature.

Example 13.2: Consider a column extending throughout the atmosphere at a location on the equator at noon on the equinox (such that the surface is perpendicular to the orientation of the sun's radiation). Determine the change in temperature experienced by that column of air during one minute assuming that all of the solar energy impinging upon it goes toward heating it up with no emission of radiation.

Answer 13.2: Using 9.8 N/kg for g, 1367.6 W/m^2 for F_{net} (since there is emission and the zenith angle is zero), 1004.67 J kg^{-1} K^{-1} for c_{p} and 1.01325×10^5 Pa for ΔP (i.e., the pressure difference from the bottom to the top of the atmosphere), one gets

$$\begin{aligned} \frac{dT}{dt} &= (gF_{\text{net}})/(c_{\text{p}}\Delta P) \\ &= 1.32 \times 10^{-4} \text{ K/s}. \end{aligned}$$

Multiply by the time (60 s) to get 0.0079 K. This is how much the temperature of that column would increase in one minute if there was no emission of radiation (which there is).

Check Point 13.10: Why is the temperature tendency inversely related to ΔP?

In the project, you will simulate the diurnal variation of temperature, warming up during the day and cooling down during the night, by comparing the amount of radiation that is absorbed each minute with the amount of radiation that is emitted each minute.

We'll do this for the lowest 10 mb of the atmosphere. We'll assume that there is no radiation absorbed during the night (from 6 pm to 6 am) and we'll assume the radiation during the day goes from zero at sunrise to a maximum of 547.8 W/m^2 at noon and then back to zero at sunset.

Why 547.8 W/m^2?

Remember that the solar constant (1367.6 W/m^2) is the solar irradiance incident upon the outside the atmosphere. By the time the radiation reaches the Earth's surface, we need to subtract out the reflected amount (i.e., due to the Earth's albedo) and the amount absorbed by the atmosphere.

As mentioned above, about 30.6% is reflected. It turns out that an additional 18% or 19% is absorbed by the atmosphere. So, only about 51% of the original radiation is actually absorbed by the Earth. Instead of the original 1367.6 W/m^2, we only have 697 W/m^2 (i.e., 51% of 1367.6 W/m^2).

> ✍ This is the maximum (at noon), not the average during the entire day. Much of this radiation gets absorbed by the Earth's surface, not the air above it

In reality, the lowest 10 mb of the atmosphere doesn't absorb radiation all of the radiation that makes it through the atmosphere. Much of this is absorbed by the ground, which warms up. However, since the ground then warms the air above it, we'll assume that all of the energy is absorbed by the 10-mb layer.

The air near the surface also absorbs radiation from the atmosphere above it (which radiates just like everything else). We'll assume this provides an additional 215.4 W/m^2. In reality, this wouldn't be constant but it helps simplify the calculations.

As for the emission, that will depend on the temperature and will be equal to σT^4, where T is whatever temperature that layer of air happens to be.

We'll start with the temperature begin 282 K at midnight (9°C). You can experiment with different values if you'd like. If we choose too low a value, the air will warm up over the course of 24 hours. If we choose too high a value, the air will cool down over the course of 24 hours.

Unless the starting temperature is really cold, however, the air should cool at midnight since there is no solar irradiance.

To calculate the change in temperature, first determine the solar irradiance, add to it the irradiance from the atmosphere and then subtract out what is emitted. Use the net irradiance with equation 13.2 to figure out the change in temperature. Then add that change to the temperature to determine the new temperature.

You'll find that, on average, the incoming radiation balances the outgoing radiation and, on average, the surface temperature is something reasonable. However, during the night the outgoing radiation is greater and so the temperature goes down. You should also find that the maximum temperature is not reached at noon but rather at some time in the afternoon. At noon, the incoming radiation is still larger than the outgoing radiation and so the temperature continues to go up.

Check Point 13.11: Should the temperature continue to rise throughout the day since there is radiation being absorbed the entire time?

13.7 Insolation

The intercepted solar irradiance is called the **insolation**.[vi]

[vi] The word "insolation" is obtained by combining parts of the words **in**-*coming* **sol**-*ar* *radi*-**ation**. Occasionally you may hear someone (even me) say "solar insolation," but technically that would be redundant.

Thus, we could say that the insolation outside of the Earth's atmosphere is S_0 or 1367.6 W/m^2. However, we can also use it to express what the irradiance is after a portion is reflected. In other words, we can say that the insolation absorbed by the Earth/atmosphere is $(1-a)S_0$.

Going further, we can say that the insolation incident on a surface whose normal is oriented at an angle Z relative to the radiation direction is $S_0 \cos Z$ or that the average insolation absorbed over the entire globe is $S_0/4$.

How was the last one determined?

If you go back to the discussion where we identified the effective temperature of the entire Earth (see page 176), we found that the average irradiance over the entire globe is one-quarter of what it would be otherwise because it considers the entire globe, even the half that is in the dark. This is where the "4" came from in the following expression:

$$T = \left(\frac{(1-a)S_0}{4\sigma} \right)^{1/4}$$

which can be rearranged as follows:

$$T = \left(\frac{S_0}{4} \frac{(1-a)}{\sigma} \right)^{1/4}$$

in order to make it more clear that the irradiance is one-quarter of what it would be otherwise.

> ✍ By the same reasoning, we can say that the average insolation incident at the equator during the equinox is S_0/π (see page 181).

In each case, if we want the *absorbed* insolation rather than just the *incident* insolation, we multiply by $(1-a)$.

Check Point 13.12: Why is the average insolation over the entire globe less than the solar constant value of 1367.6 W/m^2?

Project

In this project, you will simulate the evolution of the temperature throughout the day, warming when the sun comes up and cooling when the sun goes down.

Project 13.1: Create a spreadsheet.
(a) Title the first column "Time" (A1) and set the first cell (A2) equal to "12:00 AM" (without the quotes). Set the format for the entire column as "hh:mm AM/PM" (click on or highlight the entire column and then choose Format/Cells, click on the Number tab, and choose the Time category and the "1:30 PM" type).
(b) Title the second column "Temp" (B1) and set the first cell (B2) as 282 (Kelvin). This column will contain the temperatures at various times throughout the day.
(c) Title the third column "Fsolar" (C1) to represent the radiation from the sun absorbed at the Earth's surface. We want to set the solar radiation to zero at night and then follow a sinusoidal pattern (peaking at 547.8 W/m^2 *at noon) during the day. So, for the first cell (C2), use an "IF" statement where the value is 0 if the hour is less than 6 or greater than 17, and otherwise the value is*

$$547.8 \sin[(H - 6 + M/60) * \pi/12]$$

where H *is the hour and* M *is the minute. To get the hour and minute of cell A2 use "HOUR(A2)" and "MINUTE(A2)". Use "PI()" to get the value of* π*.*
(d) Title the fourth column "Fatmos" (D1) to represent the radiation from the atmosphere absorbed at the Earth's surface. We will assume this is constant at 215.4 W/m^2 *throughout the day so put 215.4 into the first cell (D2).*
(e) Title the fifth column "Fin" (E1) to represent the total radiation absorbed at the Earth's surface. Set the first cell (E2) equal to the sum of the solar and atmospheric contributions (C2+D2).
(f) Title the sixth column "Fout" (F1) to represent the radiation emitted by the Earth's surface. This can be calculated via the Stefan-Boltzmann law (equation 13.1):

$$F_{\mathrm{out}} = \sigma T^4$$

where $\sigma = 5.67 \times 10^{-8}$ $W{\cdot}m^{-2}{\cdot}K^{-4}$*. Use the temperature in column B for* T*.*
(g) Title the seventh column "Delta T" (G1) to represent the temperature change that results from the imbalance between "Fin" and "Fout". In the first cell (G2) use equation 13.2

$$\frac{dT}{dt} = \frac{gF_{\mathrm{net}}}{c_{\mathrm{p}}\Delta P}$$

to calculate the change in temperature. Set the time increment to be 1 minute (remember to convert to seconds), use 9.806 N/kg *for* g, $1004.67\ \mathrm{J \cdot kg^{-1}K^{-1}}$ *for* c_{p}*, use the difference between "Fin" and "Fout" for* F_{net} *(subtract "Fout" from "Fin" so that we get warming if "Fin" is greater) and* 10 mb *for* ΔP *(remember to convert to Pascals; this represents the layer of the atmosphere near the surface of the Earth). If you did everything correctly, you should get a temperature change of about* -0.084 *K in the cell (G2).*

Project 13.2: Now that you have the first row of data, you will calculate the temperature for each minute during the day, using the temperature from the previous minute and the change in temperature due to the difference in radiation calculated in the previous row.
*(a) Set the next cell (A3) in the "Time" column as the previous cell (A2) plus 1 minute. To do this, add "1/(24*60)" to the value of the A2 cell. You divide by 24 times 60 because the time is in units of days.*
(b) Set the next cell (B3) in the "Temp" column as the previous cell (B2) plus the change in temperature calculated previously (G2).
(c) All of the remaining cells in columns C though G can be copied from the previous row. If you did everything correctly, you should again get a temperature change of about -0.084 *K in cell G3.*

Project 13.3: Copy the third row as many times as necessary to get a record for every minute of the day up to midnight.
(a) What is the average temperature (i.e., average of column B)? Is this reasonable (i.e., based upon your experience)? Why or why not?
(b) At what time is the temperature a minimum? Is this reasonable (i.e., based upon your experience)? Why or why not?
(c) At what time is the temperature a maximum? Is this reasonable (i.e., based upon your experience)? Why or why not?

Project 13.4: Create a graph showing the variation in "Fin" and "Fout" during the day (highlight columns A, E and F, and then create a "Line" chart).

Problems

Problem 13.1: (a) From the measurement of the solar constant, determine the total power emitted by the sun. Note: we are about 1.496×10^{11} *m away*

from the sun and the surface area of a sphere is $4\pi R^2$.
(b) From (a), calculate the solar irradiance at a point on the sun's surface. Note: the radius of the sun is about 6.96×10^8 *m.*

Problem 13.2: In problem 13.1, you calculated the solar irradiance at a point on the sun's surface. Use (13.1) to calculate the temperature of the sun. Compare your answer to that provided by a reference.[vii]

Problem 13.3: The solar constant for Venus, being closer to the sun than the Earth, is 2613.9 W/m^2 *but, being covered in clouds, has an albedo of 0.75.*
(a) Show that the blackbody temperature of Venus is 231.7 K.
(b) Why is this so much colder than the surface temperature of Venus[viii] *(i.e., 737 K)?*

Problem 13.4: Calculate the intercepted insolation at noon on the equator for the (a) vernal equinox (sun directly overhead) and (b) summer solstice (sun 23.45° from directly overhead).

Problem 13.5: Calculate the daily average intercepted insolation for the following locations and times.
(a) the equator during the vernal equinox (see section 13.5.2),
(b) the north pole during the winter solstice (hint: where is the sun during the winter solstice at the north pole?)
(c) the north pole during the summer solstice (hint: at the north pole during the summer solstice, the sun is always 23.45° above the horizon)

Problem 13.6: (a) Suppose the air had a temperature of 273 K. What is the irradiance emitted by the air? (see Stefan-Boltzmann law, equation 13.1)
(b) How does that compare to the average absorbed insolation, $(1-a)S_0/4$*?*

Problem 13.7: Why does the minimum daily temperature occur a little bit after sunrise, not right at sunrise or just before sunrise?

[vii] See, for example, the NASA Fact Sheet on the sun at http://nssdc.gsfc.nasa.gov/planetary/factsheet/sunfact.html.

[viii] See http://nssdc.gsfc.nasa.gov/planetary/factsheet/venusfact.html (NASA Fact Sheet on Venus).

14. Greenhouse Effect

14.1 Introduction

In the previous chapter, we examined how objects can warm up by absorbing more radiation than they emit. For an object to be in equilibrium (i.e., neither warm up nor cool down), there must be a balance between the emission and the absorption.

We found that, when in radiation balance, the atmosphere will have a blackbody temperature of 254 K. This is colder than the average surface temperature on Earth (288 K) because most of the radiation emitted to space comes from the atmosphere, not the Earth, and the atmosphere is colder than the Earth.

It turns out that the mechanism behind this temperature contrast (i.e., the surface being warmer than the atmosphere) relies on the fact that the atmosphere is a better absorber of terrestrial radiation (i.e., radiation emitted by the Earth) than solar radiation (i.e., radiation emitted by the sun).

As you are probably aware, the atmosphere is mostly **transparent** to visible wavelengths of light (that is why, for example, we can see through it). What you may not be familiar with is that the atmosphere is mostly **opaque** to particular frequencies outside of the visible range. When I say "opaque", I mean that radiation emitted at those wavelengths will not be able to propagate through the entire depth of the atmosphere without getting absorbed.[i]

In other words, how much the atmosphere absorbs depends upon the frequency of the radiation. Thus, to explain why this leads to a surface temperature that is higher than the equilibrium temperature, we need to first explore how the solar and terrestrial irradiance depends on frequency (in section 14.3). In section 14.5, we'll explore how this frequency dependence leads to the warmer surface temperature, a phenomena sometimes referred to as the greenhouse effect.

---

[i] For short distances, like a couple hundred meters or so, air is still relatively transparent to most wavelengths.

14.2 Types of radiation

As mentioned in section 6.4.3, electromagnetic waves can have different frequencies. Each frequency corresponds to a particular wavelength λ, according to the wave equation (equation 6.1):

$$v = f\lambda$$

For our purposes, we will be interested in the frequency range from about 10^{12} Hz to about 10^{16} Hz because this encompasses most of the radiation emitted by the sun. This range is broken down into three regions: **ultraviolet**, **visible**, and **infrared**.

14.2.1 Visible light

The human body is not built to detect electromagnetic waves except for a narrow range from about 5×10^{14} to 10×10^{14} Hz. In a vacuum, this corresponds to wavelengths from about 300 nm to 700 nm. We detect these waves through our eyes and so we call this range **visible light**. The shorter wavelengths are interpreted as blue and violet colors and the longer wavelengths are interpreted as red and orange colors.

WHAT WAVELENGTH IS WHITE LIGHT?

White light is really the combination of all of the colors.[ii]

Check Point 14.1: Which color is associated with a higher frequency of electromagnetic waves: blue or red?

14.2.2 Infrared

We use the term **infrared** for wavelengths longer than what the eye can detect (i.e., frequencies lower than the eye can detect). In other words, these

---

[ii] Our eyes consist of three different types of cones, each able to sense a different frequency range. The frequency ranges roughly correspond to the red, green and blue regions. When all three types of cones are excited, we see white. Consequently, we can "trick" our eyes into see white by shining red, green and blue light together. Indeed, your computer or television screen takes advantage of this and only uses red, blue and green light.

waves are “beyond red”.

The existence of infrared radiation was first detected by taking light from the sun and passing it through a prism. The light disperses into a band where smaller wavelengths (like blue) are at one end and longer wavelengths (like red) are at another. By taking a thermometer and measuring the temperature at each end, it was found that enhanced readings were present beyond the red, where there was no visible light.

As seen in the experiment, one property of infrared radiation is that it tends to increase the temperature of objects around us. As discussed in section 13.3, these objects also tend to emit infrared radiation.

14.2.3 Ultraviolet

We use the term **ultraviolet** for wavelengths shorter than what the eye can detect (i.e., frequencies higher than the eye can detect). In other words, these waves are “beyond violet”. Like infrared, this region is invisible to the eye.

Ultraviolet radiation can be harmful. The part of light that causes tanning and sunburn, for example, is the ultraviolet part of the sunlight.

Check Point 14.2: Which type of radiation is associated with a higher frequency of electromagnetic waves: ultraviolet or infrared?

14.3 Spectral irradiances

If both the sun and the atmosphere emit radiation, how come I can only see the radiation emitted from the sun?

Because the two are different. Radiation from the sun (which we call **solar radiation**) is mainly in the visible portion of the spectrum whereas radiation from the atmosphere (which we call **terrestrial radiation**) is mainly infrared.

Why are the frequencies different?

Because the frequency of the emitted radiation depends upon the emitter's temperature. The sun, being much hotter, emits radiation with not only a larger irradiance (see Stefan-Boltzmann equation in section 13.3) but also extends the range into the higher frequencies (compared to the radiation emitted by the Earth).

The **peak emission** frequency or wavelength can be predicted from **Wien's law**:

$$\lambda_{\text{peak}} = \frac{2897\mu\text{m} \cdot \text{K}}{T} \tag{14.1}$$

The peak emission wavelength depends on the temperature of the object. The hotter the object, the shorter the peak emission wavelength (i.e., the higher the peak emission frequency).

As you will demonstrate in problem 14.1, the peak emission wavelength for the sun is around 500 nm, which corresponds to yellow-green light.

Doesn't the sun also emit ultraviolet and infrared?

Yes. As mentioned in section 6.4.3, solar radiation is not at one frequency but rather covers a range of frequencies, the bulk of which covers the ultraviolet, visible and infrared parts of the electromagnetic spectrum.

The actual frequency distribution of solar radiation is shown as the black curve in figure 14.1 (values correspond to observations just outside the atmosphere). As you can see, there is a peak in the distribution, which as predicted by Wien's law is in the visible part of the spectrum.

What is being plotted? Is it irradiance?

No. It is the **spectral irradiance**.

What is the spectral irradiance?

To explain what it is, let's first look at how the plot in figure 14.1 was obtained. To measure the variation with frequency, they measure the irradiance for a very small range of wavelengths (say, those between 500 nm and 501 nm). The smaller the range, the less irradiance that is received. Consequently, they can't plot the irradiance because then it depends on how small the range is.

Instead, they plot the **spectral irradiance**, which is the irradiance divided by the range used (e.g., the range would be 1 nm, if the irradiance is measured

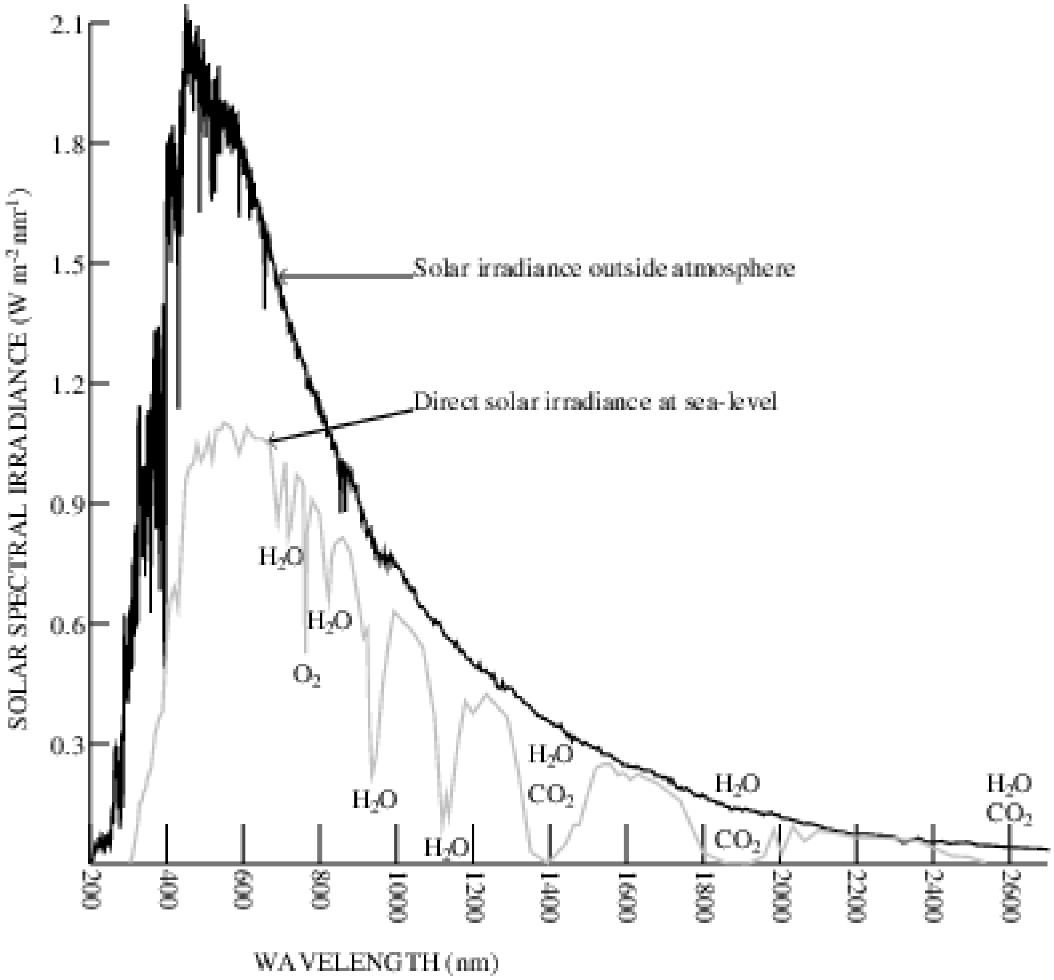

Figure 14.1: Solar spectral irradiance incident on the top of the atmosphere (black curve) and transmitted through the atmosphere to the Earth's surface (gray curve), which principal gases responsible for atmospheric absorption noted. Data obtained from the National Renewable Energy Laboratory (rredc.nrel.gov/solar).

in 1-nm bins). Consequently, the result has units of W/m^2 per nm (or, W/m^2 per μm, whichever is more convenient).

For our purposes, the symbol F_λ will be used to represent the spectral irradiance. Just remember that F_λ has different units than F.

Is there an equation to predict this distribution?

Yes. For a blackbody, the spectral irradiance is given by **Planck's law**:

$$F_\lambda = \frac{c_1}{\lambda^5 \left[\exp\left(\frac{c_2}{\lambda T}\right) - 1\right]} \tag{14.2}$$

where $c_1 = 3.74 \times 10^8 \text{ W} \cdot \text{m}^{-2} \cdot \mu\text{m}^4$ and $c_2 = 1.44 \times 10^4 \ \mu\text{m} \cdot K$ (λ is assumed to be units of μm). This expression was used in the project to calculate the spectral irradiance of the sun and the Earth.

Does the spectral irradiance, like the irradiance, depend on how far you are from the object?

Yes. The spectral irradiance in equation 14.2 corresponds to the spectral irradiance observed at the object's surface.

Thus, if you want to determine the solar spectral irradiance at the sun's surface, you use the surface temperature of the sun, as you did in project 14.1. Since the irradiance decreases as the radiation propagates away from the sun and spreads out in space, the solar spectral irradiance by the time it reaches the Earth will be much smaller (see project 14.2).

Why doesn't the solar spectral irradiance at the Earth's orbit, as calculated in project 14.2, equal the observed solar spectral irradiance in figure 14.1 (black line)?

The values do not exactly match the observed values because the sun is not a perfect blackbody.[iii] Still, the total irradiance, which you can obtain by integrating the spectral irradiance (as in project 14.3), is similar to the observed **solar constant**.

If the sun emits infrared as well as visible, does the atmosphere emit visible as well as infrared?

[iii] Remember that to properly compare the two the units must match. You will have to convert from W m^{-2} μm^{-1} to W m^{-2} nm^{-1}.

No. The atmosphere is much colder than the sun and, as such, does not emit radiation with as high a frequency as the sun. The atmosphere has temperatures that are similar to that of the Earth, and the plot of Earth's spectral irradiances (see project 14.4) show that at those temperatures practically no radiation in the visible or ultraviolet is emitted.

> ✍ Although the sun emits more infrared radiation than the Earth or atmosphere (see projects), we are much further from the sun than the atmosphere. Consequently, for our purposes, the atmosphere (and Earth) are the major contributors of infrared. The sun, though, is still the major contributor of visible and ultraviolet. Indeed, in the atmosphere, there is a nearly complete absence of overlap between the wavelength range of solar irradiance and wavelength range of atmospheric irradiance. For this reason, it is common to refer to solar radiation as **shortwave** and atmospheric radiation as **longwave**.

Check Point 14.3: Which would be larger, the irradiance in W/m^2 per nm or the irradiance in W/m^2 per μm?

14.4 Atmospheric windows

Now that we've examined the dark black line in figure 14.1, let's look at the light gray line, which indicates the solar irradiance received at the surface of the earth.

First of all, we notice that the gray curve isn't the same as the black curve. This is because some of the solar irradiance is reflected or absorbed by the atmosphere. When we compare the peak at 500 to 600 nm, we find that the solar spectral irradiance starts out at about 1.85 W/m^2 per nm outside the atmosphere but is only 1.06 W/m^2 per nm by the time it reaches the Earth's surface.[iv] This means that only 58% or so of the solar radiation reaches the Earth's surface (divide 1.06 by 1.85).

Since about 30% of the solar radiation is reflected (due to the Earth's albedo) and since some of the albedo is due to reflection off the Earth's surface, you

---

[iv] Values are obtained by taking the average spectral irradiance between 500 and 600 nm.

might expect that at least 70% of the solar radiation should reach the Earth's surface, not 58%.

The difference is due to absorption by the atmosphere. As mentioned in section 13.6, the Earth only absorbs about 51% of the solar radiation that is incident outside the atmosphere. Subtracting out the portion that is reflected (**albedo**), we get that the atmosphere absorbs about 18 to 19 percent of the solar radiation. This is why the radiation received at the Earth's surface is less than we would expect due to reflection alone.[v]

A closer examination of the gray curve reveals that things are a little more complicated, however. The amount of radiation that reaches the surface depends on frequency.

For example, below 300 nm or so, no solar radiation reaches the Earth's surface. This is partially due to the sun not emitting much at those frequencies but it is also because the atmosphere is very effective at absorbing that radiation. This absorption is due principally to the presence of ozone (O_3) between the stratosphere and mesosphere. It is ozone that prevents much of the ultraviolet light from reaching the Earth's surface. That is a good thing, as ultraviolet light can cause sunburn and skin cancer.[vi]

The ultraviolet region, however, is not the only region of enhanced atmospheric absorption. At 1400 nm, 1900 nm and 2600 nm, for example, hardly any solar radiation passes through the atmosphere. These wavelengths are in the infrared region. Unlike the absorption in the ultraviolet, which is due to ozone, the absorption in the infrared is due mainly to water vapor and carbon dioxide.[vii]

---

[v] It is worth noting that this implies that most of the Earth's albedo is due to atmospheric reflection, not reflection of the surface of the Earth. In fact, using the measured irradiances, 7% of the solar irradiance is reflected from the Earth's surface, 23% is reflected from the atmosphere and 19% is absorbed by the atmosphere. According to NASA, however, the atmosphere reflects 26% (20% of which is due to clouds) and the Earth's only reflects 4%. The difference may be due to the way we restricted our analysis to only the radiation between 500 and 600 nm.

[vi] This is the concern scientists have about **ozone depletion** (e.g., due to **chlorofluorocarbons**, **CFC**'s). Fortunately, we have more than enough ozone at mid-latitudes to absorb the ultraviolet radiation. However, above the south pole, the circulation is such that ozone depletion is enhanced, producing periodic areas of lower ozone concentrations (every spring).

[vii] Depending on the structure of the molecule, certain frequencies can be absorbed while other frequencies are transmitted. That is why the absorption is restricted to only certain

Between these locations, there are regions where almost all of the solar radiation gets through. Such regions are called **atmospheric windows**.[viii]

Check Point 14.4: What would be different if the atmospheric window was mainly in the ultraviolet instead of the visible light range?

14.5 The greenhouse effect

As one can see by figure 14.1, the absorption by water vapor and carbon dioxide is not significant for solar radiation because most of the solar radiation is in the visible bands, with significant drop-off in the infrared. However, as seen in project 14.4, terrestrial radiation is entirely infrared and longer wavelengths. And, though it is not shown in figure 14.1, atmospheric absorption is also significant in the infrared wavelengths greater than 3000 nm, with 100% absorption around 4200-4400 nm (mainly due to CO_2), 5500-7300 nm (mainly due to H_2O) and 14000-16000 nm (mainly due to CO_2), with atmospheric windows in between.[ix]

As such, a greater percentage of **terrestrial radiation** (compared to **solar radiation**) is absorbed by the atmosphere.[x] This imbalance leads to the surface being warmer than the atmosphere.

To explore how this imbalance leads to a surface (288 K in our case) that is warmer than the atmosphere (254 K in our case), compare the two models in figure 14.2.

The top model (which I'm calling the zero-layer model) was used to determine the black body temperature of the atmosphere. Radiation is incident from

---

ranges of frequencies.

[viii] There is a significant atmospheric window within the visible portion of the solar radiation. This is pretty nice – if air wasn't transparent to visible light we wouldn't be able to see through it.

[ix] There is significant, but not total absorption around 9500 nm due to ozone (O_3).

[x] Since water vapor is such a good absorber of infrared radiation, most of the infrared radiation emitted by the Earth doesn't get through clouds. Consequently, the infrared radiation that is emitted typically comes from the region at the cloud top, as discussed in section 6.5.2.

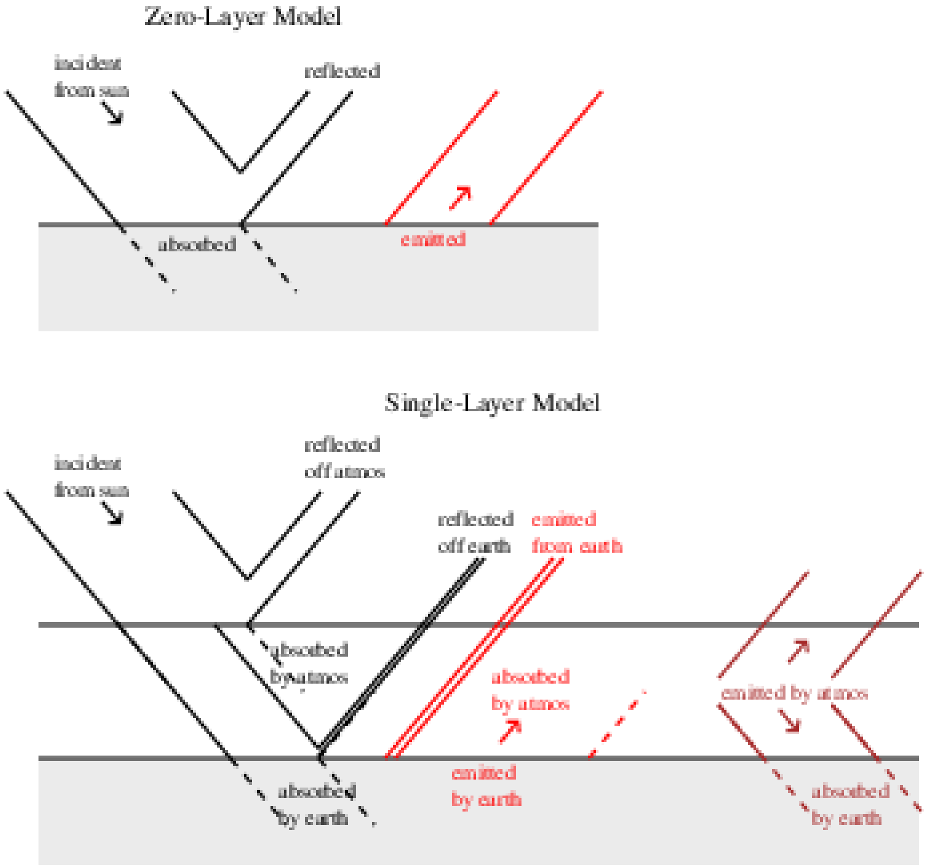

Figure 14.2: (top) Earth's radiation budget [top] considering the Earth/atmosphere as a single entity (with a single temperature) and [bottom] considering the atmosphere as being a single layer (with a single temperature) separate from the Earth. In both cases, the incident radiation (from the sun) is represented by the beam on the left. At each interface, some radiation may be reflected, absorbed or transmitted. Thickness of each beam is scaled to the corresponding amount of radiation. Source of relative amounts: NASA Langley Research Center (CERES: Clouds and the Earth's Radiant Energy System).

the sun and whatever doesn't reflect off the Earth/atmosphere gets absorbed. This absorption is balanced by the emitted radiation (see right side of top figure), which simply goes off into space.

By balancing the absorbed radiation with the emitted radiation, we get a black body temperature of 254 K.

A more accurate picture is represented by the bottom figure (which I'm calling the single-layer model). The input and output to the entire Earth/atmosphere system is the same. However, there are some significant differences with the zero-layer model.

First, let's look at the incident radiation from the sun. Here we see that part of the incident radiation is reflected off the atmosphere and part is reflected off the Earth.[xi] Similarly, some of the incident radiation is absorbed by the atmosphere and some is absorbed by the Earth.

Comparing the two models, the Earth absorbs less of the incident solar radiation in the single-layer model than in the zero-layer model.

However, this is not the only thing going on. The atmosphere also emits radiation (see far right of figure). Thus, the Earth also gets radiation from the atmosphere.

HOW DOES THE TOTAL RADIATION RECEIVED FROM BOTH THE SUN AND THE ATMOSPHERE COMPARE TO WHAT THE EARTH WOULD RECEIVE IF THERE WAS NO ATMOSPHERE?

The total could be greater than or less than what the Earth would receive if there was no atmosphere. It depends on how much the atmosphere absorbs of the terrestrial radiation (see center of figure). The more it absorbs, the warmer it will get and, subsequently, the more radiation it will emit back to the Earth.

In our case, the atmosphere absorbs a large portion of the terrestrial radiation and, as such, the Earth is warmer than it otherwise would be.

This effect (of the atmosphere leading to warmer ground temperatures than would otherwise be) is commonly referred to as the **greenhouse effect** since a greenhouse keeps the interior warmer than it would otherwise be. Note, however, that a greenhouse keeps the interior warmer due to a decrease in

---

[xi] There are additional cascading reflections off the atmosphere and Earth again but these aren't shown since they are considered to be small.

wind advection (much like your sweater) more so than an absorption of outgoing infrared radiation.

Check Point 14.5: According to figure 14.2, which provides more radiation to the Earth's surface: solar radiation that passes through the atmosphere or radiation emitted by the atmosphere?

Project

In this project you will explore how solar radiation differs from terrestrial radiation in terms of the frequency.

Project 14.1: Create a spreadsheet. In the first column, create a list of wavelengths 0.1 μm to 100 μm in 0.01 μm increments (almost 10000 points). Then, in a second column, calculate the solar spectral irradiance (at the Sun's surface) for each wavelength (first column) using equation 14.2:

$$F_\lambda = \frac{c_1}{\lambda^5 \left[\exp\left(\frac{c_2}{\lambda T}\right) - 1\right]}$$

where

$$\begin{aligned} \lambda &= \text{wavelength (in } \mu\text{m)} \\ T &= \text{solar blackbody temperature (see problem 13.2)} \\ c_1 &= 3.74 \times 10^8 \ \text{W} \cdot \text{m}^{-2} \cdot \mu\text{m}^4 \\ c_2 &= 1.44 \times 10^4 \ \mu\text{m} \cdot K \end{aligned}$$

Project 14.2: The solar spectral irradiance decreases as one gets further from the Sun because it spreads out in three dimensions as it propagates away from the Sun.[xii] *Knowing this, create another column listing the solar spectral irradiance just outside the Earth's atmosphere.*

[xii] The surface area of a sphere is $4\pi R^2$, where R is the radius of the sphere. At the surface of the sun, the energy is spread over an area equal to $(4\pi R^2_{\text{sun}})$. By the time the energy has reached the orbit of the Earth, it is spread over an area equal to $(4\pi R^2_{\text{Earth's orbit}})$. This means that the solar spectral irradiance has decreased by the ratio of $(4\pi R^2_{\text{sun}})$ to $(4\pi R^2_{\text{Earth's orbit}})$.

Project 14.3: Create another column where you multiply the solar spectral irradiance (just outside the Earth's atmosphere) by the wavelength increment you used (which should be 0.01 μm). Show that if you add up all of numbers in that column, you get a number similar to the solar constant (see appendix B.4).

Project 14.4: In another column, calculate the terrestrial spectral irradiance at the Earth's surface using the Earth's surface temperature (use the U.S. Standard Atmosphere surface temperature) instead of the Sun's surface temperature.

Project 14.5: Create the following two plots:
(a) the terrestrial spectral irradiances (at the Earth's surface) from project 14.4, and
(b) the solar spectral irradiance (at the sun's surface) from project 14.1.
Are there any wavelengths for which the atmospheric spectral irradiance (at the Earth's surface) is greater than the solar spectral irradiance (at the sun's surface)? If so, which ones?

Project 14.6: Create the following two plots:
(a) the terrestrial spectral irradiances (at the Earth's surface) from project 14.4, and
(b) the solar spectral irradiance just outside the Earth's atmosphere from project 14.2.
Are there any wavelengths for which the atmospheric spectral irradiance (at the Earth's surface) is greater than the solar spectral irradiance (just outside the Earth's atmosphere)? If so, which ones?

Problems

Problem 14.1: (a) Calculate the peak emission wavelength for the sun assuming a temperature equal to that calculated in problem 13.2. In which part of the electromagnetic spectrum does it lie (e.g., infrared, visible or ultraviolet)?
(b) Calculate the peak emission wavelength for the atmosphere assuming a temperature equal to its average surface temperature (from the U.S. Standard Atmosphere). In which part of the electromagnetic spectrum does it lie?

Problem 14.2: Which (sun or atmosphere) contributes more to the infrared irradiance impinging upon the Earth?

Part D

Equilibrium and Newton's Second Law

15. Vertical Balance

Way back on page 8, a simple forecast method was described. The first step of that method involved a short prediction of temperature. The physics involved in that were discussed in part C.

The second step is to identify how the pressure changes based upon the new temperature structure. We'll examine that in this chapter.

The third step is to identify how the new pressure structure changes the winds. We'll examine that in the remaining chapters of this part of the book.

As discussed in chapter 10, the wind influences the temperature, and so the forecast cycle starts over again. It is in this way, with repeated very short forecasts, that a numerical weather forecast is made.

15.1 Hydrostatic balance

In this chapter, we'll look at how the pressure depends upon the temperature. In the process, we'll explain why the pressure decreases with height in the way described in chapter 4 (i.e., roughly exponential).

The reason why the pressure decreases with height has to do with vertical balance. By vertical balance I mean that the decrease in pressure (with height) must produce an upward force (from high to low pressure) that exactly balances the downward force due to the gravitational attraction with the earth.

From Newton's second law, if the forces are in balance, the object doesn't accelerate. It may be moving. It just won't be speeding up or slowing down.

Is this a reasonable assumption to make for the atmosphere?

In general, yes. There are, of course, times when the vertical forces are not in balance, such as in the middle of a thunderstorm. However, for the purposes of our simple model, it is a very good assumption, as setting those two forces equal results in an expression that looks a lot like the exponential expression in equation 4.2 (see section 3.6 on page 50).

To quantify the balance, then, we need to have an expression for each force acting on the air. In this case, we have an upward force due to the vertical pressure gradient and we have a downward force due to the Earth's gravity.

> ✍ Rather than deal with force, per se, we will instead deal with the force *per volume*. In this way, we don't need to worry about how big the volume of air happens to be (this is the same reason why we use density instead of mass). Even though I am dealing with a force per volume instead of a force, I will still tend to use F as the variable abbreviation. It has many of the same properties as force but the units are different.

The gravitational force per volume is

$$\vec{F}_{\mathrm{z,g}} = -\rho g \hat{k} \tag{15.1}$$

where $\hat{k}$ is upward, so the negative sign is used to indicate that the force is downward.[i] The variable g represents the gravitational field and is approximately 9.8 N/kg.

The air doesn't accelerate downward because, countering the gravitational force, there is an upward force due to the decrease with pressure with height. The force due to this vertical pressure gradient is

$$\vec{F}_{\mathrm{z,PGF}} = -\frac{\partial P}{\partial z}\hat{k} \tag{15.2}$$

We call this force the *pressure gradient force* (that is what "PGF" stands for).

---

[i] Since the earth is rotating, there is a small centrifugal force that makes the "downward" direction slightly equatorward of "toward the center of the earth." For our purposes, we are combining the gravitational force with the centrifugal force to get an apparent gravitational force. This apparent gravitational force is what leads to the "vertical" acceleration of an object that is in "free fall". We take g to represent this apparent vertical acceleration.

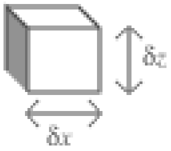

Figure 15.1: A representation of a box of very small volume.

> ✍ You can derive the expression in equation 15.2 by considering a tiny cube of air of dimensions δx, δy and δz as in Figure 15.1. One side has infinitesimal length δx, another side has infinitesimal length δy (not labeled in figure) and the third side has infinitesimal length δz. There are four vertically oriented faces (two of area $\delta x \times \delta z$ and two of area $\delta y \times \delta z$). There are two horizontally oriented faces (of area $\delta x \times \delta y$). When there is a vertical pressure gradient, the force per volume pushing up on the bottom ($P_{\text{bottom}} \delta x \delta y / \delta x \delta y \delta z$) is greater than the force pushing down on the top ($P_{\text{top}} \delta x \delta y / \delta x \delta y \delta z$) and the difference is $\delta P / \delta z$. We use the partial derivative ($\partial P / \partial z$) because we want the change at a particular fixed location (i.e., a box of infinitesimal size).

These two forces are opposite in direction (i.e., the gravitational force is directed downward while the pressure gradient force is directed upward). Consequently, they counteract.

I UNDERSTAND THAT THE GRAVITATIONAL FORCE IS ALWAYS DIRECTED DOWNWARD. IS THE PRESSURE GRADIENT FORCE ALWAYS DIRECTED UPWARD?

Yes. In a sense, the pressure gradient force is set up as a result of gravity. Gravity pulls the air downward. The air would pile up at the bottom except that such a "piling up" would cause a great pressure gradient force directed upward.

DO THESE TWO FORCES COUNTERACT EXACTLY?

For our purposes, we will assume that the two forces are exactly the same. It is left as an exercise (see problem 15.2) to show that these two forces are observed to be comparable in magnitude in general.

It is by assuming the two forces are exactly equal, however, that we can derive an expression for how the pressure must decrease with height. If the two forces counteract exactly then $F_{\mathrm{z,PGF}} + F_{\mathrm{z,g}}$ equals zero and, using equations 15.1 and 15.2, we get the following expression

$$\frac{\partial P}{\partial z} = -\rho g \tag{15.3}$$

which we will call the equation of **hydrostatic balance**.

Why is the right-hand side negative?

Because the pressure decreases with height.

If the air is not accelerating vertically, how can there be vertical motions?

Keep in mind that hydrostatic balance does not imply that the air cannot move up or down. Hydrostatic balance just means that the air is not *accelerating* up or down. In other words, the air can be moving upward or downward. It just isn't speeding up or slowing down.

How often is the atmosphere in hydrostatic balance?

As it turns out, for the most part the atmosphere is in hydrostatic balance. And, even when the atmosphere is not in hydrostatic balance, it turns out that it is usually pretty close. Indeed, the approximation is so good that it can be used to predict a structure of the atmosphere that is pretty close to the actual structure.

Check Point 15.1: If the atmosphere is in hydrostatic balance, does that mean that the pressure necessarily *has to decrease with height?*

15.2 Pressure decrease with height

Hydrostatic balance (equation 15.3) can be used to show that the pressure must decrease exponentially with height.

If we take the hydrostatic relationship (equation 15.3)

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

and replace ρ by P/RT (from the ideal gas law), we get that the change in pressure is proportional to the negative of its value (i.e., $dP \propto -P$). This implies exponential decay. Essentially, when the pressure is large (near the surface), it decreases rapidly with height. When the pressure is less (higher up), it decreases more slowly with height.

The actual relationship is as follows (see derivation in appendix C.5):

$$P(z_2) = P(z_1) \exp\left(-\frac{g}{R\bar{T}}\Delta z\right) \tag{15.4}$$

where $\bar{T}$ is the layer-mean temperature of the layer of depth $\Delta z = z_2 - z_1$ (note that $\exp x$ is equivalent to e^x). A version of this is used in project 15.1 to predict the pressure at some height z_2 given the pressure at another height z_1.

> ✍ Since there is no single value of ρ between $P(z_1)$ and $P(z_2)$, there is also no single value of the temperature. This is why the layer-mean temperature must be used.

Check Point 15.2: Why does the expression for hydrostatic balance have $\bar{T}$ instead of just T?

Note that the basic assumption of hydrostatic balance (along with the ideal gas law) forces the pressure to decrease exponentially with height. In this way, it is similar to equation 4.2:

$$P(z) = P(0)e^{-z/H}$$

where H represented the height at which the pressure falls to $1/e$ of its initial value (remember that H is called the **scale height**). The difference between that expression and the one we've derived here from hydrostatic balance is that hydrostatic balance expression has $R\bar{T}/g$ instead of H.

What this means is that in a hydrostatic atmosphere, scale height is proportional to the temperature as follows:

$$H = R\bar{T}/g \tag{15.5}$$

Since the temperature varies with height, this means that the scale height varies with height also.

Check Point 15.3: According to equation 15.5, if the temperature of a layer increases, what happens to the scale height?

15.3 Applications

With hydrostatic balance, we can do a lot of things we weren't able to do before. We'll examine several in this section.

15.3.1 Mass of atmosphere

WHAT DOES THE MASS OF THE ATMOSPHERE HAVE TO DO WITH HYDROSTATIC BALANCE?

It turns out we can use the hydrostatic balance equation to estimate the mass of the atmosphere. We'll start by providing the mass of the atmosphere, which is about 5.136×10^{18} kg (source: Handbook of Chemistry and Physics).

HOW IS THE MASS OF THE ATMOSPHERE OBTAINED?

As implied above, it isn't measured. Instead, it is inferred from the sea-level pressure and the hydrostatic assumption. From the hydrostatic equation one can get the following equation (see derivation in appendix C.6):

$$M = \frac{A}{g} P(0) \tag{15.6}$$

where A is the surface area, g is the gravitational acceleration and $P(0)$ is the pressure at the surface.

Thus, to get the total mass of the atmosphere, we estimate the surface pressure and then use the total surface area of the Earth (recall that the surface area of a sphere is $4\pi R^2$).

Check Point 15.4: According to equation 15.6, how is the surface pressure related to the mass of the atmosphere?

Although equation 15.6 uses the pressure at the surface, $P(0)$, we could use the equation to obtain the mass above any point as long as we are given the pressure at the point. In other words, the atmospheric pressure that you measure at any altitude is the weight per area of all the air molecules above you: $P = Mg/A$.

We can generalize further to obtain the mass of the atmosphere between any two levels. The pressure difference between any two levels in the atmosphere equals the weight per area of the air between the two levels:

$$\Delta P = \frac{Mg}{A}. \tag{15.7}$$

Check Point 15.5: According to equation 15.7, if the mass in a 1-km deep layer of the atmosphere increases and the pressure at the bottom remains the same, what happens to the pressure at the top?

15.3.2 Warm advection and pressure falls

In chapter 10, we saw how advection can cause the air temperature to change. We've now seen that the pressure depends upon the temperature. It follows, then, that advection can also lead to pressure changes, especially in surface pressure.

To explain how this occurs, let's go back to equation 15.4, which shows how the pressure decreases with height.

$$P(z_2) = P(z_1)\exp\left(-\frac{g}{R\bar{T}}\Delta z\right)$$

Let's rewrite this in the following way. Let's make $P(z_2)$ equal to the pressure at the tropopause and $P(z_1)$ equal to the pressure at the surface. Solving for the surface pressure, we get:

$$P_{\text{sfc}} = P_{\text{trop}}\exp\left(\frac{g}{R\bar{T}}z_{\text{trop}}\right) \tag{15.8}$$

Let's further assume that the height of the tropopause is fixed[ii] and that the pressure at the tropopause doesn't change (i.e., the mass above the tropopause doesn't change).

Given those assumptions, the only way the surface pressure can change is if the mean temperature $\bar{T}$ of the troposphere changes. If that mean temperature goes up then the surface pressure must go down (since $\bar{T}$ is in the denominator).

Let's see what this means in terms of the polar front model that was discussed in section 1.2.2. It was mentioned that the center of storms tend to be areas of low surface pressure, with the wind flowing counter-clockwise around the center in the northern hemisphere.

That means there will be warm advection east of the center (where the wind is coming from the equator) and cold advection west of the center (where the wind is coming from the poles). This increases $\bar{T}$ east of the center and decreases it west of the center.

Based on the expression above, that means the surface pressure decreases east of the center and increases west of the center, acting to move the system as a whole toward the east.

Check Point 15.6: If warm advection is occurring at a location, what is your short-term forecast for the surface pressure?

15.3.3 Sea-level pressure

We have seen that surface maps, such as that shown in figure 4.1 on page 48, show the sea-level pressure, not the actual pressure measured on the surface.[iii] The sea-level pressure that is plotted is not measured (i.e., they

---

[ii] This height isn't fixed. However, the stratosphere can be thought of as a "lid" of sorts (see chapter 17) and so the tropopause won't change significant over the time scale we are looking at.

[iii] As mentioned before, the reason why surface pressure isn't plotted is because we are interested in the horizontal pressure variations and the surface pressure distribution would reveal more about vertical pressure variations (i.e., low pressure at higher elevations) than it would about horizontal pressure variations. Thus, it is important to see the pressure values from the same height.

don't dig a hole down to sea-level and drop down a barometer). Rather, they use the hydrostatic balance relationship.

Basically, a "representative" temperature is used and, from that, the pressure that would exist at sea-level is calculated.[iv]

Check Point 15.7: On a surface map, is the plotted pressure typically greater than, equal to, or smaller than the observed surface pressure?

15.3.4 Thickness

In the hydrostatic relationship, Δz represents the depth of the layer for which the average temperature is $\bar{T}$. For this reason, Δz is known as the **thickness**[v]. As you can see from equation C.6, the height difference (or thickness) between two pressure levels depends upon the layer-mean temperature. The warmer the temperature, the larger the height difference (i.e., the thicker the layer).

In other words, the thickness can be used to get a sense of the layer-mean temperature. Warm layer-mean temperatures will be related to larger thicknesses.

By convention, the temperature of the air in the lower troposphere is gauged by the 1000-500 mb thickness (i.e., the difference in height between where the pressure is 1000 mb and where the pressure is 500 mb).

Check Point 15.8: If the atmosphere warms up and the pressure at the surface remains 1000 mb, what happens to the height where the pressure is 500 mb?

15.4 Vertical gradient in density

UNDER HYDROSTATIC BALANCE, THE PRESSURE MUST DECREASE WITH

[iv] The representative temperature is based somewhat on the air temperature at the surface.

[v] Don't confuse "thickness" of the layer with whether the air itself is "thick" or "thin." In this case, thickness refers to how deep the layer is rather than the density of the air.

HEIGHT (IN ORDER FOR THE PGF TO BALANCE GRAVITY). WHAT ABOUT THE DENSITY? MUST IT ALSO DECREASE WITH HEIGHT?

Most people expect the **density** to *decrease* with height. After all, the air gets thinner the higher one goes up a mountain. And, most of the time, the density does decrease with height.

However, under certain circumstances, the density can *increase* with height. To do so, however, the temperature has to decrease drastically with height. By the ideal gas law, if P decreases then so does the product ρRT. If ρ doesn't decrease then T must.

It isn't sufficient for T to decrease a little. It must decrease a lot. If fact, using the hydrostatic balance, we can show that temperature must decrease with height at a rate greater than g/R:

$$\frac{\partial \rho}{\partial z} > 0 \quad \text{if} \quad \frac{\partial T}{\partial z} < -\frac{g}{R}. \tag{15.9}$$

Since g/R is approximately 34 K/km, this means that the lapse rate has to be greater than 34 K/km (i.e., temperature gradient must be less than -34 K/km). That rarely happens, except perhaps during the summer when the air just above a black surface gets very hot compared to the air further up. So, normally the density will decrease with height.

HOW IS THIS RELATIONSHIP OBTAINED?

The hydrostatic relationship (equation 15.3) provides a link between height and pressure. To get a relationship between height and density, instead, we can convert the pressure in equation 15.3 to density using the ideal gas law (equation 7.2).

$$P = \rho RT$$

Taking the derivative of each side with respect to z, we get

$$\frac{\partial P}{\partial z} = RT\frac{\partial \rho}{\partial z} + \rho R\frac{\partial T}{\partial z}$$

From the hydrostatic relationship, the first term is just $-\rho g$. Solving for $\partial \rho/\partial z$, we get

$$\frac{\partial \rho}{\partial z} = -\frac{\rho}{T}\left(\frac{\partial T}{\partial z} + \frac{g}{R}\right) \tag{15.10}$$

Thus, the density will decrease with height (i.e., $\partial\rho/\partial z < 0$) as long as the lapse rate ($-\partial T/\partial z$) is less than g/R.

Check Point 15.9: If the atmosphere is in hydrostatic balance, does that mean that the density necessarily *has to decrease with height?*

Project

Project 15.1: For this project, you will create a graph of the pressure profile of the U.S. Standard Atmosphere from sea-level up to 86 km, but using a different technique than what was used in the previous project. In this project, you are to use equation 15.4 (see footnote[vi] *for assumptions) to determine the pressure at each level:*

$$P(z_2) = P(z_1)\exp\left(-\frac{g}{R\bar{T}}\Delta z\right)$$

where $\bar{T}$ *is the average temperature of the layer between* z_1 *and* z_2*,* R *is the gas constant, and* g *is the gravitational acceleration (see appendix B.4). See the footnote*[vii] *for a suggestion as to how to do this.*

Project 15.2: (a) How does your graph compare to the one you created in the previous project that assumes a strictly exponential function? Describe any differences.

[vi] As we will see in this chapter, this expression assumes the atmosphere is in vertical balance, meaning that the vertical pressure gradient is such that there is an upward force due to the pressure that exactly balances the weight of the air.

[vii] Since this expression depends upon the temperature, and the temperature varies with height, you need to do this in steps. Create a new column in the spreadsheet of the U.S. Standard Atmosphere used in the previous projects. In that column, for each level above the surface, estimate an average temperature of the layer between that level and the level below. You can use a straight average of the two temperatures (a better method would be to weight each temperature by $\ln P$ but we don't know the pressure). Then, at each level, calculate the pressure at that level, $P(z_2)$, using the average temperature of the layer just below it (between z_1 and z_2), the appropriate values of g and R, the difference in height between the two layers, Δz, and the pressure at the previous layer $P(z_1)$. At the surface, set the pressure to be 1013.25 mb.

(b) Would a different scale height make the match any better? If not, why not? If so, what scale height is the best match and why?

Project 15.3: (a) What is the pressure at the U.S. Standard Atmosphere tropopause (where the temperature no longer decreases with height)? Your answer should be about 226 mb.
(b) If the temperature of the air was warmer, what would happen to your answer in (a)? Try it out.
(c) Based on your answer to (a), where would the height of the 500-mb level be higher: where it is warm or where it is cold?

Project 15.4: Print out the following two maps (corresponding to the same time): (a) 500 mb heights and (b) 1000-500 mb thickness.
(a) According to the 500-mb map, where on your map is the layer below that height warmest? How can you tell?
(b) The assumption in (a) is that the surface pressure is the same everywhere. That is not necessarily the case. To avoid having to make this assumption, we can examine the difference in height between the 1000 mb level and the 500 mb level. We call this the 1000-500 mb thickness. Based on the thickness field, where on your map is the layer from 1000 to 500 mb warmest? How can you tell?
(c) According to the thickness field, is the layer-mean temperature warmer east of your location or west of your location (or is it the same)? Indicate your location on the map and explain your answer.

Problems

Problem 15.1: Based on the information obtained in project 15.1, the pressure at about 5.5 km above sea-level is about half that at sea-level. From those pressure values, determine the vertical pressure gradient in Pa/m. Compare your answer to the maximum horizontal pressure gradient measured on your surface map (see figure 4.1 on page 48). Which is greater?

Problem 15.2: (a) Estimate the gravitational force (per volume) using equation 15.1 between sea-level and 5.5 km. A rough order-of-magnitude estimate of the density is sufficient.
(b) Show that the units of gravitational force per volume (N/m^3) is the same as the units of pressure gradient (Pa/m).

(c) How does the value obtained in (a) compare to the typical vertical pressure gradient force (per volume; see problem 15.1)? Should they be equal or should one be greater? If one should be greater, which one?

Problem 15.3: (a) In problem 15.1, it was mentioned that the pressure at about 5.5 km above sea-level is about half that at sea-level. Estimate the percentage of the atmosphere's mass that is below 5.5 km.
(b) What is the air pressure at a height where one-half of the atmosphere (by mass) lies below you?

Problem 15.4: In order for the pressure gradient force to balance the gravitational force, the pressure gradient force must be directed upward. Why does this necessarily imply that pressure decreases with height rather increase?

Problem 15.5: Standard sea-level pressure is assumed to be 1013.25 mb. Examine the sea-level pressures as indicated on a surface map. How does the 1013.25 mb figure compare to those plotted on the surface map? Why aren't all of the values equal to 1013.25 mb? What about the average value?

Problem 15.6: (a) For dry air, what is the value of g/R? Convert your answer to units of °C/m.
(b) In a mirage, the density of air increases with height. What vertical temperature gradient would be associated with an increase in density with height?
(c) For the U.S. standard atmosphere (check the table you've downloaded), what is the vertical temperature gradient for the layers near the surface?
(d) Compare the answers from (a) and (c). Does density decrease with height near the surface of the U.S. standard atmosphere? Explain how you know.
(e) Repeat (c) for the skew-T shown in figure 3.2 (page 39).
(f) Compare the answers from (a) and (e). Does density decrease with height at Little Rock as observed in figure 3.2? Explain how you know.

Problem 15.7: Suppose we warm an air parcel at 700 mb to a temperature 1°C above its surroundings, which is at -10°C.
What is the buoyant force (per volume) on the air parcel? Is it directed upward or downward? Note: assume that the pressure of the air parcel is the same as that around it.

Problem 15.8: If the atmosphere had a constant density equal to its density at the surface, how thick would it be such that its total mass is 5.136×10^{18} *kg?*

Problem 15.9: Use equation 15.6 to obtain the mass of the atmosphere. Why is it not equal to 5.136×10^{18} *kg?*

Problem 15.10: Consider that the wind blows around a low pressure center in a counter-clockwise manner in the northern hemisphere.
(a) Assuming a north-south temperature gradient with colder temperatures toward the north, what kind of temperature advection occurs north, west, south and east of the low pressure center (i.e., warm advection, cold advection or no temperature advection).
(b) Based upon the relationship between warm advection and pressure falls discussed in section 15.3.2 and the pattern identified in part (a), in which direction should the low pressure center move (east, west, north or south) to be consistent with that relationship? Explain.

Derivations

Show-me 15.1: Show that the net force per volume due to the pressure gradient exerted on a box of infinitesimal size is equal to the vertical pressure gradient (dP/dz), in agreement with equation 15.2.

Show-me 15.2: Show how equation 15.4 can be obtained from equation C.6.

Show-me 15.3: Show how equation 15.9 is obtained from equation 15.10.

Show-me 15.4: To derive equation 15.6 on page 293, we had to assume that ΔP equals $P(0)$ if M is the entire mass of the atmosphere (see last step of the derivation in appendix C.6). Why does ΔP equal $P(0)$ when M is the entire mass of the atmosphere?

16. Lapse rates

16.1 Environmental vs. parcel lapse rate

Recall that the **lapse rate** is the rate at which the temperature decreases with height (see section 3.6.2). For example, the U.S. Standard Atmosphere describes an average atmosphere where the temperature in the troposphere decreases at a rate of 6.5°C every km. That means that if you had really tall ladder, the temperature of the air at one would be 6.5°C warmer than the air 1 km higher.

In a real atmosphere, the lapse rate is not a constant 6.5°C every km throughout the troposphere. That is just an idealized average.

In any event, in chapter 17 we will look at how the lapse rate at a particular level can tell us something about the potential for vertical forcing.

In this chapter, we don't look at the lapse rate of the atmosphere. Instead, we look at how an *individual parcel* changes its temperature when it ascends or descends in the atmosphere.

In other words, the temperature we are trying to predict is not the observed temperature. Rather, it is the temperature that a particular **parcel** would have if we moved it up or down in the atmosphere.

WHAT IS A PARCEL?

A parcel is a small region of air that we assume is not mixing with the air around it. Thus, as it rises or falls, we assume it will not warm or cool due to its mixing with the surrounding air, even if the air around it is vastly warmer or colder than the parcel itself.

IF IT DOESN'T MIX WITH THE ENVIRONMENT, WHY WOULD IT WARM OR COOL?

The parcel warms or cools because it is being compressed or expanded as it encounters different pressure.

The air pressure is lower at higher altitudes. Thus, a rising parcel encounters lower pressure as it rises. As it does so, it expands and cools. Conversely, a descending parcel encounters higher pressure as it descends. As it does so, it is compressed and warms.

The variable abbreviation for the lapse rate is Γ. To distinguish between (a) the variation in temperature experienced by the air parcel as is ascends or descends and (b) the vertical variation in temperature that already exists at the location, I'll use subscripts. For example, I'll use Γ_{env} to represent the lapse rate of the environment through which the parcel passes. As I mentioned before, I'll look at this in chapter 17.

In this chapter, I will look at the lapse rate experienced by the rising or falling parcel. This lapse rate will depend on whether the parcel is saturated or not. If it is not saturated, it will cool or warm because of expansion or compression. If the air is saturated, then any cooling will result in condensation, which releases latent heat. Consequently, a saturated parcel will cool less slowly when ascending.

> ✍ A descending parcel warms. Consequently, there is no difference between a saturated and unsaturated parcel upon descending, since no condensation occurs either way.

Check Point 16.1: If a parcel of air rises in the atmosphere, does its lapse rate necessarily equal the environmental lapse rate?

16.2 Adiabatic lapse rate

In this section, we examine how an individual parcel's temperature decreases as it is raised in the atmosphere. The reason why the parcel's temperature decreases is because it expands as its pressure decreases (in order to be in balance with the air around it).

We've already identified how the temperature of a parcel changes with a change in *pressure* (equation 11.1):

$$dT = \frac{1}{c_{\mathrm{p}}\rho} dP$$

To find the parcel's lapse rate, we need to find how the parcel's temperature changes with a change in *height*.

It is here where the hydrostatic relationship comes in from the last chapter. That relationship, after all, describes how the air *pressure* decreases with *height* (equation 15.3):

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

In particular, we use the hydrostatic relationship to replace dP in the adiabatic expansion expression with $-\rho g dz$. This gives the following relationship:

$$\frac{\partial T}{\partial z} = -\frac{g}{c_\mathrm{p}} \tag{16.1}$$

The particulars of the derivation is left as an exercise for the reader (see Show-me 16.1).

The **lapse rate**, Γ, is defined as the negative of the vertical temperature gradient: $\Gamma = -\partial T/\partial z$. Consequently, the **adiabatic lapse rate** is:

$$\Gamma_\mathrm{d} = \frac{g}{c_\mathrm{p}}$$

where I've used Γ_d to indicate the adiabatic lapse rate.

Since g and c_p are independent of height, temperature and pressure, the adiabatic lapse rate is likewise independent of height, temperature and pressure. Its value can be found by plugging in the values of g and c_p:

$$\begin{aligned}\Gamma &= \frac{g}{c_\mathrm{p}} \\ &= \frac{9.8\ \mathrm{N/kg}}{1004.67\ \mathrm{J \cdot kg^{-1}K^{-1}}} \\ &\sim 10\ \mathrm{K/km}.\end{aligned}$$

Thus, the adiabatic lapse rate (see equation 16.1) is about $+10\ \mathrm{K \cdot km^{-1}}$. This means that an insulated parcel of dry air will warm (due to compression) 10 K for every km it descends (in a hydrostatic atmosphere) and cool (due to expansion) 10 K for every km it ascends.

WHY USE A SUBSCRIPT "d" FOR THE PARCEL'S ADIABATIC LAPSE RATE?

Because we're assuming that there isn't any condensation as the parcel cools. Technically, if the parcel cools below its dew point there will be condensation and this changes the parcel's lapse rate. Thus, the "d" stands for "dry".

Check Point 16.2: In the U.S. Standard Atmosphere, the temperature 1 km above the surface is 6.5° C colder than the temperature at the surface (8.5° C compared to 15° C). What temperature would that air have if it was brought down to the surface? Is that air warmer, or cooler, than the air that was already at the surface?

16.3 Moist adiabatic lapse rate

As discussed in chapter 5, condensation occurs when the air cools to its dew point and, as we know from chapter 11, rising motion causes the air to cool. Such cooling can lead to condensation (and clouds) and this introduces latent heating that we cannot ignore.

So as the air rises, it doesn't cool at the adiabatic lapse rate of 10°C per km?

Not if it is saturated. Only if the air is unsaturated will it cool at the adiabatic lapse rate as it rises.

The end result is that, if saturated, one kilometer of lifting no longer results in 10°C of cooling. Rather, it may only be five or six °C of cooling.

This new rate is called the **moist adiabatic** lapse rate and we'll indicate it as Γ_{m}. To distinguish this from the non-saturated rate (see equation 16.1), we'll call the non-saturated rate the **dry adiabatic** lapse rate (indicated as Γ_{d}).

Can we predict exactly what the new cooling rate is?

The new cooling rate will depend on how much moisture condenses and that depends on what the temperature is (since more moisture is present if the temperature is warmer). While an equation can be derived for the moist adiabatic lapse rate, we will instead identify the lapse rate via the skew-T log-P chart (see next section).

WHAT IF THERE IS SOME MOISTURE IN THE AIR BUT IT ISN'T SATURATED?

For air parcels that have some water vapor but are unsaturated, the air parcel will cool according to the dry adiabatic lapse rate until it becomes saturated. Technically,. as the air parcel cools, the saturation mixing ratio drops and may eventually reach the actual mixing ratio (which remains the same as the parcel rises). Once the air parcel becomes saturated, the vapor condenses and we get latent heating.

Check Point 16.3: Suppose we have air at 1000 mb and 0°C with a vapor pressure of 6 mb. Estimate the temperature of the air after being lifted 1 km.

16.4 Skew-T log-P diagrams

16.4.1 Dry adiabatic lapse rate

In section 3.6.3, the skew-T log-P graph was introduced. We can now discuss the series of lines that run vertically like temperature but are skewed the other way (oriented perpendicular to the temperature lines. An example is shown in figure 16.1 (the lines are colored light red in the electronic version of this textbook).

These lines have a slope equal to the adiabatic lapse rate. They are placed on the diagram so that we can quickly determine the temperature a parcel would have if it were raised or lowered in the atmosphere. This process is illustrated in the following example.

Example 16.1: From the skew-T log-P diagram in figure 16.1, determine the temperature an air parcel would have if it initially had a temperature of −50°C at 400 mb and then was brought down to 600 mb.

Answer 16.1: First we identify the initial temperature and pressure of the parcel. This is indicated by the heavy black circle in figure 16.1, which is at the intersection of where the −50°C temperature line (solid straight lines going from lower left to upper right) intersects the 500 mb pressure line (horizontal solid straight lines).

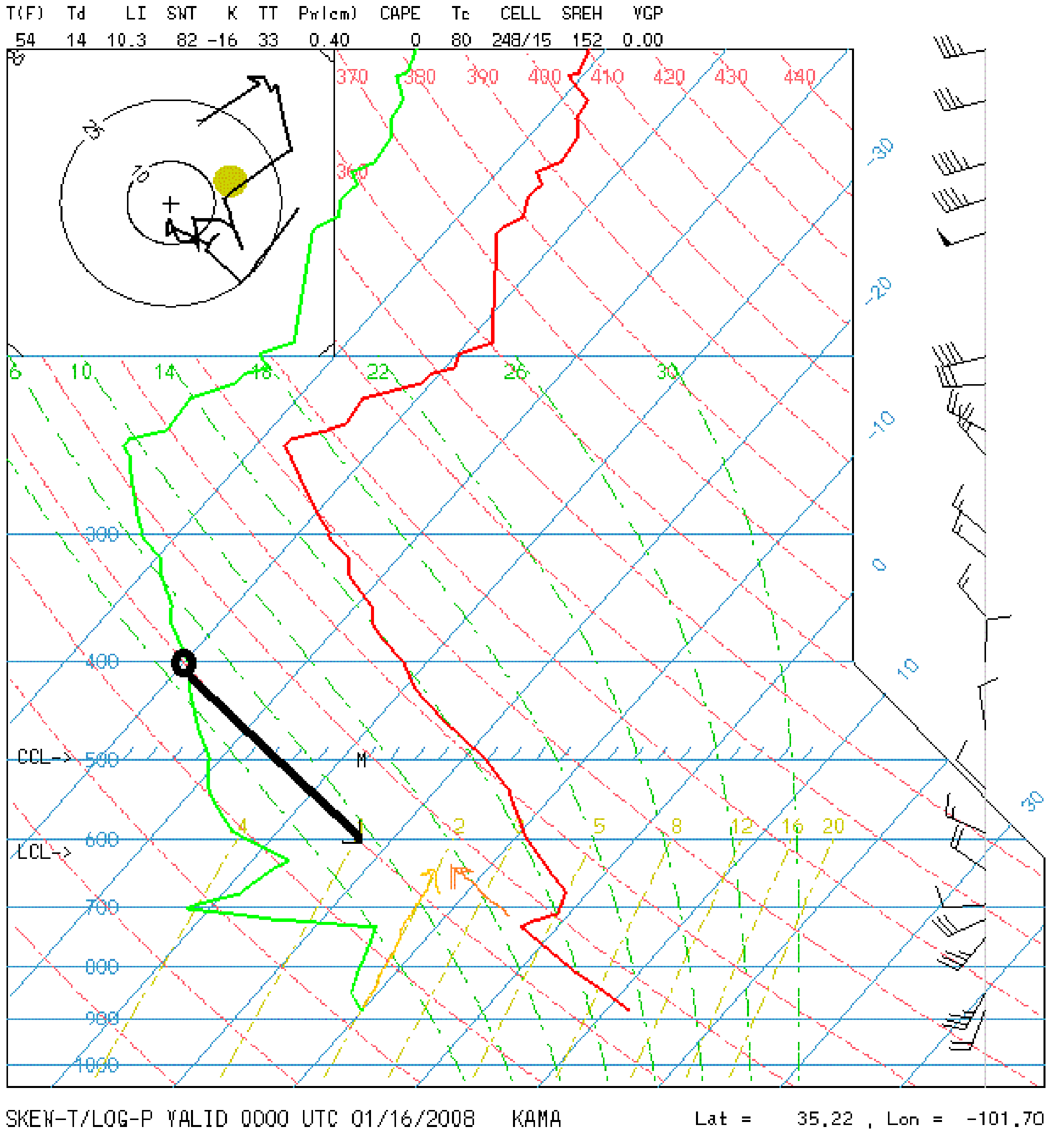

Figure 16.1: A skew-T log-P graph for Amarillo, Texas. Data is from 0000 UTC (7 PM local the previous day), 16 Jan, 2008. Horizontal solid lines are pressure (in mb). Skewed, solid lines are temperature (in °C). Circle and heavy black line are discussed in the text.

Next, starting from this location, we follow the direction of the dry adiabatic lapse rate line (dashed curves going from upper left to lower right) toward the 600 mb pressure line. When we reach that line, we stop and look at what temperature that point corresponds to.

In this case, that point corresponds to a temperature of −23°C. This is warmer than it was before, which makes sense since it is descending and, as it descends, it should be compressed and warm up.

Check Point 16.4: Which of the following is represented by the dry adiabatic lapse rate curve on the skew-T graph?
(a) The temperature of a dry parcel (as a function of height),
(b) The temperature of the actual atmosphere (as a function of height),
(c) The temperature of the U.S. Standard Atmosphere (as a function of height).

16.4.2 Potential temperature

On the Skew-T diagram, the pressure lines are drawn every 100 mb and the temperature lines are drawn every 10°C. You may wonder, then, what values the dry adiabatic lapse rate curves correspond to.

To answer this question, look at where the dry adiabatic lapse rate curves intersect the 1000 mb pressure line. You should find that the lines intersect the 1000 mb pressure line every 10°C, with −23°C being the coldest and 47°C being the warmest.

At first glance, those may seem strange, but they make more sense when you write them in kelvin instead of celsius. In kelvin, the values are 250 K for the coldest curve and 320 K for the warmest, with a curve drawn every 10 K. Indeed, if you look at the top of the skew-T diagram, you see that these curves are indeed labeled in terms of kelvin, with 10 K increments.

Notice, for example, that the parcel in the previous example (−50°C at 400 mb) lies along the 290 K dry adiabatic curve. This mens that this particular parcel would have a temperature of 290 K if it was compressed to 1000 mb.

We call that temperature value the **potential temperature**. This is the same potential temperature discussed on page 156, where it was indicated as θ and defined to be the temperature a parcel would have if compressed (or expanded) to 1000 mb adiabatically.[i].

Any parcel that has a temperature and pressure that puts it along the 290 K dry adiabatic curve has the same potential temperature value (290 K in this case). In other words, as a dry parcel rises or falls, its potential temperature must remain constant. Its *temperature* changes (due to adiabatic warming or cooling) but its *potential* temperature does not (i.e., the temperature it would have if brought to 1000 mb is the same regardless of which path it takes to get there).

✍ Instead of calling these lines the "dry adiabatic lapse rate curves" we could instead call them the "potential temperature curves."

Check Point 16.5: Suppose we have dry air at 1000 mb and 0° C.
(a) What is its temperature and potential temperature in kelvin?
(b) Suppose the air is then lifted 1 km. What happens to the parcel's temperature and potential temperature as it rises?

16.4.3 Moist adiabatic lapse rate

On a skew-T log-P diagram you may have noticed that there are curved lines that are almost vertical at the bottom of the diagram but curve toward the upper left, becoming parallel with the potential temperature (i.e., dry adiabatic lapse rate) curves at near the top of the chart.

These represent the moist adiabatic lapse rate curves.[ii] Recall that the moist adiabatic lapse rate is the rate at which the temperature changes when the saturated air cools. It cools at a slower rate than dry air because the cooling is somewhat countered by latent heating (as the moisture condenses).

---

[i] By convention, potential temperatures are given in kelvin.

[ii] Meteorologists tend to call the dry and moist adiabatic lapse rate curves the dry and moist **adiabats**.

> ✍ For Amarillo, Texas, as shown in figure 16.1 on page 224, the observed temperature between 620 mb and 450 mb just happens to be roughly parallel to the moist adiabatic lapse rate curves.

Why do they match the dry adiabatic lapse rates at low pressures?

Technically, it has to do with the *temperature* being low in that region, not the *pressure*. At low temperatures, there is relatively little moisture present, even if saturated. Since there is not much moisture to condense, there isn't much heating. Consequently, at high altitudes the moist adiabatic lapse rate is roughly the same as the dry adiabatic lapse rate (i.e., about 10 K/km).

Why do they differ at high pressures?

Again, it has to do with the temperature. At higher temperatures, there is more moisture available at saturation. The moist adiabatic rate is the rate at which an air parcel would cool if it was saturated – the moisture that condenses would lead to latent heating, thus slowing down the rate of cooling.

Wouldn't the moist adiabatic lapse rate then be less than the dry adiabatic lapse rate?

Yes. At low altitudes, the moist adiabatic lapse rate is about 6 K per kilometer.

Why, then, are the moist lines more vertical than dry lines?

Because of the way the skew-T log-P chart is made. Remember that the temperatures are skewed. Low temperatures are to the upper left. By sloping toward the left, the dry adiabatic lines indicate rapid cooling with height. The moist adiabatic lines, on the other hand, slope less slowly toward the left, indicating less rapid cooling with height.

So if the parcel is saturated and rising, it doesn't cool as rapidly as it would if it were dry – does this mean the air parcel's potential temperature changes?

Yes. The potential temperature increases due to latent heating as moisture is condensed. This can be seen in the skew-T diagram, as the moist adiabatic lines shift toward higher potential temperature lines as one goes higher.

It the parcel is returned back to its original height, does the potential temperature return to its original value?

That depends on what happens to the condensed water. If it is in the form of cloud droplets then they could evaporate when the air is compressed, and evaporative cooling exactly balances out the condensational warming that occurred during ascent, thus returning the potential temperature to its original value.

On the other hand, if the condensed water precipitates out of the parcel, then there is no liquid water available to evaporate as the parcel warms. Consequently, the parcel will warm back up to a temperature that is higher than when it started (because of the added latent heat from the condensation that occurred during the ascent).

Just as each dry adiabatic curve corresponds to a particular potential temperature, each moist adiabatic curve corresponds to a particular **equivalent potential temperature** (indicated as θ_e) which represents the temperature the parcel would have if all of its moisture were condensed out and the parcel then compressed to 1000 mb. The equivalent potential temperature, then, remains the same regardless of whether the air parcel is saturated or not.

> ✍ If you look across the 200 mb line of the skew-T diagram, you see that the moist adiabatic curves are labeled in terms of celsius, with 4°C increments. If you trace these curves down to the 1000 mb line, you'll see that they correspond to the temperature there.

Check Point 16.6: Suppose we have air at 1000 mb and $0^\circ C$ with a vapor pressure of 6 mb. The air is then lifted 1 km.
(a) What happens to the parcel's temperature as it rises?
(b) What happens to the parcel's potential temperature as it rises?
(c) What happens to the parcel's equivalent potential temperature as it rises?

16.5 Lifting condensation level (LCL)

What happens if an air parcel is neither dry nor saturated (i.e., the mixing ratio is not zero but neither is it equal to the saturation mixing ratio)?

In that case, the parcel's temperature cools at the dry adiabatic lapse rate until saturation is reached. Then, the parcel's temperature cools at the moist adiabatic lapse rate.

How do we know when saturation is reached?

If no condensation occurs (i.e., saturation has not been reached) then the parcel's *mixing ratio* will remain the same. Consequently, follow the dry adiabatic curve until the saturation mixing ratio at that point equals the initial mixing ratio (i.e., identify the point where the potential temperature and mixing ratio lines cross).

Further raising would lead to condensation (i.e., a cloud). This level is thus called the **lifting condensation level** or **LCL**, for short.

To illustrate this process, consider the skew-T log-P chart for Amarillo in figure 16.1 (page 224). Let's suppose we take the air at the surface. According to the diagram, that air has a mixing ratio of 2 g/kg and a potential temperature of around 296 K (this is an estimate, as it lies between the 290 K and 300 K lines).

If we raise this parcel, its potential temperature will remain at 296 K. Meanwhile, its mixing ratio will remain at 2 g/kg.

Until it reaches saturation, that is.

That point is indicated by the intersection of the two curves (see the two arrows; gold and red in the electronic version), which is about 630 mb. Any further lifting will result condensation and the temperature will follow the moist adiabatic curve.

Via this process, we can estimate where the bottom of the cloud is, which is the location of the lifting condensation level (630 mb in this case).

> ✍ This particular skew-T diagram also provides an estimate for where the LCL is, which it indicates it is just above my estimated location.[iii]

[iii] This particular skew-T log-P chart calculates the LCL using values from a "mixed" layer near the surface rather than just the values at the surface. This is somewhat more realistic (since the real atmosphere gets mixed somewhat during lifting). In figure 16.1, the person who created the chart used the lowest 110 mb layer. Since the potential temperature is pretty uniform throughout that layer (in figure 16.1), mixing doesn't change the temperature. However, it does push the mixing ratio a little lower, which pushes the LCL a little higher than you might otherwise calculate.

Check Point 16.7: (a) As an air parcel is raised, but before the lifting condensation level, what happens to the parcel's mixing ratio and potential temperature?
(b) What happens to the parcel's mixing ratio and potential temperature as it is raised above the lifting condensation level?

16.6 Vertical mixing

We've assumed that air parcels don't mix but sometimes there is quite a bit of mixing that goes on in the atmosphere, particular near the ground. Thus, it is useful to examine what happens in such cases.

You might think it would be very difficult to predict the mixed temperature in such a situation, since we now have to worry about the added difficulty of adiabatic warming/cooling due to different pressures.

However, it turns out to be relatively easy if we focus on the potential temperature. It turns out that when two parcels from different heights mix together, the final **potential** temperature will be an average of the initial **potential** temperatures.

In fact, a well-mixed layer will have the same potential temperature throughout the layer (as you'll see in section 17.5, we call this a layer of **neutral** stability).

✍| Such a well-mixed layer is shown near the surface in figure 16.1.

Although the temperature usually *decreases* with height, the potential temperature usually *increases* with height. Thus, mixing the air near the surface with the air above it will result in the surface air warming and the air above it cooling.[iv] The end result is a layer that is colder on top than on the bottom, with a lapse rate equal to g/c_p (i.e., about 10 K/km).

[iv] In fact, citrus growers take advantage of this when the temperature dips below freezing. By setting up huge fans, they mix the air with the air higher up, hoping that the temperature near the surface (where the citrus trees are) will warm up.

Compare that to a typical troposphere (as represented by the U.S. standard atmosphere). The lapse rate of the standard atmosphere is 6.5 K/km. The reason for this is that it is somewhat well-mixed, which makes the lapse rate somewhat less than 10 K/km. We also have moisture in the atmosphere and the condensation of water vapor gives off heat, which affects the lapse rate (see section 16.4.3).

If a well-mixed atmosphere has a lapse rate of 10 K/km, why is a house warmer on the top floors than the bottom floors?

The house differs from the troposphere in two ways. First, the source of heat in the troposphere is the ground (due to solar radiation being absorbed there and making the ground warm) whereas that is not typically the case in a house (either the heating system warms the entire house or the house is warmed by the sun hitting the roof). Second, the air in a house is typically not well-mixed.

Doesn't hot air rise?

As will be discussed in chapter 17, hot air rises only if it is hotter than the surroundings. And then, as it rises, it cools due to adiabatic expansion (as it encounters lower pressure). It is for this reason that the troposphere tends to get cooler with height (i.e., the top of a mountain is colder than the valley).

Check Point 16.8: What is it about the layer of air near the surface as shown in figure 16.1 that implies the layer is well-mixed?

Project

In this project, you will determine how the temperature of the air can change when it is compressed, as will happen when a parcel of air is moved from a region of low pressure to a region of high pressure. You will do this three ways: (1) using an expression that depends upon pressure, (2) using an expression that depends upon the difference in height associated with the difference in pressure (in a hydrostatic atmosphere) and (3) using a skew-T log-P chart.

Project 16.1: From the **U.S. standard atmosphere** *that you downloaded into Excel, you have previously added the air pressure at each level assuming* *hydrostatic balance. Identify the air temperature and pressure at a height of 2 km.*

Project 16.2: Suppose you take a parcel of air at 2 km (with temperature as indicated by the **U.S. standard atmosphere***) and bring it down to the surface, whereupon it is compressed by the higher pressure air that is there. The parcel will warm up. Calculate the new temperature of that air parcel by using equation 11.5 on page 155. In that expression,* T_0 *and* P_0 *correspond to one set of temperature and pressure (at 2 km for our case) and* T *and* P *correspond to another set of temperature and pressure (at the surface in our case).*

Project 16.3: Instead of using equation 11.5 to determine the new temperature of the air parcel (if brought down to the surface), instead use equation 16.1 on page 221. Note that equation 16.1 gives you dT/dz*. You need to multiply this by the change in height to get the change in temperature.*

Project 16.4: A third way to determine the new temperature of the air parcel (if brought down to the surface) is to use a skew-T log-P diagram.
(a) Print out a blank skew-T log-P diagram[v] *and identify the point on it that corresponds to the temperature and pressure of an air parcel at a height of 2 km according to the U.S. Standard Atmosphere.*
(b) Identify the straight lines on the skew-T log-P diagram that have a slope opposite those of the temperature. Highlight the one that comes nearest the point indicated in (a).
(c) Identify the point on the graph where the highlighted line in (b) crosses the horizontal line representing the surface pressure (1013.2 mb).
(d) Identify the temperature corresponding to the point indicated in (c).

Project 16.5: (a) Do all three methods predict the same temperature?
(b) How does the new temperature compare with the temperature the parcel had prior to being brought down to the surface?
(c) How does the new temperature compare with the temperature of the air around it (i.e., what is typically found at the surface according to the U.S. Standard Atmosphere)?

---

[v] See, for example, http://meteora.ucsd.edu/wx_pages/stuff/Blank_Skew-T.pdf

Problems

Problem 16.1: (a) In projects 16.2 and 16.3, you were asked to calculate the temperature an air parcel would have if it was lowered 2 km to the surface. Both assumed the air was in hydrostatic balance but one used equation 11.5

$$T = T_0 \left(\frac{P}{P_0}\right)^{R/c_\mathrm{p}}$$

and the other used equation 16.1

$$\frac{dT}{dz} = -\frac{g}{c_\mathrm{p}}.$$

Should they both give the same answer? Why or why not?
(b) Should the answer in project 16.4 be the same as the other two? Why or why not?

Problem 16.2: (a) Consider a dry air parcel with temperature 5° C and pressure 800 mb. Using the skew-T log-P diagram, what temperature would the air parcel have if it were expanded to 500 mb?
(b) Suppose the parcel was initially saturated. Using the skew-T log-P diagram, what temperature would the air parcel have it were expanded to 500 mb?

Problem 16.3: Go to http://weather.rap.ucar.edu/upper/ and print off a skew-T log-P chart for a location where the dew point and temperature differ by at least 10° C. Trace over the mixing ratio curve corresponding to the mixing ratio observed at the surface. Also trace over the potential temperature curve corresponding to the potential temperature at the surface.
(a) At what pressure do these two traces cross?
(b) How does this compare to where the LCL is indicated on the skew-T log-P?

Problem 16.4: A north-south oriented mountain range is experiencing westerly winds. In the valley on the west side, the temperature T and dew point T_d are both 20° C.
(a) Estimate T and T_d at the top of the mountain range (2 km above the valley).
(b) Estimate T and T_d in the valley on the east side.
(c) Where, if at all, do you expect rain (assume all moisture that condenses

falls out as rain)?
(d) Why is it warmer and drier in the valley downwind of the mountain range (an effect called the **rain shadow effect***)?*

Problem 16.5: In the U.S. Standard Atmosphere, the temperature 1 km above the surface is 6.5° C colder than the temperature at the surface. What would the temperature of the surface be if it is mixed with some air that was brought down from 1 km above the surface?

Problem 16.6: (a) From the skew-T log-P diagram in figure 16.1 (on page 224), estimate the temperature that would exist at the surface if all the air below 720 mb was well mixed. Explain how you obtained your answer.
(b) How would your answer to (a) change if the depth of the mixed air was extended to 700 mb? Would the surface be warmer or cooler? Explain.

Derivations

Show-me 16.1: Derive equation 16.1 by combining the hydrostatic balance relationship (equation 15.3) with the first law of thermodynamics (equation 11.3) for adiabatic processes ($dQ = 0$).

17. Stability

17.1 Buoyancy

To get clouds and precipitation, we need to cool the air to the dew point. In a storm, this is accomplished by having air ascend and expand against the lower pressure environment into which it is ascending.

For this reason, a crucial piece of a forecast has to do with predicting the vertical motion. The question, then, is how do we predict the vertical motions?

The answer is pretty simple, really: hot air rises and cold air sinks.

Well, this isn't exactly true. More accurately, air accelerates upward if it is warmer than the surrounding air, and accelerates downward if it is colder than the surrounding air.

WHY?

As discussed in section 15.1, there are two forces acting on an air parcel: the gravitational force pulling downward and the pressure gradient force pushing upward. We assumed these were in balance. However, there are situations when it is not imbalance, mainly when the air is warmer or colder than the surrounding air.

To see why, we'll assume that the surrounding air is in hydrostatic balance. The gravitational force per mass is ρg (downward), where ρ is the density of the *surrounding* air. The vertical pressure gradient force per mass is $\partial P/\partial z$ (upward). Since the surrounding air is in hydrostatic balance,

$$\frac{\partial P}{\partial z} = \rho g$$

Now let's consider our air parcel, which has a different temperature but same pressure as the surrounding air. By the ideal gas law, that means the parcel's density is different than the surrounding air. That means the gravitational force per mass is different. Yet, the pressure gradient force is the same (since it has the same pressure). That means the two forces are no longer in balance.

In particular, if the parcel's temperature is lower, the density is higher (since $P = \rho RT$). A higher density means the gravitational force per mass is greater and thus the parcel accelerates downward (in the direction of the gravitational force).

Conversely, if the parcel's temperature is higher, the density is lower. A lower density means the gravitational force per mass is less and thus the parcel accelerates upward (in the direction of the pressure gradient force).

How do we know the pressure of the parcel is the same as the surrounding air?

If the pressure was different, the parcel would expand or contract such that the pressures would equalize. That is why we can say that the vertical pressure gradient is the same.

> ✍ Although we used air in our analysis, we could've used any fluid. The same formulation could be used, for example, for air bubbles in water. Whenever the immersed object (whether it is a bubble or an air parcel) has a density less than the fluid it is immersed in, the object will accelerate upward.

Check Point 17.1: Does warm air always rise?

Mathematically, we can write the gravitational force per mass on the parcel as ρ'. That means the force imbalance on the air parcel is $\partial P/\partial z - \rho^{p}rimeg$. Since the surrounding air is in balance, we can replace $\partial P/\partial z$ with ρg. Making this replacement we have that the force imbalance on our parcel is as follows:

$$\vec{F}_{\text{buoyancy}} = g(\rho - \rho')\hat{k} \tag{17.1}$$

where the force imbalance $\vec{F}_{\text{buoyancy}}$ is known as the **buoyancy** force. In this form, one can see that an air parcel that is less dense than its surroundings will experience an positive (upward) buoyant force.

> ✍ The buoyant force is really the *net* force resulting from two separate forces: the gravitational force pulling downward and the pressure gradient force pushing upward.

Check Point 17.2: Why is it said that equation 17.1 implies that an air parcel that is less dense than its surroundings will experience an positive (upward) buoyant force.

17.2 Equilibrium level

In the previous section, we saw how we can determine whether an air parcel will rise or fall, assuming it starts at rest, based upon whether its temperature is warmer or cooler than the environment. In this section, we show how we can determine *how far* it will rise or fall until it reaches an environment where the buoyant force is zero. That level is called the **equilibrium level**.

The equilibrium level will depend upon the temperature of the environment. To obtain that temperature, we will use a sounding, as plotted on a skew-T log-P chart.

As the parcel rises or falls, it expands or compresses, which changes its temperature adiabatically. How it changes can also be obtained via a skew-T log-P chart.

Consequently, our determination of the equilibrium level will utilize a skew-T log-P chart.

To illustrate, consider the sounding from Omaha, Nebraska, on 1200 UTC 20 May 2004, as plotted on the skew-T log-P chart in figure 17.1.

There are a couple of circles, diamonds and other figures plotted on the chart in addition to the temperature and dew point. I put those there to illustrate a couple of examples.

For the first example, consider what would happen to a −30°C parcel that we “introduce” at 400 mb. This parcel is indicated by the thick black circle on the diagram.

As we know from the previous section, this parcel will sink since its temperature is colder than the environment at that level (about −23°C according to the chart).

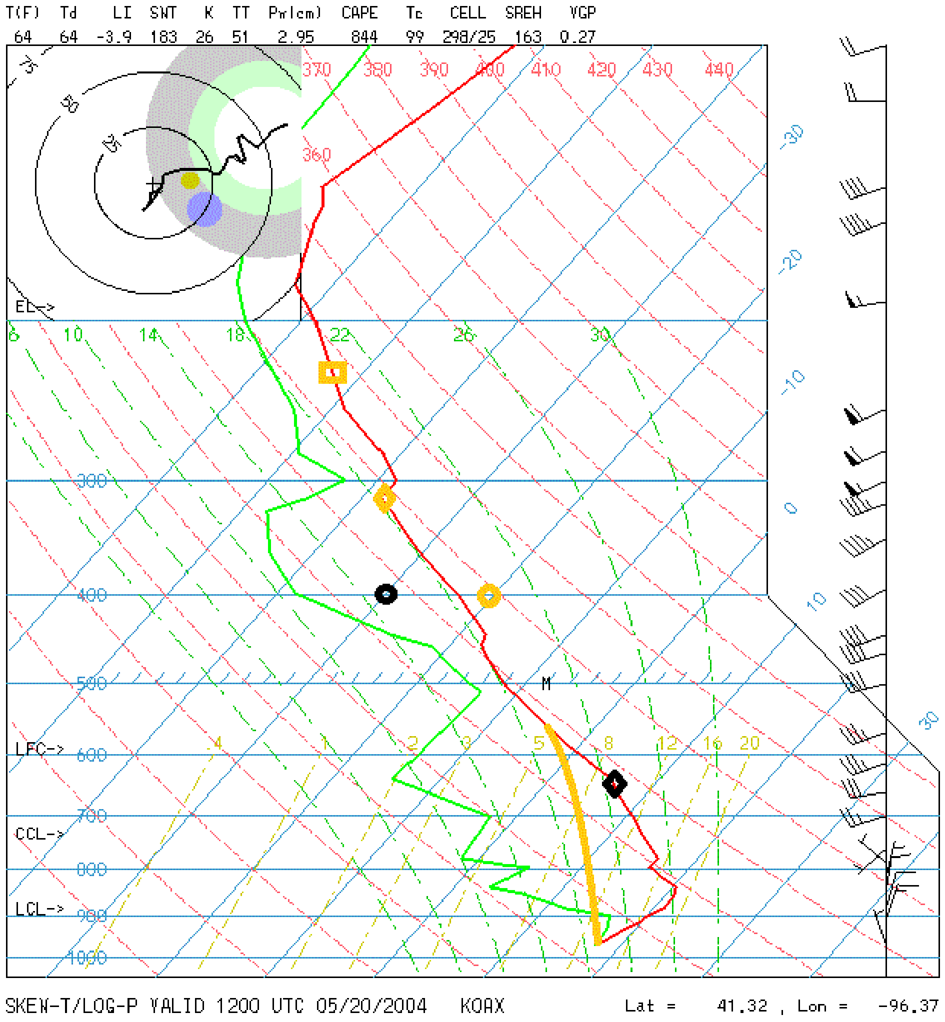

Figure 17.1: A skew-T log-P graph for Omaha, Nebraska. Data is from 1200 UTC (7 AM Central Daylight Time), 20 May, 2004. Squares, circles and diamonds are described in the text.

As it sinks, it warms adiabatically according to the dry adiabatic lapse rate (keeping its potential temperature the same). This rate is indicated by the light red lines that slope toward the upper left of the chart.

Looking at the chart, the black circle is located between the 310 K and 320 K potential temperature curves.[i] As the parcel descends, it maintains this potential temperature and follows the dry adiabatic lapse rate curves until it encounters air that is the same temperature.

As can be seen on the skew-T chart, this happens at around 640 mb, which I've indicated with the thick black diamond. This represents the equilibrium level for this particular parcel.

What happens if the parcel is originally warmer than the environment, rather than cooler?

To illustrate what happens, consider a −20°C parcel "introduced" at 400 mb. This parcel is indicated by the thick gold circle on the diagram.

This parcel will rise since its temperature is warmer than the environment at that level. As it rises, it cools adiabatically according to the dry adiabatic lapse rate. Looking at the chart, the gold circle is located between the 320 K and 330 K potential temperature curves.

As the parcel ascends, it maintains this potential temperature and follows the dry adiabatic lapse rate curves until it encounters air that is the same temperature. The equilibrium level for this particular parcel is around 310 mb, which I've indicated with the thick gold diamond.

What happens if the parcel is originally saturated?

If it is saturated then as it rises it cools according to the moist adiabatic lapse rate, keeping its equivalent potential temperature the same. In this case, the equivalent potential temperature is about 20°C (look across the 200 mb line of the skew-T diagram for the labels). Tracing the moist adiabatic curves upward, one eventually reaches the point marked with the thick gold square at around 230 mb, which indicates the equilibrium level for this particular parcel.

---

[i] This can be determined by either looking at the labels at the top of the skew-T chart and incrementing backwards until one reaches the black circle, or by tracing the potential temperature (dry adiabatic lapse rate) curves to 1000 mb and reading off the temperature there.

Check Point 17.3: Suppose a $-20^\circ C$ parcel is introduced at 500 mb over Omaha, Nebraska, at 1200 UTC 20 May 2004. Where is its equilibrium level?

17.3 Level of free convection

How do we get the air to be warmer than the surroundings?

For air near the surface, the air can warm simply by being next to the ground, which warms up during the day as it absorbs solar radiation. The problem is that it can't simply be warmer – it must be warmer than the surroundings. And, if all of the air near the surface warms then there won't be any buoyant force. Usually, however, there are enough differences in the surface that the air may experience more warming at one location than another.

Another way to get the air warmer than the environment is to cause it to ascend. That requires some other force. For example, if the wind send the air into a mountain, the ground may force the air to move upward. Or, when two air masses meet (as with a front[ii]), some air may be pushed up at the boundary.footnoteConsider, for example, how squeezing some clay in your hand would make some may ooze out through the gaps in your fingers.

Doesn't the air cool when it ascends?

Yes. However, it doesn't matter if the air is cooler than it was. What matters if it is warmer than the air in which it ascends into.

For example, consider the air at the surface in Omaha, Nebraska, as shown in the skew-T of figure 17.1. That air is saturated (i.e., the dew point equals the temperature). If we raise that air by some mechanism, it will cool according to the moist adiabatic rate. I've indicated that rate on the diagram with a bold gold curve.

Up until 580 mb or so, the air we've pushed up remains colder than the environment. At 580 mb, though, it now equals the air that is there at 580 mb. And, any further lifting results in it being *warmer* than the environment

---

[ii] The two air masses could have different temperatures, which may introduce air of a different temperature into a region.

air at those altitudes. If we release the air parcel at a pressure less than 580 mb, it will experience an upward buoyancy force.

This height is called the **level of free convection**. It is the height (or pressure) at which any further lifting results in an upward buoyancy force.

That results in further ascent, bringing it past the $-20°\text{C}$ temperature at 400 mb that was explored in the previous section (gold diamond in the skew-T chart) and eventually to the 230 mb equilibrium level (gold square in the skew-T chart).[iii]

> ✍ On the left side of the skew-T diagram, they've indicated the level of free convection and its equilibrium level as "LFC" and "EL", respectively. There levels are a little different than the levels obtained by us because they first assume that the lowest levels are mixed. Mixing the air would increase the surface temperature since the air above the surface has a higher potential temperature (see section 16.6) while the mixing ratio would remain roughly the same (since the mixing ratio is pretty much the same throughout the lowest 100 mb or so). With the surface air no longer saturated, you'd have to raise the surface air higher to reach condensation (higher LCL), lower LFC and higher EL.

Check Point 17.4: Not all soundings have a level of free convection. Given the sounding for Omaha, Nebraska, shown in figure 17.1 (on page 238), how much colder would the surface have to be for there to be no level of free convection? Assume the dew point continues to match the temperature (i.e., the relative humidity remains at 100%).

17.4 The nature of stability

In this section, we examine **stability**.

To understand stability, consider the following two situations. One situation has a marble in the bottom of a bowl. When the marble is moved to one side,

---

[iii] This is just because I had chosen the initial parcel characteristics with the intention of matching this situation.

it rolls back to the bottom. This is considered a **stable** situation because the marble returns to its initial position when displaced slightly from it.

Now consider the reverse situation. In this situation, we invert the bowl and place the marble on top of the inverted bowl. It might stay there, but if the marble is moved to one side the marble will roll off the inverted bowl. This is considered an **unstable** situation because the marble continues to move away from its initial position when displaced slightly from it.

When applying stability to the atmosphere, we are considering air parcels, not marbles. And what forces the air parcel back toward its initial position or away from its initial position is the effect of buoyancy, not the shape of a bowl.

To illustrate what I mean, consider an air parcel that is initially at the same temperature as the air around it. With no difference between the air parcel temperature and the surrounding (environment) temperature, the air parcel would not naturally feel a force forced upward nor downward. We say that the air parcel is in **equilibrium**.

Then suppose that the air parcel is given a little push so that it rises a small amount (see previous section). When the air parcel rises, it cools a small amount. It cools because it expands as it encounters an environment where the pressure is less.

As we've seen in the previous section, just because the air parcel cools that does not mean it is colder than the environment at that new level. The temperature of the air at the new level could be anything, since the environmental lapse rate can be almost anything.

> ✍ Being higher up, the new environment is likely colder than the prior environment, but there is no reason why it *has* to be colder.

If the air parcel happens to be colder than its new environment, it will sink and return to its initial position. We call this situation a **stable equilibrium** because a small displacement will force the air parcel back toward its initial position.

On the other hand, if the air parcel ends up being warmer than its new environment, it will rise, moving farther from its initial position. We call this situation an **unstable equilibrium** because a small displacement will force the air parcel away from its initial position.

WHAT HAPPENS IF THE AIR PARCEL IS INITIALLY PUSHED DOWNWARD INSTEAD OF UPWARD?

If pushed downward, the reverse happens. The air warms, rather than cools. As before, however, the key is whether the air parcel is warmer or cooler than the new environment. If it is cooler, it continues falling and we have an unstable equilibrium. If it is warmer, it rises and returns to where it started, which is called a stable equilibrium.

Check Point 17.5: (a) If an air parcel, initially the same temperature as the environment, is raised a small amount, will it cool or warm?
(b) Suppose, after being raised, it is then warmer than the air that surrounds it at that level. Is that associated with a stable equilibrium or an unstable equilibrium? Why?

17.5 Stable, neutral and unstable layers

Of course, it isn't too useful knowing whether the air is unstable or stable *after* it has been slightly displaced upward or downward. What we really want to know is whether it is unstable or stable *without* having to move the air parcel.

To do this, we go back to what makes the air stable. It is stable if, when we displace an air parcel upward, it becomes colder than the environment. A dry parcel cools at the dry adiabatic rate. Consequently, for it to become colder than the environment, the environment must have a lapse rate *less* than the dry adiabatic rate.

WHAT IF WE LOWER THE AIR PARCEL?

The same thing happens but in reverse.

It is stable if, when we displace an air parcel downward, it becomes warmer than the environment. A parcel warms at the dry adiabatic rate. Consequently, for it to become warmer than the environment, the environment must have a lapse rate *less* than the dry adiabatic rate.

In general, then, a stable environment is one that has a lapse rate less the dry adiabatic lapse rate. Mathematically that means that a stable environment has a lapse rate *Gamma* that is less than g/c_p or about 10 K/km:

For example, the sounding for Omaha, Nebraska, as shown in figure 17.1 (page 238), is stable throughout. Since the dry adiabatic rate is indicated on the skew-T log-P charts as the lines curving from bottom right to upper left, one can quickly identify whether a particular layer is stable or not by comparing the observed temperature trace with the dry adiabatic lines (i.e., potential temperature lines).

Actually, there is a layer from 620 to 570 mb or so, which has a lapse rate very similar to the dry adiabatic rate, as shown by the temperature profile being parallel to the potential temperature curve. Such a layer is called a **neutral** environment. It means that you can move a parcel from within that layer to any other location within that layer and it will remain there.

In most situations, the lapse rate in the troposphere is less than the dry adiabatic lapse rate and thus is stable.

For example, consider the U.S. standard atmosphere as representative of the typical state of the atmosphere. In the lower 11 km of the U.S. standard atmosphere, the lapse rate is 6.5 K/km, which is less than the dry adiabatic lapse rate. This means that the U.S. Standard Atmosphere is considered to be stable.

Check Point 17.6: What is it about the air between 620 mb and 570 mb Omaha, Nebraska, in figure 17.1 that indicates the layer is neutral?

Can the atmosphere ever be unstable?

For the atmosphere to be unstable, it would have to have a lapse rate greater than the dry adiabatic lapse rate. In other words, it isn't sufficient to just have a vertical temperature gradient with warm air below cold air. The air below must be at least 10°C warmer than the air a kilometer above it. One situation where this can happen is in the afternoon, when the surface has been heated by the sun.

> ✍ Note that, contrary to what people say about warm air rising, warm air will only rise if it is warmer than the air on either side. An air parcel can be warmer than the air above it (or colder than the air below it) and still be in a perfectly stable situation. Even it were unstable, the air won't rise unless something pushes it out of its equilibrium state (albeit an unstable one).

If an air parcel gets pushed out of equilibrium, will it rise forever?

No. Eventually it will encounter air that is warmer than it is because of the warm air associated with ozone absorption in the stratosphere. To find the equilibrium level, follow the discussion in section 17.2.

When the temperature is warmer at higher elevations, we have a negative lapse rate. Such a situation is called an **inversion** (since it is opposite what is normally expected in the lower atmosphere). Such a lapse rate is associated with very strong stability, as a negative lapse rate is much smaller than the dry adiabatic lapse rate of 10 K/km. Such strong stability above the tropopause is responsible for the lack of clouds in the stratosphere. After all, clouds are formed when the air temperature cools to the dew point. The main mechanism for that cooling is the cooling associated with adiabatic expansion when parcels rise. The strong stability in the stratosphere inhibits rising motion[iv] and thus cloud formation.[v]

> ✍ The tropopause is around 200 mb. As seen in the Omaha, Nebraska, sounding the lapse becomes close to zero above 200 mb. A zero lapse rate is called **isothermal** (since the temperature is the same with height).

Check Point 17.7: (a) If the atmosphere matched the profile of the U.S. Standard Atmosphere, and an air parcel near the surface of that atmosphere were pushed upward a kilometer, by how much would the air parcel have cooled?
(b) How much warmer or cooler would the air parcel be than the air around it at that level?
(c) Based on this, is the U.S. Standard Atmosphere stable, neutral or unstable?

[iv] It is for this reason that the top of a thunderstorm (cumulonimbus cloud) is characterized by a flat area called an anvil (because it looks like an anvil). This “flattening out” corresponds to the stability of the atmosphere above the tropopause. Once the parcel reaches the stratosphere, further upward motion is inhibited. Consequently, the upper levels clouds spread out horizontally (since upward motion is suppressed).

[v] A common misconception is that there just isn’t enough water available above the tropopause to produce clouds at that level. While it is true that the air is so cold that there is very little moisture, the lack of clouds has more to do with the fact that there is no mechanism for the cooling.

17.6 Conditional instability

If a rising saturated air parcel cools at a different rate than an unsaturated air parcel would, how does this affect the stability?

The general idea is the same but the conditions required for instability are a little different.

If the air is saturated, we need to compare the environmental lapse rate with the *moist* adiabatic lapse rate (about 6 K/km near the surface), not the *dry* adiabatic lapse rate (about 10 K/km). In other words, if the environmental lapse rate is greater than 6 K/km (near the surface) and the environment is saturated, it will be unstable. The environmental lapse rate need not reach 10 K/km.

What if moisture is present but the environment isn't saturated?

It turns out that if the environmental lapse rate is between the dry and moist adiabatic lapse rates, whether it is unstable or not depends on the amount of moisture that is present. Such a condition is known as **conditional instability**. Only if the environmental lapse rate is greater than the dry adiabatic lapse rate will instability be guaranteed (such a situation is known as **absolute instability**).

Meteorologists look for conditions where the atmosphere is stable if dry but unstable if moist. If there is a mechanism present to provide enough lift to bring a parcel to its lifting condensation level, then the resulting instability in the atmosphere will do the rest – the parcel will rise until it reaches a region that is stable (typically the stratosphere).

Check Point 17.8: Does having water vapor in the air make the air more stable or less stable?

Problems

Problem 17.1: (a) From the skew-T log-P chart in figure 16.1 (page 224), is the air between 700 mb and 720 mb stable or unstable? How can you tell?

(hint: compare with adiabatic lapse rate lines)
(b) In the same skew-T log-P chart, is there any level above 700 mb where the air is unstable? If so, where? How about neutral?

Problem 17.2: Consider the air above the surface as indicated by the skew-T log-P diagram in figure 16.1. Suppose we take a parcel of air at 850 mb and heat it to 7° C above the surrounding air.
(a) Use the adiabatic lapse rate lines on the skew-T to determine how high it will rise. Provide the pressure level and explain how you used the skew-T to get the answer.
(b) Using equation 11.6 (Poisson's equation), determine the potential temperature of the parcel, both at 850 mb and at the pressure level you indicated in (a).
(c) Compare the two potential temperatures. Are they the same? If so, explain why they should be. If not, explain why they aren't.

Problem 17.3: In the early morning, the air near the ground can cool a great deal, much more than the air above. The result is a situation where the vertical temperature gradient is positive (negative lapse rate), known as an **inversion***. Is this a region of strong stability or strong instability? Explain.*

Problem 17.4: In the U.S. standard atmosphere, the lapse rate above 11 km or so (the tropopause) goes to zero (i.e., it is **isothermal***).*
(a) Is this isothermal region one of strong stability or strong instability? Explain.
(b) To have clouds, you need upward motions so that the air cools to the dew point. Is your answer in (a) consistent with the observation that clouds typically don't exist much above the tropopause?

Problem 17.5: Given the sounding for Omaha, Nebraska, shown in figure 17.1 (on page 238), what would happen to the lifting condensation level (LCL), level of free convection (LFC) and equilibrium level (EL) if the surface temperature was increased to 30° C (while keeping the same dew point)?

Problem 17.6: Go to http://weather.rap.ucar.edu/upper/ and print off a skew-T log-P chart for a location where the "EL" and "LFC" are indicated on the chart (note that charts will not have these marked unless there is a possibility of a level of free convection). Check their positions by determining the level of free convection on your own, by doing the following. First determine the lifting condensation level as in problem 16.3. If the parcel, on the way up to the LCL, ever encounters air that is colder than itself, it will continue

rising on its own. That point would be the level of free convention (LFC). Otherwise, continue up from the LCL following the moist adiabat (since the air is now saturated) until the parcel encounters air that is colder than itself. The first occurrence is the level of free convection.
(b) Determine how high the parcel will rise before it first encounters air that is warmer than itself (i.e., equilibrium level).

Problem 17.7: (a) What is the potential temperature at the surface for Amarillo Texas, as shown in figure 16.1 on page 224?
(b) What is the potential temperature at 800 mb for Amarillo Texas, as shown in figure 16.1?
(c) What is it about the air between the surface and 800 mb for Amarillo Texas that indicates the layer is neutral?

18. Horizontal Wind Direction

18.1 Wind and pressure

In this chapter, we examine the mechanisms responsible for horizontal motion and use those mechanisms to predict the wind direction.

> ✍ Horizontal motion is referred to as the **wind**. Vertical motion is usually referred to as **updrafts** or **downdrafts**.

As you probably already recognize, the wind is intimately tied to the air pressure. So, we will spend a lot of time focusing on the pressure pattern and those aspects of the pressure pattern that impact the wind direction.

Over small time scales, the wind flows from high pressure to low pressure. I'll illustrate this by examining a phenomenon called the sea breeze.

Over large time scales, particularly far from the frictional effects associated with the surface, the wind flows parallel to the isobars. This means it flows *around* the low and high pressure centers, with the direction being opposite in opposite hemispheres (i.e., northern vs. southern).

Near the surface, where there is friction, the wind still flows around the low and high pressure centers for the most part, but the wind direction is somewhat toward the low pressure regions and away from the high pressure regions (i.e., it has a component that is across the isobars).

> ✍ Meteorologists use the word **cyclonic** to describe the direction that the air spins around a center of low pressure. Conversely, meteorologists use the word **anticyclonic** to describe the reverse direction – the direction that the air spins around a center of high pressure.

This leads to several questions:

- Why doesn't the wind from high to low pressure in all cases?
- Why is the direction opposite in the southern hemisphere vs. the northern?

These questions will be answered in this chapter.

18.2 Land/sea breezes

The **land/sea breeze** is a common phenomenon along the coast. Actually, it is called the **land breeze** when the breeze blows from land out toward the sea and it is called the **sea breeze** when the breeze blows in from the sea toward land.

It turns out that the sea breeze tends to blow during the day (when the land is warmer than the sea) and the land breeze tends to blow at night (when the land is colder than the sea).

The reason for the land breeze is because at night there tends to be higher pressure over the land and lower pressure over the sea. That produces a horizontal force on the air, pushing it from the land (higher pressure) toward the sea (lower pressure). The opposite situation occurs during the day, with the higher pressure over the sea.

WHY IS THERE A PRESSURE DIFFERENCE?

The pressure difference is a consequence of the temperature difference.

During the day, the land heats up more than the ocean.[i] Conversely, at night, the land cools down more quickly than the ocean. Thus, the land tends to be colder at night and warmer during the day.

HOW DOES THE TEMPERATURE DIFFERENCE SET UP A PRESSURE DIFFERENCE?

Based on the ideal gas law, one might expect the pressure to *increase* in a region where the temperature increases. To see why the reverse is true, we have to look at what is going on step-by-step.

Consider two points on the surface straddling the shore line, one on water and one on land (indicated as x_1 and x_2, respectively, in figure 18.1). We'll assume

---

[i] The main reason for this is because the land absorbs solar radiation within a very shallow layer of land whereas the sea absorbs solar radiation over a much deeper layer (since the water is mostly transparent). Another reason has to do with the specific heat of land vs. water. The specific heat of water is higher, which means it doesn't warm up as much for the same amount of absorbed energy.

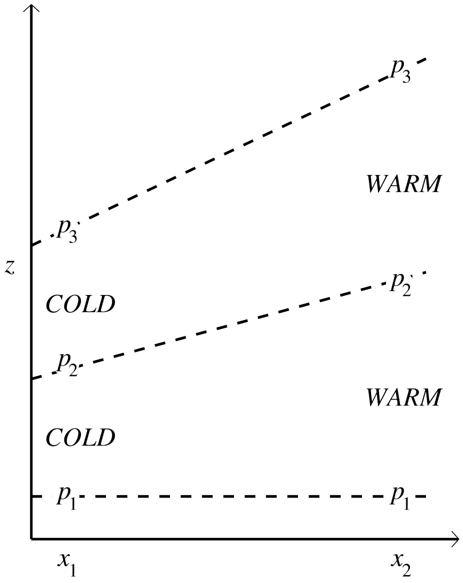

Figure 18.1: A representation of a vertical slice of the atmosphere, with coordinates z and x. The surface pressure at locations x_1 and x_2 are the same (indicated by p_1). The height at which the pressure equals p_2 is higher at x_2 than it is at x_1. In the figure, the pressure decreases with height (i.e., $p_1 > p_2 > p_3$).

that initially the atmosphere is the same over each point. As discussed in section 15.3.1, this means there will be just as much mass over each point and the pressure at each point will be the same (indicated by the p_1 at each location).

Now let's see what happens during the day. During the day, the land (location x_2) heats up more than the ocean. The pressure must remain the same, since there has been no change to the *amount* of air above that location. Consequently, according to the ideal gas law, as the temperature increases over land, the density over land decreases.

The only way that the density can decrease without changing the amount of mass above x_2 is if the column increases in height (i.e., the thickness increases). This is indicated by the higher height at which the pressure p_2 is found.

At upper levels, then, there is a horizontal pressure gradient. In particular, the pressure over the ocean (x_1 in the figure) is *lower* than the pressure over land (x_2 in the figure). Consider, for example, that p_3 is less than p_2 but at the height z in the figure they are roughly at the same altitude. Thus, the air at that altitude will be forced from p_2 to p_3.

Consequently, at upper levels, the air moves from land to ocean (from right to left in the figure).

What about at the surface, which is where most of us live?

Well, due to the movement of air from over land to over water, the amount of air over the land decreases and the amount of air over the water increases. Under hydrostatic balance, the surface pressure must subsequently decrease over land and increase over the ocean. This causes a pressure gradient at the surface as well.

It is this *resulting* surface pressure gradient that pushes air from the sea to the land.

This is called the **sea breeze** and it occurs mainly during the day (when the land is warmer).

The reverse process occurs at night and is called the **land breeze**.

Check Point 18.1: When the land warms up during the day, what happens to the 1000-500 mb thickness over land? If the surface pressure doesn't change, what happens to the horizontal pressure gradient at around 500 mb along the shore line?

18.3 Large scale pressure systems

As mentioned in the introduction, when we get to larger scale pressure systems, the wind is *not* directly toward the lowest pressure but rather *around* the center of lowest pressure.

In fact, the direction is counter-clockwise in the northern hemisphere and clockwise in the southern hemisphere (with the directions reversed around high pressure centers).

> ✍ As mentioned earlier, this is really only true for special circumstances. However, very close to what we observe at high altitudes, where friction doesn't have much of an effect. And, even when it does, like near the surface, the motion is still roughly parallel to the isobars.

To explain why this is, I will use a conceptual approach. Predicting the wind speed, on the other hand, is more mathematical (see chapter 19).

To start, we recognize that air is forced from high pressure toward lower pressure.

However, if this was the only thing forcing the air, the air would accelerate toward regions of lower pressure, meaning that the wind should be directed toward lower pressure (and get faster and faster as it moves).

This prediction is not supported by our observations. Instead, the wind is directed parallel to the isobars (i.e., *around* regions of high and low pressure).

To explain why, we can apply Newton's laws. Since the earth is rotating (and the winds are also observed to be characterized by circular flow) we will use a form of Newton's laws known as conservation of angular momentum. To refresh your memory, I will review the basic concept of this, as well as the concept of frames of reference, before showing how it can be used to explain the direction of the wind.

Check Point 18.2: Describe the motion of the air, relative to the isobars, at high altitudes.

18.3.1 Conservation of angular momentum

To illustrate the meaning of **conservation of angular momentum**, consider the case of the spinning skater. When the skater brings his[ii] arms toward his body, the skater spins. Why?

First, we notice that the skater doesn't start at rest. In other words, if we look carefully, the skater is already spinning. By bringing his arms in toward his body, the skater doesn't *start* to spin but rather starts to spin *faster*. The

---

[ii] For ease of reading, I am using "his" as a generic term for the skater, rather than the more appropriate "his/her".

direction of the spin is always the same as what it was before the arms were brought in.

In physics language, we say that as the arms are brought in, the skater's **angular velocity** increased. The angular velocity had to increase because, by **Newton's second law**, the skater's **angular momentum** has to remain the same as long as there is no net **torque** acting on the skater, which is the case for the skater.[iii]

An object's angular momentum about an axis is the product of the object's **rotational inertia**[iv] about the axis and the object's **angular velocity** about the axis. If one goes down, the other must go up in order to keep the product the same (assuming no torques acting upon the system).

WHY DOES BRINGING THE ARMS IN REDUCE THE ROTATIONAL INERTIA?

The rotational inertia depends upon how the mass is distributed about the axis of rotation.[v] The further the mass from the axis, the greater the rotational inertia (i.e., the more **torque** needed to change the object's angular velocity). Conversely, when most of the mass is near the axis of rotation, the rotational inertia is smaller and it is easier to change the object's rotation rate.

By bringing the arms in toward the body, the skater's rotational inertia about the axis of rotation decreases. Since angular momentum must remain the same, the skater's angular velocity must increase.

Check Point 18.3: When a skater brings his arms toward his body, he spins faster. What happens to his angular momentum during this time?

18.3.2 Frames of reference

HOW DOES THIS APPLY TO THE CASE OF WIND AROUND A LOW PRESSURE?

[iii] Torques are associated with forces applied to the object and directed around the axis of rotation. By definition, torque is equal to the product of r and F_{T}, where r is how far the force is applied from the axis of rotation and F_{T} is the component of the force about the axis.

[iv] Or, **moment of inertia**.

[v] The rotational inertia is defined as $m(r^2)$avg, where (r^2)avg is the mass-weighted average square radius of the object.

The air, like the skater, is already rotating because it is on the Earth, which is rotating. Then, when the air is brought in toward the low pressure center, the air spins faster, just like the skater when he brings his arms in.

The problem with the air is that we are observing the air on a rotating frame of reference (the Earth) and so the air does not appear to be rotating initially. But it is.

To illustrate what is going on, let's replace the skater with a person on a turntable that is free to rotate. If the person is already rotating and then extends his arms, his rotational inertia (about the rotating axis of the turntable) will increase and his angular velocity (about the rotating axis of the turntable) will conversely decrease. In other words, by extending the arms, the person's rotation rate slows down. Conversely, by bringing the arms in, the person's rotation rate speeds up.

Key point #1: If the person is not spinning to begin with, the person's rotation rate will remain zero regardless of whether the arms are brought in or out.

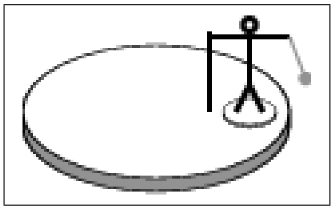

Now consider the situation illustrated in the figure. The person and the rotating turntable are placed on an even larger rotating platform (like the rotating platforms found at many playgrounds). What happens if the larger platform is rotating?

If the person (who is standing on the turntable which is in turn on the platform) is holding onto something fixed to the platform, the person will *not* appear to be rotating *from the perspective* of someone on the platform. From the perspective of someone on the ground, however, the person *is* rotating, at the same rate as the larger platform.

WHAT HAPPENS WHEN THE PERSON ON THE TURNTABLE BRINGS HIS ARMS IN TOWARD HIS BODY?

Since the person (being on the turntable) is free to rotate, he will rotate *faster* in the direction he was rotating before. Since the person is *already* rotating with the large platform, he will rotate *faster* in the direction the large platform is rotating.

Key point #2: From the point of view of someone on the larger platform, the person is not spinning to begin with – the person's rotation rate appears to change rather than remain zero when the arms are brought in or out.

In physics terms, we say that the apparent discrepancy between key points #1 and #2 is that someone on the larger platform is viewing the situation from a **non-inertial reference frame**. The reference frame of someone on the large platform is considered to be *non-inertial* because it is accelerating (anything moving in a circle must be accelerating; otherwise its direction would remain constant). Our laws of physics (like the conservation of angular momentum that was used to get key point #1) apply only to **inertial reference frames**. The reference frame of someone on the ground is considered to be *inertial* because it is not accelerating. Thus, someone on the ground sees the situation correctly (i.e., the person on the turntable rotates because he/she was *already* rotating).

So, what does this have to do with the wind?

Since we are on a rotating earth, we are observing the situation from a non-inertial frame of reference. Air that appears to be stationary to us is actually rotating with the earth. When looking down from above the northern hemisphere, the earth is rotating counter-clockwise. Consequently, air that appears to be stationary is actually rotating counter-clockwise (as seen by someone looking down from above the northern hemisphere).

Since angular momentum is conserved, the air will spin faster as it moves toward the center of lowest pressure. Consequently, the air will appear to "develop" a counter-clockwise rotation (in the northern hemisphere). The closer the air gets to the center of rotation, the faster it goes. The speed of the wind around a low can get very high as a result (as in a tornado).

> ✍ As mentioned earlier, meteorologists use the word **cyclonic** to describe the direction that the air spins around a center of low pressure. Large low pressure systems, like hurricanes and cyclones, will spin cyclonically (hence its name).

What about around a high pressure center?

A similar process is at work for wind around a center of high pressure. Since stationary air is rotating cyclonically, it spins slower cyclonically as the air moves away from the high pressure center. The slower rotation rate then appears as an apparent rotation that is anticyclonic (clockwise in the northern hemisphere and counter-clockwise in the southern hemisphere). Note that there is a limit to how fast the anticyclonic rotation can be. It can't be any faster than the rotation rate of the earth. If it was faster, it would *really* be rotating anticyclonically (from the perspective of the stars). Consequently, the wind speeds around high pressure regions tend to be less than the wind speeds around low pressure regions.

What happens in the southern hemisphere?

In the southern hemisphere, the direction of rotation will be clockwise only because the earth spins clockwise as seen by someone looking down from above the southern hemisphere.

> ✍ The direction is actually the same in both hemisphere – it is just the vantage point that is different.

What if there is no high pressure or low pressure center?

If there is no high pressure or low pressure, we can still determine what the wind direction should be (given our assumptions) but noting the orientation of the isobars. Since the wind should go around the low or high pressure center, that is equivalent to saying that the wind goes parallel to the isobars. So, once you find the orientation of the isobars, the wind direction should be parallel to that.

In the northern hemisphere, the wind flows parallel to the isobars such that low pressure is on the left (as you face the direction the wind is going in). In the southern hemisphere, the low pressure would be on the right.

Check Point 18.4: In which direction does the air in a hurricane spin in the northern hemisphere: the same direction as the rotation of the earth, or opposite that direction? What about in the southern hemisphere?

18.4 Size of effect

Is this the reason why water always goes down the drain in a certain direction?

Not really. First of all, a simple experiment can show that the water doesn't always go down the drain counter-clockwise in the northern hemisphere and clockwise in the southern hemisphere. Either it is somewhat random or it is being forced to go one way (the same way in both hemispheres) because of the make-up of the sink or toilet (most toilets have water entry points that are oriented such that the water is forced to move in a certain direction).

Still, why is it that the effect is only seen on large scales, such as with low pressure systems? From the discussion above, it seems as though it should apply to all scales, even to the sink drain (and the land/sea breeze).

The answer is that it *would* apply *if* the fluid (air or water) is initially perfectly still relative to the earth. This was assumed in the discussion above. If the fluid is perfectly still then it would definitely have an initial cyclonic rotation relative to the stars (because the earth has a cyclonic rotation relative to the stars). Consequently, no matter what the size of the situation, bringing the fluid toward some center (like a drain) will make the fluid spin faster cyclonically (counter-clockwise in the northern hemisphere) to conserve angular momentum.

However, more likely, the fluid is not perfectly still. In such situations, it is possible that it initially had an anticyclonic rotation, not only relative to the earth's surface but also relative to the stars. Thus, upon moving toward some center, it will spin faster anticyclonically. As it turns out (see below), the likelihood of such a situation is higher for smaller systems.

To illustrate why, we need to compare the angular velocity of the wind around the low pressure center with the angular velocity of the Earth around its axis. So, we first need to determine the angular velocity of the Earth around its axis.

Check Point 18.5: Does the water always go down the drain in a particular direction north of the equator?

18.4.1 Angular velocity of the Earth

An object's **angular velocity** is the same as its **rotation rate** (or rotational speed).[vi] To indicate the angular velocity, one provides the angle the object

---

[vi] Technically, this is really the **angular speed** since the angular velocity is a vector and thus includes the direction. To indicate the direction of a rotation, one usually indicates the direction of the axis of rotation, according to the "right-hand" convention. One can also indicate the direction as clockwise or counter-clockwise, but then one also has to indicate the position of the observer.

rotates in a given amount of time. So, for instance, an object that completes one revolution in time T has an angular speed of $(360°)/T$ or $(2\pi \text{ radians})/T$.

A radian is a unit of angle, just like degrees or revolutions. Whereas there are 360 degrees in one revolution, there are 2π radians in one revolution (where the Greek letter π is used to indicate the value 3.1415927).[vii] You can use whatever unit you want unless you are using an equation in which one or the other unit was assumed. In this chapter, I will introduce equations that relate the angular velocity of the Earth with the velocity of a point on the Earth's surface. Those equations are derived assuming radians as the unit of angle. Thus, to use those equations, you need to use radians as the unit of angle.

> ✍ A lower-case omega (ω) will be used to indicate the angular velocity of an object, except in the case of the Earth, where a capital omega (Ω_e) will be used to emphasize its importance.

So what is the angular velocity of the Earth?

The Earth rotates once every 24 hours.

Since angular velocity is the angle the object rotates divided by the time to do so, the Earth's angular velocity is one revolution per day or, using unit abbreviations, 1 rev/d.

What about in units of revolutions per second?

In units of revolutions per second, that would be

$$\frac{1\ \text{rev}}{1\ \text{d}} = \frac{1\ \text{rev}}{(24\ \text{h}) \times (3600\ \text{s/h})}$$

or 1.157×10^{-5} rev/s.

What about in units of radians per second?

As mentioned above, we will eventually encounter expressions where we must use radians as our unit of angle. Thus, it becomes important to express the Earth's angular velocity in units of radians.

In units of rad/s, the Earth's angular velocity would be

$$(1.157 \times 10^{-5})\frac{1\ \text{rev}}{1\ \text{s}} = (1.157 \times 10^{-5})\frac{2\pi\ \text{rad}}{1\ \text{s}}$$

[vii] This value of π has been rounded. More precisely, it is the value of a circle's circumference to its diameter.

or 7.272×10^{-5} rad/s.

> ✍ Notice that the rotation rate of the Earth is slower than the hour hand of a clock (2 revolutions per day).

Before moving on, I'd like to modify the value in two ways.

First, I will write the units as s^{-1} instead of rad/s. The reason for this is that a radian is really a ratio of two similar length units (i.e., length of arc divided by length of radius). So, it is typically written without a unit.

Second, I will replace the "7" with a "9", so that the angular velocity is given as follows:

$$\Omega_{\mathrm{e}} = 7.292 \times 10^{-5}\ \mathrm{s}^{-1}$$

This is the value that you'll find in appendix B.4.

Why is the value equal to $7.292 \times 10^{-5}\ \mathrm{s}^{-1}$ instead of $7.272 \times 10^{-5}\ \mathrm{s}^{-1}$?

This apparently insignificant difference can be accounted for by recognizing that the Earth does not really rotate every 24 hours.

Huh? Why not?

Well, it is correct that twenty-four hours is the **solar day** and is the time it takes for the sun to "move" from its zenith (highest point above the horizon), set, rise and return to its zenith. This apparent motion is due to the earth spinning on its axis.

The earth, however, does not stay in one place as it spins. The earth is also orbiting the sun. We must take that motion into account. The time it takes for the earth to rotate once relative to the fixed stars is called the **sidereal day** (pronounced: sigh-DIR-ee-al). In problem 18.6, you will calculate the actual angular speed of the earth and show that it actually equals the value given in the appendix.

> ✍ Most people are more comfortable with specifying speed in terms of distance per time. However, when dealing with solid objects that are rotating, describing the rotation is more convenient than describing the linear motion because every point on the Earth shares the same rotation rate (see problem 18.4).

Check Point 18.6: The rotation rate of the Earth is given as $7.292 \times 10^{-5}\ \mathrm{s}^{-1}$.

How does this compare to the rotation rate of the hour hand on a clock?

18.4.2 Angular velocity of the air

The purpose of identifying the angular velocity of the Earth was so that we could investigate why big things, like hurricanes, always rotate cyclonically but little things, like drains, do not.

The key has to do with our initial assumption: that the fluid has to be stationary relative to the Earth before it is drawn inward toward the central point. If it is stationary, then relative to the stars the fluid has the same rotation rate as the Earth, which is cyclonic (and, as we've seen above, is close to 1 revolution per day).

WHAT IF THE FLUID IS NOT STATIONARY?

As long as the fluid is not moving too much, it would still have roughly the same rotation rate as the Earth and so would still be rotating cyclonically relative to the stars. As such, when it is drawn to some central point, it will still start to rotate cyclonically relative to the Earth.

HOW MUCH IS TOO MUCH?

Too much would be if the fluid was rotating *anticyclonically* relative to the stars. For this to occur, then *relative to the Earth* the fluid would have to have an angular velocity that was anticyclonic with a magnitude at least as large as Ω_e (to counter the rotation of the Earth).

Only then would motion toward a center result in greater anticyclonic motion, both relative to the stars and to the Earth.

HOW FAST IS THAT?

As we've already seen, that is equal to 1 revolution per day.

IS THAT HARD TO DO?

It depends on the size.

Suppose we have a low pressure center of radius 1000 km. The circumference is $2\pi r$ or about 6280 km. If the air in this system was rotating anticyclonically at a rate of 1 revolution per day, that would mean the air around its edge would have to travel 6280 km each day. In units of meters per second, that would be

$$\frac{6.28 \times 10^6 \text{ m}}{(24 \text{ h}) \times (3600 \text{ s/h})}$$

or 72.92 m/s.

This is very fast (over 160 mph) and it assumes this is the speed *before* the air starts to move toward the low center. We just don't see such ambient wind speeds and aren't likely to. More likely, the ambient wind speeds are calm, with magnitudes less than 2 m/s.

This means that we aren't likely to find a large (1000 km radius or us) region of air that has an anticyclonic rotation. And that means we aren't like to have it rotate anticyclonically when the air is drawn toward a center location.

Let's compare this to a region with a diameter of 50 cm (about the size of a sink). The circumference is πD or about 1.57 m. Again, for the air to have an anticyclonic angular velocity equal to 1 revolution per day, that would mean the wind at the edge would have to travel 1.57 m in a day. In units of meters per second, that would be

$$\frac{1.57 \text{ m}}{(24 \text{ h}) \times (3600 \text{ s/h})}$$

or 18×10^{-6} m/s.

The speed is so slow that imperceptible motions are sufficient to create rotation rates around the drain equal in magnitude to the earth's rotation rate around its axis.

Since we are very likely to find such speeds along the perimeter of our sink, it is not unusual to have anticyclonic motion when the fluid (air or water) moves toward the center.

Check Point 18.7: (a) If the air is calm, what is its angular velocity relative to the stars? (b) What is its angular velocity relative to the Earth?

Project

In this chapter, you examined the relationship between wind direction and pressure. Thus, for this chapter's project, you will look at maps that have both pressure and wind barbs plotted. In particular, you will have two maps (see course home page), both for the same observation time.

1. A surface map with isobars and wind (plotted as wind barbs)
2. A map at 18,000 ft, with pressure indicated at observation stations and wind barbs

Project 18.1: Observe the surface map. Roughly speaking, which of the following best describes the way the wind blows?
(a) Away from high pressure regions and toward low pressure regions (across the isobars).
(b) Away from low pressure regions and toward high pressure regions (across the isobars).
(c) Clockwise around high pressure regions and counter-clockwise around low pressure regions (parallel to the isobars).
(d) Counter-clockwise around high pressure regions and clockwise around low pressure regions (parallel to the isobars).
(e) Both (a) and (c).
(f) Both (b) and (d).

Project 18.2: Observe the map at 18,000 ft. The pressure is indicated by the three-digit number in the upper-right portion of each station model. Draw a rough analysis of the isobars, every 4 mb, indicating the regions of high and low pressure with H's and L's, respectively. Which of the following best describes the way the wind blows in the surrounding region?
(a) Away from high pressure regions and toward low pressure regions (across the isobars).
(b) Away from low pressure regions and toward high pressure regions (across the isobars).
(c) Clockwise around high pressure regions and counter-clockwise around low pressure regions (parallel to the isobars).
(d) Counter-clockwise around high pressure regions and clockwise around low pressure regions (parallel to the isobars).

(e) Both (a) and (c).
(f) Both (b) and (d).

Problems

Problem 18.1: Describe an observation you can make other than the direction of winds around low pressure that supports the contention that the earth rotates counter-clockwise when looking down above the northern hemisphere. For example, consider the apparent movement of the stars and sun.

Problem 18.2: (a) Why do we say that the earth rotates clockwise when looking down above the southern hemisphere?
(b) Based upon the discussion in the text, which way should the winds flow around low pressure in the southern hemisphere: clockwise or counter-clockwise?
(c) Test your prediction by examining a surface map (see course web page) for the southern hemisphere. Was your prediction correct?

Problem 18.3: Determine the speed of a point on the surface of the earth (relative to the earth's axis) for a point on each of the following latitudes:
(a) on the equator,
(b) at 45° N, and
(c) at the North Pole.
(d) Compare your answers in (a)-(c) with a typical surface wind speed (relative to the earth's surface) found on a surface map (e.g., 10 m/s). Which is bigger/smaller?

Problem 18.4: Determine the angular speed (in radians per second) of a point on the surface of the earth (relative to the earth's axis) for the following latitudes assuming the earth rotates once every 24 hours:
(a) on the equator,
(b) at 45° N, and
(c) at the North Pole.

Problem 18.5: (a) Suppose the u-component of the wind at 40° N is 10 m/s. Convert this to an angular velocity about the axis of the earth.
(b) How does this compare to the angular velocity of the earth?

Problem 18.6: (a) Calculate the number of solar days in one full orbit around the sun. Hint #1: It isn't 365, since there is an extra day every leap year.

Hint #2: If a leap year occurred every fourth year, that would mean there was 365.25 solar days in one full orbit. However, every 100 years, there is no leap year unless the year is divisible by 400.
(b) The number found in (a) is the time it takes to go completely around the sun. How many times does the earth rotate during that time? Hint: Add one to the number in (a) in order to account for the one additional rotation the earth makes with respect to the stars as it orbits the sun.
(c) The value in (a) is the time it takes for the earth to make the number of rotations obtained in (b). The ratio of (b) to (a) is the angular speed in rotations per day. Convert this ratio to units of radians per second. Compare your answer to the angular speed given in appendix B.4.

Problem 18.7: If the pressure increases in the positive x-direction then according to Newton's Second Law in what direction is the air accelerating (assuming no other forces are acting)?

Problem 18.8: Calculate the wind speed necessary to create an anticyclonic rotation relative to the surface for air at the edge of the following weather phenomena at 40 degrees north:
(a) A tornado that forms within a region of diameter equal to 20 km.
(b) A hurricane that forms within a region of diameter equal to 2000 km.
(c) What do these results imply about the likelihood of the motion being anticyclonic at these scales?

Problem 18.9: Draw a figure, like that in Figure 18.1, which illustrates the surface and upper-air pressures over water (x_1) and land (x_2) at night. Explain how your figure leads to the wind coming off the land at night.

19. Horizontal Wind Speed

19.1 Geostrophic balance

In this chapter, we examine the process for predicting the wind speed given the pressure distribution.

We'll start with an expression that gives the wind speed under a special condition called **geostrophic balance**. It is a lot easier to start with this and modify it for real situations than start with a really complicated relationship and then modify it for simplified situations.

The relationship states that the horizontal wind speed is proportional to the pressure gradient as follows:

$$V_{\mathrm{g}} = \frac{1}{f\rho}\nabla P \tag{19.1}$$

where f is a parameter known as the **Coriolis parameter**, which we will discuss later. The other variables you should already be familiar with. The density and pressure gradient are indicated as ρ and ∇P. The subscript "g" on the wind speed V is used to indicate the special circumstance under which the relationship holds.

This leads to several questions:

- Where does equation 19.1 come from?
- How should equation 19.1 be modified if there is friction of if the flow isn't straight?
- What is the meaning and value of f, the coriolis parameter?

Let's address those questions now.

First, the expression comes from an assumption that the pressure gradient force, which is pushing the air toward lower pressure, is balanced by the tendency for the air to "spin," being how the air is on a spinning Earth. Mathematically, we take into the account that the Earth is spinning via "extra" force terms. For example, Newton's second law, when expressed on a rotating reference frame and applied to the atmosphere, is given by equation D.2 in appendix D:[i]

$$\frac{dv}{dt} = -\frac{1}{\rho}\frac{\partial P}{\partial y} - F_y - (2\Omega_e u \sin\phi) - \frac{u^2 \tan\phi}{a} - \frac{vw}{a}$$

This looks complicated but let's identify the terms.

First, the left-hand side (dv/dt) represents the acceleration. By Newton's second law, the acceleration must equal the net force per mass. So, all of the terms on the right represent some kind of forcing.

The first term on the right-hand side represents the force due to the pressure gradient. It is negative because if $\partial P/\partial y$ is positive then the pressure is higher toward the north, which pushes the wind toward negative y.

The second term on the right-hand side represents the friction. To simplify our analysis, we will be ignoring that term.

The remaining terms on the right-hand side are all a consequence of the spinning Earth. As discussed in the previous chapter, relative to our reference frame the air appears to "start" to spin when it is brought toward the center of a low pressure system when, in actuality, it was already spinning. The extra terms on the right express the "forcing" that appears when the air moves. This is why those terms have u, v and/or w in them – they depend on how fast the air is moving.

In any event, we will ignore all of those extra terms except for the first one, since the others are so much smaller (they have a, the radius of the Earth, in the denominator). As mentioned above, we will also ignore friction. This results in the following simplification:

$$\frac{dv}{dt} = -\frac{1}{\rho}\frac{\partial P}{\partial y} - (2\Omega_e u \sin\phi)$$

The first term on the right means that a north-south pressure gradient will push the wind toward low pressure. The second term means that the rotation

---

[i] This expression is just for the y component. The expressions for the other components look similar and the differences are not crucial to our discussion at the moment.

of the Earth will cause the wind to turn in a circle. For example, motion toward the east (positive u) results in an acceleration toward the south (a negative dv/dt). It depends on the rotation rate of the Earth Ω_e, the speed of the wind u and the latitude ϕ.

The next simplification we will make is to assume that the air has been flowing long enough to be in balance. In other words, we'll assume that the air is no longer accelerating. That makes dv/dt equal to zero. Moving terms, we then have:

$$(2\Omega_\mathrm{e} u \sin\phi) = -\frac{1}{\rho}\frac{\partial P}{\partial y}$$

This assumption is called geostrophic balance.

Our final steps are to replace $2\Omega \sin\phi$ with f, the Coriolis parameter,

$$f = 2\Omega_\mathrm{e} \sin\phi. \tag{19.2}$$

and solve for the wind speed u. This gives the geostrophic balance expression (equation 19.1), which is written in vector form rather than just for one component.

> ✍ It is unfortunate that the convention is to use f to indicate this parameter since it is so similar to the f that is used to indicate a force per mass. Oh well. I will try to remember to use subscripts with the f to indicate a force per mass. You can then assume that an "f" with no subscripts means the Coriolis parameter.

What happened to the negative sign?

The negative sign just impacts the direction of the wind. Since we already know from the previous chapter what direction the wind flows, I removed it for simplicity.[ii]

What happens if there is friction?

If there is friction then the F term has to be included. That will lead to a slower speed. With that slower speed, the pressure gradient is able to bring the air in toward the low pressure and the wind is no longer parallel to the isobars.

[ii] The expression for u has the negative sign while the expression for v does not. This is because positive x is to the right of y, but negative y is to the right of positive x.

WHAT HAPPENS IF THE WIND IS MOVING AROUND A CENTER (LIKE A LOW PRESSURE CENTER)?

We've assumed no acceleration (i.e., dv/dt is zero). The only way the wind can move in a circle is if there is a force imbalance (i.e., non-zero acceleration). That means that the coriolis term no longer balances the pressure gradient term. Either the wind speed is greater than geostrophic (if it is turning away from low pressure) or less than geostrophic (if it is turning toward low pressure).

Check Point 19.1: Does the Coriolis parameter depend upon latitude? If so, is it greater near the poles or near the equator?

19.2 Determining the geostrophic wind speed

We use equation 19.1,

$$V_{\mathrm{g}} = \frac{1}{f\rho}\nabla P$$

to obtain the geostrophic wind speed.

Thus, to find the geostrophic wind speed, first determine the pressure gradient ∇P then divide that by the density of the air and the Coriolis parameter f valid at that latitude.

Example 19.1: Suppose the pressure changes by 4 mb over a distance of 300 km. What is the geostrophic wind speed near the surface if the latitude is 40 degrees?

Answer 19.1: First we calculate the Coriolis parameter from equation 19.2:

$$\begin{aligned} f &= 2\Omega_{\mathrm{e}} \sin\phi \\ &= 2(7.292 \times 10^{-5}\ \mathrm{s}^{-1}) \sin 40^\circ \end{aligned}$$

which gives a value of $9.4 \times 10^{-5}\ \mathrm{s}^{-1}$.

Assuming a density of 1.22 kg/m^3, we apply equation 19.1 to get

$$V_{\mathrm{g}} = \frac{1}{f\rho}\nabla P$$

$$= \frac{1}{(9.4 \times 10^{-5}\ \mathrm{s}^{-1})(1.22\ \mathrm{kg/m}^3)} \frac{(4 \times 10^2\ \mathrm{Pa})}{(300 \times 10^3\ \mathrm{m})}$$

which gives a value of 11.7 m/s.

> ✍ At 43 degrees latitude, the Coriolis parameter has a value close to $10^{-4}\ \mathrm{s}^{-1}$. For mid-latitudes (30 to 50 degrees), the Coriolis parameter varies from about $7 \times 10^{-5}\ \mathrm{s}^{-1}$ to about $11 \times 10^{-5}\ \mathrm{s}^{-1}$. In problem 19.1, you will determine the value of the Coriolis parameter at your location (using equation 19.2) but it should be around $10^{-4}\ \mathrm{s}^{-1}$.

Check Point 19.2: If the pressure gradient increases, what happens to the geostrophic wind speed?

19.3 Determining the geostrophic wind direction

While it is possible to determine the wind direction using the geostrophic wind equation, we don't need to – we've already figured out how to get the wind direction using the conceptual approach outlined in section 18.3.

First, we identify the direction of the pressure gradient force. That points toward low pressure. In the northern hemisphere, the wind flows counterclockwise around the low pressure center. That means it must flow with the lower pressure to its left. In the southern hemisphere, the wind flows with the lower pressure to its right.

Example 19.2: Suppose the pressure increases toward the east. What is the direction of the geostrophic wind in the northern hemisphere?

Answer 19.2: If the pressure increases toward the east, that means that the lower pressure is toward the west. In the northern hemisphere, the wind flows with the lower pressure to its left. That means it must be flowing toward the south.[iii]

[iii] In the northern hemisphere. In the southern hemisphere, the wind would be directed toward the south.

✍ The geostrophic should move parallel to the isobars rather than across.

What happens if the air encounters a region with a larger pressure gradient?

If the pressure gradient increases, then for a short time the air will turn toward the low pressure and speed up until it once again reaches geostrophic balance. During this time of acceleration, it is no longer under geostrophic balance and so equation 19.1 won't apply.

Check Point 19.3: In the northern hemisphere, if the geostrophic wind is toward the east, in which direction is the pressure gradient force?

19.4 Using constant-pressure charts

One of the simplifying assumptions in geostrophic balance is that there is no friction. This means that winds near the surface tend not to be geostrophic.

On the other hand, at higher elevations there is relatively little friction and geostrophic balance is more common (see, for example, project 19.2).

Unfortunately, it is difficult to calculate the geostrophic wind from upper-level maps using equation 19.1

$$V_{\mathrm{g}} = \frac{1}{f\rho}\nabla P$$

because the pressure typically isn't plotted on upper-level maps (and so one can't get the pressure gradient ∇P).

Instead of pressure, they plot the height. This is because most upper-air charts are at a particular pressure (e.g., 500 mb) instead of a particular height.[iv] In fact, rather than listing the pressures at a particular height, they list the heights of a particular pressure.

[iv] The reason for plotting pressure maps instead of height maps is historical. As a weather balloon rises, it measures the temperature and pressure. Consequently, it is easy to take all of the temperatures measured by various balloons when they are at a particular

Won't the pressure be the same everywhere on a pressure map?

Yes, but we don't plot the pressure. We plot the height at which that pressure is found.

How does that change the map?

For the most part, not at all. The overall pattern will be the same[v] whether we plot pressure on a constant height surface or height on a constant pressure surface. In fact, the lines of constant pressure on a constant-height map look almost exactly the same as the lines of constant height on a constant-pressure map, assuming the same approximate level is being plotted.

> Lines of constant height are called **contours**. Lines of constant pressure are called **isobars**.

How do we calculate the horizontal pressure gradient from such a map if the pressure is the same everywhere on the map?

We don't. Or, more precisely, we could but we won't.

Then how do we determine the geostrophic wind speed?

If we are given a constant pressure map, we calculate the geostrophic wind speed by rewriting equation 19.1 in terms of the height gradient $\nabla_p Z$ instead of the pressure gradient ∇P:

$$V_\mathrm{g} = \frac{g}{f}\nabla_p Z \tag{19.3}$$

Where does equation 19.3 come from?

To obtain equation 19.3, I replaced the pressure gradient ∇P in equation 19.1

$$V_\mathrm{g} = \frac{1}{f\rho}\nabla P$$

with $\rho g \nabla_p Z$.

---

pressure. To plot the temperature at a particular height requires that the balloons have a radar that detects the distance to the ground or someone to calculate the height by the process followed in problem 4.2. Before computers, this would be a long, involved process.

[v] In other words, the "lows" are in the same location and the "highs" are in the same location.

This is a valid replacement as long as the atmosphere is hydrostatic, since under hydrostatic balance

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

which was listed previously as equation 15.3.

You can see that the hydrostatic relationship looks very similar to our replacement

$$\nabla P = \rho g \nabla_p Z.$$

The difference is that the hydrostatic relationship does not involve the horizontal derivative. However, from calculus[vi], we know that

$$\left.\frac{\partial P}{\partial x}\right|_z = \left.\frac{\partial P}{\partial z}\right|_x \left.\frac{\partial z}{\partial x}\right|_P \tag{19.4}$$

where P (at constant z) and z (at constant P) are assumed to vary only along x. We then use the hydrostatic relationship to replace $\partial P/\partial z$ with $-\rho g$.

Thus, if you have a constant pressure map and you want to calculate the geostrophic wind, simply use equation 19.3 instead of equation 19.1.

What about the direction of the geostrophic wind?

Remember how the geostrophic wind direction is parallel to the isobars on a constant height map? Well, on a constant pressure map, the geostrophic wind direction is still parallel to the isolines except that the isolines now represent a particular height, not a particular pressure.

> ✍ Since g is constant with height whereas ρ is not, it turns out that it is easier to calculate the geostrophic wind from a constant pressure map, as opposed to a constant height map.

Check Point 19.4: If the pressure increases toward the east when plotted on a constant height map, where is the height highest on a constant pressure map (for roughly the same height)?

[vi] If you don't recognize this from your calculus class, ask your instructor to derive it.

19.5 Jet stream

With what we know now, we can explain the band of very high wind speeds that occur in the upper-troposphere coincident with the boundary between cold polar air and warm tropical air. This band of high winds is called the **jet stream**.

To explain the jet stream, first recall that the height of a pressure surface depends upon the mean temperature of the layer (see section 15.3.4). Thus, at the boundary between cold polar air and warm tropical air, there is a height gradient at upper levels between the low heights over the cold air and the high heights over the warm air.

Associated with this height gradient is a wind, with a speed that can be determined by the geostrophic wind equation (equation 19.3) and direction that runs parallel to the contours.[vii]

Check Point 19.5: Which way does the jet stream typically flow: toward the east or toward the west?

Project

In this chapter, you examined the relationship between wind direction and pressure. Thus, for this chapter's project, you will look at maps that have both pressure and wind barbs plotted. In particular, you will have two maps (see course home page), both for the same observation time.

1. A map at 18,000 ft, with pressure indicated at observation stations and wind barbs
2. A map of 500-mb heights, with wind barbs

[vii] You may recall that there is also a height gradient associated with the **sea breeze**, which develops due to temperature differences between land and sea (see section 18.2). Unlike the jet stream, the sea breeze tends to be directed across the contours (i.e., toward land or out to sea). This is because the domain of the sea breeze is too small for Coriolis force to have much of an effect (although there will be some effect which will make the winds have a small component parallel to the shoreline).

Project 19.1: Compare the height pattern on the 500-mb map you obtained with the pressure pattern you drew on the 18,000-ft map. Is the overall pattern the same?

Project 19.2: (a) Choose a wind observation on the 500-mb map. What is the speed (in m/s) and direction of the wind at that location?
(b) Measure the height gradient at that location.
(c) Using equation 19.3, determine the geostrophic wind speed corresponding to the height gradient calculated in part (b). How does that correspond to the observed wind speed?
(d) What is the direction of the geostrophic wind at that location. How does that correspond to the observed wind direction?
(e) Is the observed wind consistent with the analysis on page 270 (i.e., greater than geostrophic if going around high pressure and less than geostrophic if going around high pressure)?

Problems

Problem 19.1: Calculate the value of the Coriolis parameter, f, for your location.

Problem 19.2: For the same horizontal pressure gradient, how might the observed wind near the surface be different than the observed wind at higher altitudes? Consider the presence of friction near the ground.

Appendices

A. Lists of Abbreviations

A.1 Variable Abbreviations

$\vec{\nabla}$	horizontal gradient, $\hat{\imath}(\partial/\partial x)+\hat{\jmath}(\partial/\partial y)$
Δx	displacement
$\partial/\partial t$	rate of change at a particular location
η	viscosity
ϵ	ratio of water vapor molar mass to dry air molar mass ($\mu_{\mathrm{v}}/\mu_{\mathrm{d}}$)
Γ	lapse rate
λ	wavelength
λ	longitude
μ	molar mass (also molar, molecular or atomic weight)
ϕ	latitude
Φ	gravitational potential energy
π	ratio of circle circumference to diameter
ρ	density
σ	Stefan-Boltzmann constant
θ	potential temperature
ω	angular velocity
Ω_e	Earth's angular velocity

a	acceleration
a	albedo
A	area

b	a constant used in Teten's formula
b	absorptivity of solar radiation

c	specific heat (or specific heat capacity)
c	absorptivity of terrestrial radiation
c_{p}	specific heat of dry air (constant pressure)
c_{v}	specific heat of dry air (constant volume)
c_1	constant used in Planck's law
c_2	constant used in Planck's law
C	heat capacity

d/dt	rate of change following a parcel
e	vapor pressure
e_{s}	saturation vapor pressure
f	Coriolis parameter
f	frequency
F	force, force per mass or force per volume
F	irradiance (radiant flux density)
F_λ	spectral irradiance
g	gravitational force per mass
g_0	gravitational acceleration (at sea-level)
G	gravitational force constant
H	geopotential height
$\mathbf{i}$	eastward unit vector
I	moment of inertia
$\mathbf{j}$	northward unit vector
L	latent heat
L_{v}	latent heat of vaporization
L_{s}	latent heat of sublimation
L_{f}	latent heat of fusion
m	mass
m_{d}	mass of dry air
m_{T}	mass of total air
m_{v}	mass of water vapor
n	number of moles
N_A	Avogadro's number
P	pressure (total air)
P	power (or luminosity for radiation)
P_{d}	pressure of dry air
q	specific humidity ($m_{\mathrm{v}}/m_{\mathrm{T}}$)
Q	energy transferred (via heat processes)

r	mixing ratio ($m_\mathrm{v}/m_\mathrm{d}$)
r	distance between two objects or from axis
$\bar{r}$	mass-weighted average radius
R	radius
R	gas constant
R_e	radius of the earth
R_d	gas constant for dry air
R_v	gas constant for water vapor
R^*	universal gas constant
RH	relative humidity (e/e_s)

S_0	solar constant (average solar irradiance just outside Earth's atmosphere)
S_avg	solar insolation (solar irradiance averaged over surface of the Earth)

t	time
T	period
T	temperature
T_1	a constant used in Tetens formula
T_2	a constant used in Tetens formula
T_a	temperature of the atmosphere
T_e	temperature of the earth

u	eastward or x-component velocity
u_g	(eastward or x-component) geostrophic wind speed
U	internal energy

v	velocity or speed
v	northward or y-component velocity
v_g	(northward or y-component) geostrophic wind speed
$\vec{v}$	horizontal velocity vector
V	horizontal speed
V	volume

w	vertical velocity
W	work (thermodynamic)

x	eastward distance

y	northward distance

z	vertical distance
Z	geometric (or radar) height

A.2 Unit Abbreviations

atm	atmosphere (101.324 kPa)	pressure
bar	bar (100 kPa)	pressure
°C	degrees Celsius	temperature
h	hour (60 min)	time
hPa	hectoPascals (10^2 Pa)	pressure
Hz	Hertz (cycle/s)	frequency
J	Joules ($kg \cdot m^2/s^2$)	energy
K	Kelvin	temperature
kg	kilogram	mass
knot	nautical mile per hour	speed
m	meter	length
mb	millibars (10^{-3} bar)	pressure
mi	mile (1609 m)	distance
min	minute (60 s)	time
mol	mole	$6.02214199 \times 10^{23}$ molecules
N	Newtons ($kg \cdot m/s^2$)	force
nm	nautical mile (111.12 km = 60 nm)	distance
Pa	Pascals (N/m^2)	pressure
ppm	parts per million	ratio
ppb	parts per billion	ratio
rad	arc length per radius	angle
s	second	time
W	Watts (J/s)	power
%	parts per hundred	ratio
°lat	One degree of latitude (60 nm)	distance

B. Reference Information

B.1 General Constants

Avogadro's number (N_A)	$6.02214199 \times 10^{23}$ molecules
gravitational force constant (G)	6.67390×10^{-11} N · m^2/kg^2
Stefan-Boltzmann constant (σ)	5.67×10^{-8} W m^{-2} K^{-4}
speed of light in vacuum (c)	2.998×10^{8} m/s
universal gas constant (R^*)	8.3145 J mol^{-1} K^{-1}

B.2 Air Properties

gas constant for dry air (R_d)	287.06 J kg^{-1} K^{-1}
mass of the atmosphere	5.136×10^{18} kg
molar mass of dry air	28.964 g mol^{-1}
specific heat for dry air, constant pressure (c_p)	1004.67 J kg^{-1} K^{-1}
standard sea-level pressure	1013.25 mb or 1.01325×10^5 Pa
standard surface air density	1.217 kg/m^3
viscosity of air	about 2×10^{-5} Pa s
scale height of atmosphere	8.5 km

B.3 Water Properties

gas constant for water vapor (R_v)	461.52 J kg^{-1} K^{-1}
boiling point of water	100°C
freezing point of water	0°C
latent heat of vaporization for water at 0°C(L_v)	$\sim 2.5 \times 10^6$ J/kg
latent heat of sublimation for water at 0°C(L_s)	$\sim 2.83 \times 10^6$ J/kg
latent heat of fusion for water at 0°C(L_f)	$\sim 3.35 \times 10^5$ J/kg
molar mass of water vapor	18.015 g mol^{-1}
water density at 0°C	1000 kg/m^3

B.4 Astronomical Properties

angular velocity of the earth	$7.2921151467 \times 10^{-5}\text{s}^{-1}$
gravitational field (at sea-level)	9.80665 N/kg
mean earth-moon distance	3.80×10^{8} m
mean earth-sun distance	1.4959789×10^{11} m
mean radius of the earth	6.371×10^{6} m
mean radius of the sun	6.96×10^{8} m
mass of the earth	5.9736×10^{24} kg
mass of the sun	1.9891×10^{30} kg
solar constant	$1367.6\ \text{W/m}^2$

C. Derivations

C.1 First Law of Thermodynamics

The **first law of thermodynamics** provides a relationship between the changes in three things: the internal thermal energy U, the heating Q and the work W. It is usually written generically as follows (see equation 11.2):

$$dU = dQ - dW$$

In terms of pressure P and temperature T, it can be expressed as follows (see equation 11.3):

$$c_\mathrm{p} dT = \frac{dQ}{m} + \frac{dP}{\rho}$$

where $c_\mathrm{p} = 1004.67$ J kg^{-1} K^{-1} for dry air (no moisture).

In this appendix, I derive the latter from the former.

At first glance, it appears as though we just need to replace dU by $mc_\mathrm{p}dT$ and replace dW by $-VdP$. Then we can divide everything by m.

However, it isn't that simple. If we use the definitions of work and pressure[i], we find that dW equals PdV, not $-VdP$. If we make this replacement, we get

$$dU = dQ - PdV \tag{C.1}$$

This isn't what we want but it is still important because I'll be using this later.

However, this at least gives us a hint as to what to do. By the chain rule, we know that $PdV = d(PV) - VdP$. And, by the ideal gas law[ii], $PV =$

---

[i] The definition of work is $dW = Fdx$. Since $P = F/A$ (from the definition of pressure), we can replace F by $P \cdot A$. That means that dW is equal to $PAdx$. Since Ax is the volume, we have that $dW = PdV$.

[ii] There are several ways we can write the ideal gas law. The way we've used it here is $P = \rho RT$. Multiply both sides by V to get $PV = mRT$ (and recognize that $\rho = m/V$).

mRT and so $d(PV) = d(mRT) = mRdT$. Thus, dW can be replaced by $mRdT - VdP$.

Making this replacement, we get

$$\begin{aligned} dU &= dQ - dW \\ &= dQ - mRdT + VdP \end{aligned} \tag{C.2}$$

Although closer to what we want, this is still not what we need.

To complete the derivation and go from equation C.2 to equation 11.3, we need to learn what is meant by c_p. Thus, I'm going to take a short detour to explain what c_p represents. That will allow us to rewrite the expression appropriately.

- First, we define **heat capacity** C as dQ/dT, the amount of heat needed to raise the temperature (per degree of temperature).
- Next, we define the **specific heat**[iii] as the heat capacity per mass (i.e., C divided by m). We'll indicate the specific heat as c. Thus, $c = C/m = (1/m)(dQ/dT)$. Solve for dQ to get $dQ = mcdT$.
- The specific heat c will depend on how volume and pressure are allowed to change during the process. Usually, two values are determined; one obtained for constant pressure c_p and one obtained for constant volume c_v. These are very closely related but will have different values[iv].
- Suppose we use the version of the first law expressed in equation C.2 ($dU = dQ - mRdT + VdP$) and assume *pressure* is kept constant. In that case, $dP = 0$ and $dQ = mc_\mathrm{p}dT$ (where I used the p subscript on c because the volume is constant). The first law then simplifies to

$$dU = mc_\mathrm{p}dT - mRdT$$

Now that we have a version of the expression that involves c_p, it is relatively simple to complete the derivation.

[iii] This is sometimes called the **specific heat capacity**.

[iv] For dry air, c_p, c_v, and R form simple ratios, i.e., $c_\mathrm{p} : c_\mathrm{v} : R = 7{:}5{:}2$. This is a consequence of the diatomic make-up of the majority of atmospheric molecules (i.e., nitrogen and oxygen).

Basically, since we know that dU is equivalent to $(mc_{\mathrm{p}}dT - mRdT)$ (see last step above), we can replace dU in equation C.2

$$dU = dQ - mRdT + VdP$$

to get

$$mc_{\mathrm{p}}dT - mRdT = dQ - mRdT + VdP.$$

This simplifies to equation 11.3 if we cancel the $-mRdT$ on both sides, solve for dQ and divide by m.

C.2 Specific heats

During the derivation of the first law in the previous section, it was mentioned that there are two specific heat values: one at constant pressure (c_{p}) and one at constant volume (c_{v}). As you might expect, there is a rather simple relationship between them. You might surprised, however, to see just how simple it is.

To derive the relationship, note from before that we can use the version of the first law expressed in equation C.2 ($dU = dQ - mRdT + VdP$) and assume *pressure* is kept constant to get

$$dU = mc_{\mathrm{p}}dT - mRdT.$$

Alternately, we can instead use the version of the first law expressed in equation C.1 ($dU = dQ - PdV$) and assume *volume* is kept constant. In that case, $dV = 0$ and $dQ = mc_{\mathrm{v}}dT$ (where I used the v subscript on c because the volume is constant). The first law then simplifies to

$$dU = mc_{\mathrm{v}}dT. \tag{C.3}$$

These two expressions both give dU except that one is in terms of c_{p} and the other is in terms of c_{v}. Equating the two expressions gives us the relationship between c_{p} and c_{v}:

$$\begin{aligned}
dU &= dU \\
mc_{\mathrm{p}}dT - mRdT &= mc_{\mathrm{v}}dT \\
(c_{\mathrm{p}} - R)\cancel{mdT} &= c_{\mathrm{v}}\cancel{mdT} \\
(c_{\mathrm{p}} - R) &= c_{\mathrm{v}}
\end{aligned}$$

or, rearranging terms,

$$c_{\mathrm{p}} = c_{\mathrm{v}} + R.$$

Like I said, this is a rather simple relationship.

C.3 Poisson's equation

In this section, I derive equation 11.5,

$$T = T_0 \left(\frac{P}{P_0}\right)^{R/c_{\mathrm{p}}}$$

from

$$c_{\mathrm{p}} \frac{dT}{T} = R \frac{dP}{P}.$$

First, we recognize that

$$d \ln x = \frac{dx}{x}.$$

In other words, the derivative of the natural log of x is the inverse of x.

Since $dx/x = d \ln x$, we can replace dT/T and dP/P with $d \ln T$ and $d \ln P$, respectively. This gives

$$c_{\mathrm{p}} d \ln T = R d \ln P. \tag{C.4}$$

In other words, a small change in $\ln P$ results in a small change in $\ln T$. We no longer need to worry about what T and P are, as long as we consider the small changes in $\ln P$ and $\ln T$ (instead of dP and dT).

The large change in $\ln T$ (i.e., $\Delta \ln T$) is just the summation (i.e., integration[v]) of all of the small changes in $\ln T$ (i.e., $\int d \ln T$):

$$\begin{aligned}
\Delta(\ln T) &= \int d \ln T \\
&= \int \frac{R}{c_{\mathrm{p}}} d \ln P \\
&= \frac{R}{c_{\mathrm{p}}} \int d \ln P \\
&= \frac{R}{c_{\mathrm{p}}} \Delta(\ln P)
\end{aligned}$$

---

[v] As we did with equation 15.4 before.

We now take advantage of the fact[vi] that $\ln T - \ln T_0$ is equivalent to $\ln(T/T_0)$, where T and T_0 are the final and initial temperatures. Rewriting the expression in that form gives:

$$\ln\left[\frac{T}{T_0}\right] = \frac{R}{c_\mathrm{p}} \ln\left[\frac{P}{P_0}\right].$$

This can be further simplified by using the fact[vii] that $a \ln x$ is equal to $\ln x^a$. This allows us to write it as follows:

$$\ln\left[\frac{T}{T_0}\right] = \ln\left[\left(\frac{P}{P_0}\right)^{R/c_\mathrm{p}}\right] \tag{C.5}$$

Finally, we take the exponential[viii] of each side and then multiply by T_0 to get equation 11.5:

$$T = T_0 \left(\frac{P}{P_0}\right)^{R/c_\mathrm{p}}$$

where T_0 and P_0 are the old temperature and pressure and T and P are the new temperature and pressure.

C.4 Temperature tendency due to radiation

To obtain equation 13.2,

$$\frac{dT}{dt} = \frac{gF}{c_\mathrm{p}\Delta P}$$

we go back to the first law of thermodynamics (equation 11.3):

$$\frac{dQ}{m} = c_\mathrm{p} dT - \frac{dP}{\rho}.$$

As we did in section 12.3.4, we'll assume that the change in temperature is not accompanied by a corresponding change in pressure (i.e., the parcel

---

[vi] See footnote on page 291.

[vii] To see why this is, represent the number x as an exponential of the form e^b. In that case, $\ln x$ would be equal to b. That means that $a \ln x$ would be equal to ab. However, the same result would be obtained if we raised x to the power of a first. Raising e^b to the power of a gives e^{ab}. Taking the natural logarithm of that gives ab, just as before.

[viii] See footnote on page 292.

simply expands or contracts in order to maintain the same pressure with the surroundings). Consequently, $dP = 0$ and we have

$$\frac{dQ}{m} = c_\mathrm{p} dT.$$

Here we are interested in the time tendency, so we rewrite the expression as follows:

$$\frac{dT}{dt} = \frac{1}{c_\mathrm{p} m} \frac{dQ}{dt}.$$

In this case, dQ/dt represents the energy per time contributed by the solar radiation. Since F represents an energy per time that impinges upon a certain area, dQ/dt can be replaced by FA, where A is the area of our parcel:

$$\frac{dT}{dt} = \frac{FA}{c_\mathrm{p} m}.$$

Unfortunately, this expression is in terms of mass and area, which we'd like to get rid of. From equation 15.6, the mass per area (under hydrostatic equilibrium) is

$$\frac{m}{A} = \frac{\Delta P}{g}.$$

Making this replacement, we get equation 13.2.

C.5 Pressure decrease with height

To obtain equation 15.4,

$$P(z_2) = P(z_1) \exp\left(-\frac{g}{R\overline{T}} \Delta z\right)$$

use the ideal gas law (equation 7.2)

$$P = \rho R T$$

to replace ρ by P/RT in the hydrostatic relationship (equation 15.3)

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

Doing so and rearranging terms, we get:

$$\frac{dP}{P} = -\frac{g}{RT}dz$$

Since $dx/x = d\ln x$, we have

$$d\ln P = -\frac{g}{RT}dz$$

This means that if we look at a very tiny slice of atmosphere, of depth equal to dz, the product $(gdz)/(RT)$ for that tiny slice must equal the tiny change in $\ln P$ that exists between the top and bottom of that layer.

We could make the layer dz as deep as we want except that the temperature T changes as we go up in the atmosphere. So, there is no single value of T that is appropriate for the entire layer — the value of T is really only appropriate for a very thin layer. Anything thicker will not have a single temperature associated with it.

We can get around this problem by assuming we know the average temperature of a deep layer of depth Δz. In that case,

$$\Delta\ln P = -\frac{g}{R\bar{T}}\Delta z$$

where $\bar{T}$ is the average temperature of the layer.

We now set the top and bottom of the layer such that the top is at height z_2 and the bottom is at height z_2. The equation then becomes

$$\ln P(z_2) - \ln P(z_1) = -\frac{g}{R\bar{T}}\Delta z$$

Since the difference in two logs is equal to the log of the ratio[ix], we have

$$\ln\left(\frac{P(z_2)}{P(z_1)}\right) = -\frac{g}{R\bar{T}}\Delta z \tag{C.6}$$

Take the exponential of both sides[x] to get equation 15.4.

[ix] This can be seen by interpreting the meaning of the natural logarithm. Basically, it represents the exponent when the number is written in terms of e. For example, suppose two numbers A and B are equal to e^a and e^b, respectively. Then the natural logarithm of each would be a and b, respectively. The difference in the natural logarithms would be $a-b$. However, the same result would be obtained if we first divided A by B. That would give us e^a/e^b or e^{a-b}. Taking the natural logarithm of that would give us $a-b$ also. So, the log of A minus the log of B is equivalent to the log of A/B.

[x] The natural logarithm of e^a is a. Taking the exponential of that gives e^a, exactly

C.6 Mass of the atmosphere

To obtain the mass of the atmosphere, M, as expressed in equation 15.6,

$$M = \frac{A}{g} P(0)$$

we start with a mathematical representation of the total mass of the atmosphere:

$$M = \int_0^\infty dm$$

where dm is the infinitesimal mass of a slice of atmosphere dz thick and the $\int$ symbol means that we add up all of the tiny pieces of mass dm to get the total mass M. This is known as an integral.

The integral is from zero to infinity. The zero and infinity refer to the height in the atmosphere, we need to first convert the infinitesimal mass dm to be in terms of an infinitesimal height dh (so that we can do the integration). This can be done by using the definition of density:

$$\rho = dm/(Adz)$$

where A is the area of the slice (so that Adz is like an infinitesimal volume).

Using this expression to replace dm with ρAdz, the integral becomes

$$M = \int_0^\infty \rho Adz$$

The problem with this expression is that it has ρ in it and ρ varies with height. It would be better to replace it with something that does not depend on z.

To solve this problem, we use the hydrostatic relationship (equation 15.3). Instead of integrating over height, we can integrate over pressure. The hydrostatic relationship, then, provides the conversion from height to pressure.

At the surface, the height equals zero and the pressure equals $P(0)$. At the top, the height equals infinity and the pressure equals zero. Using the hydrostatic relationship, dz is equivalent to $dP/(\rho g)$.

---

what we had at the start. So, taking the exponential of a natural logarithm gives you the initial value (i.e., the two cancel each other out).

Making these replacements, we get

$$M = -\int_{P(0)}^{0} \frac{A}{g} dP$$

The area of each slice A is the surface area of the earth $4\pi R_\mathrm{e}^2$ and is constant. Assuming g to be constant also, we can take both A and g out of the summation to get

$$M = -\frac{A}{g}\int_{P(0)}^{0} dP$$

The sum of all the dP's just gives ΔP, which gives

$$M = -\frac{A}{g}\Delta P$$

It is left to the reader (see problem 15.4) to determine why ΔP is equal to $P(0)$ when M is the total mass of the atmosphere.

D. Equations of Motion

In this appendix, I derive the three equations of motion, one equation for each component direction.

$$\frac{du}{dt} = -\frac{1}{\rho}\frac{\partial P}{\partial x} - F_x - 2\Omega w\cos\phi + 2\Omega v\sin\phi - \frac{uw}{a} + \frac{uv\tan\phi}{a} \tag{D.1}$$

$$\frac{dv}{dt} = -\frac{1}{\rho}\frac{\partial P}{\partial y} - F_y - (2\Omega u\sin\phi) - \frac{u^2\tan\phi}{a} - \frac{vw}{a} \tag{D.2}$$

$$\frac{dw}{dt} = -\frac{1}{\rho}\frac{\partial P}{\partial z} - g - F_z + (2\Omega u\cos\phi) + \frac{u^2+v^2}{a} \tag{D.3}$$

These equations look very complicated but they are obtained by applying Newton's second law to the atmosphere, using a coordinate system where u is toward the east, v is toward the north and w is upward.

The terms on the right-hand-side of each equation can be organized into three types:

Type	Examples	Explanation
Forces (per mass)	$(1/\rho)(\partial P/\partial x)$	pressure gradient force
	F_x	friction
	g	gravity
Coriolis terms	$2\Omega v\sin\phi$	Due to rotating reference frame
Curvature terms	uw/a	Due to rotating reference frame Usually less significant than Coriolis terms

In my derivation, I will first focus on the Coriolis and curvature terms. These are obtained by writing Newton's second law in spherical coordinates for a rotating reference frame.

As you should already be aware, Newton's second law relates the forces acting on an object with the motion of the object.

$$\vec{a} = \frac{\vec{F}_{\text{net}}}{m}$$

We will first assume there are no external forces acting on our object (the air parcel).

$$\vec{a} = 0$$

This isn't true, of course, and so we'll add in the forces later.

If the acceleration is zero, you would think that the resulting motion of the object is pretty straightforward in that the velocity must remain constant. That is true. However, since the Earth is a sphere (or close to it), we will be using spherical coordinates and in spherical coordinates the velocity along each spherical component need not be constant, even if the total vector acceleration is zero.

> ✍ In spherical coordinates, there are three components: two angular coordinates (latitude and longitude) and a radial coordinate.

To see why this is, consider a stone being swung in a horizontal circle via a string. If the string is suddenly released, the object will move in a straight line. However, in spherical coordinates, the radial position of the object will be seen to increase with time as the object moves away from the center. Conversely, the angular motion of the object will slow.

So, we must see what $\vec{a} = 0$ means in terms of our spherical coordinates.

I will do this two ways. First I will use two versions of Newton's second law that are used with circular motion. I'll do this because you are probably already familiar with these versions and their application to spherical coordinates is straightforward. My second method will use a more rigorous, mathematical approach.

> ✍ For the derivation, we define our three component directions as follows. The $\hat{\imath}$ direction is directed along a line of latitude toward the east (i.e., toward increasing longitude λ). The $\hat{\jmath}$ direction is directed along a line of longitude toward the north (i.e., toward increasing latitude ϕ). The $\hat{k}$ direction is directed upward, away from the center of the Earth (i.e., toward increasing altitude r).

D.1 Using circular motion expressions

In circular (or polar) coordinates, there are two components: an angular (or tangential) component and a radial coordinate. The radial distance is indicated as r and the velocity in that direction is v_r. The angular component is indicated as θ (angular measure) or $r\theta$ (tangential measure) and the velocity in that direction is indicated as ω (angular measure) or v_t (tangential measure).

If there are no radial forces acting on the object, the radial component of the velocity will change as follows:

$$\frac{dv_r}{dt} = \frac{v_t^2}{r} \tag{D.4}$$

where v_r is the radial velocity and v_t is the tangential velocity (see section D.6 for derivation). This is essentially Newton's second law applied in the radial direction, recognizing that a force per mass equal to v_t^2/r directed toward the rotation axis is necessary in order to keep the object in circular motion (where v_r equals zero). Conversely, with no net force applied, the object will accelerate away from the axis with an acceleration equal to v_t^2/r.

If there are no tangential forces acting on the object, the tangential component of the velocity will change as follows:

$$\frac{d(rv_t)}{dt} = 0 \tag{D.5}$$

This is essentially Newton's second law applied to rotation (net torque equals the moment of inertia times the angular acceleration). With no net force applied, angular momentum (rv_t) is conserved (see section D.6 for derivation).

Since there are two angular coordinates (latitude and longitude) in spherical coordinates, each expression needs to be applied twice (once for each angular coordinate) if we are to use these relationships in spherical coordinates. This I will do in the next two sections. First I will apply it around a line of latitude (where longitude varies) and then I will apply it around a line of longitude (where latitude varies).

D.2 Motion around line of latitude

When applied around a line of latitude, the Earth's axis is the coordinate axis. We'll use r as the distance from the object to the axis. For our purposes,

r is equal to $a \cos \phi$ (within a very small error[i] where a is the radius of the Earth and ϕ is the latitude.

The tangential direction is $\hat{\imath}$. Therefore, $v_t = u + \Omega r$, where I've added the motion of the Earth, since u is measured relative to the Earth's rotation.

We will assume there is no net force acting on our object. We'll add in the appropriate forces later.

D.2.1 Radial motion

When applied to this situation, equation D.4 is written as follows:

$$\frac{dv_r}{dt} = \frac{(\Omega r + u)^2}{r}$$

and, expanding, we get

$$\frac{dv_r}{dt} = (\Omega^2 r) + (2\Omega u) + \frac{u^2}{r}.$$

Since the coordinate axis is the Earth's axis, the radial direction has two components: $\hat{k}$, which is the vertical direction, and $-\hat{y}$, which for Northern latitudes (positive ϕ) is toward the south. Therefore, $w = v_r \cos \phi$ and $v = -v_r \sin \phi$.

In addition, to a very good approximation, $r = a \cos \phi$.

We can then obtain an expression for w:

$$\begin{aligned} dw/dt &= \partial(v_r \cos \phi)/\partial t \\ &= \Big((\Omega^2 a \cos \phi) + (2\Omega u) + (u^2)/(a \cos \phi)\Big) \cos \phi \\ &= (\Omega^2 a \cos^2 \phi) + (2\Omega u \cos \phi) + (u^2/a) \end{aligned} \tag{D.6}$$

---

[i] The error comes about because of two things. First, the Earth is not a perfect sphere. Consequently, the distance from the center of the Earth to the Earth's surface is not always equal to the same number. Second, we are studying the atmosphere and, as such, we are not restricted to just the Earth's surface. In both cases, the variation is small compared with the average radius of the Earth.

And a similar expression for v:

$$\begin{aligned} dv/dt &= \partial(-v_r \sin\phi)/\partial t \\ &= -\Big((\Omega^2 a\cos\phi) + (2\Omega u) + (u^2)/(a\cos\phi)\Big)\sin\phi \\ &= -(\Omega^2 a \sin\phi\cos\phi) - (2\Omega u\sin\phi) - (u^2\tan\phi/a) \qquad \text{(D.7)} \end{aligned}$$

D.2.2 Tangential motion

When applied to this situation, equation D.5 is written as follows:

$$\frac{d(\Omega r^2 + ur)}{dt} = 0$$

so

$$\Omega\frac{dr^2}{dt} + r\frac{du}{dt} + u\frac{\partial r}{\partial t} = 0$$

or

$$2r\Omega v_r + r\frac{du}{dt} + uv_r = 0$$

or

$$\frac{du}{dt} = -uv_r/r - 2\Omega v_r$$

Again, the radial direction has two components: $\hat{k}$, which is the vertical direction, and $-\hat{y}$, which for Northern latitudes (positive ϕ) is toward the south.

However, whereas before we had that $w = v_r\cos\phi$ and $v = -v_r\sin\phi$, in this case we have that $v_r = w\cos\phi - v\sin\phi$. The roles are reversed because we are now examining how the vertical motion influences the tangential motion – we now need the radial portions of w and v. Before, we needed the vertical and horizontal portions of v_r.

Again, to a very good approximation, $r = a\cos\phi$.

With these substitutions, we have

$$\frac{du}{dt} = -\frac{uw\cos\phi}{a\cos\phi} + \frac{uv\sin\phi}{a\cos\phi} - 2\Omega w\cos\phi + 2\Omega v\sin\phi$$

or

$$\frac{du}{dt} = -\frac{uw}{a} + \frac{uv\tan\phi}{a} - 2\Omega w\cos\phi + 2\Omega v\sin\phi \qquad \text{(D.8)}$$

D.3 Motion around line of longitude

When applied around a line of longitude, the center of the circle is the center of the Earth. Consequently, the distance r, which represents the distance from the object to the axis is equal to a, the radius of the Earth, within a very small error.

The radial direction in this case is $\hat{k}$, the vertical direction, and the tangential direction is $\hat{y}$, which is toward the north. Therefore, $w = v_r$ and $v = v_t$.

> ✍ Because this application only involves w and v, u will not appear in the two expressions we obtain.

D.3.1 Radial motion

When applied to this situation, equation D.4 is written as follows:

$$\frac{dw}{dt} = \frac{v^2}{a} \tag{D.9}$$

where I've assumed all other vertical forces are in balance (e.g., hydrostatic balance).

D.3.2 Tangential motion

When applied to this situation, equation D.5 is written as follows:

$$\frac{d(rv)}{dt} = \frac{rF_t}{m}$$

With no tangential forces, we have that

$$\frac{d(rv)}{dt} = 0$$

so

$$v\frac{dr}{dt} + r\frac{dv}{dt} = 0$$

or

$$\frac{dv}{dt} = -\frac{vw}{a} \tag{D.10}$$

D.4 Acceleration in spherical coordinates

By applying the two expressions, one for radial motion (equation D.4) and one for tangential motion (equation D.5), to the two coordinates (latitude and longitude), we obtained five expressions. These five expressions can be combined into three expressions, one for each component direction.

Equation D.8 gives us the acceleration in the $\hat{\imath}$ direction:

$$\frac{du}{dt} = -2\Omega w \cos\phi + 2\Omega v \sin\phi - \frac{uw}{a} + \frac{uv\tan\phi}{a} \tag{D.11}$$

Combining equations D.7 and D.10 gives us the acceleration in the $\hat{\jmath}$ direction.

$$\frac{dv}{dt} = -(\Omega^2 a \sin\phi\cos\phi) - (2\Omega u \sin\phi) - \frac{u^2\tan\phi}{a} - \frac{vw}{a} \tag{D.12}$$

Combining equations D.6 and D.9 gives us the acceleration in the $\hat{k}$ direction.

$$\frac{dw}{dt} = (\Omega^2 a \cos^2\phi) + (2\Omega u \cos\phi) + \frac{u^2+v^2}{a} \tag{D.13}$$

> ✍ The terms that are divided by a are known as the **curvature** terms. The terms that include Ω^2 (the square of the rotation of the Earth) are called the **centrifugal** terms.[ii] The remaining terms, which include Ω, are known as the **Coriolis** terms.

D.5 Mathematical Derivation

A more straightforward, but less intuitive approach, is to write $\vec{a}$ in spherical coordinates.

As before, we define our three component directions as follows. The $\hat{\imath}$ direction is directed along a line of latitude, toward increasing longitude λ.

[ii] The word "centrifugal" is from the Latin "fugere", which means "to flee" (as in "fugitive"). Knowing the etymology of the word is useful because it helps us remember the direction of a force that acts centrifugally (i.e., it is "center fleeing" or directed away from the center of rotation). The opposite direction, toward the center, is called the "centripetal" direction. The word "centripetal" is from the Latin "petere", which means "to go to" or "seek."

The $\hat{\jmath}$ direction is directed north along a line of longitude, toward increasing latitude ϕ. The $\hat{k}$ direction is directed away from the center of the Earth, toward increasing altitude r.

> Mathematical textbooks tend to use $\hat{\phi}$, $\hat{\theta}$ and $\hat{r}$ instead of $\hat{\imath}$, $\hat{\jmath}$ and $\hat{k}$ (and ϕ and θ instead of λ and ϕ).

In our coordinate system, the velocity is expressed as follows:

$$\vec{v} = u_* \hat{\imath} + v\hat{\jmath} + w\hat{k}$$

where I expressed the velocity along the latitude direction as u_* because u is used to represent the motion of the object relative to the reference frame. In our case, the reference frame (the Earth) is moving and so u_* includes the motion of the Earth (equal to $\Omega r \cos\phi$, where Ω is the rotational velocity of the Earth).

Even if the acceleration is zero (when the net force is zero), that does not mean that u, v and w must each remain constant. That is because in time the component directions $\hat{\imath}$, $\hat{\jmath}$ and $\hat{k}$ change. For example, if w is zero (so that the object remains on the surface of a sphere), a non-constant u or v will lead to a latitude and/or longitude change which, in turn, changes the $\hat{k}$ direction. Mathematically, then, we have

$$\begin{aligned}
d\vec{v}/dt &= d(u_*\hat{\imath})/dt + d(v\hat{\jmath})/dt + d(w\hat{k})/dt \\
&= (du_*/dt)\hat{\imath} + u_*(d\hat{\imath}/dt) + (dv/dt)\hat{\jmath} + v(d\hat{\jmath}/dt) + (dw/dt)\hat{k} + w(d\hat{k}/dt) \\
&= (du_*/dt)\hat{\imath} + (dv/dt)\hat{\jmath} + (dw/dt)\hat{k} + u_*(d\hat{\imath}/dt) + v(d\hat{\jmath}/dt) + w(d\hat{k}/dt)
\end{aligned}$$

and since we are assuming no net force acting on the object, we have

$$0 = (du_*/dt)\hat{\imath} + (dv/dt)\hat{\jmath} + (dw/dt)\hat{k} + u_*(d\hat{\imath}/dt) + v(d\hat{\jmath}/dt) + w(d\hat{k}/dt).$$

or

$$(du_*/dt)\hat{\imath} + (dv/dt)\hat{\jmath} + (dw/dt)\hat{k} = -u_*(d\hat{\imath}/dt) - v(d\hat{\jmath}/dt) - w(d\hat{k}/dt). \quad \text{(D.14)}$$

What we need, then, are relationships for how $\hat{\imath}$, $\hat{\jmath}$ and $\hat{k}$ depend upon time. There are several ways to do this. However, in each case, it involves first identifying how these component directions depend upon ϕ (latitude) and λ

(longitude). There is no dependence on r.

$$\begin{aligned}
\partial\hat{\imath}/\partial\phi &= 0 \\
\partial\hat{\imath}/\partial\lambda &= -\hat{k}\cos\phi + \hat{\jmath}\sin\phi \\
\partial\hat{\jmath}/\partial\phi &= -\hat{k} \\
\partial\hat{\jmath}/\partial\lambda &= -\hat{\imath}\sin\phi \\
\partial\hat{k}/\partial\phi &= \hat{\jmath} \\
\partial\hat{k}/\partial\lambda &= \hat{\imath}\cos\phi
\end{aligned}$$

These relationships can be obtained algebraically or geometrically (not shown).

The next step is to expand each time derivative as follows:

$$\begin{aligned}
\frac{d\hat{\imath}}{dt} &= \frac{d\hat{\imath}}{d\phi}\frac{d\phi}{dt} + \frac{d\hat{\imath}}{d\lambda}\frac{d\lambda}{dt} \\
&= 0 + (-\hat{k}\cos\phi + \hat{\jmath}\sin\phi)\dot{\lambda} \\
&= \sin\phi\dot{\lambda}\hat{\jmath} - \cos\phi\dot{\lambda}\hat{k} \\
\frac{d\hat{\jmath}}{dt} &= \frac{d\hat{\jmath}}{d\phi}\frac{d\phi}{dt} + \frac{d\hat{\jmath}}{d\lambda}\frac{d\lambda}{dt} \\
&= -\hat{k}\dot{\phi} - \hat{\imath}\sin\phi\dot{\lambda} \\
&= -\sin\phi\dot{\lambda}\hat{\imath} - \dot{\phi}\hat{k} \\
\frac{d\hat{k}}{dt} &= \frac{d\hat{k}}{d\phi}\frac{d\phi}{dt} + \frac{d\hat{k}}{d\lambda}\frac{d\lambda}{dt} \\
&= \hat{\jmath}\dot{\phi} + \hat{\imath}\cos\phi\dot{\lambda} \\
&= \cos\phi\dot{\lambda}\hat{\imath} + \dot{\phi}\hat{\jmath}
\end{aligned}$$

We now rewrite the variables where

$$\begin{aligned}
\dot{\lambda} &= u_*/(r\cos\phi) \\
\dot{\phi} &= v/r \\
\dot{r} &= w
\end{aligned}$$

This gives us the following expressions:

$$\frac{d\hat{\imath}}{dt} = (u_*\tan\phi/r)\hat{\jmath} - (u_*/r)\hat{k}$$

$$\begin{aligned}
\frac{d\hat{j}}{dt} &= -(u_* \tan\phi/r)\hat{\imath} - (v/r)\hat{k} \\
\frac{d\hat{k}}{dt} &= (u_*/r)\hat{\imath} + (v/r)\hat{\jmath}
\end{aligned}$$

Plugging in to what we had before, the right side of equation D.14 can be written as

$$-u_*((u_* \tan\phi/r)\hat{\jmath} - (u_*/r)\hat{k}) - v(-(u_* \tan\phi/r)\hat{\imath} - (v/r)\hat{k}) - w((u_*/r)\hat{\imath} + (v/r)\hat{\jmath})$$

or, rearranging terms,

$$(u_* v \tan\phi/r - wu_*/r)\hat{\imath} + (-u_*^2 \tan\phi/r - wv/r)\hat{\jmath} + (u_*^2/r + v^2/r)\hat{k}).$$

Since the two sides of equation D.14 must be equal for each component direction, this leads to three equations of motion:

$$\begin{aligned}
du_*/dt &= (u_* v \tan\phi/r) - (u_* w/r) \\
dv/dt &= -(u_*^2 \tan\phi/r) - (vw/r) \\
dw/dt &= (u_*^2/r) + (v^2/r)
\end{aligned}$$

The final piece of the puzzle is to replace u_* with $u + \Omega r \cos\phi$. I'll do this in two steps. First, I'll replace it in du_*/dt to get

$$\begin{aligned}
du_*/dt &= d(u + \Omega r \cos\phi)/dt \\
&= du/dt + (\Omega \cos\phi \dot{r}) - (r\Omega \sin\phi \dot{\phi}) \\
&= du/dt + (\Omega w \cos\phi) - (\Omega v \sin\phi)
\end{aligned}$$

Then I will make the replacement elsewhere. The three expressions then can be written as follows:

$$\begin{aligned}
du/dt &= -(\Omega w \cos\phi w) + (\Omega v \sin\phi) + (u + \Omega r \cos\phi)(v \tan\phi/r) - (u + \Omega r \cos\phi)(w/r) \\
&= -(\Omega w \cos\phi) + (\Omega v \sin\phi) + (uv \tan\phi/r) + \Omega v \sin\phi - (uw/r) - (\Omega w \cos\phi) \\
&= -(2\Omega w \cos\phi) + (2\Omega v \sin\phi) + (uv \tan\phi/r) - (uw/r) \\
dv/dt &= -((u + \Omega r \cos\phi)^2 \tan\phi/r) - (vw/a) \\
&= -(u^2 \tan\phi/r) - (2\Omega u \sin\phi) - (\Omega^2 r \sin\phi \cos\phi) - (vw/a) \\
dw/dt &= ((u + \Omega r \cos\phi)^2/r) + (v^2/a) \\
&= (u^2/a) + (2\Omega u \cos\phi) + (\Omega^2 r \cos^2\phi) + (v^2/a)
\end{aligned}$$

Replacing r with a (the radius of the Earth) leads to the same equations as before (see equations D.11, D.12 and D.13).

D.6 Polar Coordinates

For reference, in this section I derive equations D.4 and D.5, which represent the acceleration in polar (circular) coordinates given no net force.

First, we define our two component directions: $\hat{\theta}$ and $\hat{r}$.

In our coordinate system, the velocity is expressed as follows:

$$\vec{v} = v_t\hat{\theta} + v_r\hat{r}$$

Even if the acceleration is zero (when the net force is zero), that does not mean that v_t and v_r must each remain constant. That is because in time the component directions $\hat{\theta}$ and $\hat{r}$ change. For example, if v_r is zero (so that the object remains on the surface of a circle), a non-constant v_t will lead to a change in the direction of motion (changes in $\hat{\theta}$ and $\hat{r}$). Mathematically, then, we have

$$\begin{aligned} d\vec{v}/dt &= d(v_t\hat{\theta})/dt + d(v_r\hat{r})/dt \\ &= (dv_t/dt)\hat{\theta} + v_t(d\hat{\theta}/dt) + (dv_r/dt)\hat{r} + v_r(d\hat{r}/dt) \\ &= (dv_t/dt)\hat{\theta} + (dv_r/dt)\hat{r} + v_t(d\hat{\theta}/dt) + v_r(d\hat{r}/dt) \end{aligned}$$

and since we are assuming no net force acting on the object, we have

$$0 = (dv_t/dt)\hat{\theta} + (dv_r/dt)\hat{r} + v_t(d\hat{\theta}/dt) + v_r(d\hat{r}/dt).$$

What we need, then, are relationships for how $\hat{\theta}$ and $\hat{r}$ depend upon time. There are several ways to do this. However, in each case, it involves first identifying how these component directions depend upon θ. There is no dependence on r.

$$\begin{aligned} \partial\hat{\theta}/\partial\theta &= -\hat{r} \\ \partial\hat{r}/\partial\theta &= \hat{\theta} \end{aligned}$$

These relationships can be obtained algebraically or geometrically (not shown).

The next step is to expand each time derivative as follows:

$$\begin{aligned} \frac{d\hat{\theta}}{dt} &= \frac{d\hat{\theta}}{d\theta}\frac{d\theta}{dt} + \frac{d\hat{\theta}}{dr}\frac{dr}{dt} \\ &= -\hat{r}\dot{\theta} \\ \frac{d\hat{r}}{dt} &= \frac{d\hat{r}}{d\theta}\frac{d\theta}{dt} + \frac{d\hat{r}}{dr}\frac{dr}{dt} \\ &= +\hat{\theta}\dot{\theta} \end{aligned}$$

We now rewrite the variables where

$$\begin{aligned}\dot{\theta} &= v_t/r \\ \dot{r} &= v_r\end{aligned}$$

This gives us the following expressions:

$$\begin{aligned}\frac{d\hat{\theta}}{dt} &= -(v_t/r)\hat{r} \\ \frac{d\hat{r}}{dt} &= (v_t/r)\hat{\theta}\end{aligned}$$

Plugging in to what we had before, we get

$$(dv_t/dt)\hat{\theta} + (dv_r/dt)\hat{r} = (v_t^2/r)\hat{r} - (v_r v_t/r)\hat{\theta}.$$

Since the two sides of the expression must be equal for each component direction, this leads to two equations of motion:

$$\begin{aligned}dv_r/dt &= v_t^2/r \\ dv_t/dt &= -(v_r v_t/r)\end{aligned}$$

The first expression is the same as equation D.4 with no net force. The second expression can be seen to be the same as equation D.5 (with no net force) by first rewriting it as

$$r\frac{dv_t}{dt} + v_r v_t = 0$$

then replacing v_r by dr/dt,

$$r\frac{dv_t}{dt} + v_t\frac{dr}{dt} = 0$$

and finally using the chain rule to get

$$\frac{d(rv_t)}{dt} = 0$$

which is conservation of angular momentum.

D.7 Adding in appropriate forces

In this section, we add the forces acting on the air. If you compare the equations we've obtained so far for the acceleration in spherical coordinates (equations D.11 to D.13):

$$\begin{aligned} du/dt &= -(2\Omega w \cos\phi) + (2\Omega v \sin\phi) + (uv\tan\phi/a) - (uw/a) \\ dv/dt &= -(\Omega^2 r \sin\phi\cos\phi) - (2\Omega u \sin\phi) - (u^2\tan\phi/a) - (vw/a) \\ dw/dt &= (\Omega^2 r\cos^2\phi) + (2\Omega u\cos\phi) + (u^2/a) + (v^2/a) \end{aligned}$$

with the equations of motion introduced at the beginning of the appendix (equations D.1 to D.3):

$$\begin{aligned} \frac{du}{dt} &= -\frac{1}{\rho}\frac{\partial P}{\partial x} - F_x - 2\Omega w\cos\phi + 2\Omega v\sin\phi - \frac{uw}{a} + \frac{uv\tan\phi}{a} \\ \frac{dv}{dt} &= -\frac{1}{\rho}\frac{\partial P}{\partial y} - F_y - (2\Omega u\sin\phi) - \frac{u^2\tan\phi}{a} - \frac{vw}{a} \\ \frac{dw}{dt} &= -\frac{1}{\rho}\frac{\partial P}{\partial z} - g - F_z + (2\Omega u\cos\phi) + \frac{u^2+v^2}{a} \end{aligned}$$

you'll see that they differ into two ways.

The first difference is that the equations of motion include the force terms, which makes sense since the acceleration equations were derived assuming no forces were acting. The second difference is that the equations of motion lack the centrifugal terms (i.e., the terms with Ω^2) that are present in the acceleration equations.

The second difference can be explained as follows.

If we look at the three acceleration equations, we see that when an object is at rest relative to the Earth ($u = v = w = 0$), we still have an acceleration upward and toward the south:[iii]

$$\frac{\partial u}{\partial t} = 0$$

[iii] This acceleration is traditionally associated with an apparent **centrifugal force** that shows up when an object is observed from a rotating reference frame. It is called a centrifugal force because there appears to be a force acting on the object that is pushing the object away from the axis or **centrifugally**. Keep in mind that this force only appears when making the observation from the rotating reference frame. If viewed from a stationary frame or one that is not accelerating, such a force will not appear because the motion can be explained by Newton's laws without utilizing this force.

$$\begin{aligned}\frac{\partial v}{\partial t} &= -(\Omega^2 a \sin\phi \cos\phi) \\ \frac{\partial w}{\partial t} &= (\Omega^2 a \cos^2\phi)\end{aligned}$$

To interpret what this means, let's suppose that we dropping an object, like a rock, from rest. In that case, there is a gravitational force acting on the rock, acting in the $-\hat{k}$ direction with a magnitude equal to Gm_{earth}/a^2. According to the acceleration equations, countering that force (per mass) in the $-\hat{k}$ direction is an acceleration in the $+\hat{k}$ direction equal to $(\Omega^2 a \cos^2\phi)$ and an acceleration in the $-\hat{j}$ direction equal to $(\Omega^2 a \sin\phi \cos\phi)$.

In other words, the rock will not accelerate in the $-\hat{k}$ direction with an acceleration equal to Gm_{earth}/a^2. Instead, it will accelerate in a direction that is slightly to the $-\hat{j}$ side of $-\hat{k}$, with a magnitude slightly less than Gm_{earth}/a^2.

> The actual acceleration will have a direction that deviates from $-\hat{k}$ by tne inverse tangent of
>
> $$\frac{(\Omega^2 a \sin\phi \cos\phi)}{(Gm_{\text{earth}}/a^2) - (\Omega^2 a \cos^2\phi)}$$
>
> and has a magnitude equal to the square root of
>
> $$(\Omega^2 a \sin\phi \cos\phi)^2 + [(Gm_{\text{earth}}/a^2) - (\Omega^2 a \cos^2\phi)]^2$$
>
> Plugging in appropriate values for the Earth, one gets a magnitude of 9.803 m/s^2 and a deviation of 0.1° at a latitude of 41°N.

Since we define "downward" as being the direction that a rock, released at rest, accelerates in free-fall, it turns out that our "downward" is slightly shifted from $-\hat{k}$, which we defined as being directed radially in toward the center of the Earth.

The way we deal with this problem is by redefining $\hat{k}$ to be aligned with our local perception of vertical. It involves a small rotation to our coordinate system (our $\hat{j}$ will also be shifted slightly) and it allows us to combine the gravitational and centrifugal terms (which would normally show up in both the dv/dt and dw/dt expressions) into one term that only shows up in the dw/dt expression. The resulting term is called the **effective gravity** and is indicated as g. The standard value of the effective gravity (at sea level) is 9.80665 m/s^2 (see appendix B.4).

Using our realigned coordinate system and the effective gravity, the acceleration equations become:

$$\begin{aligned}
du/dt &= -(2\Omega w\cos\phi) + (2\Omega v\sin\phi) + (uv\tan\phi/a) - (uw/a)\\
dv/dt &= -(2\Omega u\sin\phi) - (u^2\tan\phi/a) - (vw/a)\\
dw/dt &= -g + (2\Omega u\cos\phi) + (u^2/a) + (v^2/a)
\end{aligned}$$

and, adding in the other forces (pressure gradient force and friction), we get the equations of motion.

> ✍ Since $\hat{k}$ and $\hat{\jmath}$ are slightly shifted from the spherical coordinates we used in the derivation of the acceleration equations, a more accurate derivation would use the shifted coordinates, not the spherical coordinates. However, the derivation would follow along the same lines, with the velocity still expressed as $u_*\hat{\imath} + v\hat{\jmath} + w\hat{k}$ (since v and w are measured relative to our shifted coordinates). The only difference is that $\partial\hat{\jmath}/\partial\phi$ and $\partial\hat{k}/\partial\phi$ will depend upon ϕ (latitude) slightly. Exactly how this impacts the expressions is unclear but is assumed to be negligible.

Index

of protons
(atomic number)

Name 0 Symbol molar mass

Source: Atomic weights of the elements 2007 (IUPAC Technical Report)
Pure and Applied Chemistry 81:2131-2156, 2009
Note: 2011 data is different for some elements; 2009 data is used because that is what NIST uses
Molar masses in g/mol; * = no stable isotope

Hydrogen 1 H 1.00794(7)																		Helium 2 He 4.002602(2)
Lithium 3 Li 6.941(2)	Beryllium 4 Be 9.012182(3)												Boron 5 B 10.811(7)	Carbon 6 C 12.0107(8)	Nitrogen 7 N 14.0067(2)	Oxygen 8 O 15.9994(3)	Fluorine 9 F 18.9984032(5)	Neon 10 Ne 20.1797(6)
Sodium 11 Na 22.98976928	Magnesium 12 Mg 24.3050(6)												Aluminum 13 Al 26.9815386(8)	Silicon 14 Si 28.0855(3)	Phosphorus 15 P 30.973762(2)	Sulfur 16 S 32.065(5)	Chlorine 17 Cl 35.453(2)	Argon 18 Ar 39.948(1)
Potassium 19 K 39.0983(1)	Calcium 20 Ca 40.078(4)		Scandium 21 Sc 44.955912(6)	Titanium 22 Ti 47.867(1)	Vanadium 23 V 50.9415(1)	Chromium 24 Cr 51.9961(6)	Manganese 25 Mn 54.938045(5)	Iron 26 Fe 55.845(2)	Cobalt 27 Co 58.933195(5)	Nickel 28 Ni 58.6934(4)	Copper 29 Cu 63.546(3)	Zinc 30 Zn 65.38(2)	Gallium 31 Ga 69.723(1)	Germanium 32 Ge 72.64(1)	Arsenic 33 As 74.92160(2)	Selenium 34 Se 78.96(3)	Bromine 35 Br 79.904(1)	Krypton 36 Kr 83.798(2)
Rubidium 37 Rb 85.4678(3)	Strontium 38 Sr 87.62(1)		Yttrium 39 Y 88.90585(2)	Zirconium 40 Zr 91.224(2)	Niobium 41 Nb 92.90638(2)	Molybdenum 42 Mo 95.96(2)	Technetium 43 Tc *	Ruthenium 44 Ru 101.07(2)	Rhodium 45 Rh 102.90550(2)	Palladium 46 Pd 106.42(1)	Silver 47 Ag 107.8682(2)	Cadmium 48 Cd 112.411(8)	Indium 49 In 114.818(3)	Tin 50 Sn 118.710(7)	Antimony 51 Sb 121.760(1)	Tellurium 52 Te 127.60(3)	Iodine 53 I 126.90447(3)	Xenon 54 Xe 131.293(6)
Cesium 55 Cs 132.9054519	Barium 56 Ba 137.327(7)	†	Lutetium 71 Lu 174.9668(1)	Hafnium 72 Hf 178.49(2)	Tantalum 73 Ta 180.94788(2)	Tungsten 74 W 183.84(1)	Rhenium 75 Re 186.207(1)	Osmium 76 Os 190.23(3)	Iridium 77 Ir 192.217(3)	Platinum 78 Pt 195.084(9)	Gold 79 Au 196.966569(4)	Mercury 80 Hg 200.59(2)	Thallium 81 Tl 204.3833(2)	Lead 82 Pb 207.2(1)	Bismuth 83 Bi 208.98040(1)	Polonium 84 Po *	Astatine 85 At *	Radon 86 Rn *
Francium 87 F *	Radium 88 Ra *	‡	Lawrencium 103 Lr *	Rutherfordium 104 Rf *	Dubnium 105 Db *	Seaborgium 106 Sg *	Bohrium 107 Bh *	Hassium 108 Hs *	Meitnerium 109 Mt *	Darmstadtium 110 Ds *	Roentgenium 111 Rg *	Ununbium 112 Uub *	Ununtrium 113 Uut *	Ununquadium 114 Uuq *	Ununpentium 115 Uup *	Ununhexium 116 Uuh *		Ununoctium 118 Uuo *

†	Lanthanum 57 La 138.90547(7)	Cerium 58 Ce 140.116(1)	Praseodymium 59 Pr 140.90765(2)	Neodymium 60 Nd 144.242(3)	Promethium 61 Pm *	Samarium 62 Sm 150.36(2)	Europium 63 Eu 151.964(1)	Gadolinium 64 Gd 157.25(3)	Terbium 65 Tb 158.92535(2)	Dysprosium 66 Dy 162.500(1)	Holmium 67 Ho 164.93032(2)	Erbium 68 Er 167.259(3)	Thulium 69 Tm 168.93421(2)	Ytterbium 70 Yb 173.054(5)
‡	Actinium 89 Ac *	Thorium 90 Th *232.03806(2)	Protactinium 91 Pa *231.03588(2)	Uranium 92 U *238.02891(3)	Neptunium 93 Np *	Plutonium 94 Pu *	Americium 95 Am *	Curium 96 Cm *	Berkelium 97 Bk *	Californium 98 Cf *	Einsteinium 99 Es *	Fermium 100 Fm *	Mendelevium 101 Md *	Nobelium 102 No *