An introduction to the
PHYSICS
of
WEATHER PREDICTION

Robert A. Cohen
Physics Department
East Stroudsburg University

August 17, 2025
Version 11.2

Introduction

The focus of this textbook is on the physics used to make weather forecasts. In particular, the textbook focuses on the relationships and mechanisms affecting heat, moisture and air motion that are used in numerical models, which are the computer programs that simulate the atmosphere and used by weather forecasters to help predict the weather. This is not a textbook on how to use the output of such simulations to make weather forecasts, nor is it about using other observations to make weather forecasts, like discriminating between various types of precipitation or clouds and what they suggest about future weather.

How to use this book

This textbook is designed to be used as a PDF, as many items are hot-linked to either other locations in the textbook or to websites. In addition, many of the graphs and figures are best viewed electronically, either because of the color or to allow you the ability to zoom in on important features.

This textbook is also designed to be read differently than most textbooks. For this text you must actively try to evaluate your understanding of the concepts as you read through the text. In support of this, there are questions, which I call **checkpoints**, embedded throughout the readings.

These checkpoints are designed to provide an opportunity for you to check your understanding of the reading before moving on. Just because you can answer the checkpoint correctly doesn't necessarily mean that you fully understand the reading but it at least provides one way of self-assessing. You are responsible for constantly evaluating what you are reading.

On the other hand, if you can't answer the question without looking at the checkpoint solutions (a separate document), chances are you didn't fully comprehend the readings. If that is the case, do not continue, thinking that if you are confused about something you can always get it addressed later when you have more time. Such an approach will not work here because the readings build upon one another. So if you don't understand one section, future readings will not build properly on your previous understanding. Whatever the weakness, whatever you fail to address will only come back to haunt you later. In addition, when you find you are unable to understand something later, it will be all that more difficult to figure out where the weakness is because by that time you will have so many weaknesses you won't know where to start.

You will be left with two options: trying to start all over again or simply memorizing everything in the hope that that will get you through the semester. If you chose the latter approach, don't be surprised if you get more and more frustrated by the textbook's approach. After all, it wasn't designed for such an approach.

Projects vs. Questions vs. Problems vs. Show-me's

There are four types of homework tasks included in this textbook.

One type of task is called a **Project**. Not all chapters include a project and some projects can be completed prior to reading the chapter. The projects introduce techniques such a graphing, map reading and Excel skills that are important. If you have trouble with these skills, you may be tempted to just skip the project. Please do not succumb to that temptation.

Projects general require you to interpret real observations, like current maps and such. The reason for using real observations is so that you can relate the material to what is going on in the world right now. However, sometimes the real world isn't always in the "proper" state to clearly illustrate what is intended. For that reason, it is important that you work with the instructor to make sure the observations are conveying the important idea.

A second type of task are **questions**. Questions typically do not involve numbers. Instead, they ask you to explain or interpret some phenomenon.

A third type of task, called **problems**, typically involve numbers and ask you to calculate something.

Sometimes a problem can be difficult to answer because it requires a skill (like algebra or calculus) that you are weak at. The problems are designed such that later chapters use progressively more advanced mathematical techniques. If you can't manage a particular calculation, get it addressed so that it doesn't cause additional problems later on.

When doing a problem, make sure you can *interpret* the answer and not just *get* an answer. In other words, make sure the answer makes sense. If you cannot tell whether an answer makes sense or not, chances are you are missing some crucial concept. For example, suppose the problem asks you to predict how long it takes for a stone to drop one meter in free fall (ignoring air resistance). If you obtain an answer of 3600 seconds (1 hour), you should be able to realize that this is not correct. If you are not able to recognize whether an answer is reasonable or not, check with the instructor.

The fourth type of task is called a **Show-Me**. These are tasks where you are asked to derive a relationship. As with the problems, Show-me's are designed such that later chapters use progressively more advanced mathematical techniques.

Pay particular attention to the derivations within the chapter readings. Although the Show-me's won't ask you to repeat the same derivations, the procedures for the derivations are very similar. In addition, the derivations can give you insight that you might otherwise miss.

Overview

After a short chapter examining what is meant by a computer model of the atmosphere, the book is broken up into four parts.

Part A examines properties of the air that makes up the atmosphere. Much of part A should be a review from your introductory meteorology class, chemistry class, calculus class and physics class. At the same time, it is hoped that the reader gets a feeling for what the atmosphere is like so that he/she can interpret what comes later.

Part B investigates the processes that influence temperature, namely adiabatic expansion and compression, the first law of thermodynamics, and radiation emission and absorption.

Part C applies Newton's second law to atmospheric motion, both vertically and horizontally. The atmosphere tries to be "in balance" and moves in an attempt to restore the balance.

Finally, part D looks moisture and how it impacts the atmosphere.

Necessary background

I assume you have already completed introductory courses in meteorology, physics, calculus and chemistry. Indeed, one of the purposes of this textbook is to show how physics, chemistry and calculus are used to study the atmosphere.

Naturally, many students wonder if they really have the background necessary to use this textbook. To guide you, I outline here those concepts that I assume you have already been exposed to.

Meteorology: You should be able to do the following.

- Distinguish between different types of weather, like snow, sleet, rain and fog.
- Read a surface and upper-air weather map
- Construct isobars and isotherms
- Recall typical values of barometric pressure, temperature and wind speeds at the surface
- Read and interpret vertical profiles of temperature and dew point, ideally via a skew-T log-P graph

Chemistry: You should be able to do the following.

- Recall the definition of density and pressure, and the approximate densities of water and air
- Interpret a periodic table and recall the approximate molar masses of elements like carbon, nitrogen and oxygen

Physics: You should be able to do the following.

- Recall and apply Newton's second law of motion

- Describe vectors via components
- Recall the definitions of velocity and acceleration (and units)
- Recall and apply the first law of thermodynamics.
- Apply the law of gravitation and recall that a body in free fall near Earth's surface accelerates at about 9.8 m/s^2
- Recall the approximate radius of Earth
- Utilize metric units such as meter, second, kilogram, Newton, Joule and Pascal
- Utilize concepts of energy conservation and linear and angular momentum conservation
- Describe angular motion in terms of angular quantities

Calculus: You should be able to do the following.

- Take the derivative of a function of the form ax^n with respect to x
- Recall that the derivative of the function $\ln x$ with respect to x is x^{-1}
- Interpret the slope of a line as one application of the derivative of a function
- Distinguish between time and space averages
- Distinguish between average and instantaneous values

Acknowledgments

Several students of mine have made informal comments that have since been incorporated into this text. Students who have reviewed very early versions are Jocelyn Ross, Michael Salaski, Joel Geisinger and John Hearon.

Discussions with, and support from, my wife Maria were also instrumental. I thank her for her patience.

Contents

1. Weather Forecasting

Introduction

The purpose of these readings is to show how physics (and also chemistry and math) can be used to describe and make predictions about 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 air properties, in general, like what happens to the air if we cool or warm it, compress or expand it, move it, or whatever. All three types of predictions will be periodically mentioned in these readings. The physics (and chemistry and math) is the same, regardless.

Before showing how we apply physics (and chemistry and math) to make predictions about the atmosphere, we first need to recognize that there are several ways to predict the weather. I'll go over four in this chapter. Of the four, only the third accurately represents a method that uses the physics we'll explore in this book. It turns out that the third method is also the main way that meteorologists predict the weather and for good reason, too, since it traditionally has been the most reliable.

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

The *local forecast method* was the predominant method of weather forecasting until about two hundred years ago. Nowadays, meteorologists do not use this method for day-to-day forecasting, as it is very unreliable compared to the other two techniques I will 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 plotting 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 from above.

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]

Checkpoint 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, where the time zone is EST (Eastern Standard Time)?
(b) What local time would it be in Greenwich, England?
(c) What local time would it be in San Francisco, California, , where the time zone is PST (Pacific Standard Time)?

---

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

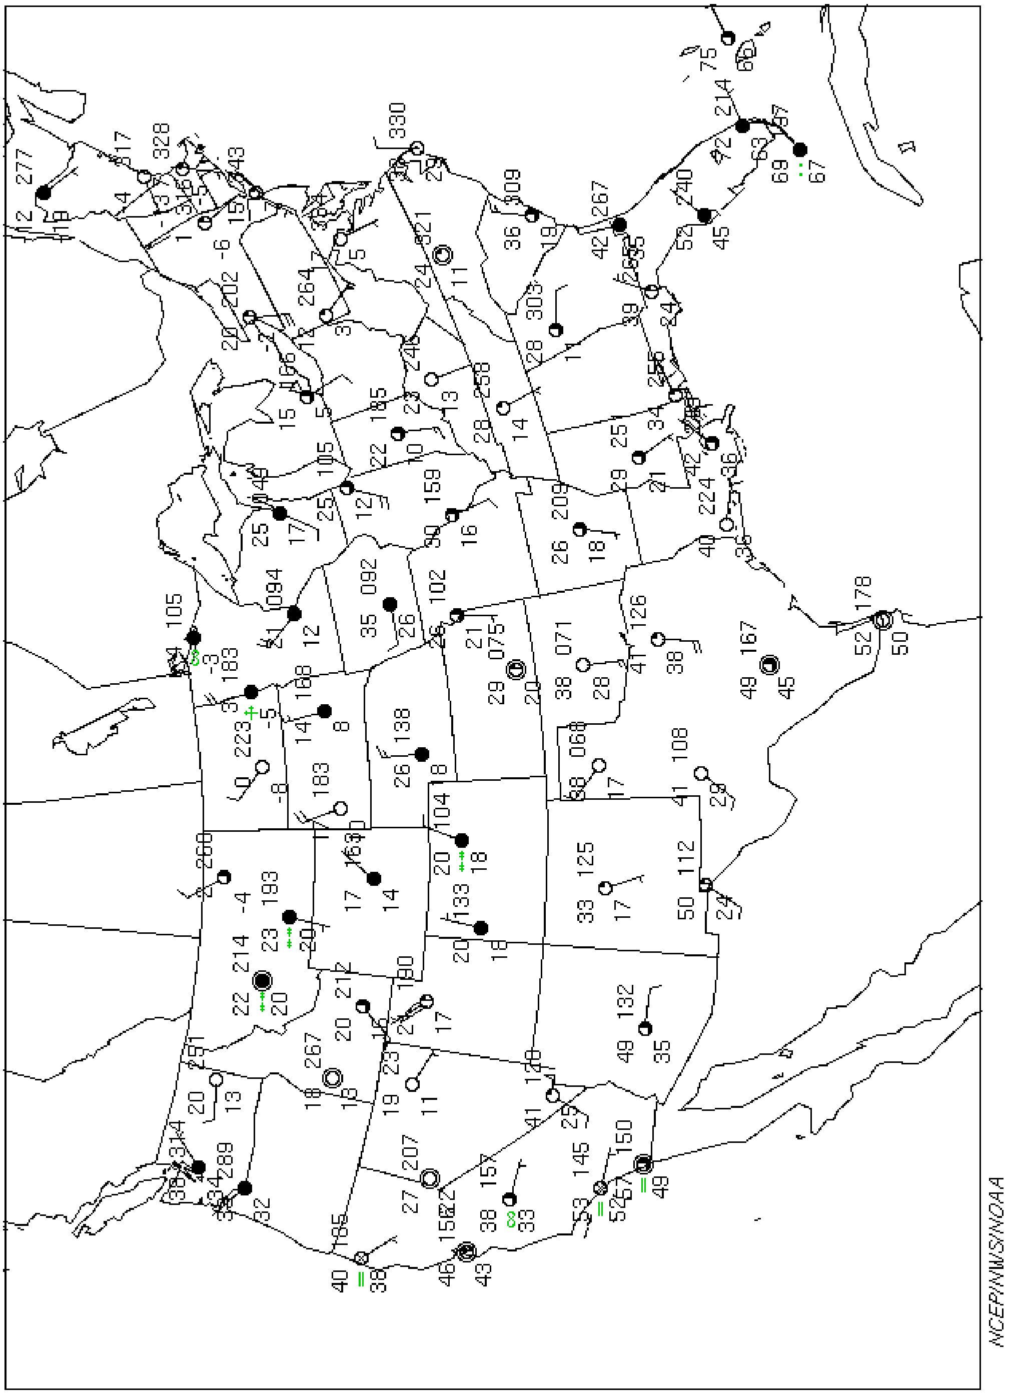

Figure 1.1: A map of surface observations for 14 UTC on 04 January 2014 (source: University of Illinois WW2010 ProjectDataStreme).

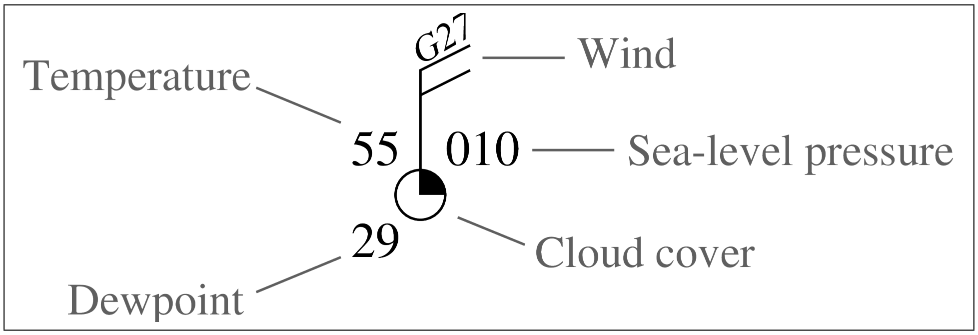

Figure 1.2: A sample station plot.

1.2.2 METAR data

The data for the surface map (see Figure 1.1) is obtained from hourly observations across the globe (with higher frequency observations made when significant weather occurs). Most observations are now being done automatically, via **ASOS** (Automated Surface Observation Stations) stations.

As illustrated in Figure 1.2, each station provides the temperature (in degrees Fahrenheit), dew point (in degrees Fahrenheit), cloud cover (in fraction of the sky), sea-level pressure (more on that in chapter 4) and wind (discussed in chapter 5). Precipitation, if there is any, is indicated by a symbol to the left of the station model (see, for example, some of the sites in California in Figure 1.1).

Weather information is typically encoded in a METAR[ii] format. Although it isn't necessary to be able to read the METAR format, it helps to be a little bit familiar with the meteorological conventions used as we'll occasionally need to check our theory with actual weather data.

An example METAR observation is the following:[iii]

KDHT 182053Z AUTO 36020G27KT 10SM FEW085 SCT110 13/M02 A2963 RMK AO2 PK WND 35027/2053 SLP010 T01281017 55005

[ii] Meteorological Aerodrome Report

[iii] METAR observations can be obtained from weather.rap.ucar.edu/surface/ and aviationweather.gov/metar.

Each set of characters represents a specific piece of information:

KDHT indicates the station location, which in this case is Dalhart Municipal Airport, Texas.

182053Z indicates the time, which in this case is 2053 UTC (see discussion above) on the 18th of the month.

AUTO indicates that the observations are from an automated station.

36020G27KT indicates that the wind is from 360 degrees at 20 knots, with gusts to 27 knots. We'll explain how to interpret the direction in chapter 5.

FEW085 SCT110 indicates the cloud cover and height. Each set of characters represents one cloud layer. In this case, there are two, with the lower layer being "FEW" and the upper layer being "SCT" (scattered).

13/M02 indicates the temperature (13°C in this case) and dew point (−2°C in this case, where the "M" means "minus").

A2963 indicates the altimeter reading in inches of mercury (29.63 inches in this case). We'll discuss this later in chapter 9.

The other characters are not important except for the set that includes "SLP", which stands for sea-level pressure. The meaning of this is discussed in chapter 4.

Precipitation, if observed, would be indicated before the wind information.

Checkpoint 1.2: Suppose the following observation was made (in METAR format): KTUL 270753Z 11004KT 10SM SCT200 21/20 A2988 RMK AO2 SLP110 T02110200. What is the time of the observation, and what is the temperature and dew point?

1.2.3 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 (where 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, in the polar front model, the air pressure is lowest where the weather system is.[iv] As part of the model, then, the air pressure at our location should decrease as a weather system approaches and increase as the weather system moves away from us. Tracking the trend of the local air pressure allows us to anticipate and track the coming and going of weather systems.

Another part of the model states that the wind blows in a direction that is around, rather than toward, the weather system center. The direction is counter-clockwise in the northern hemisphere. The model also states that weather systems tend to move from west to east. The combination of these two aspects means that in the northern hemisphere the wind will be coming from the south prior to the arrival of the system. One can therefore use the direction of the wind to determine when a weather system is approaching.

Not surprisingly, 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 approaching.

Finally, 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.

Checkpoint 1.3: Which of the following would be expected 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

---

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

1.3 Numerical weather prediction

It is useful to know the characteristics and structure of weather systems (as described by the polar front model; see previous section) because it provides a useful context for applying the physics and math that we'll encounter in this course. However, you don't need to understand much physics to use the polar front model to make weather predictions. As such, I won't be discussing it much (beyond what I've already discussed in the previous section).

Instead, I'll focus on the final forecast method, which is called **numerical weather prediction**. The reason 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 are based upon this technique.[v] 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.[vi]

1.3.1 Initialization process

To start, we need to know the temperature, humidity, wind and pressure throughout the entire atmosphere.

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.

As mentioned above, there are observing stations scattered throughout the globe. Unfortunately, those only give us information about the air near the surface. To get information above the surface, we release balloons[vii]

---

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

[vi] Technically, I'm referring to **deterministic** numerical models as opposed to **probabilistic** numerical models. The difference is described here.

[vii] These balloons are launched twice a day, at 00Z and 12Z. There are over 800 weather balloon sites worldwide (source). Once the balloon pops, it just falls to the ground. If you happen to find one, instructions on what to do can be found online. If you are really interested, check out the radiosonde museum.

(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 4.5, 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 all of these 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 (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.

1.3.2 The prediction process

Once we have the temperature, humidity, pressure and wind at each grid point, the next step is to predict how these variables 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 update process can be broken down into the following three steps:

1. The temperature is updated, based on the wind and processes like the heating during the day or cooling at night.
2. The air pressure is updated, based on the new temperature values.
3. The wind is updated, based on the new pressure values.

Once the predictions are made for the first time step, we then repeat the process, moving forward in time until we reach the forecast time.

1.3.3 The physics

WHERE IS THE PHYSICS IN THIS PROCESS?

The physics is used to make the predictions during each time step. For example, let's consider the three steps considered earlier.

1. Temperature
   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 conservation of energy (i.e., the first law of thermodynamics), which is the focus of part B of the readings.
2. Air pressure
   The air pressure depends on the air temperature. Consequently, once the temperature has been updated the pressure must likewise be updated. This is done using a relationship called the ideal gas law, which is discussed in part A of the readings.
3. Wind
   The wind depends on the air pressure. So, once the pressure has been updated the wind must likewise be updated. This is done using a relationship based upon Newton's second law, which is discussed in part C of the readings.

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.

1.3.4 Actual weather forecasts

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; see part D), 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[viii]

[viii] Each degree of latitude is equal to 111.12 km. Each degree of longitude is equal to 111.12 km times the cosine of the latitude.

(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[ix], 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[x] (i.e., the ones used by the weather services) are the **GFS** (Global Forecast System), **NAM** (North American Model), **RAP** (Rapid Refresh) 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 Service (**NWS**). Models like the GFS 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.

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

[x] UCAR (University Corporation for Atmospheric Research) provides information and the most common numerical models as well as various places you can access the output from numerical models.

Research models (i.e., those used by research organizations like universities), tend to have smaller domains. 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.

Checkpoint 1.4: 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?

1.4 Artificial intelligence

Whereas numerical prediction models rely on physics, and the mathematical equations that represent the physics, artificial intelligence (AI) uses patterns it has gleaned from loads of past observations. In this way, AI models are like the local forecast method mentioned at the beginning of this chapter, the difference being that the local forecast method uses past observations made at the location in question, whereas the AI models uses the full set of observations across the globe or region. While AI forecasts are relatively new, early results seem to show that they are comparable in accuracy to those made by numerical models. Since AI forecasts require a lot less computation, they are more efficient, and so we expect AI forecasts to grow in popularity.

Questions

Question 1.1: A degree of latitude is defined as 60 nautical miles, which is equal to about 69 statute miles. It turns out that a degree of latitude is also equal to 111.12 km. Based on this information, which is biggest: a nautical mile, a statute mile or a kilometer? Which is smallest? Provide your reasoning.

Question 1.2: According to the polar front model, which way would the wind likely be blowing in the northern hemisphere after a weather system passes? Explain.

Problems

Problem 1.1: The most common operational or research models used by the National Weather Service are RAP, NAM and GFS. Relevant information about each one is provided here:

	GFS	*NAM*	*RAP*
horizontal domain	*global*	*north america*	*north america*
horizontal grid spacing	*0.25-2.5 deg*	*12 km*	*13 km*
maximum height	*0.27 hPa*	*2 mb*	*10 mb*
# of vertical levels	*64*	*60*	*51*
output time step	*6 h*	*3 h*	*1 h*
Length of simulation	*240 h*	*84 h*	*21 h*

Assuming they all use the same time integration step (not actually true), which of the models do you think would require the greatest number of computational steps to make the entire forecast? Why?

Problem 1.2: Suppose you have a synoptic map using data in Chicago, IL, from 10 PM on December 31, 2013. What would be the date (include time) of the map in UTC? How do you know?

Problem 1.3: 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? Make whatever approximations you think are reasonable. Describe the reasoning in each step of your calculation.

Part A

Air Properties

2. Ideal Gas Law

Introduction

The physics of the atmosphere can be summarized by three relationships: the ideal gas law[i], Newton's second law of motion, and the first law of thermodynamics. 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). These four properties are the focus of part A of this textbook.

2.1 Properties of an ideal gas

WHY IS THE RELATIONSHIP CALLED THE IDEAL GAS LAW?

It is called "the ideal" gas law, and not just "the" gas law, because technically it only works for "ideal" gases (which don't exist in the real world). Fortunately for us, air is pretty close to an ideal gas.

An ideal gas 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 most of these postulates hold true 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.

---

[i] The ideal gas law is also known as the **equation of state**.

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 (so the average *velocity* is zero but the average *speed* is not). 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, the resulting volume would be just a minute fraction of the original 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 (meaning they have no excess positive or negative charge), this seems like a reasonable approximation.

The fourth postulate means that the collisions don't impact the average kinetic energy of the molecules. If the average kinetic energy was impacted by collisions, the temperature (which is a measure of the average kinetic energy) would increase or decrease as a result of each collision. Again, it seems reasonable to assume this won't happen (unless we happen to be doing something else to the air, like compressing it or heating it).

Such postulates define what is called an **ideal gas**. Although the gases in the atmosphere are not *exactly* ideal, they are very close. 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 several remarkable things about this. One is that molecules must be incredibly tiny to get so many in such a space (22.4 liters is less than one cubic foot)! Another is that it is possible to know the exact volume of so many molecules regardless of the molecules involved. The third is that we only need the ideal gas law to show this (see problem 3.1).

Checkpoint 2.1: Can we consider air to be an ideal gas?

2.2 Meteorological version

You may be familiar with the ideal gas law from chemistry. It is usually written as $PV = nRT$ and relates how the pressure (P), volume (V) and temperature (T) of a gas are related.

Air is close to an ideal gas, so we can use the ideal gas law to describe the atmosphere, but air doesn't have a definite volume, so we can't use *this* version of the ideal gas law. The version you learned in chemistry is well-suited for chemistry because they deal with gases in a container that has a specific volume. For our purposes, we'll use a version that uses density instead of volume (we'll derive this from $PV = nRT$ in chapter 3):

$$P = \rho RT \tag{2.1}$$

where P, ρ and T are the pressure, density and temperature, respectively. 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 2.2.3).

WHAT IS R?

We call R the **gas constant**. Different gases have different values of R because the value depends on the composition of the gas. For dry air (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 chapter 3.

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 (abbreviated as K).[ii] Since

---

[ii] The convention is to use lower-case names for unit names, even when the unit is someone's name, like in this case, where the kelvin is named for William Thomson, better known as Lord Kelvin, a British mathematician and physicist (1824-1907) who did work in thermodynamics. As far as I know, however, the corresponding abbreviation is always capitalized when corresponding to someone's name.

kelvin is not part of the units for pressure, R must have units of inverse kelvin (K^{-1}) to cancel out the units of kelvin (K) in the temperature.

How do we use the ideal gas law?

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).

Checkpoint 2.2: 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?

2.2.1 Pressure

Although we are not as sensitive[iii] 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. One example of this is that if the pressure is decreasing, that typically indicates that a weather system is approaching (see section 1.2.3).

Pressure vs. force

How do we know there is air around us?

We can tell there is air around us because 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 on us is the same in all directions (see next section), which

---

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

means the *net* force on us is zero. As we'll explore in section 2.4, we can only sense the air when the pressure is *different* on opposite sides of an object. Such a situation can be obtained by holding your hand out a car window while the car is in motion. Doing so will compress the air in front of your hand, increasing the pressure there, with the opposite occurring behind your hand, decreasing the pressure there. You'll experience a force pushing your hand backward, due to the difference in air pressure.

If we sense the force of the air, why do we measure the pressure instead of force?

Certainly force *could* be used. However, the force exerted by the air depends on 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 the surface and the larger the force exerted on that surface.

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

$$P = \frac{F}{A} \tag{2.2}$$

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

Because we divide out the area, pressure is independent of the area and thus can be assigned to a single point in space. It is similar to how we refer to the price per gallon when comparing how expensive the gas is at different gas stations. By looking at the price *per gallon*, we remove the dependence on how many gallons we are actually purchasing.

> ✍ We will use the same reasoning for defining density as the mass per volume (see section 2.2.3).

Since the pressure per area is independent of the area, we can talk about the pressure at a single point. Even though a single point has an **infinitesimal** area, with the force exerted by the air on that area also infinitesimal, the *ratio* of the two is not infinitesimal.[iv] It is for this reason that we usually use pressure instead of force when dealing with fluids like air.

[iv] Something infinitesimal would be infinitely small, or so close to zero that it is equal to zero for all practical purposes. More is said on this in section 2.3.

Isn't air a gas?

Yes, air is a gas, but it is also a fluid. 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.

Checkpoint 2.3: 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?

Isotropy

As mentioned above, 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. At a particular location, then, the air pushes up as much as it pushes down or horizontally. In other words, if could enclose an extremely tiny volume of air in a cube, the pressure on each face of the cube would be the same. This is called **isotropy**.

✍ Isotropy doesn't mean the pressure is the same everywhere. Certainly, the pressure is different on top of a mountain than down in a valley.

Checkpoint 2.4: Suppose we orient a piece of paper horizontally, parallel to the ground. If the pressure is isotropic, is the pressure exerted on one side always equal to the pressure on the other side? Explain.

2.2.2 Temperature

Next to precipitation, day-to-day variations in temperature are probably what concerns most people. This is because changes in temperature can be easily perceived. Most people are relatively unaware of the day-to-day changes in pressure and density.

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.

> ✍ The average kinetic energy is equivalent to the total kinetic energy of the gas divided by the number of molecules. By dividing out the number of molecules, we get rid of the dependence on how many molecules we measure. This is similar to what we did with pressure, where we defined pressure as the force exerted by the air divided by the area. By dividing out the area, we get rid of the dependence on the area.

2.2.3 Density

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

$$\rho = \frac{m}{V} \tag{2.3}$$

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, by dividing out the volume, we get rid of the dependence on the volume.

Checkpoint 2.5: The temperature and density are defined as the kinetic energy per molecule and the mass per volume, respectively. If the temperature and density is uniform throughout a particular region, does the temperature and/or density value depend on the number of molecules or how large a volume we use?

2.3 Typical values and units

As mentioned in the previous section, pressure, temperature and density are defined such that the values do not depend on the size or volume of the gas being measured. This means that each single point in the gas can be assigned a particular value of pressure, temperature and density.

WOULDN'T THE VOLUME BE ZERO AT A SINGLE POINT?

Yes, but that is okay. For example, density is the ratio of mass to volume. The volume of a point is technically infinitesimal, but so is the mass of the gas within that infinitesimal volume. Although the mass and volume of a point are both infinitesimal, the *ratio* of the two is not.

Indeed, this is the idea behind derivatives. A derivative is essentially a ratio of infinitesimal values, with the ratio being some finite value. In a similar way, density, pressure and temperature have finite values at a single point because they are ratios.

In this section, we examine typical values for the pressure, temperature and density of air.

Checkpoint 2.6: At a single point, of infinitesimal size, the air has a particular non-zero temperature and a particular non-zero pressure. Does it also have a particular non-zero density at that point, or do you need a non-infinitesimal region in order to have a density?

2.3.1 Pressure

Pascals

Since pressure is force per area, and the SI units of force and area are newtons (N) and square meters (m^2), respectively, pressure has SI units of N/m^2. The average air pressure at sea level is 101,325 N/m^2 (or 1.01325×10^5 N/m^2).[v]

Since pressure is used so often, the units (N/m^2) are frequently replaced with a single SI 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.

---

[v] We will examine in section 4.1.2 where I obtained the value for the average air pressure at sea level.

Since a typical sea-level pressure is such a large number of pascals, a metric prefix is often used. For example, instead of pascals, the pressure might be given in **kilopascals**. So, the average sea-level pressure would be 101.325 kPa.

In meteorology, the convention is to use **hectopascals** (abbreviated as hPa) instead of pascals or kilopascals. In hectopascals, the average sea-level pressure is 1013.25 hPa. See appendix A.3 for a description of the metric prefixes.

Millibars

As the "hecto" prefix is rarely used, you might wonder why meteorologists use hectopascals.

The reason is because a hectopascal is equivalent to a **millibar**, which was the meteorological convention prior to hectopascals. In fact, many meteorologists and maps still use millibars, so it is a useful exercise to explain what a millibar is and why we used it.

To explain the millibar, we first have to explain the **bar**. One bar is set equal to 10^5 Pa. It is defined as such because then one bar is approximately the average sea-level pressure at Earth's surface, which happens to be 1.01325 bars (divide 101,325 Pa by 10^5 Pa/bar).

> ✍ Another unit of pressure is the *atmosphere* (abbreviated as "atm"), defined as 1.01325 bars, which is the average sea-level pressure at Earth's surface.

The bar is a convenient unit to use for the air pressure near Earth's surface but it isn't convenient for the pressure changes from day-to-day because the air pressure only changes by a fraction of a bar over the course of a day. For this reason, meteorologists used millibars instead of bars. The unit abbreviation for millibars is "mb," so the average sea-level pressure can be written as 1013.25 mb.

> ✍ The millibar is not an SI unit, which is why many people, particularly in Europe, use hectopascals (hPa) instead of millibars (mb). They are equivalent, though. You can use either one.

Checkpoint 2.7: Express the average sea-level pressure in pascals, kilopascals and millibars.

Station model

It was mentioned that the average sea-level pressure is 101,325 Pa or 1013.25 mb.

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 "010" indicate a pressure of 1001.0 mb.

Checkpoint 2.8: Suppose the station model on a synoptic map of the surface indicates the pressure as "981". What is the surface pressure in millibars?

2.3.2 Temperature

Three common temperature units are degrees **Fahrenheit** (°F), degrees **Celsius**[vi] (°C) and **kelvin**[vii] (K). 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 units for some typical situations.

Freezing point of water	32°F	0°C	273.15 K
Cool autumn day	50°F	10°C	283.15 K
Typical room temperature	68°F	20°C	293.15 K
Normal body temperature	98.6°F	37°C	310.15 K
Boiling point of water	212°F	100°C	373.15 K

[vi] Before 1948 this unit was known as **centigrade**.

[vii] Unlike the other common temperature units, the name kelvin is not preceded by the "degrees" term (nor is the abbreviation preceded by the degrees symbol) because it is an absolute scale, not a relative scale (the Rankine temperature scale is sometimes written similarly for the same reason, while other times it is written with the degrees symbol due to convention). The name kelvin is also not capitalized. In general, unit names are not capitalized. The "C" in Celsius and the "F" in Fahrenheit are capitalized only because those words are not at the beginning of the unit name (so they act as modifiers of the unit, not the name of the unit itself).

Note that zero does not correspond to the same situation in all three temperature units. Zero degrees Celsius corresponds to the temperature water freezes (as with °C) but zero kelvin corresponds to the temperature where no thermal energy exists (not shown).[viii]

Do we need to use kelvin for temperature?

It depends on the situation.

For the temperature in the ideal gas law (equation 2.1), we use kelvin because a zero temperature in that relationship means each side must be zero, and the left side is zero only when the pressure is zero. Whenever we need the temperature *change*, on the other hand, we can use *either* °C or K. This is because the **increment** is the same in both, meaning a *change* of 1 °C is the same a *change* of 1 K.[ix] Be careful, though. A temperature of 1°C is not equivalent to a temperature of 1 K. Only the *increment* is the same. Looking at the table, one can see that the *values* always differ by 273.15.

> ✍ Just as most of our constants will use units such as meters, kilograms and seconds, so will most of our constants use kelvin.

Checkpoint 2.9: (a) In kelvin, what is 10° C?
(b) In kelvin, what is a change of 10° C?

As mentioned above, we need to use kelvin when using the ideal gas law because, according to the ideal gas law (equation 2.1), 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 2.1. A temperature scale must be used that equals zero only when the pressure exerted by the gas goes to zero.[x] The kelvin scale is designed such that the temperature, in kelvin, is zero when the pressure is zero.

[viii] In comparison, it seems no one really knows what zero degrees Fahrenheit is supposed to correspond to, although one common theory is that it corresponds to the lowest temperature one can get with a mixture of salt and ice.

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

[x] 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).

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

Checkpoint 2.10: 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 the ideal gas law (equation 2.1)?

How are degrees Celsius and degrees Fahrenheit related?

The conversion between degrees Celsius and degrees Fahrenheit is a bit more complicated because the two scales use different increments. The Fahrenheit scale has 180 degrees between freezing and boiling whereas the Celsius scale has only 100 degrees between freezing and boiling. That means that an increment of 1°C is a greater temperature difference than an increment of 1°F. No matter which two points you choose, there will be a greater number of °Fbetween the two points than °Cbetween the two points. For example, the freezing point and normal body temperature are 0°C and 37°C, for a difference of 37°C. In the degrees Fahrenheit scale, the temperatures are 32°F and 98.6°F, for a difference of 66.6°F..

It turns out that a difference of 1°C is equivalent to a difference of 1.8°F.[xi] You can verify this by taking the ratio of 66.6°F to 37°C and seeing that the ratio is equal to 1.8°F/°C. You can use any two points and you'll find the same ratio.

Knowing the ratio of the increments, you can find the relationship between the two scales by first determining how far each value is from the associated reference point and then setting the ratio of those values to the ratio of the increments. This is easier to show than it is to explain. For example, we know that the freezing point of water is 0°C and 32°F. If T_C and T_F represent the temperatures in degrees Celsius and degrees Fahrenheit, respectively, then $(T_\mathrm{C} - 0°\mathrm{C})$ and $(T_\mathrm{F} - 32\ \mathrm{F})$ represent how far the temperature is from freezing in the two scales.

The ratio of these two is equal to

$$\frac{(T_\mathrm{C} - 0°\mathrm{C})}{(T_\mathrm{F} - 32°\mathrm{F})} = \frac{1°\mathrm{C}}{1.8°\mathrm{F}}$$

[xi] Some people prefer to write that a temperature difference of 5°C is equivalent to a temperature difference of 9°F.

If you want, you can then rearrange the expression to solve for one of the temperatures, like as follows:[xii]

$$T_{\mathrm{C}} = \left(\frac{1^{\circ}\mathrm{C}}{1.8^{\circ}\mathrm{F}}\right)(T_{\mathrm{F}} - 32^{\circ}\mathrm{F}) \tag{2.4}$$

> ✍ We could have used the same process for finding the relationship between degrees Celsius and kelvin. The result is $T_{\mathrm{K}} = (1\ \mathrm{K}/^{\circ}\mathrm{C})T_{\mathrm{C}} + 273.15\ \mathrm{K}$.

Checkpoint 2.11: The boiling point of water is 212°*F in the Fahrenheit scale and* 100°*C in the Celsius scale. Suppose we had used the boiling point as our reference instead of the freezing point. Would the relationship in equation 2.4 be any different? If so, how? If not, why not?*

Why does the Fahrenheit scale have a different increment?

As mentioned above, 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 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 and so he used the freezing and boiling points of water and set them as 0°Cand 100°C, for a difference of 100°C.[xiii]

Checkpoint 2.12: Why is it said that the Fahrenheit scale has 180 degrees between freezing and boiling when boiling is 212 degrees, not 180 degrees?

[xii] Derivation 2.3 asks you to derive an expression for T_{F} rather than T_{C}.

[xiii] 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 (see Show-me 2.4), perhaps to avoid using negative numbers. It is not clear why he decided to use 100 increments.

2.3.3 Density

A typical density of the air (near sea-level) is between 1.2 and 1.3 kg/m^3. What this means is that a cubic meter of air typically contains between 1.2 kg and 1.3 kg of air molecules. A cubic meter is the volume of a cube one meter on each side. At higher altitudes, the density is less (to be explored in chapter 4).

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

Checkpoint 2.13: (a) What is the density of air near sea level?
(b) What is the density of water?
(c) Which is larger?

2.4 Measurements

Density

To measure density, 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 directly like this. Rather, we first measure the temperature and pressure and then use the ideal gas law to determine the density based upon the temperature and pressure.

Temperature

Temperature is measured using a **thermometer**. Although temperature reflects the average kinetic energy of the molecules, 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 that is measured instead.

[xiv] 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.

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.

Pressure

With all of the units one can use with pressure (see section 2.3.1), you might be surprised to learn that until recently the most popular unit used by TV meteorologists (and still used by pilots) was 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 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.

Since the air is exerting its pressure all around us, you might think it should be relatively straightforward to measure the air pressure. What makes it difficult to do so, though, is that the air is exerting its pressure *all around us*.

Suppose we hold a piece of paper by its edge. There are two sides of the paper. The air pressure on both sides are equal. Consequently, the air exerts the same force on each side of the piece of paper, but the direction of the two forces are opposite. Consequently, the *net* force on the paper is zero. It doesn't matter what the air pressure is.

The only way we could measure the air pressure, then, would be to physically remove all of the air on one side. The force needed to keep the paper stationary would then equal the force exerted by the air on the other side. The air pressure would then be that 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."

This method is impractical because is very difficult to entirely remove the air from only one side of a piece of paper.

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).

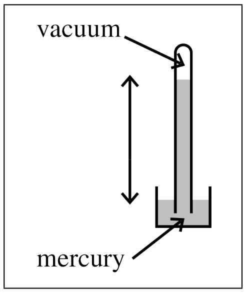

When we do this, we find that, at sea-level, the height of that mercury column (that balances the force of the air pushing it up) is typically about 760 mm high (about 30 inches). That means that the 760-mm high column of mercury has the same "weight" as the atmosphere. The mercury barometer essentially 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.

Checkpoint 2.14: What would weigh more: a container of mercury 30 inches deep or a container of air 30 inches deep?

Questions

Question 2.1: Of density, temperature and pressure, which do we typically measure (rather than calculate via the ideal gas law)? Why?

Question 2.2: 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 measurement that has a finite value even though it is technically a ratio where the numerator and denominator are both infinitesimal.

Problems

Problem 2.1: (a) Using the fact that the density of liquid 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.

Problem 2.2: (a) Determine the weight[xv] *of a mercury column of height 760 mm and area 1 cm^2. Assume a density of 13.5951 g/cm^3. Check units!*
(b) Using the value obtained in (a), 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 24 for the typical sea-level pressure. Explain the difference, if 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 2.3: 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 24?

Problem 2.4: Estimate the total number of drops in the ocean by first estimating the volume of Earth's oceans (from an estimate of the surface area and average depth) and the comparing that to an estimate volume of a single drop of water. Make sure your units work out.

Problem 2.5: Using the fact that there is 1 mole of atoms in 22.4 liters[xvi]*, calculate the average distance between each atom.*[xvii] *Look up the average*

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

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

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

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

Derivations

Show-me 2.1: One version of the ideal gas law is written as follows: $PV = mRT$*, where* m *is the mass of the gas. Show how this version can be derived from equations 2.3 and 2.1.*

Show-me 2.2: (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 2.3: Using equation 2.4, one can convert the temperature in °F to °C. Derive an equation that can be used to convert temperature in °C to °F (i.e., the reverse) and use your equation to show that a typical **room temperature** *of 20°C is equivalent to 68°F.*

Show-me 2.4: 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[xix]*.*
(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.*
(b) Use your relationship to determine what room temperature was in the original Celsius scale.

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

[xix] 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).

3. What is in the Air?

Introduction

As mentioned in chapter 2, the ideal gas law can be expressed as follows:

$$P = \rho RT$$

The ideal gas law relates the four atmospheric variables of temperature, pressure, density and composition. We've already discussed the first three. In this chapter, we examine composition, which is reflected in the variable R.

3.1 Gas constant

As mentioned in chapter 2, R is called the **gas constant**. The value depends on the gases that are present.

Of all of the gases that make up air, water vapor is the only gas that varies significantly from day to day and from location to location. To simplify things, then, we'll focus on either the gas constant for dry air[i] (the air without the water vapor)

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

or the gas constant for water vapor alone[ii] alone

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

Then, if we want to use the ideal gas law for air that contains some 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 14). Consequently, the gas constant for air will likely be close to the R_{d} value.

---

[i] See appendix B.2.

[ii] See appendix B.3.

> ✍ 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.

Checkpoint 3.1: 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?

3.2 The universal gas constant

Before going into detail as to how the gas constant R reflects the composition, I first want to point out that the R in our ideal gas law ($P = \rho RT$) is <u>not</u> the same R^* that is present in the version you probably learned in chemistry class ($PV = nR^*T$). Both are normally written as R, so it is easy to confuse the two. I am writing the latter as R^* to distinguish it from the R used in *our* version of the ideal gas law.

The main difference is that, whereas R depends on the composition of the gas, R^* does not depend on the make-up of the gas and has the same value regardless of what gas is present. For that reason, R^* is known as the **universal gas constant**. Its value is given as follows:

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

Notice that the units of R^* are similar to the units of R except that R^* uses “mol” (moles) instead of “kg” (kilograms). The meaning of “mol” is examined in the next section.

Checkpoint 3.2: (a) Does the value of R depend on whether there is water vapor present or not?
(b) Does the value of R^* depend on whether there is water vapor present or not?

3.3 Molar mass

To see how the air composition influences the value of R, it helps to derive our version of the ideal gas law ($P = \rho RT$) from the chemistry version ($PV = nR^*T$).

> ✍ Do not just skip over the derivation, or any derivation, thinking that only the final relationships are important. Derivations are not just fun academic exercises for your brain. There is physics within the derivation that will provide insight into the meaning of the different variables.

First, we need to recognize what n represents, namely 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".

> ✍ The number of atoms or molecules in one mole, 6.0221415×10^{23}, is called **Avogadro's number** (N_A).

If we divide both sides of the chemistry version by V, we get

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

The quantity n/V represents the number of atoms or molecules per volume. This quantity is like the mass density discussed in the previous chapter but in terms of number, not mass.

To get the ideal gas law in a form that has the mass density, we can multiply n/V by m/m, where m is 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 $m \times n$).

Making this replacement in the ideal gas law, we get the following:

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

The quantity n/m 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”:[iii]

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

For example, the molar mass of oxygen atoms is 15.9994 g/mol.[iv]

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

The m/n value is an *average* because some of the molecules may have a greater mass than others. The quantity m is the total mass of all the molecules all together. By dividing by n, the total number of moles, we get the average mass of a mole full of molecules.[v]

Replacing m/n with μ in the ideal gas law, we get the following:

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

As a final step in our derivation, then, we define the **gas constant** R as the universal gas constant R^* divided by the molar mass:

$$R = \frac{R^*}{\mu} \tag{3.1}$$

Note that whereas R^* does not depend upon the gas, R does depend upon the gas. In particular, the greater the average molar mass, μ, the smaller the gas constant. Since water vapor is lighter than air, the gas constant for water vapor is greater than the gas constant for dry air.

With this replacement, we have our version of the ideal gas law: $P = \rho R T$ (equation 2.1).

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

[iv] Given the units, this is technically the molar mass density, as it represents the mass per mole, much like a volume mass density represents the mass per volume. However, everyone calls it the molar mass, as though it is the average mass of a mole of particles.

[v] The meaning of the average will be examined further in section 3.4.

Checkpoint 3.3: The constants R and R^* have slightly different units. Is that consistent with the relationship between the two given in equation 3.1?

3.4 Molar mass of dry air

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 3.1), $R = R^*/\mu$, 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 for dry air is 28.964 g/mol.[vi]

Let's compare this value to some elements, as listed in the periodic table (see inside back cover; the molar mass is provided below the letter abbreviation for each element[vii]). The element with a molar mass closest to that of air is silicon[viii], and the periodic table lists the molar mass of silicon as being 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.

[vi] According to the CRC Handbook (1983), the official value is 28.96443 g/mol, which with rounding is the same as what we have obtained).

[vii] Even if all of the atoms are the same element, there might still be some variations in the mass because of the different number of neutrons that may be present in the nucleus of each atom. The number of protons define the element, so all atoms of the same element have the same number of protons. However, the number of neutrons may be different, which will impact the mass. Such variations of a single element are called **isotopes**. The molar mass value listed on the periodic table for a particular element is actually the average molar mass for all of the isotopes of that element. This is discussed in more detail in appendix E.

[viii] By convention, element names are not capitalized.

> 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 periodic table lists the molar masses of nitrogen and oxygen as 14.0067 g/mol and 15.9994 g/mol, respectively. 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).

Checkpoint 3.4: What does it mean for a molecule to be diatomic?

Examining the molar mass of air a bit closer reveals that it has a value that 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 $\mu_{\rm d}$ to represent the molar mass of dry air:

$$\begin{aligned}\mu_{\rm d} &= (28.0134\ \text{g/mol}) \times (0.78) + (31.9988\ \text{g/mol}) \times (0.21) \\ &= 21.85045 + 6.71975\end{aligned}$$

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, when weighting the two values (in determining the average), 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 argon makes up most of the remaining 1% and argon has a molar mass of 39.948 g/mol, which is higher than that of nitrogen and oxygen. Including argon in our calculation, we get:

$$\begin{aligned}\mu_{\mathrm{d}} &= (28.0134\ \mathrm{g/mol}) \times (0.78) + (31.9988\ \mathrm{g/mol}) \times (0.21) \\ &\quad +(39.948\ \mathrm{g/mol}) \times (0.01) \\ &= (28.85045 + 6.71975 + 0.39948)\ \mathrm{g/mol} \\ &= 28.9697\ \mathrm{g/mol}\end{aligned}$$

As expected, this is even closer to the actual value. It is still a little off, though, because the relative abundance values were rounded and did not include additional gases that are present in tinier amounts. The actual fractions of each gas in dry air are listed in table 3.1 along with the molar masses.

Checkpoint 3.5: If the average molar mass of air was calculated using only the nitrogen and oxygen, with values given in table 3.1, one gets a value of 28.8564 u. That is less than the accepted value of 28.96443 u. Why?

3.5 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

Table 3.1: Molar masses (from physics.nist.gov) and proportion (by number; from CRC 1983) of molecule type in dry air (at sea level), except for carbon dioxide, which uses the 2025 data from the Mauna Loa Observatory (314 ppm in 1958). 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	0.04275	0.4275	4.275
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

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 in Table 3.1 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 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 2) 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 (meaning the ratios are the same everywhere).

Checkpoint 3.6: 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 or parts per billion. An increase of one part per million is equivalent to an increase of one-thousandth of a part (10^{-3}) per thousand or one ten-thousandth of a part (10^{-4}) per hundred. Since one part per hundred is equivalent to one percent, that means that one part per million is equivalent to 10^{-4} percent (or 0.0001%). One part per billion would be even smaller.

Consequently, the day to day variations in pollutants don't really change the overall molar mass. 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.

What about longer-term changes? Certainly the atmosphere has evolved over time, right?

Although the atmosphere certainly evolved over the liftime of Earth (billions of years), we can consider it to be fairly constant over a timescale of decades, centuries or even millenia. See Appendix F for more information on the mechanisms for those longer-term changes.

What about the increase of carbon dioxide that I hear about? Is that changing the molar mass of dry air?

Carbon dioxide is definitely increasing,[ix] from about 315 ppm in 1960 to about 430 ppm in 2025.[x] Since carbon dioxide, CO_2, has a molar mass of about 44 g/mol (add twice the molar mass of oxygen to the molar mass of carbon), an increase in carbon dioxide concentration would push the molar mass of dry air higher.

However, “ppm” means “parts per million”, and so an increase of 100 pm (for carbon dioxide) would be equivalent to an increase of 0.01%, still a tiny percentage compared to the percentages for gases like nitrogen, oxygen and argon.

Why isn’t water vapor included in the list of constituents listed in table 3.1?

Because the amount of water vapor, unlike the others, is both highly variable and likely to reach appreciable percentages of the total air (as much as 4%). Because of the importance of water vapor, I discuss it in a separate part of the book (see part D).

Wouldn’t the presence of water vapor change the average molar mass of air?

[ix] The increase in carbon dioxide comes from the burning of **fossil fuels**. Whenever one burns a hydrocarbon (e.g., the gasoline in a car), oxygen is utilized with the hydrocarbon to produce water vapor and carbon dioxide. A similar process occurs during respiration except with carbohydrates. This is why carbon dioxide is exhaled when animals and plants undergo respiration.

[x] This is the concentration observed at Mauna Loa in Hawaii. The concentration tends to be more between 30°N and 40°N.

Yes. That is why everything we've done to this point is accurate only for **dry air**, as indicated by the subscript "d."

Checkpoint 3.7: What does it mean for the air to be dry?

Project

Project 3.1: 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 3.1. In the second column type in the relative proportion of each gas, in percent.

Project 3.2: Add two additional columns to your spreadsheet, one containing the concentrations in ppm (parts per million) and one 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.[xi]

Project 3.3: Add a new column to your spreadsheet. In it, copy the average molar mass for each gas from table 3.1. In a separate cell, use a spreadsheet equation to calculate the average of all of the molar masses, remembering to weight the values by their relative proportions.[xii]

Questions

Question 3.1: How do you think the relative abundances of N_2, O_2, *Ar*, CO_2, H_2 *and* He_2 *might change over the next one thousand years. Based on your predictions, will the molar mass of the atmosphere be significantly different than what it is now? In what way?*

[xi] 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 zeros. For those, it is better to use units of ppm (parts per million). Indeed, when examining the concentration of pollutants, it is not unusual to have the concentration reported in ppb (parts per billion).

[xii] You can check your answer by comparing it with the accepted value of 28.96443 u (source: CRC Handbook, 1983).

Problems

Problem 3.1: Use the chemistry version of the ideal gas law to determine the volume per mole (V/n) of an ideal gas at standard pressure (1013.25 hPa) and temperature (273.15K). Compare your answer to the value of 22.4 liters per mole provided on page 18.

Problem 3.2: Use the information in the periodic table (inside back cover) to determine the molar masses of N_2, O_2, *carbon dioxide* (CO_2) *and water vapor* (H_2O).

Problem 3.3: Given the molar mass of water vapor (H_2O) *in comparison to the molar masses of diatomic nitrogen and oxygen (see problem 3.2), should moist air have a molar mass less than, or greater than, the value for dry air (28.96443 g/mol)? Explain.*

Problem 3.4: (a) Determine the value of the gas constant for dry air by dividing the universal gas constant by the molar mass of dry air.
(b) Determine the value of the gas constant for water vapor, R_v, *by dividing the universal gas constant by the molar mass of water vapor.*
(c) Compare your results with the values of the gas constant for dry air and the gas constant for water vapor that are given in appendix B.2.

Problem 3.5: Suppose you land on an unknown planet and find that the pressure at the surface is 7.5 hPa with a temperature of $-20^\circ C$ *and a density of 0.0155* kg/m^3. *Based on these measurements and assuming the atmosphere is made up of a gas that follows the ideal gas law, what is your guess for the gas that makes up most of this atmosphere?*

Derivations

Show-me 3.1: 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 2.1 can be derived from this expression if the R in equation 2.1 is equal to the universal gas constant* R^* *divided by the average molar mass density.*

Show-me 3.2: 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.

4. Spatial Variations

Introduction

To forecast how air pressure, temperature and density vary with *time*, we must have an idea how those things vary in *space*, meaning the vertical and horizontal. That is the focus of this chapter.

4.1 Vertical profile

4.1.1 Pressure vs. height

The value given in chapter 2 for the average sea-level pressure was 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. For the time being, we'll just refer to the sea-level pressure as P_0, where the subscript zero refers to an elevation of zero (sea level).

As one might expect, the higher one goes in the atmosphere the lower the pressure. In chapter 9, we'll examine why this is (hint: it has to do with Newton's second law). For now it is sufficient to model the decrease as an exponential decay.

WHAT DOES IT MEAN TO DECREASE EXPONENTIALLY WITH HEIGHT?

When the pressure decreases exponentially with height, its value P can be modeled by the following equation:[i]

$$P = P_0 e^{-z/H} \tag{4.1}$$

---

[i] The letter e represents a number (about 2.718282).

where z is the height[ii] of the pressure value and H is the height at which the pressure falls to $1/e$ of its initial value.

> Relationships are exponential when the rate of change of something is proportional to the value of that something. Population growth, for example, tends to be exponential since the growth rate is greater when the population is greater. The vertical variation of pressure tends to be exponential because the pressure decreases more rapidly with height where the pressure value is larger (near Earth's surface) than where the pressure is smaller (higher up).

When z is equal to zero, $e^{-z/H}$ equals one (since e^0 is equal to one). By multiplying by P_0 (the sea-level pressure) we ensure that the pressure equals the sea-level pressure when z equals zero. The negative sign in equation 4.1 ensures that the pressure *decreases* exponentially with height (rather than increases exponentially).

The value of H in equation 4.1 is called 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.

According to NASA, the scale height of Earth's atmosphere is 8.5 km, but that is just an average value, as the pressure variation with height changes from day to day. Similarly, equation 4.1 is just a rough representation of the actual air pressure, which varies from day to day.

Checkpoint 4.1: Why is it said that the scale height represents the height at which the pressure is about one-third the pressure at the surface?

4.1.2 Temperature vs. height

Like pressure, the temperature also varies from day to day. To provide a reference against which we can compare, an average temperature profile has

---

[ii] We are using z for altitude to be consistent with the convention of using x, y and z for the three directions in three-dimensional space (x and y for eastward and northward, respectively).

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

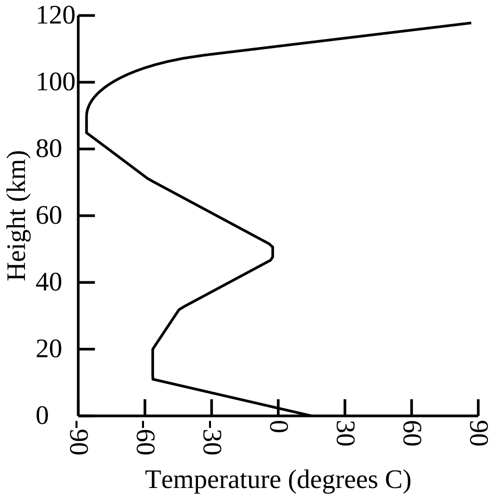

Figure 4.1: A graph showing the U.S. standard temperature profile.

been identified by the United States Committee on Extension to the Standard Atmosphere (COESA)[iv]. This average temperature profile is known as the **U.S. Standard Atmosphere.**[v] The U.S. Standard Atmosphere is sort of like an ideal average.

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 profile from the surface to about 120 km is shown in Figure 4.1.

> ✍ Notice that Figure 4.1 has the dependent variable (temperature) on the horizontal axis rather than the more traditional vertical axis. This is done so that the vertical axis mirrors the vertical direction, with the bottom of the graph corresponding to the bottom of the atmosphere.

Checkpoint 4.2: On any particular day, will the temperature of the real atmosphere at various levels equal the temperature given in the U.S. Standard Atmosphere?

[iv] COESA represented 29 U.S. scientific and engineering organizations.

[v] Presented by COESA on 15 October 1976. 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.

4.1.3 Density vs. height

Like the air pressure and temperature, the density of the air is not the same throughout the atmosphere, nor does it remain the same with time.

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

Does the density decrease with altitude?

The density doesn't *always* decrease with altitude but, as you can show in Show-me 4.2, the density does decrease with height under normal conditions. That is why it is harder to breathe the higher one goes. The density of the air at the cruising altitude of planes (12-15 km), for example, is typically only a third or a quarter of the density of the air near Earth's surface.

When does the density not decrease with altitude?

Whether the density decreases with height or not depends on how the *temperature* changes with height (see section 9.4 and problem 9.4). Optical effects called **mirages** can occur when the density *increases* with height.

Checkpoint 4.3: The peak of Mt. Everest is almost 9 km above sea-level. At that altitude, do you expect the density to be greater than 1 kg/m^3 *or less than 1 kg/m*^{3}*?*

4.2 Vertical gradients

4.2.1 Temperature (lapse rate)

Although the U.S. Standard Atmosphere is an average, and not a picture of the real atmosphere on any given day, it contains certain features that are usually present in the real atmosphere.

For example, as one can see from Figure 4.1, the U.S. Standard Atmosphere can be broken down into layers by examining how quickly the temperature changes with height. These layers are also present in the real atmosphere, although theboundaries between the layers won't be as "clean" as in the U.S. Standard Atmosphere.

In the U.S. Standard Atmosphere, the lowest layer has a temperature that decreases by 6.5°C every kilometer (see project 4.1). In the real atmosphere, it varies a great deal, not only from place to place but also at different levels of the atmosphere. Sometimes the temperature actually *increases* with height near the surface. However, on average, the general trend near the surface is to decrease with height, as with the U.S. Standard Atmosphere.

> Since degrees Celsius has the same increment as kelvin, a decrease of 6.5°C per kilometer is equivalent to a decrease of 6.5 K per kilometer.

In addition, just as with the U.S. Standard Atmosphere, in the real atmosphere there is a height when the temperature stops decreasing with height, and above that point the temperature either remains the same with height or increases. However, the height of that transition varies. In the U.S. Standard Atmosphere, that transition occurs around 11 km whereas in the real atmosphere it typically varies between 10 and 15 km.

The **troposphere** is the name given to this lowest layer of the atmosphere, where the temperature gets colder with height. The level above the troposphere, where temperatures remain the same with height and/or get warmer with height, is called the **stratosphere** and extends from about 10 km to about 50 km.

Since variations in temperature are such an important piece of weather forecasting, we need to review the language we use to describe these variations, both in English and in mathematics. In particular, the vertical temperature **gradient** is the rate at which the temperature changes with altitude and, mathematically, the vertical temperature gradient can be written as a derivative with respect to height: $\partial T/\partial z$.[vi]

> A zero $\partial T/\partial z$ value corresponds to a vertical line (infinite slope) on the graph in Figure 4.1. This may seem strange if you are used to the derivative being the slope of a line. The derivative is the slope only <u>if</u> one creates a graph and <u>if</u> the variable in the numerator is on the vertical axis and the variable in the denominator is on the horizontal axis. That is not the case for the temperature graph in Figure 4.1, since the numerator, T, is plotted on the horizontal axis, not the vertical axis. That means $\partial T/\partial z$ does not correspond to the slope of the line on our graph.

[vi] As mentioned earlier in this chapter, we are using z for the vertical position (altitude).

As noted above, the temperature typically decreases with height in the troposphere, which means the temperature gradient is typically negative in the troposphere. Rather than deal with negative numbers all the time when discussing the vertical temperature gradient in the troposphere, meteorologists define the **lapse rate** as the *negative* of the vertical temperature gradient:

$$\Gamma = -\partial T/\partial z \tag{4.2}$$

where the Greek letter Gamma (Γ) is used to represent the lapse rate. In the troposphere, then, the lapse rate is typically positive.

> ✍ Meteorologists tend to use special names for vertical gradients. Besides the lapse rate (negative of the vertical temperature gradient), there is also the **wind shear**, which is the vertical gradient of the wind (to be discussed in section 13.6).

Checkpoint 4.4: 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?

4.3 Notation

Is $\partial T/\partial z$ related to the derivative of T notation we learned in calculus?

The ratio $\partial T/\partial z$ is known as the derivative of T with respect to z and, as such, it is pretty much the same derivative you learned in calculus class as dT/dz.

The difference is mostly irrelevant to much of what we will cover. However, you'll notice that sometimes I'll write the derivative with a curly "∂" and sometimes I'll write it as a standard "d" so it might help to know the difference.

First of all, $\partial T/\partial z$ is called the *partial* derivative while dT/dz is called the *total* derivative. The partial derivative is used when a variable (like T) depends upon several things (like both the position and the time) and I only want to refer to how it depends on *one* of those things. In those cases, a derivative like $\partial T/\partial z$ means that things like the time is held fixed and only the position has changed.

The total derivative is used when no variables 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.

Checkpoint 4.5: 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?

4.3.1 Pressure

As we know, the pressure is less the higher one is in the atmosphere. That means the vertical pressure gradient is negative.

This can also be seen in equation 4.1, $P = P_0 e^{-z/H}$ by taking the derivative with respect to z:

$$\begin{aligned} \frac{\partial P}{\partial z} &= \frac{\partial (P_0 e^{-z/H})}{\partial z} \\ &= -\frac{P_0}{H} e^{-z/H} \\ &= -\frac{P}{H} \end{aligned}$$

Here I have taken advantage of a special property by the number represented by e, namely that $d(e^x)/dx = e^x$. In other words, when the number e is raised to a power of x, a change in x by dx results in a change of e^x equal to e^x times dx.

What this means is that as the pressure gets smaller and smaller, at higher and higher altitudes, the vertical pressure gradient likewise gets smaller and smaller. 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 other words, if you were in an 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 (if the elevator could actually go up that high).

Another consequence of the exponential decay in pressure is that the vertical pressure gradient $\frac{\partial P}{\partial z}$ is proportional to the pressure. This means that if you measure the pressure at two different levels, the ratio of the pressures depends only on the height difference. For example, if you measured the pressure at height z_1 and compared it to the pressure at height z_2, the ratio of the two pressures would be:

$$\frac{P_1}{P_2} = \frac{P_0 e^{-z_1/H}}{P_0 e^{-z_2/H}}$$

This can be simplified by rearranging as follows:

$$\frac{P_1}{P_2} = \frac{P_0}{P_0} e^{(z_2-z_1)/H} = e^{\Delta z/H}$$

In other words, in our ideal elevator moving upward at a constant speed, you'd move from 1000 mb to 750 mb (a decrease of 25%) in roughly the same amount of time it would take you to move from 400 mb to 300 mb (also a decrease of 25%).

Checkpoint 4.6: Suppose the surface pressure is 1000 mb and the pressure 1.5 km above the surface is 800 mb.
(a) How far above 800 mb would we likely find the air pressure to be 600 mb: less than 1.5 km, about 1.5 km or greater than 1.5 km?
(b) How far above 800 mb would we likely find the air pressure to be 640 mb: less than 1.5 km, about 1.5 km or greater than 1.5 km?

4.4 Horizontal maps

4.4.1 Temperature

The air temperature near Earth's surface typically varies between 0°C (the freezing point of water) and 37°C (normal body temperature). As mentioned before, room temperature is typically around 20°C.

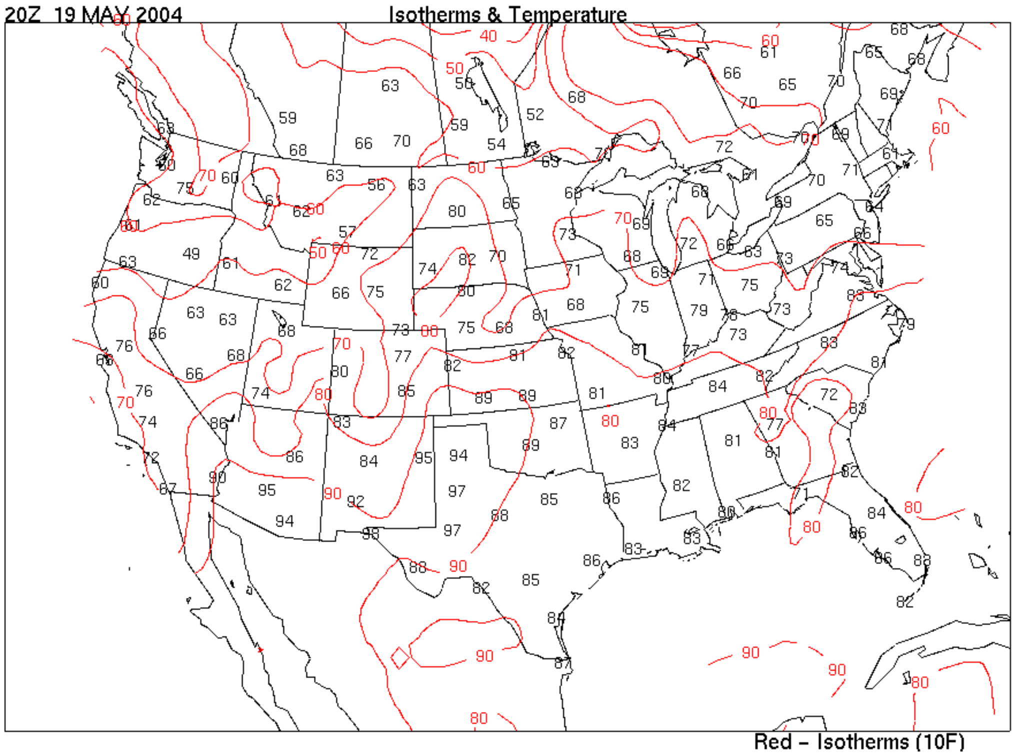

Figure 4.2: A map of surface temperature for 20 UTC on 19 May 2004 (source: DataStreme).

To examine how the temperature varies in the horizontal, we can examine a synoptic map (see chapter 1). An example of such a map is shown in Figure 4.2, 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 curves that connect places that have similar values of temperature. The generic name for such curves are **isolines** because each place along the line has the same (*iso*) value. Depending on what the parameter is, they can also be called something specific to the parameter, such as **isobars** ("bar" = pressure) or **isotherms** ("therm" = temperature).[vii] In this particular map, the lines connect locations that have the same temperature and so the isolines in this case are isotherms.

[vii] 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**.

By convention, U.S. surface maps tend to have temperatures in degrees Fahrenheit, and that is the case for the map in Figure 4.2, with isotherms plotted every 10°F.

Checkpoint 4.7: On the surface temperature shown in Figure 4.2, one can see a temperature of 68°F in southern Iowa and another temperature of 81°F along the border between Iowa and Nebraska. Given those two temperatures, why would one expect two isotherms to pass between the two? Why do you think no isotherms are drawn between those two numbers?

4.4.2 Pressure

As mentioned before, the actual sea-level pressure varies from place to place and from time to time. Typical values can vary from, say, 950 mb to about 1050 mb. This can be seen on the map of sea-level pressures in Figure 4.3.

In this case, the sea-level pressure values are in millibars (which are equivalent to hectapascals). Notice how the sea-level pressure values are all around the standard sea-level pressure of 1013.25 mb. Also notice the isobars, which are plotted every 4 mb. The isolines in this case are called isobars because they connect locations that have the same pressure.

Important: The values depicted in figure 4.3 are not those of the air pressure at the surface but rather the estimated pressure *at sea level* (more on that in section 9.3.4).

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.

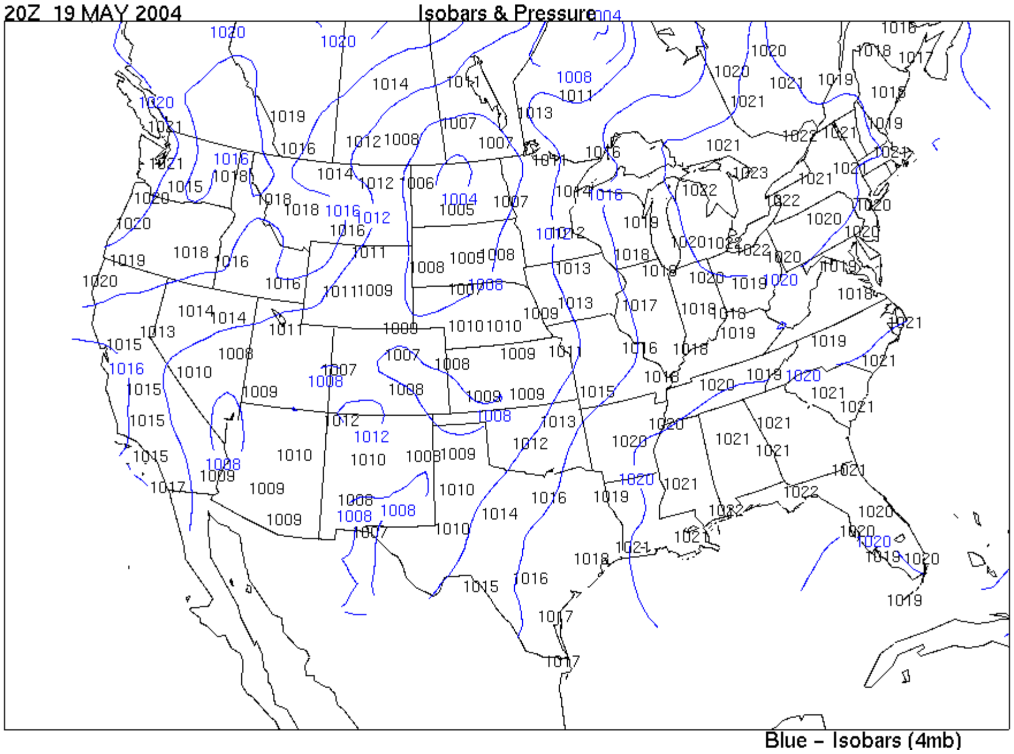

Figure 4.3: A map of sea-level pressure for 20 UTC on 19 May 2004 (source: DataStreme). This map is referenced in Problem 4.1.

Checkpoint 4.8: 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 Earth? If not, which is greater: the sea-level pressure or the pressure at the surface?

4.4.3 Horizontal gradients

Isolines are valuable because, 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 pressure or temperature is higher and lower. They can also be used to determine the value at any location as well as the *horizontal* gradient.

The horizontal gradient, like the vertical gradient, represents how much a quantity changes per distance. The difference is that the horizontal gradient tells us how the quantity varies in the horizontal, as opposed to the vertical.

The horizontal gradient is greatest where the isolines are closest together. For example, if we plotted a topographic map with contours (lines of constant height), the height gradient would be greatest where the contours are close together, which represented regions with large slopes, like a cliff or side of a mountain.

In the same way, when we plot a pressure map with isobars (lines of constant pressure), the pressure gradient is greatest where the isobars are close together. Likewise, a strong temperature gradient exists where the temperature is very different at two nearby locations.

The horizontal pressure gradient is a useful thing to identify because it tells us how the pressure varies from one location to another. As we'll find out in part C, we need to know the horizontal pressure gradient in order to predict the wind. And, as we'll find out in chapter 5, the horizontal temperature gradient can tell us something about how much the temperature is likely to change.

Checkpoint 4.9: 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?

HOW DO WE DETERMINE WHAT THE HORIZONTAL GRADIENT IS?

The horizontal gradient is obtained by identifying how close the isolines are. For example, consider the isobars drawn to the right. Four isobars are drawn, from 1020 mb to 1026 mb.

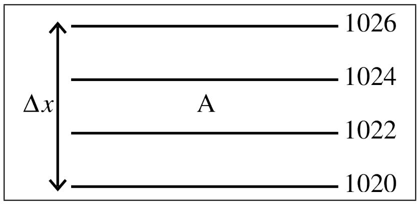

Let's suppose that the distance from the 1026 mb isobar to the 1020 mb isobar is 100 km (the distance is indicated as Δx in the figure). That means that the pressure changes by 6 mb over a distance of 100 km. That would correspond to a pressure gradient of (6 mb)/(100 km) or 0.06 mb/km. Basically, we divide the pressure difference by the distance between the two points to get the pressure gradient.

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

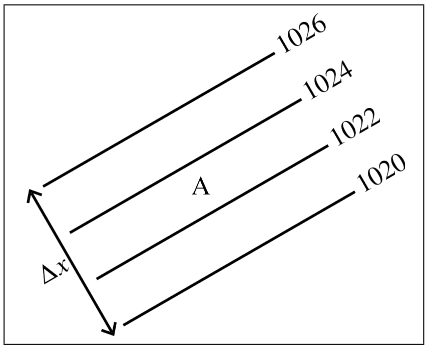

While it is relatively easy to use a synoptic map to find the pressure difference between two points, it is a little trickier to find the distance between the two since distance is not usually indicated on synoptic maps.

To find the distance you'll need to use a reference on the map for which you know the distance. For example, if you know the size of Nebraska (333 km north to south)[viii], then you can use that as a reference. Measure the distance from one isobar to the next and then compare that distance to the reference distance.

Checkpoint 4.10: Suppose isobars are drawn every 4 mb. At a particular location two isobars are separated by a distance measured to be twice the distance from the northern to southern boundary of Nebraska. What is the pressure gradient at that location?

4.4.4 General notation

To determine the horizontal gradients, the examples utilized **finite** differences, which are represented mathematically by the capital Greek letter "delta" (Δ), as in Δx. While all measurements use finite differences, in theory we do not need to use a finite difference. We instead estimate an

---

[viii] 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).

infinitesimal[ix] difference and, in so doing, calculate the pressure gradient at a point. This is particularly useful if the pressure gradient varies from one place to the next.

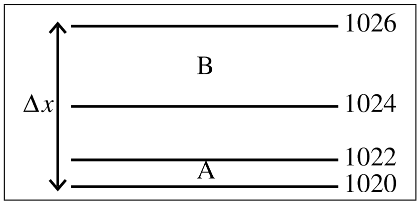

For example, in the figure to the right the pressure gradient is greater at A than at B because the isobars are closer together at A than at B. In this case, what gradient is represented by $\Delta P/\Delta x$, which uses the total difference in pressure and distance?

The answer is that it gives the *average* pressure 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", and so would be represented as $\partial P/\partial x$ rather than $\Delta P/\Delta x$.

In this way, it is the same as what we did when examining the vertical gradients of temperature, pressure and density (as with $\partial T/\partial z$). There is just one additional complication when dealing with horizontal gradients, and that has to do with there being two directions, x and y, in the horizontal. As such, horizontal gradients really should involve two gradients, one for the x direction and one for the y direction.

In practice, meteorologists use x and y to indicate the two perpendicular directions of east and north, so $\partial P/\partial x$ would only represent the gradient in the eastward direction. If there is only a gradient in the northward direction then $\partial P/\partial x$ would be zero.

To address this, we can either write the horizontal gradient as two derivatives, one with respect to x and one with respect to y, or we can write the horizontal gradient as a single derivative, in the direction perpendicular to the isobars, which is how I did it with the tilted isobars on page 59. The latter is represented mathematically with the "∇" symbol. For example, the horizontal gradient of pressure would be indicated as ∇P. Notice that the symbol used for the gradient is not the symbol "Δ", which is represents a finite difference. It actually uses the "∇" symbol, which is an inverted "Δ".

[ix] 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.

✍ Technically, ∇ represents the total gradient in three dimensions, but we'll use it to indicate just the horizontal gradient.

Checkpoint 4.11: Suppose you are examining a region of a synoptic map where the isobars are close together. Is ∇P large or small?

4.5 Skew-T graphs

As mentioned in chapter 1, 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. Another is by showing you several synoptic maps, each at a different level. A third way is via a single graph.

I'm going to show you the third way, since that is the most common and, perhaps, simplest. However, indicating the vertical temperature profile via a graph is not as straightforward as you might think.

Let's start by considering a graph much like the one shown in Figure 4.1 for the U.S. Standard Atmosphere except for *real* data and only for the lowest 15 km of the atmosphere where most of the weather takes place (see chapter 6 to learn why). Figure 4.4 shows two ways of plotting the temperature profile for Flagstaff, Arizona. As was the case when we plotted the U.S. Standard Atmosphere (figure 4.1), the dependent variable (temperature) is on the horizontal axis rather than the more traditional vertical axis so that the vertical axis mirrors the vertical direction, with the bottom of the graph corresponding to the bottom of the atmosphere.

The difference is that the left graph has the tempeature plotted vs. height, as we did with the U.S. Standard Atmosphere, while the right graph has the temperature plotted vs. pressure.

WHY WOULD WE WANT TO PLOT THE TEMPERATURE VS. PRESSURE?

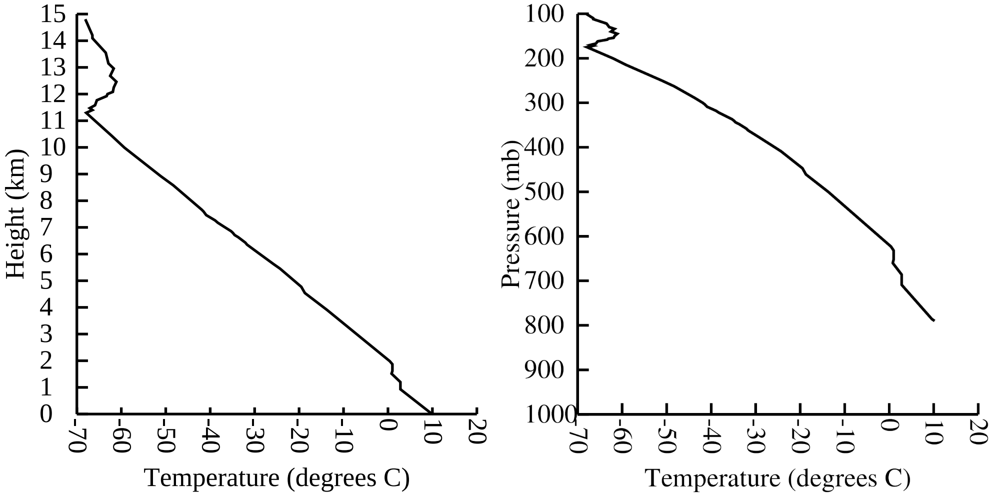

Figure 4.4: [left] A graph showing the temperature for the lowest 15 km of the atmosphere over Flagstaff, Arizona, on 0000 UTC, 20 Jan 2019. [right] The same profile but with pressure instead of height.

The problem with plotting vs. height is that, until the advent of GPS, it was difficult to determine the height of the radiosondes as they ascended. Consequently, it wasn't possible to produce the left graph in Figure 4.4. The radiosonde did have a barometer, though, and so we knew the temperature at various *pressures*, which meant we *could* produce a graph of the temperature vs. *pressure*, as in the right graph.

Notice that the pressure values are arranged "backward", with lowest pressure at the top rather than the bottom. This is because then the top of the graph mirrors the pressures at higher elevations (since pressure is less at higher elevations).

While the right graph in Figure 4.4 is sufficient to represent the vertical temperature profile, there are some issues with it that make it a bit cumbersome to use.

First, we know from earlier in this chapter that the pressure doesn't decrease linearly with height, as the right graph in Figure 4.4 seems to suggest. Instead, it decreases exponentially. Second, the temperatures in the bottom left of the graph (-70°C at 1000 mb) and the top right of the graph (20°C at 100 mb) are almost never observed. In fact, temperatures at the surface rarely get colder than -30°C or so and the temperature at 100-200 mb rarely gets warmer than −40°C or so.

These two issues are fixed by using a **skew-T/log-P** graph, as shown in Figure 4.5, which is the same as the right graph in Figure 4.4 except for two important differences.

First, the pressure lines are not equally spaced (shown in blue, with labels 1000, 900, 800, etc.). This is because the vertical coordinate is actually linear with the logarithm[x] of pressure. In other words, if we wrote the logarithm of pressure on the chart (instead of the pressure), you'd see that the logarithm of pressure values would be equally spaced (i.e., the lines corresponding to 10^5 Pa, 10^4 Pa, 10^3 Pa, if plotted, would be equally spaced). One can show this mathematically, which is left as an exercise for you in Show-me 4.1.

Checkpoint 4.12: 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 distance 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)?

The second difference between the right graph in Figure 4.4 and that in Figure 4.5 is that the temperature lines in Figure 4.5 (shown in blue with labels 30, 20, 10, etc.) are skewed at an angle (i.e., not horizontal). In this way, the graph doesn't include the rarely observed cold surface temperatures and warm upper-air temperatures.

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.

On a skew-T log-P, there are two thick squiggly lines. The one on the right (red in the figure) is the measured temperature of the atmosphere and corresponds to the vertical temperature profile shown in Figure 4.4. The squiggly

---

[x] 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.

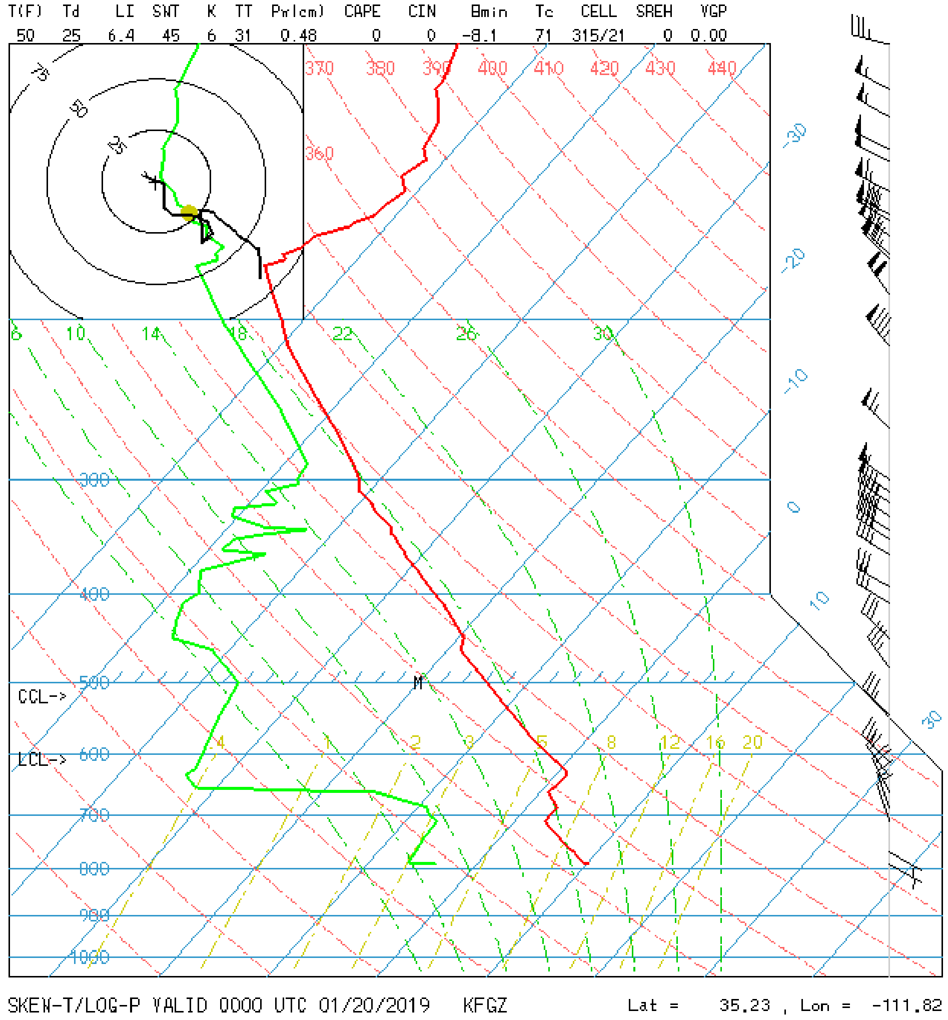

Figure 4.5: A skew-T log-P graph for Flagstaff, Arizona, for 0000 UTC, 20 Jan 2019. Horizontal solid blue lines are pressure (in mb). Skewed, solid blue lines are temperature (in °C). The solid red somewhat jaggedly line is the bserved temperature profile. This figure is also used in Problem 4.4.

line on the left (green in the figure) represents the amount of moisture measured in the atmosphere (this is discussed in chapter 14).

Checkpoint 4.13: According to the skew-T log-P graph shown in Figure 4.5, what is the temperature at a height where the pressure is 500 mb?

Project

Project 4.1: (a) Download the temperatures of the **U.S. Standard Atmosphere** *and import them into a spreadsheet.*
(b) Create a graph of the U.S. Standard Atmosphere temperature profile from sea-level up to 90 km such that height is on the vertical axis. Note that the file gives two sets of heights: geopotential height and geometric height. The difference will be discussed in chapter 9. They are very similar, however, and for this project you'll be using the geopotential heights (since the U.S. Standard Atmoshere below 90km is based on geopotential heights) even though the geometric heights are the actual, radar heights.[xi]
(c) How many different layers of constant lapse rate regions are there? Identify the lapse rate of each of those layers (use change in temperature divided by change in geopotential height for the purpose of this activity). Note: the bottom layer corresponds to the troposphere.

Project 4.2: (a) In the spreadsheet, use equation 4.1 ($P = P_0 e^{-z/H}$) to calculate the pressure at every height from sea-level up to 120 km (where H is 8500 m and $P_0 = 1013.25$ mb). If you do this correctly, you'll find that the pressure at about 12 km (typically the altitude where airplanes fly) is approximately 250 mb.
(b) Create a graph showing the pressure profile up to 120 km. Note: In Excel, one can use "EXP(A1)" to determine the value of e raised to the value stored in cell A1. As with project 4.2, plot height on the vertical axis.

Project 4.3: (a) From the results in project 4.2, how high does one have to go to reach a pressure one-half that at the surface?

[xi] The file uses H and Z as variable abbreviations for the geopotential and geometric heights, respectively. This usage of H is different than what we've used before for H, which was the scale height.

(b) Plug the height value obtained in (a) as z in Equation 4.1 ($P = P_0 e^{-z/H}$)*.*
What pressure do you get?
(c) Should your answer in (b) be the half the surface pressure? Explain.
(d) At a height equal to the scale height H*, is the pressure half that at the surface? Explain.*

Project 4.4: (a) From the results in project 4.2, how high does one have to go to reach a pressure equal to $1/e$ *the value at the surface, where e is equal to 2.718282?*
(b) Plug the height value obtained in (a) as z in Equation 4.1 ($P = P_0 e^{-z/H}$)*.*
What pressure do you get? (c) Should your answer in (b) be $1/e$ *times the surface pressure? Explain.*
(d) At a height equal to the scale height H*, is the pressure equal to* $1/e$ *times the surface pressure? Explain.*

Questions

Question 4.1: Why is a skew-T log-P graph called a skew-T log-P graph?

Problems

Problem 4.1: Figure 4.3 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.
(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 4.2: Figure 4.2 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.
(c) What is the range of temperatures that are observed? Include units.

(d) Calculate the maximum horizontal temperature gradient that is present (see the region of Southern California, with 90°F along the border with Arizona and the 67°F along the Pacific coast near the border with Mexico).
(e) Which is greater (absolute value): the maximum horizontal temperature gradient observed in figure 4.2 or a typical vertical temperature gradient in the troposphere (see U.S. Standard Atmosphere).

Problem 4.3: From problems 4.1 and 4.2 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 2.1) and the value of R for dry air (see appendix B.2). Compare your value of density to the range mentioned in section 2.3.3 and comment on any differences.

Problem 4.4: (a) According to Figure 4.5, what is the temperature at a height where the pressure is 750 mb? (b) In the figure, does the temperature tend to decrease or increase with height?

Problem 4.5: The following expression is derived in Show-me 4.2:

$$\frac{1}{\rho}\frac{\partial \rho}{\partial z} = -\frac{1}{H} + \frac{\Gamma}{T}$$

Given this expression and assuming a tropospheric lapse rate of 6.5 K/km and a scale height equal to 8.5 km, should the density typically decrease near the surface, where the temperature is around 300 K (i.e., 27 C)? What about near the top of the troposphere, where the temperature is around 220 K (i.e., −53 C)? Provide the rationale for your answers.

Problem 4.6: The following expression is derived in Show-me 4.3:

$$\frac{V}{V_0} = \left(1 - \frac{\Gamma Z}{T_0}\right) e^{Z/H}$$

According to the expression, by what percentage does the volume of a balloon increase if the balloon ascends by 1000 m from the surface? Assume a typical surface temperature and a typical tropospheric lapse rate of 6.5 K/km.

Derivations

Show-me 4.1: Mathematically, the following relationships are true.

$$\begin{aligned} \log(AB) &= \log A + \log B \\ \log(A^B) &= B \log A \end{aligned}$$

Using these two properties, derive the following by applying the logarithm to equation 4.1.[xii]

$$\log P = [-1/(H \ln 10)]z + \log P_0$$

Note: The difference between $\log A$ *and* $\ln A$ *is that* $\log A$ *represents the logarithm with base 10 (i.e., by what factor does 10 have to be raised in order to equal A), so* $\log 10 = 1$, *whereas* $\ln A$ *represents the logarithm with base e (i.e., by what factor does e have to be raised in order to equal A), so* $\ln \mathrm{e} = 1$. *Both follow the two rules described above.*

Show-me 4.2: Mathematically, the vertical density gradient is represented by the derivative of the density ρ *with respect to the height* z, *and is indicated in mathematical notation as* $\partial\rho/\partial z$. *Derive the following expression for the vertical density gradient:*

$$\frac{1}{\rho}\frac{\partial \rho}{\partial z} = -\frac{1}{H} + \frac{\Gamma}{T}$$

where H *is the scale height,* Γ *is the lapse rate* $(-\partial T/\partial z)$ *and* T *is the temperature. Do this by using the ideal gas law* ($P = \rho RT$; *solved for density) and the vertical pressure equation* ($P = P_0 \mathrm{e}^{-z/H}$; *see equation 4.1). Also use the fact that the gas constant* R *does not vary with height, so* $\partial R/\partial z = 0$.

Show-me 4.3: Suppose we raise a balloon a distance Z. *Assuming the temperature of the air inside the balloon changes to match the temperature of the surrounding air, show that the balloon will expand (rather than shrink) such that the new volume* V *divided by the old volume* V_0 *is equal to the following expression:*

$$\frac{V}{V_0} = \left(1 - \frac{\Gamma Z}{T_0}\right) e^{Z/H}$$

where H *is the scale height of the atmosphere and* Γ *is the lapse rate. To do this, you'll need to use the ideal gas law (equation 2.1), the definition of density (equation 2.3), the scale height expression for the decrease of pressure with height (equation 4.1 on page 47), and the knowledge that the temperature decreases at a rate given by the lapse rate (see section 4.2.1).*

[xii] This shows that the logarithm of the pressure $\log P$ varies linearly with height z, with slope equal to $-1/(H \ln 10)$ and y-intercept equal to $\log P_0$.

Part B

First Law of Thermodynamics

5. Advection

Introduction to part B

In chapter 1, we described how a numerical model predicts the evolution of the atmosphere. Part of this involves predicting how the temperature changes over time.

In this part of the text, we look at the processes that might lead to a change in temperature, like the impact of radiation from the Sun (chapter 8) and the expansion and compression that occurs when air rises and descends (chapter 6).

In this chapter, we'll look at the process by which the wind brings in air that is colder or warmer.

5.1 Describing winds

Before we can predict the temperature change due to the wind, we have to first describe and quantify the wind speed and direction.

5.1.1 Wind direction

On a weather map, the wind direction is indicated by a **wind barb**. A wind barb is like an arrow but, instead of an arrow head showing the direction the wind is blowing toward, tail feathers are used to indicate the direction the wind is coming from.

Consistent with this, the meteorological convention is to indicate 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

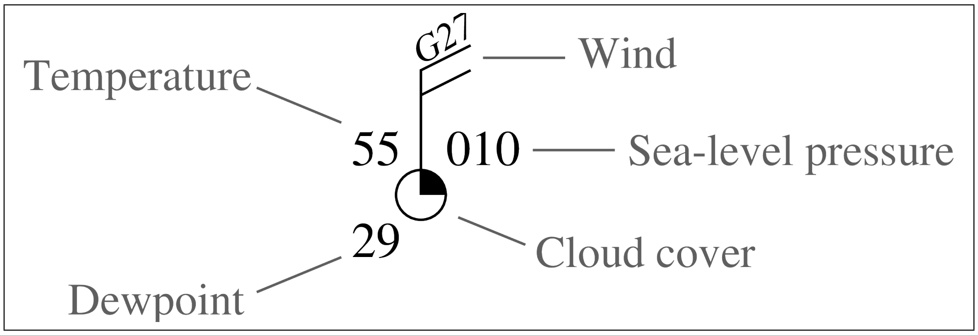

Figure 5.1: A sample station plot, repeated from Figure 1.2.

to winds coming from the west.[i]

For example, the wind barb in Figure 5.1 indicates a wind coming from the north. In comparison, on the surface map shown in Figure 1.1, the wind in Oklahoma is coming from the south.

Checkpoint 5.1: At 14 Z on 04 January 2014 (see Figure 1.1), from which direction is the wind coming over Florida?

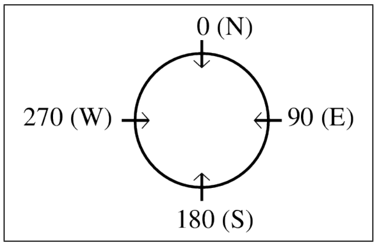

As mentioned before, meteorologists specify the wind direction *from which* the wind is blowing. To indicate the direction, we use 0° to indicate wind from the north and increase the angle clockwise around the circle (so that 90° indicates a wind from the east). Thus, a 180° wind would be coming from the south.

> ✍ You might think zero is supposed to be rightward but that happens to just be the *mathematical* convention. There is no special reason why we *must* do it that way.

Following the meteorological convention, a direction of 360° (or 0°) would mean the wind is coming from the north, as indicated in the station plot in Figure 5.1. Similarly, a wind from the south would have a direction of 180°.

---

[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**.

Checkpoint 5.2: Indicate the angle (in degrees) associated with wind coming from the west in the meteorological convention.

5.1.2 Wind speed

On a map, 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.

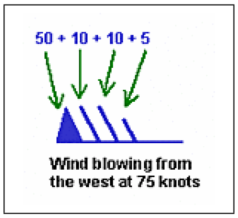

✍ When wind gusts are observed, the wind gust (in knots) is sometimes indicated next to the wind flags, as in Figure 5.1.

Checkpoint 5.3: What is the wind speed indicated by the station model in Figure 5.1? Include units.

WHAT IS A KNOT?

As noted earlier, the METAR data indicates the wind speed in **knots**, as the meteorological community traditionally has followed the nautical practice for describing the wind speed. A knot is equivalent to a **nautical mile** per hour.

WHAT IS A NAUTICAL MILE?

A nautical mile is a little larger than a **statute mile**, which is the mile they use when posting speed limits on roads.

✍ 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.

The reason for the difference is because a statute mile is for use on land while a nautical mile is for use on water. More specifically, a statute mile 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).

They are pretty similar but a statute mile is a little smaller. For example, a degree of latitude is defined as 60 nautical miles but it is equal to about 69 statute miles (actually $69.04676\bar{6}$ statute miles). That means a nautical mile is about 69/60, or 1.15, times bigger than a statute mile.

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.

Checkpoint 5.4: Which is faster: 10 mph or 10 kts?

5.2 Wind chill

Before getting too far into how we calculate the temperature advection, we need to recognize 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.[ii]

> ✍ 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.

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 16, evaporation is a cooling process.

Checkpoint 5.5: When the wind blows, does the temperature necessarily decrease? What about how it feels?

[ii] Convection can change the temperature when the wind is mixing up the air with air from another level (see section 10.5).

5.3 Warm vs. cold advection

Temperature advection refers to what happens when the wind is blowing from a colder or warmer location. If the wind is blowing from a warmer location then it is bringing warmer air to the region, a process we call **warm advection**, and leading to temperatures warming. Conversely, if the wind is bringing air that is cooler, we call that **cold advection**.

> ✍ Many properties can be advected, including moisture, momentum and temperature.

Checkpoint 5.6: If warm advection is occurring at a particular location, what would be your short-term forecast for temperature at that location: warming, cooling or neither?

For temperature advection to occur, three things need to happen:

- The temperature must be different somewhere else – this means that there must be a temperature gradient present (i.e., $\nabla T \neq 0$),[iii]
- The wind must be blowing (i.e., $\vec{v} \neq 0$), and
- The wind must be directed, at least partially, along the temperature gradient (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 or 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 5.2, 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 a region with warmer temperatures to a region with colder temperatures. This represents warm advection. Not only is the wind blowing across the isotherms, but the wind is blowing from warmer to colder areas.

[iii] Remember that $\vec{\nabla}T$ is short-hand for $(\partial T/\partial x)$, where x is directed across the isotherms.

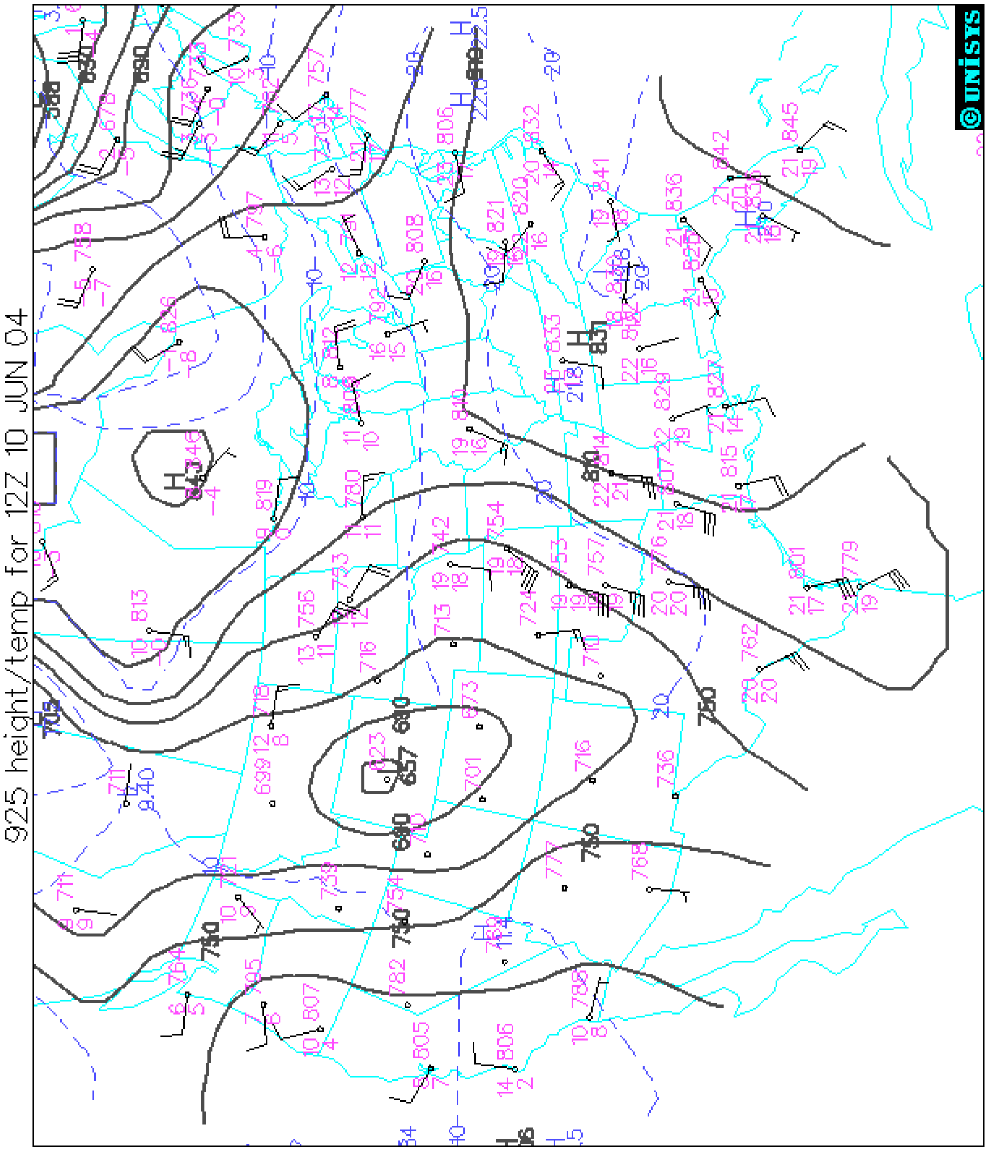

Figure 5.2: Heights (solid lines; every 30 meters) and temperatures (dashed lines; every 5°C) at 925 mb for 12 UTC on 10 June 2004. Heights, temperatures and wind bars at sounding sites are also provided.

Based on this, we would expect the temperature to increase in the central United States. This is indeed what happened, as can be seen the following day (see Figure 5.3).

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. In other words, the wind would be blowing from a region that had the same temperature.

The figure below illustrates the difference between cold advection, warm advection and no temperature advection.

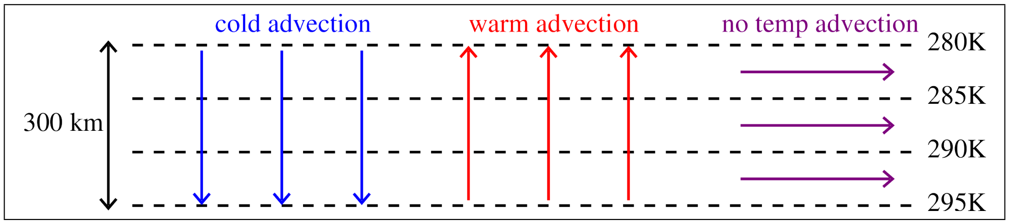

WHAT IF THE WIND HAD BEEN BLOWING FROM SOME OTHER DIRECTION?

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 the temperature advection would depend on whether the wind was blowing toward the east or toward the west. With an east-west temperature gradient, a north or south wind would produce no temperature advection.

Checkpoint 5.7: 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?

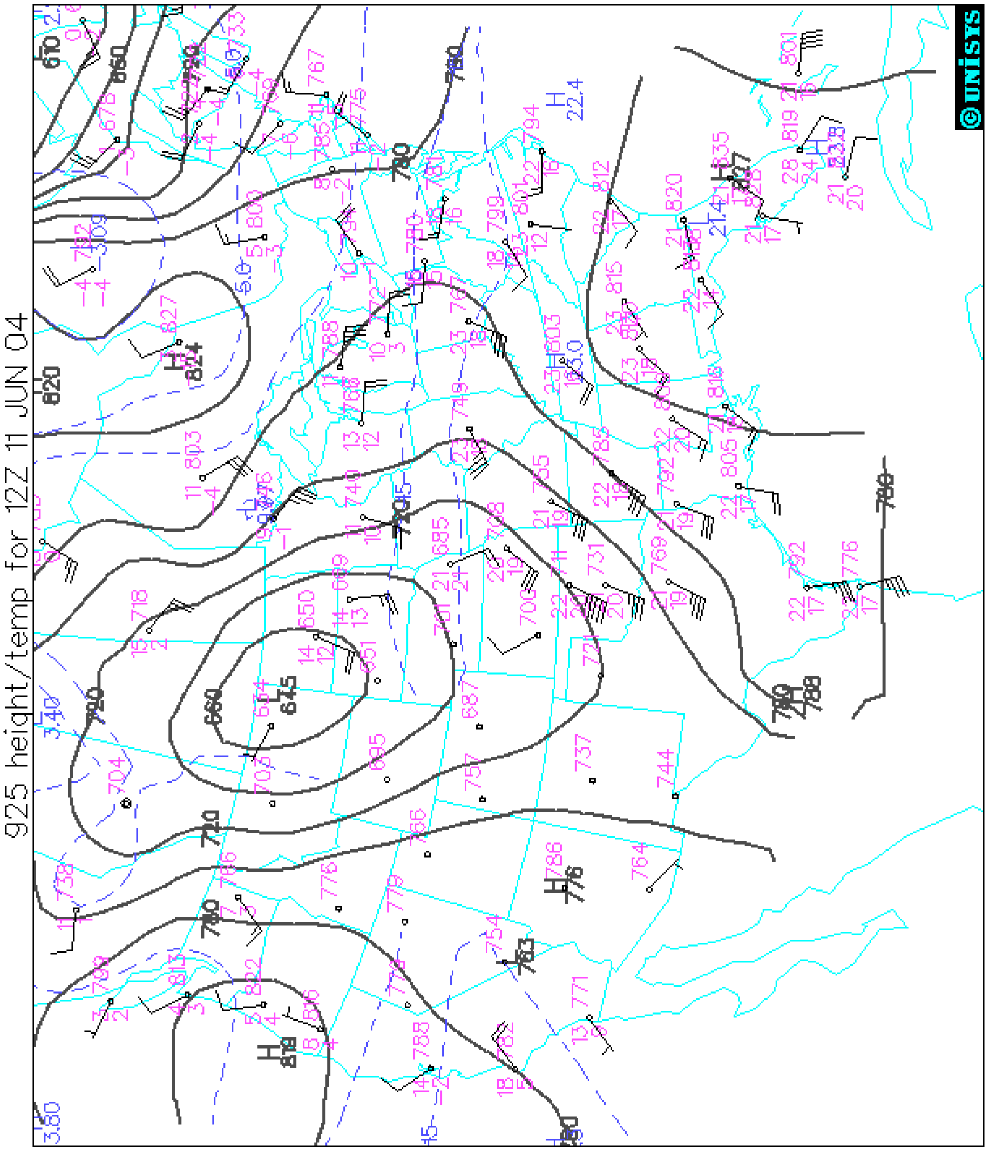

Figure 5.3: Same as in figure 5.2 except for 12 UTC on 11 June 2004 (24 hours later).

5.4 Time tendency equation

At this point, we can predict whether the temperature is going to warm or cool (or neither) by comparing the wind direction with the temperature gradient. Let's now examine how to predict the actual *amount* of warming or cooling. To do that, we need to first recognize that the amount of warming or cooling 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: the **temperature tendency**.

> ✍ In meteorology, any rate is called a **tendency**. For example, the pressure tendency is the rate at which the pressure changes. However, sometimes the word is also used to refer to the amount the quantity changes in a particular period of time. In the former case, the pressure tendency would have units of pressure per time. In the latter case, the pressure tendency would have units of pressure.

Checkpoint 5.8: 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.

Checkpoint 5.9: 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?*

5.4.1 Determining the temperature tendency

To predict *how much* warmer or colder the temperature will be, we need to know how the **temperature tendency** is impacted by the temperature advection.

> ✍ There are several processes that can change the air temperature. In this chapter, we are restricting our analysis to how the temperature advection changes the temperature. Also, the temperature advection can either be horizontal (due to horizontal winds) or vertical (due to upward or downward motion). To simplify things, we are focusing on horizontal advection.

We already know that the temperature advection depends on three things: the temperature gradient ($\vec{\nabla}T$), the wind ($\vec{v}$) and the direction of the wind relative to the temperature gradient.

These three things can be combined in a single expression as follows: $-\vec{v}\cdot\vec{\nabla}T$.

To illustrate what this expression means and how we use it, consider the figure on page 78, which illustrates cold advection on the left, warm advection in the center and no temperature advection on the right.

Let's look at the expression in four parts and apply it to the three cases shown in the figure.

- First consider the temperature gradient, $\vec{\nabla}T$. In the figure, the 280 K isotherm is 300 km from the 295 K isotherm, meaning that the temperature gradient is 0.05 K/km (divide the difference in temperature, 15 K, by the distance, 300 km). The temperature gradient is directed toward warmer temperatures and, in this case, that would be toward the bottom of the page.
- Next let's look at the wind velocity, $\vec{v}$. Let's suppose the wind speed is 20 km/h in each case, but the direction is different (down the page on the left, up the page in the middle, and toward the right on the right).
- Next let's look at the dot between $\vec{v}$ and $\vec{\nabla}T$. The dot is used to indicate we want to multiply the two in a special way, using only the components that are parallel. This is known as the **dot product**.[iv] For example, on the left, the wind and the temperature gradient are both directed down the page. In that case, the dot product equals

---

[iv] Not too imaginative, is it?

the product of the two magnitudes (0.05 K/km and 20 km/h), which is equal to 1 K/h. In the middle, the wind is directed up the page, *opposite* the direction of the temperature gradient. In that case, the dot product is equal to −1 K/h, the *negative* of the product of the two magnitudes. In the right case, the wind is perpendicular to the temperature gradient and the dot product is zero.

- Finally, let's look at the presence of the negative sign. Notice that the dot product is positive in the left case (cold advection) and negative in the middle case (warm advection). The negative sign is added in order to be consistent with how we defined temperature tendency, negative for cooling (as in the left case) and positive for warming (as in the middle case).

Checkpoint 5.10: What is the sign of the temperature advection in the central United States in Figures 5.2 and 5.3? Is this a case of warm advection, cold advection or neither? Why?

5.4.2 Dealing with different units

WHAT DOES IT MEAN FOR THE TEMPERATURE TENDENCY TO BE 1 K/h OR −1 K/h?

A temperature tendency of 1 K/h means the temperature is warming at a rate of 1 K (or 1°C) every hour. A temperature tendency of −1 K/h means the temperature is cooling at a rate of 1 K (or 1°C) every hour.

WHAT HAPPENS IF THE WIND SPEED IS NOT IN KILOMETERS PER HOUR?

If the wind speed was in m/s or mph, the units of the wind or the temperature gradient would have to converted so that the units would cancel appropriately.

To illustrate how to do this, let's consider the conversion from knots (the meteorological convention[v] at the surface) or miles per hour (the U.S. convention) to meters per second (standard scientific practice).

---

[v] This convention may have come about because weather forecasting was first focused on the maritime industry.

To convert it is easiest to focus on the difference between nautical miles, statute miles and kilometers. A kilometer is closer to the other two than a meter is. Switching from kilometers to meters is relatively straightforward, as there are a thousand meters in a kilometer.

A kilometer is smaller than a statute mile, which is smaller than a nautical mile. A nautical mile, as one might guess given the name, is based on maritime navigation and is equal to 1/60 of a degree of latitude. In other words, there are 60 nautical miles in a degree of latitude.

A statute mile (the one used in the United States) is slightly smaller than a nautical mile. There are 69.04677 statute miles in a degree of latitude.

A kilometer is still smaller. There are 111.12 kilometers in a degree of latitude.

Based on those three values, one can show that one statute mile is equal to 1.609344 kilometers (or 1609.344 meters).[vi]

And, since there are 3600 seconds in one hour, one can show that one mile per hour is a little less than 0.5 m/s.[vii]

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 5.1).

Checkpoint 5.11: Which is fastest: 10 m/s, 10 mph or 10 kts? Which is slowest?

5.4.3 Dealing with wind at an angle

WHAT IF THE WIND IS NOT ALIGNED WITH THE TEMPERATURE GRADIENT?

In those cases, the temperature advection can be calculated in one of four ways:

[vi] To get the number of kilometers per mile, divide 111.12 km/°lat by 69.04677 mi/°lat.

[vii] 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.

- 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,
- Use the total speed of the wind and the total temperature gradient and then multiply that product by the cosine of the angle between them.
- Calculate the eastward and northward components of the wind (see chapter 13) and the temperature gradient, and then calculate two products, one for the eastward components and one for the northward components, then add those together.[viii]

All four ways can be used always and you'll get the same results regardless of which method you use. In addition, all four ways simplify to the technique used before when the wind and temperature gradient happen to be parallel or perpendicular.

WHICH METHOD SHOULD I USE TO CALCULATE THE TEMPERATURE ADVECTION WHEN THE WIND IS NOT ALIGNED WITH THE TEMPERATURE GRADIENT?

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 method 1: calculate the temperature gradient along the direction of the wind then multiply that with the total wind speed.

Let's consider, for example, the illustration in Figure 5.4.3. To calculate the temperature advection, use the total wind speed but use the temperature gradient only the direction of the wind.

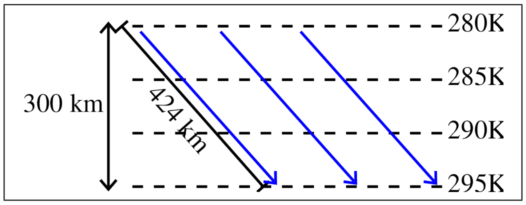

In the illustration, use a length of 424 km instead of 300 km to calculate the temperature gradient. The 424 km value refers to the distance from the 280 K isotherm to the 295 K isotherm in a direction parallel to the wind direction, rather than just straight across the isotherms.

[viii] Using u and x to indicate the wind and distance toward the east and v and y to indicate the wind and distance toward the north, the two products would be $-u\partial T/\partial x$ and $-v\partial T/\partial y$. This is the method used in numerical models because of the way the data is represented at individual points, usually in a rectangular grid.

The temperature gradient in that direction, then, is 0.035377 K/km (15 K divided by 424 km). The temperature advection, then, is 0.71 K/h (multiply the wind speed, 20 km/h, by the temperature gradient).

Checkpoint 5.12: In the example demonstration method 1, the wind direction is 45° from temperature gradient direction. Suppose we used method 3 instead. In that case, we'd use the total wind speed (20 km/h) and the total temperature gradient (0.05 K/km) and we'd multiply that product by the cosine of the angle between them. Calculate that product and compare your answer to what was found using method 1.

Project

Project 5.1: For this project, you will need two maps, 24 hours apart, that contain temperature and wind, like the archive at the Storm Prediction Center. Select two dates 24 hours apart for a particular level (850 mb tends to be good). Identify a region that, on both maps, the wind is clearly blowing (i.e., wind speed greater than 10 knots) 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.

Project 5.2: (a) At the location identified in project 5.1, 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?

Questions

Question 5.1: In which of the following cases will a cup of water freeze quickest?

(a) 35°F while winds of 35 mph are blowing.
(b) 30°F while winds are calm.
Explain your choice.

Question 5.2: For warm advection, what is the sign of the temperature advection?

Problems

Problem 5.1: On the map in Figure 1.1 (page 4), it indicates that the wind speed at Chicago, Illinois, on January 4, 2014, is 20 knots. Convert this to units of mph (miles per hour), km/h and m/s, using only the fact that there are 3600 seconds in an hour, 1000 meters in a kilometer and that one degree of latitude is equal to 69.045677 mi, 60 nm and 111.12 km. Show your work. No short-cuts.

Problem 5.2: Examine the map in Figure 5.2. Determine the following values for Albany, New York, on June 10, 2004, where the wind is 15 knots out of the NNW.
(a) Determine the wind speed in km/h.
(b) 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) 24 hours later 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 5.3?[ix]

[ix] 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.

6. Compression and Expansion

Introduction

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'll examine how changes in pressure can bring about changes in temperatures, and in chapter 7 we'll examine how radiation is related to temperature changes.

6.1 Air parcels

In this chapter, we look at how air can warm or cool 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 and the environment into which the air is moving.

For the air that is moving, imagine we envelop 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 to contain the air inside it. Such a region of air is considered an **air parcel**.

In chapter 5, we examined temperature advection, which refers to how the wind can take an air parcel from one region and transport it to another region. The assumption, at the time, was that the air parcel's temperature remains the same during the transport. If the parcel's temperature is changing en route to our location, we'll need to add that change into the temperature tendency calculation.

Mathematically, we can include that change as follows:[i]

$$\frac{\partial T}{\partial t} = -\vec{v} \cdot \vec{\nabla} T + \frac{dT}{dt} \tag{6.1}$$

In this expression, I'm using slightly different notation to distinguish between the rate the temperature changes at a particular location ($\partial T/\partial t$) and the rate at which the parcel's temperature changes as it moves (dT/dt). The temperature advection ($-\vec{v} \cdot \vec{\nabla} T$, see chapter 5) is just one way the temperature at a location can change. There can be other factors contributing to the temperature tendency, like compression, expansion, radiation, etc., and those factors are included in the dT/dt expression.

> We call dT/dt the **total derivative** of temperature and it represents a **Lagrangian** change (meaning the change following the parcel). In comparison, we call $\partial T/\partial t$ the **partial derivative** and it represents a **Eulerian** change (meaning the change at a fixed position).

In this chapter, we are examining what impacts the parcel as we follow it along its path. For that reason, I'll be using the dT notation instead of the ∂T notation.

Checkpoint 6.1: Suppose the temperature is increasing at our location and temperature advection is the only process responsible for the local change in air temperature. What is the value of dT/dt: positive, negative or zero?

6.2 Adiabatic changes in temperature

The question we want to answer is: What happens to the temperature of the air parcel as we follow it on its journey through the atmosphere?

To answer the question, we'll assume the air parcel has the following properties:[ii]

---

[i] This expression is usually written as $dT/dt = \partial T/\partial t + \vec{v} \cdot \vec{\nabla} T$.

[ii] We will also assume that its translational kinetic energy (as well as any change in gravitational potential energy) is small compared to its internal energy.

- The air parcel will *not mix* with any air it encounters. It is as though the imaginary surface surrounding the air parcel is impermeable.
- There is no conduction of heat between the parcel and its surroundings.
- 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. As long as our parcel is large enough and the time scale is small enough, this is a pretty good assumption.

For example, whenever you see a puffy, white cumulus cloud, you might notice that there is a sharp edge between the cloud and the surrounding clear air. From our perspective on the ground, the sharp edge means that the mixing is small compared to the transport of the air inside the cloud.[iii] If the air inside the cloud was mixing with the clear air outside the cloud, the cloud edge wouldn't be so sharp.

✍ | We will explore how mixing affects things in section 10.5.

How can we say that there is no conduction of heat between the parcel and its surroundings?

Air is actually a very poor conductor of heat. So, if there is no mixing, the air parcel can be considered to be thermally isolated (so any temperature changes inside the parcel are not due to heat transfer across the boundaries). This means that encountering air that is colder or warmer won't make the air parcel warmer or cooler.

If there is no heat transfer across the boundaries, how can the temperature of the parcel change?

[iii] Diffusion is not totally ignored in operational numerical models of the atmosphere, but it is small, especially when you consider that a typical numerical model (see page 12) has a resolution where each grid point represents an air parcel that has a size that dwarfs the size of an individual cloud.

There are several ways that an individual air parcel's temperature can change, even though there is no heat transfer across the boundary of the air parcel. In this chapter we examine the change in temperature that results from when an air parcel expands against the surroundings or is compressed by the surroundings. In chapter 7, we examine the change in temperature that results from when the air parcel absorbs radiation or emits radiation. In chapter 16, we examine the change in temperature that results from when water changes state.

When would an air parcel expand or be compressed and why would that change its temperature?

The parcel will expand or contract whenever it enters a region with a different air pressure, as when an air parcel rises or falls (since the pressure decreases with height, as described in section 4.1.1 and explored further in section 9.1), and consistent with 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).

When the air parcel expands and pushes against the environment, the air parcel does work on the environment and, in so doing, the air parcel loses energy and cools down. When the air parcel is compressed by the environment and contracts, 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*).

Checkpoint 6.2: 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?

For the temperature of the air parcel to change adiabatically, it must enter an environment where the pressure is different from its own. As mentioned above, this can happen when an air parcel rises or falls.

By how much does the temperature rise or fall when the air parcel compresses or expands?

As I will discuss in section 6.3, one can use conservation of energy to show that for adiabatic processes (no heat transfer) the small change in temperature dT depends upon the small change in pressure dP as follows:

$$dT = \frac{1}{c_\mathrm{p}\rho} dP \tag{6.2}$$

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"?

Exactly what this means will be discussed in section 6.3. It suffices to say at this point that it has a value of 1004.67 $\mathrm{J \cdot kg^{-1} K^{-1}}$. The phrase "of dry air" is used because it assumes there is no water vapor present.

Checkpoint 6.3: According to equation 6.2, if the pressure decreases, what happens to the temperature?

The problem with equation 6.2:

$$dT = \frac{1}{c_\mathrm{p}\rho} dP$$

is that involves the density ρ. The density will change as the pressure changes, which means that we can't use this expression for large changes in pressure, since there won't be a single value of the density that we can use.

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 (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 section 6.5), resulting in the following:

$$T = T_0 \left(\frac{P}{P_0} \right)^{R/c_\mathrm{p}} \tag{6.3}$$

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. I will call this the **adiabatic temperature equation**.[iv]

> ✍ The ratio of R to c_p is very close to 2/7, which you can show in Problem 6.2. Consequently, I will frequently write 2/7 in place of R/c_p when using this equation in examples and such.

Checkpoint 6.4: According to equation 6.3, what happens to the temperature if the pressure decreases?

Example 6.1: Suppose 10°C air at 800 mb is compressed adiabatically to 1000 mb. What would its temperature be then?

Answer 6.1: Using equation 6.3, 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 6.3, uses the absolute temperature (not the temperature relative to freezing).

[iv] Some people call this **Poisson's equation** or **the Poisson equation**, while others refer to a version of it (namely, $TP^{-\kappa}$ = constant, where $\kappa = R/c_\mathrm{p}$) as just one of three Poisson's equations (the other two being PV^γ = constant and $TV^{\gamma-1}$ = constant, where $\gamma = c_\mathrm{p}/c_\mathrm{v}$). I'm avoiding using Poisson, not only to avoid the confusion with these other expressions, but also because there is a different equation in physics that is called the Poisson's equation (namely, $\nabla^2 V = -\rho/\epsilon_0$). Besides, "adiabatic temperature equation" is so much more descriptive.

Checkpoint 6.5: When using equation 6.3, does it matter if the temperature is in kelvin or in degrees Celsius? Does it matter if the pressure is in millibars or pascals? If so, which? If not, why not?

6.3 First law of thermodynamics

On page 91, equation 6.2 was introduced to relate a parcel's change in temperature with its change in pressure (assuming no heat transfer with the environment).

$$dT = \frac{1}{c_\mathrm{p}\rho} dP$$

At the time, it was mentioned that this came from conservation of energy, but a detailed derivation wasn't provided.

While it isn't crucial to know where the expression came form to use it, knowing the derivation helps us remember the assumptions inherent in the expression and thus allows us to know the applicability of the expression. Deriving the expression also gives us insight into the meaning of c_p and why the change in temperature should depend only on the change in pressure.

Before showing you the derivation, I'd like to first discuss the major assumptions that go into the derivation.

This expression comes from the **first law of thermodynamics**. The first law of thermodynamics is essentially a statement that energy is conserved, which means that the *total* amount of energy in the universe is the same but just is transferred locally from one type to another.

When applied to a parcel of air, this means that any change in the parcel's energy is due to energy being transferred to or from it (via the surroundings).

The parcel's energy is called the **internal energy**, since it is internal to the parcel.[v] The internal energy is related to the parcel's temperature, and so I

---

[v] Technically, if the parcel is moving it also has translational kinetic energy but, as noted in the footnote on page 88, we are assuming that the parcel is moving slowly enough that its translational kinetic energy is small compared to the internal energy.

will frequently refer to it as the **thermal energy**, although technically the internal energy can be more than just the thermal energy.

Mathematically, the internal energy is indicated by U.

According to the principle of conservation of energy, the parcel's energy U can change only if energy is transferred to it from the surroundings or from the parcel to the surroundings. There are two ways energy is transferred:[vi]

Q is the energy added to the system due to *microscopic* processes (like conduction and radiation), sometimes called **heating** or **heat**, and

W is the energy subtracted from the system due to *macroscopic* processes (like expanding against the environment), called the **work** done by the system on the environment.[vii]

Mathematically, then, a small change in the parcel's internal energy, dU, must be equal to any transfer to the parcel due to heating, dQ, minus any transfer away from the parcel due to work, dW:[viii]

$$dU = dQ - dW \tag{6.4}$$

Let's interpret this relationship and make sure it makes sense. Basically, it says that for an object's temperature to increase ($dU > 0$), we have to heat it ($dQ > 0$) and/or have work done on it ($dW < 0$). Conversely, for an object's temperature to decrease ($dU < 0$), we have to cool it ($dQ < 0$) and/or have it do work on the environment ($dW > 0$).

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.

[vi] There is a fine and somewhat arbitrary line between these two mechanisms.

[vii] 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.

[viii] 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).

Checkpoint 6.6: According to the first law of thermodynamics (equation 6.4), 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 6.4 because, as written, it isn't particularly useful for our purposes. We need a version that is written in terms of variables we can measure, like temperature and pressure (dT and dP). When we do so, it looks like:

$$c_{\mathrm{p}} dT = \frac{dQ}{m} + \frac{dP}{\rho} \tag{6.5}$$

where $c_{\mathrm{p}} = 1004.67$ 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 6.2 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 6.2.

WHERE DOES EQUATION 6.5 COME FROM?

To derive this expression, we need to utilize a couple of definitions:

$dW = F dx$ 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.

Checkpoint 6.7: In the derivation of equation 6.5, four definitions are used to get it in a form involving dP and dT. What are they?

With the four definitions identified above, we can take the generic version of the first law of thermodynamics (see equation 6.4):

$$dU = dQ - dW$$

and express it in terms of pressure P and temperature T, as expressed in equation 6.5:

$$c_\mathrm{p} dT = \frac{dQ}{m} + \frac{dP}{\rho}$$

where $c_\mathrm{p} = 1004.67$ J kg^{-1} K^{-1} for dry air (no moisture). As mentioned earlier, you don't need to go through the derivation to be able to use the adiabatic temperature equation, which is based on the first law. However, by going through the derivation you gain a better sense of what is represented by c_p (the **specific heat** of dry air at constant pressure) and what it means to be "at constant pressure."

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[ix], we find that dW equals PdV, not $-VdP$. If we make this replacement, we get

$$dU = dQ - PdV \tag{6.6}$$

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[x], $PV = mRT$ and so $d(PV) = d(mRT) = mR_d T$. Thus, dW can be replaced by $mR_d T - VdP$.

Making this replacement, we get

$$\begin{aligned} dU &= dQ - dW \\ &= dQ - mR_d T + VdP \end{aligned} \tag{6.7}$$

[ix] 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 $dW = PdV$.

[x] 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 $\rho = m/V$).

Although closer to what we want, this is still not what we need.

To complete the derivation and go from equation 6.7 to equation 6.5, 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**[xi] 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.
- Suppose we use the version of the first law expressed in equation 6.7 ($dU = dQ - mR_dT + VdP$) 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.

Basically, the expression for dU above allows us to replace the left side of equation 6.7

$$dU = dQ - mR_dT + VdP$$

with $mc_\mathrm{p}dT - mRdT$, which gives us

$$mc_\mathrm{p}dT - mRdT = dQ - mR_dT + VdP.$$

This simplifies to equation 6.5 with a bit of algebra.

Checkpoint 6.8: It was stated that a bit of algebra is needed to go from the last expression above ($mc_\mathrm{p}dT - mRdT = dQ - mR_dT + VdP$) to equation 6.5. What are the algebraic steps that are needed?

[xi] This is sometimes called the **specific heat capacity**.

6.4 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 6.7 ($dU = dQ - mR_dT + VdP$) 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 6.6 ($dU = dQ - PdV$) 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:

$$dU = mc_{\mathrm{v}}dT. \tag{6.8}$$

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.

> ✍ Not only does c_{v} and R add up to c_{p} but for dry air the three form simple ratios, namely that $c_{\mathrm{p}} : c_{\mathrm{v}} : R = 7{:}5{:}2$. This is a consequence of the diatomic make-up of the majority of atmospheric molecules (nitrogen and oxygen).

Checkpoint 6.9: In this section it was stated that $dQ = m_\text{v} dT$. *However, in the derivation of our version of the first law of thermodynamics (previous section), it was stated that* $dQ = mc_\text{p} dT$. *Why was the p subscript used before and the v subscript is used here?*

6.5 Deriving the adiabatic temperature equation

In this section, I derive the adiabatic temperature equation (6.3),

$$T = T_0 \left(\frac{P}{P_0}\right)^{R/c_\text{p}}$$

from

$$c_\text{p}\frac{dT}{T} = R\frac{dP}{P}.$$

Not only does this help you see where the adiabatic temperature equation comes from but it also helps with Show-me 6.1.

First, we recognize that an infinitesimal (extremely tiny) change in a value, divided by the value, is equal to an infinitesimal change in the natural log of the value:[xii]

$$d\left(\ln x\right) = \frac{dx}{x}$$

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

$$c_\text{p} d(\ln T) = R d(\ln P). \tag{6.9}$$

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).

[xii] In other words, the derivative of the natural log of x is the inverse of x.

Although the expression just considers infinitesimal changes, it also applies to large changes (since there are no other variables undergoing a change), and so we can write the expression as follows:[xiii]

$$c_{\mathrm{p}}\Delta(\ln T) = R\Delta(\ln P)$$

or, rearranging,

$$\Delta(\ln T) = \frac{R}{c_{\mathrm{p}}}\Delta(\ln P)$$

We now take advantage of the fact[xiv] that $\Delta \ln T$ is equivalent to $\ln(T/T_0)$, where T and T_0 are the final and initial temperatures. Rewriting the expression in that form (and doing the same to $\Delta \ln P$) 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[xv] that $a \ln x$ is equal to $\ln x^a$. This allows us to write it as follows:

$$\ln\left[\frac{T}{T_0}\right] = \ln\left[\left(\frac{P}{P_0}\right)^{R/c_{\mathrm{p}}}\right] \tag{6.10}$$

Finally, we take the exponential[xvi] of each side and then multiply by T_0 to get equation 6.3:

$$T = T_0\left(\frac{P}{P_0}\right)^{R/c_{\mathrm{p}}}$$

[xiii] Technically, the large change in lnT (i.e., $\Delta \ln T$) is just the summation of all of the small changes in $\ln T$ (i.e., the integral of $d \ln T$). However, the integral of $d \ln T$ is ΔT and the integral of $d \ln P$ is ΔP.

[xiv] 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.

[xv] 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.

[xvi] The natural logarithm of e^a is a. Taking the exponential of that gives e^a, exactly what we had at the start. So, taking the exponential of a natural logarithm gives you the initial value (the two cancel each other out).

where T_0 and P_0 are the old temperature and pressure and T and P are the new temperature and pressure.

Checkpoint 6.10: In the derivation, we took

$$c_{\mathrm{p}} d \ln T = R d \ln P$$

and replaced it with

$$c_{\mathrm{p}} \Delta(\ln T) = R \Delta(\ln P)$$

Why couldn't we have done the same thing with the original expression,

$$c_{\mathrm{p}} \frac{dT}{T} = R \frac{dP}{P}$$

and replaced it with the following?

$$c_{\mathrm{p}} \frac{\Delta T}{T} = R \frac{\Delta P}{P}$$

Questions

Question 6.1: 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 at the surface. How do we know this is not because the rising water cools by expanding adiabatically, as with air parcels?

Problems

Problem 6.1: Jet aircraft normally cruise at an altitude where the pressure and temperature are typically 200 mb and -60° C. Determine the temperature of this outside air if it were adiabatically compressed to the 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 6.2: One by-product of the derivation of specific heat in section 6.4 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 readings there is a rather simple relationship between them: $c_\mathrm{p} = c_\mathrm{v} + R$.
(a) It has been stated that for dry air $c_\mathrm{p} = 1004.67\ \mathrm{J \cdot kg^{-1}K^{-1}}$. Given the value of R for dry air that we used before, calculate c_v. Which is bigger: c_p or c_v?
(b) Use the values of R, c_p and c_v to show that the ratios $c_\mathrm{p} : c_\mathrm{v} : R$ is equal to $7 : 5 : 2$ (in keeping with the note on page 98).

Derivations

Show-me 6.1: Derive the adiabatic temperature equation (equation 6.3),

$$T = T_0 \left(\frac{P}{P_0} \right)^{R/c_\mathrm{p}}$$

from equation 6.9 (see section 6.5),

$$c_\mathrm{p} d \ln T = R d \ln P.$$

7. Radiation

Introduction

So far, we have examined how compression and expansion can change the air's temperature. 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.

In this chapter and the next chapter, we examine a second process for changing the air temperature, a process that involves **radiation**. In this chapter, I'll describe the radiation. In the next chapter I will show how radiation impacts the temperature.

> ✍ In chapter 16, we examine a third process, called latent heating and cooling, which occurs when water condenses or evaporates (and is why we feel colder when water evaporates from our skin).

7.1 Describing radiation

7.1.1 What is radiation?

Radiation is an **electromagnetic wave** – 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.[i] 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.

---

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

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 the environment wants to be in 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 at one location. This "desire for the status quo" gives rise to a water wave that travels away from that location at a particular speed (which can depend on the depth of the water, among other things).

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, where the field represents its potential 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 (the speed of light).

The general term for **electromagnetic waves** is called **radiation**. As we will see in section 7.1.3, visible light is one type of electromagnetic wave. Consequently, I will tend to use the words "electromagnetic waves," "radiation" and "light" interchangeably.

Checkpoint 7.1: What is the difference between light and electromagnetic radiation?

7.1.2 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!

Mathematically, we represent the speed of light as c (as in the equation $E = mc^2$; see Appendix E). With careful measurements, though, scientists have found that in the vacuum of space this speed is close to 300,000,000 meters per second. More precisely, the speed of light in the vacuum of space is 2.99792458×10^8 m/s. 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 is 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.

Checkpoint 7.2: Does light travel in a vacuum?

7.1.3 Frequency and wavelength

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

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**)[ii] that at high doses can have an negative effect on living cells.

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

Except for visible light, which is a type of electromagnetic radiation, we (as human beings) are oblivious to electromagnetic waves. We cannot hear them, taste them or touch them. What distinguishes visible light from other types of electromagnetic radiation is its **frequency**. Other than visible light, we need specialized instruments to detect electromagnetic radiation, with each instrument designed to detect a specific range of electromagnetic frequencies.

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 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, which is the lower-case Greek letter "nu"). The frequency is the inverse of the **period** T (which is the time per cycle):

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

Just as water waves can have different frequencies, electromagnetic waves can likewise have different frequencies. Since electromagnetic waves of different frequencies can have different properties, we use names to distinguish between different frequency ranges (see Figure 7.1). The entire range of electromagnetic 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 radio wave 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

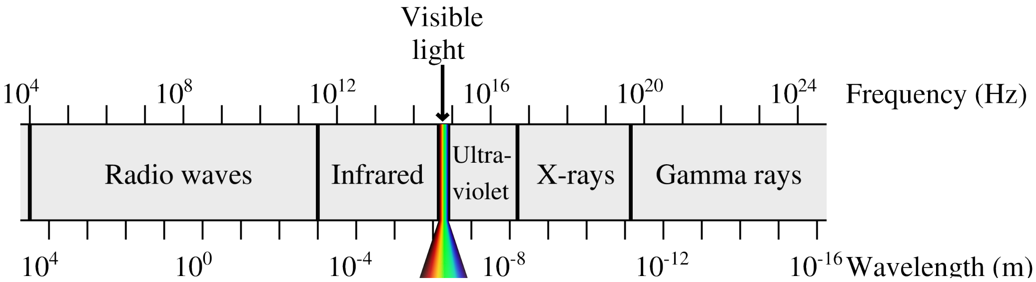

Figure 7.1: The portion of the electromagnetic spectrum from 10^4 Hz to 10^{24} Hz (wavelength range from 10^{-16} m to 10^4 m). Higher frequencies (and smaller wavelengths) are on the right.

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.[iii]

Checkpoint 7.3: 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 (see Figure 7.2. 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:

$$v = f\lambda \tag{7.1}$$

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.

[iii] 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.

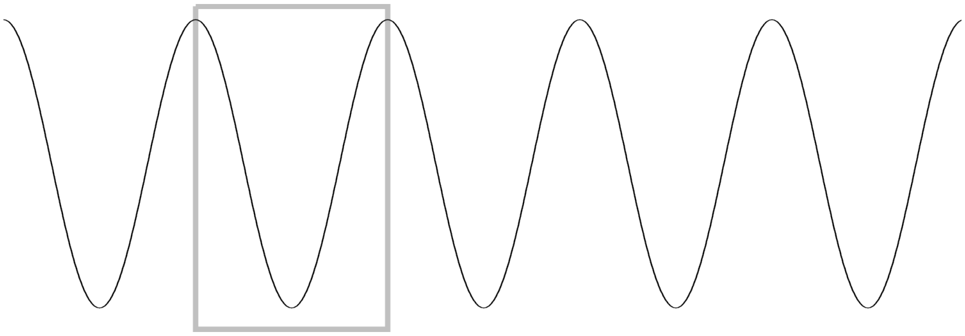

Figure 7.2: A snapshot of a wave. The box identifies a segment that, when repeated, reproduces the pattern. The width of the box is the wavelength of the wave.

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

Checkpoint 7.4: 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 frequency stays the same, the wavelength must change whenever the speed changes, and the speed depends on the medium. Thus, any reference to the wavelength depends on the medium.

On the other hand, many people are more familiar with the wavelength values than the frequency values. Plus, 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 same as those in a vacuum.

Another advantage of wavelength is that the way the radiation interacts with objects depends on the wavelength of the radiation (relative to the size of the object). For example, if the wavelength is much, much larger than the size of the object, the object will have no impact on the radiation, making the object essentially invisible to those waves. For example, FM radio waves have wavelengths about 3 meters in length. Those waves travel easily through the air, since the molecules, cloud droplets and other particles that are in 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 radiation will reflect off the object. For example, visible light reflects off us because we are much bigger than the wavelength of visible light (around half a micron or 500 nm).

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 **radar**[iv] uses 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.

An example of a radar map is shown in Figure 7.3.

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 (collocated 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 traveled 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.

[iv] Radar stands for Radio Detection And Ranging.

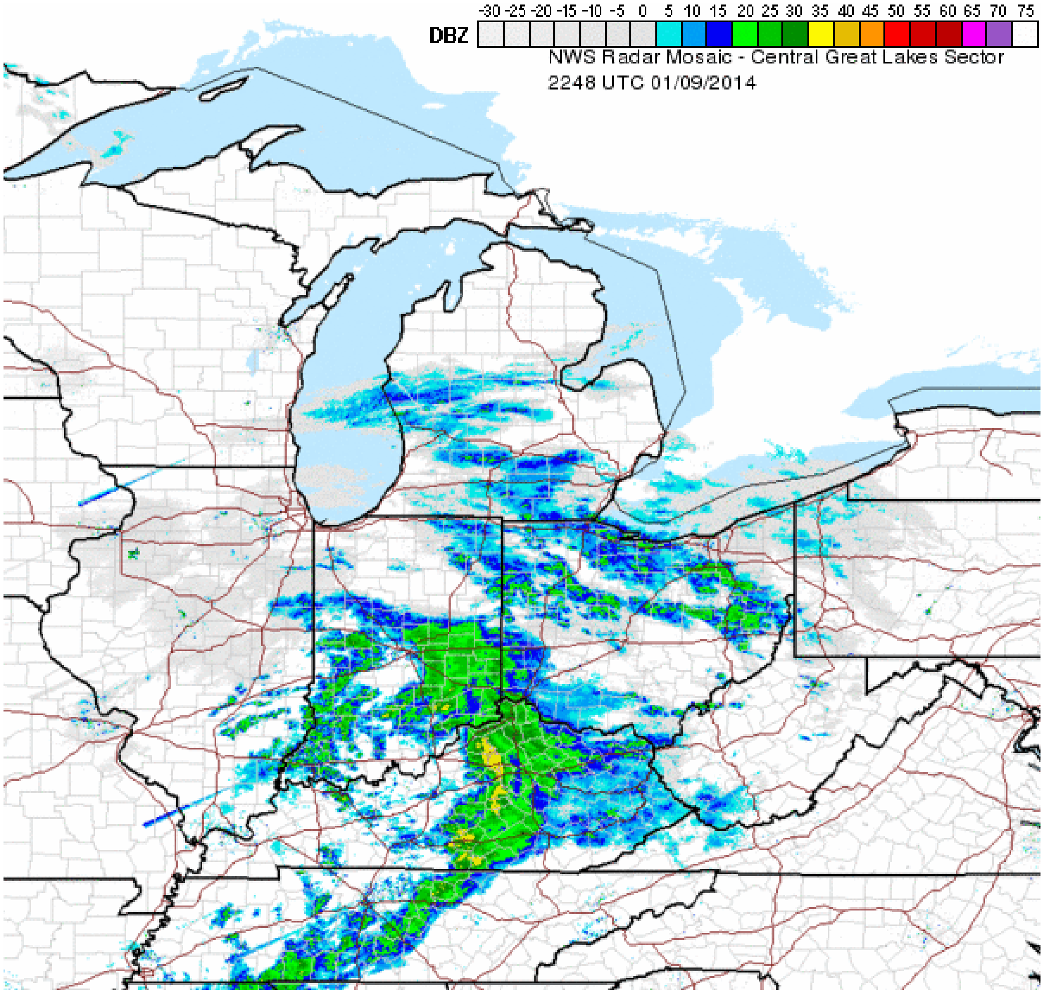

Figure 7.3: Reflectivity for the Midwest United States on January 09, 2014, at 2248 UTC.

✍ The radar map shown on page 110 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 (where "WSR" stands for 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.

Checkpoint 7.5: What is the purpose of having weather radars use a wavelength of radiation that is so much larger than cloud drops?

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**.

We use the term **ultraviolet** for wavelengths shorter than what the eye can detect (i.e., frequencies higher than the eye can detect). The shortest wavelengths our eyes can detect are around 300 nm and are interpreted as blue and violet colors. Consequently, ultraviolet radiation refers to those waves that are "beyond violet".

Ultraviolet radiation can be harmful. The part of light that causes tanning and sunburn, for example, is the ultraviolet part of the sunlight.

We use the term **infrared** for wavelengths longer than what the eye can detect (frequencies lower than the eye can detect). The longest wavelengths our eyes can detect are around 700 nm and are interpreted as red and orange colors. Consequently, infrared radiation refers to those waves that 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 will be discussed in section 7.3, these objects also tend to emit infrared radiation.

Checkpoint 7.6: Which type of radiation is associated with a higher frequency of electromagnetic waves: ultraviolet or infrared?

7.2 Satellite pictures

Strong storms tend to have clouds that reach high into the troposphere. The stronger the storm, the higher the cloud tops. Thus, it is useful to have a system that can identify how high the clouds are.

Whereas ground-based observers can determine the base of the clouds, cloud tops are determined by satellite systems. For example, Figure 7.4 displays a picture taken by a satellite high above 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 highlight 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 radar 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?

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:

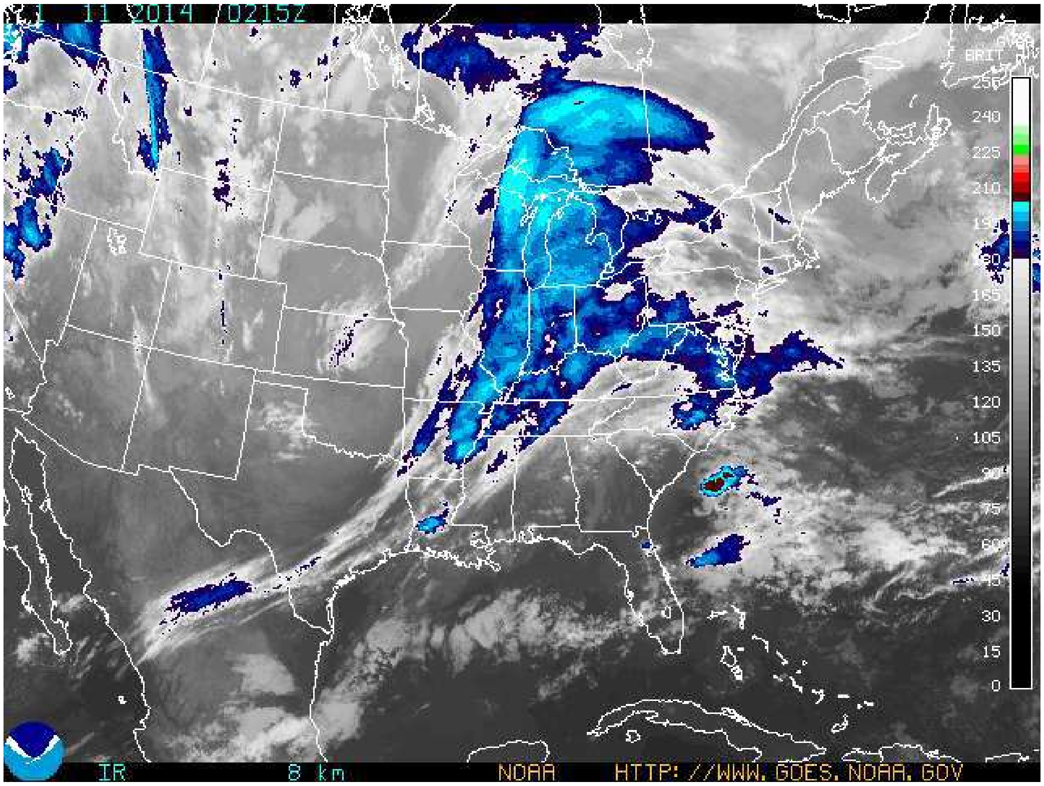

Figure 7.4: Infrared satellite picture for the Eastern Conus Sector on January 11, 2014, at 0215 UTC.

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.
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.

Checkpoint 7.7: What is the difference between a low cloud and a high cloud that allows an infrared satellite to distinguish between them?

7.3 Emission

We know from section 7.2 that all objects emit radiation, and warmer objects emit more radiation than cooler objects. It turns out that the amount of radiation depends on both the temperature of the object and its size. For example, the Sun emits almost 4×10^{26} joules every second (from NASA Sun Fact Sheet; see problem 7.1) whereas Earth only emits a little over 1×10^{17} joules every second, and the average human male emits about 1×10^{3} joules every second.

A couple of things to notice about these radiation values.

1. In a sense, radiation "carries" energy and is given as a rate, with SI units of joules per second. This combination of units is equal to a watt, abbreviated as W. Consequently, we can say that the Sun radiates at almost 4×10^{26} W.
2. The rate at which energy is emitted is called the **luminosity**. Consequently, we can say that the Sun's luminosity is almost 4×10^{26} W.
3. The luminosity depends on two things: the object's temperature and its size (or, more precisely, its surface area). The Sun has a much higher luminosity than Earth because the Sun is both hotter and bigger. Our skin temperature is warmer than Earth, on average, but our luminosity is much less than Earth's because our surface area is so much smaller.

It turns out that an object's luminosity can be determined via the following equation:

$$P = A\sigma T^4 \tag{7.2}$$

where A and T are the object's surface area and temperature, respectively, and σ is a constant that is always equal to 5.67×10^{-8} W·m^{-2}·K^{-4}. Indeed, that is how I obtained the luminosity of Earth and a human male. I simply plugged in the appropriate surface area and temperature values.

✍ The surface area of a sphere is equal to $4\pi R^2$, where R is the radius of the sphere.

Why is luminosity represented by the letter P?

That is because luminosity is a rate at which energy is transferred, and the general term for that is **power**.

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 watts.

Checkpoint 7.8: The average surface temperature of Earth is about 288 K (15°C) whereas the skin temperature of a human is about 306 K (33°C). Does that mean a human's luminosity is greater? Why or why not?

7.4 Irradiance

Can we use the luminosity to determine the air temperature?

Yes and no. A really hot star may have a very high luminosity, but if it is really far away from us it will have no impact on the temperature of our atmosphere because the energy emitted by the star spreads out in space and, by the time the light reaches us, the energy has spread out so much that the energy we receive is very little. On the other hand, a toaster may be hot and close but it is so small its luminosity is too small and so the energy we receive from the toaster is also very little.

What about the Sun?

Certainly, the Sun has an impact on air temperature. It has a high luminosity and it is relatively close. However, we don't receive the entire luminosity. To figure out the Sun's impact, we need to figure how much of the emitted solar energy is being intercepted by Earth. To do this, we'll first figure out the Sun's **irradiance**[v], which is essentially the *concentration* of the energy flow rate. To find the irradiance, divide the luminosity by the surface area the energy is incident upon:

$$F = \frac{A\sigma T^4}{A_{\text{incident}}}$$

where F is the irradiance and A_{incident} is the area through which the energy (given by $A\sigma T^4$) is flowing.

> ✍ 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 or I for irradiance. Unfortunately, E is likely to cause confusion with energy and I looks like the number one. I personally prefer Q, since F reminds me of force.

---

[v] There are many different terms floating around for this. Some examples include **radiant flux**, **radiant power**, **radiant exitance**, **radiant emittance**, **radiant flux density** and **radiative flux**. Using the term "flux" has the advantage of conveying an idea similar to **latent heat flux** and **sensible heat flux**. I'm avoiding the term "flux" because there does not seem to be a consistent definition for it. Just keep in mind 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.

For example, at the Sun's surface, the solar irradiance (concentration) is just σT_{sun}^4. This is because the area emitting the radiation A is the same area that we are using to find the concentration A_{incident} – both are equal to the Sun's surface area. In comparison, at Earth's orbit, the irradiance (concentration) is much less because the radiation spreads out in space.

Suppose we enclose the Sun with a big invisible balloon with a radius equal to the radius of the Sun. The radiation passing through that balloon's surface would be equal to the Sun's luminosity, which is almost 4×10^{26} W, as mentioned in the previous section. Now suppose we enclose the Sun with an even bigger invisible balloon, with a radius equal to the orbit of Earth. The radiation passing through that balloon would *likewise* equal 4×10^{26} W, but the *concentration* of that energy would be less because it would be more spread out, and Earth would only intercept a small fraction of the 4×10^{26} W.

To determine the solar irradiance at Earth's orbit, we divide the Sun's luminosity by the surface area of the big balloon – which has a radius equal to Earth's orbit.

$$F = \frac{A_{\text{sun}} \sigma T_{\text{sun}}^4}{4\pi R_{\text{Earth orbit}}^2}$$

Plugging in the Sun's surface area, its surface temperature and Earth's orbital radius (see appendix B.4), we find a solar irradiance equal to 1361 W/m^2.

> ✍ Notice that the irradiance has units of watts per square meter, not just watts, because it represents a "concentration" of power or a power per area.

Example 7.1: The surface of Earth has an average temperature of 288 K. What is the irradiance emitted at Earth's surface?

Answer 7.1: If we are at Earth's surface then our "balloon" happens to be the same size as Earth and we just need to divide Earth's luminosity by its surface area. Mathematically, the luminosity is $A\sigma T^4$ (see equation 7.2) so, after dividing through by A, Earth's irradiance (at its surface) would just be σT^4. Plugging in the value of σ and Earth's average temperature of 288 K, one gets 390 W/m^2.

Notice that when calculating the irradiance of an object at that object's surface, the areas cancel and we get σT^4. This relationship is known as the **Stefan-Boltzmann law**.

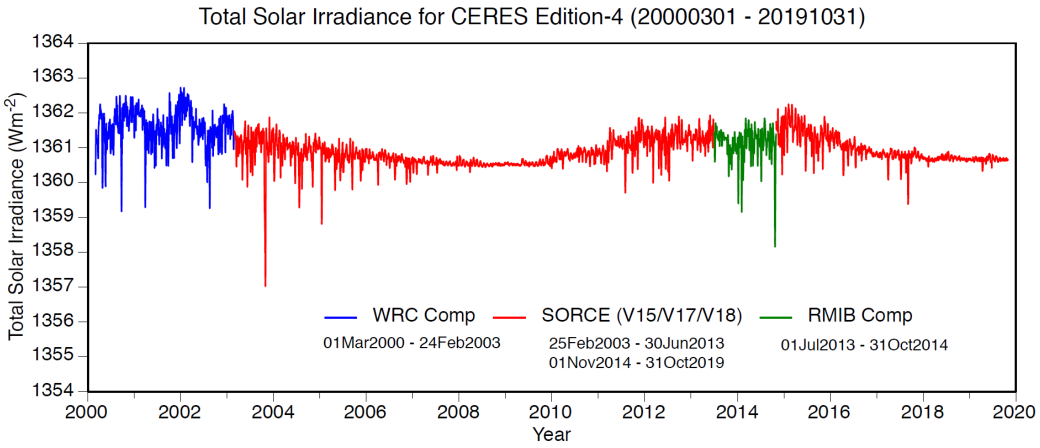

Figure 7.5: Solar irradiance outside Earth's atmosphere (from NASA).

Checkpoint 7.9: According to the Stefan-Boltzmann law, which should be greater: Earth's irradiance at its surface or the Sun's irradiance at its surface? Note: the Sun is a lot hotter than Earth.

How do we know the Sun's surface temperature?

In practice, we actually do the reverse of what was described above. Rather than measuring the Sun's surface temperature and then calculating what the solar irradiance must be at Earth's orbit, we instead measure the solar irradiance at Earth's orbit and then infer what the Sun's surface temperature must be to produce that irradiance (see Problem 7.2). The latter, of course, is much easier since we happen to be located at Earth's orbit rather than the Sun's surface. That is actually how we obtained the 1361 W/m^2 value.[vi]

The 1361 W/m^2 value is known as the **solar constant**. 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 Figure 7.5). 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.

Is the solar irradiance different outside the atmosphere than on Earth's surface?

[vi] Source: NASA Fact Sheet.

Yes. Although the extra distance to Earth's surface is insignificant, as the radiation passes through the air, some of it might reflect back out into space or get absorbed by the air (which we'll discuss in chapter 8). Consequently, the solar irradiance is less by the time the radiation gets to the surface of Earth.

Checkpoint 7.10: How does the irradiance of the Sun change as one moves away from the Sun?

7.5 Spectral irradiance

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 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 7.3) but also extends the range into the higher frequencies (compared to the radiation emitted by Earth).

The **peak emission** frequency or wavelength can be predicted from **Wien's law**:

$$\lambda_{\text{peak}} = \frac{2897\mu\text{m} \cdot \text{K}}{T} \tag{7.3}$$

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 7.3, 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 7.1.3, 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 7.6 (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 7.6 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.

Since the irradiance in a particular range depends on how small the range is, we instead 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 in 1-nm bins). Consequently, the result has units of $\mathrm{W/m^2}$ per nm (or, $\mathrm{W/m^2}$ per μm, if μm bins are used).

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**:

$$F_\lambda = \frac{c_1}{\lambda^5 \left[\exp\left(\frac{c_2}{\lambda T}\right) - 1\right]} \tag{7.4}$$

where $c_1 = 3.74 \times 10^8\ \mathrm{W} \cdot \mathrm{m}^{-2} \cdot \mu\mathrm{m}^4$ (equal to $2hc^2$, where h is Planck's constant and c is the speed of light) and $c_2 = 1.44 \times 10^4\ \mu\mathrm{m} \cdot K$ (equal to hc/k_B, where k_B is the Boltzmann constant). This expression is used in the project to calculate the spectral irradiance of the Sun and Earth. Note that λ must be in μm to have the units work out.

DOES THE SPECTRAL IRRADIANCE, LIKE THE IRRADIANCE, DEPEND ON HOW FAR YOU ARE FROM THE OBJECT?

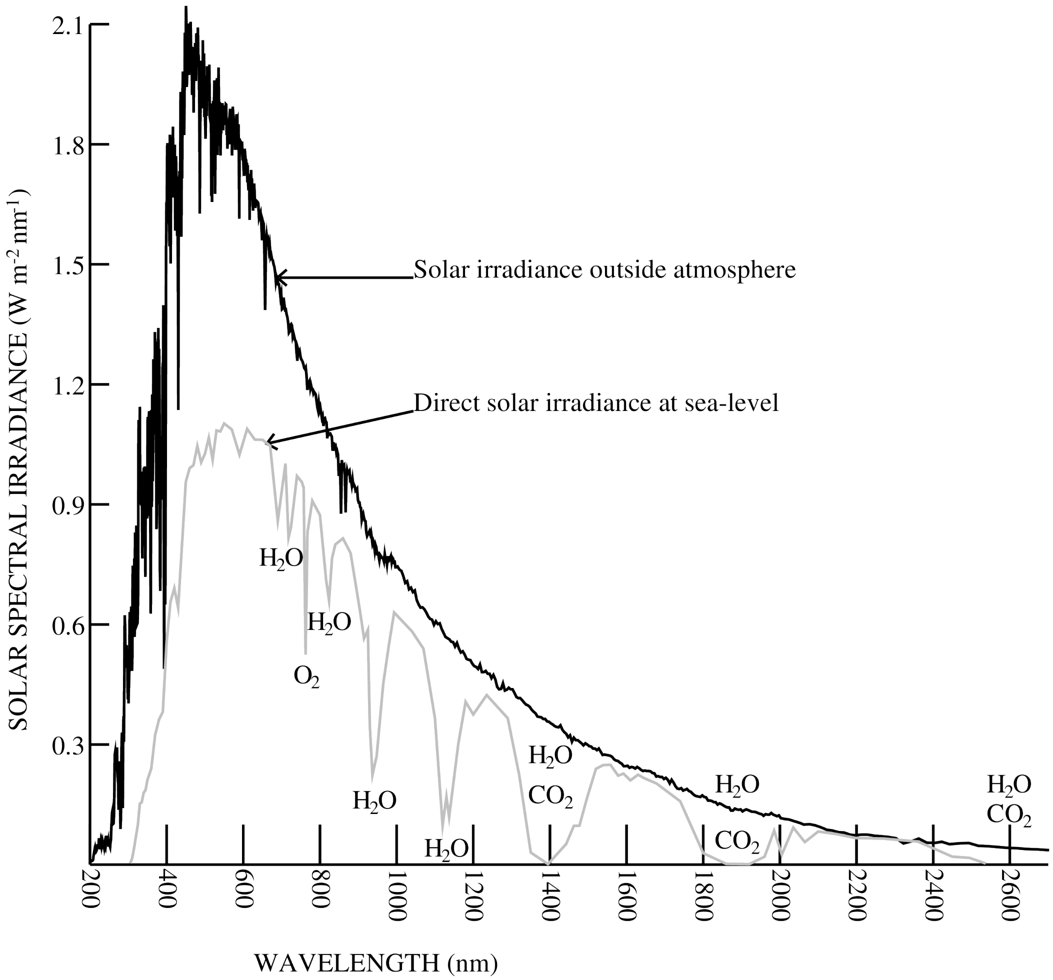

Figure 7.6: Solar spectral irradiance incident on the top of the atmosphere (black curve) and transmitted through the atmosphere to Earth's surface (gray curve), which principal gases responsible for atmospheric absorption noted. Data obtained from the National Renewable Energy Laboratory (rredc.nrel.gov/solar).

Yes. The spectral irradiance in equation 7.4 corresponds to the spectral irradiance observed at the object's surface. For example, if you use the surface temperature of the Sun, equation 7.4 will give you the solar spectral irradiance at the Sun's surface (see Project 7.1).

Since the irradiance decreases as the radiation propagates away from the Sun and spreads out in space, the solar spectral irradiance received at Earth (see Figure 7.6) is much smaller than that calculated by equation 7.4.[vii]

It is also important to note that the solar spectral irradiance calculated by equation 7.4 is smooth (see Project 7.2), whereas the observed solar spectral irradiance shown in Figure 7.6.

The difference has to do with the fact that the Sun is not a perfect blackbody. Still, the total irradiance, which you can obtain by integrating the spectral irradiance (as in Project 7.4), 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?

No. The atmosphere is much colder than the Sun and, as such, does not emit radiation over frequencies as high as the Sun. The atmosphere has temperatures that are more similar to that of Earth's surface, which does not emit any radiation in the visible or ultraviolet. That means we can treat the Sun as the only source of visible and ultraviolet radiation.

Which emits more infrared – Earth or Sun?

Since the Sun is so much hotter than Earth, the Sun emits more of all types of radiation. That means the Sun emits more infrared radiation at its surface than Earth emits at its surface.

However, all of Earth's infrared radiation is intercepted by the atmosphere whereas the atmosphere only intercepts a tiny fraction of the Sun's infrared radiation. Consequently, except for the near infrared (less than four microns wavelength or so[viii]), the infrared radiation we receive from the Sun is *less* than the infrared radiation emitted by Earth.

[vii] Remember that to properly compare the two the units must match. You will need to convert from W m^{-2} μm^{-1} to W m^{-2} nm^{-1}.

[viii] This actual overlap will depend on where we are. The solar radiation is greater outside Earth's atmosphere and when the Sun is directly overhead. However, assuming the Sun is high in the sky, the overlap will still be around four microns wavelength. Radiation less than four microns includes the IR-A and IR-B ranges, whereas the IR-C range extends to 1000 microns.

In other words, wavelengths greater than four microns are dominated by sources on Earth (and atmosphere) whereas wavelengths less than four microns are dominated by sources on the Sun. For this reason, it is common to refer to solar radiation as **shortwave** and atmospheric radiation as **longwave**.

Checkpoint 7.11: Which would be the larger number, the irradiance in W/m^2 per nm or the irradiance in W/m^2 per μm?

Project

Project 7.1: (a) 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). (b) In a second column, calculate the solar spectral irradiance (at the Sun's surface) for each wavelength (first column) using equation 7.4 (see page 120 for the meaning and values of λ, T, c_1 and c_2):

$$F_\lambda = \frac{c_1}{\lambda^5 \left[\exp\left(\frac{c_2}{\lambda T}\right) - 1\right]}$$

(c) For which wavelength is the emission for the Sun a maximum? In which part of the electromagnetic spectrum does it lie: IR, visible or UV?

Project 7.2: 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.[ix] *Knowing this, create another column listing the solar spectral irradiance just outside Earth's atmosphere.*

Project 7.3: (a) In another column, calculate the terrestrial spectral irradiance at <u>*Earth*</u>*'s surface using Earth's surface temperature (use the U.S. Standard Atmosphere surface temperature) instead of the Sun's surface temperature.*

[ix] 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 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 $(4\pi R^2_{\text{Earth's orbit}})$.

(b) For which wavelength is the emission for Earth a maximum? In which part of the electromagnetic spectrum does it lie: IR, visible or UV?
(c) Are there any wavelengths where the terrestrial spectral irradiance (at Earth's surface; third column) is greater than the solar spectral irradiance at the Sun's surface (first column)? If so, which ones?
(d) Are there any wavelengths where the terrestrial spectral irradiance is greater than the solar spectral irradiance just outside Earth's atmosphere (second column)? If so, which ones?

Project 7.4: Add up all of the spectral irradiances just outside Earth's atmosphere (created in Project 7.2) and multiply the result by the wavelength increment you used (which should be 0.01 μm). Show that the total is similar to the solar constant (see appendix B.4).

Questions

Question 7.1: A typical Helium-Neon laser produces light with a wavelength of about 632 nm. What color is this light?

Question 7.2: Does the radar map of Figure 7.3 indicate where the clouds are? If so, where are the clouds on the map? If not, why not?

Question 7.3: The infrared satellite picture in Figure 7.4 has the Gulf of St. Lawrence being brighter than the Gulf of Mexico. Why is that, given that both are at sea-level?

Question 7.4: According to Planck's law, a hot object (like the Sun) will emit more radiation than a cool object (like the atmosphere) at any frequency. For the atmosphere, however, the Sun is farther away than Earth. Given this, which source of radiation, Earth or sun, provides more infrared radiation to Earth's atmosphere, and why?

Question 7.5: Recall that the solar constant refers to the solar irradiance outside the atmosphere. However, if you sum up the values in Project 7.3 (terrestrial spectral irradiances), you get a value that is smaller than the solar constant, and that is the case if you used values of c_1 and c_2 that weren't rounded. Why is the sum smaller than the solar constant?

Problems

Problem 7.1: (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 7.2: In Problem 7.1, you calculated the solar irradiance at a point on the Sun's surface. Use equation 7.2 to calculate the temperature of the Sun. Compare your answer to that provided by a reference.[x]

Problem 7.3: (a) Calculate the peak emission wavelength for the Sun assuming a temperature equal to that calculated in Problem 7.2. 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 7.4: 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* μ*m to 300* μ*m).*

---

[x] See, for example, the NASA Fact Sheet on the sun.

8. Radiation Balance

Introduction

In the previous chapter, we examined how objects emit radiation, with the radiation amount and frequency depending on the object's temperature and size. We focused on two objects: the Sun and Earth. The Sun, being much hotter and bigger, emits much more radiation than Earth, and the radiation it emits peaks at a higher frequency and lower wavelength. In terms of wavelength, the Sun's radiation peaks around 500 nm (visible) whereas Earth's radiation peaks around 10 μm (infrared).

In his chapter, we'll take these ideas and explore how radiation causes objects to warm or cool. The basic idea is pretty simple, namely that objects warm up when they absorb more radiation than they emit and, conversely, they cool down when they emit more than they absorb.

8.1 Diurnal Variation

Everyone is probably familiar with the temperature being warmer during the day than at night. For example, consider Figure 8.1.

The left graph shows the actual observed temperatures in Chicago on March 19, 2019. Notice how the temperature reaches a minimum between 6 am and 7 am and then warms up. The temperature continues to warm up during the day, reaching a peak between 4 pm and 5 pm after which it cools down. This rise and fall of temperatures is consistent with the Sun rising in the morning and setting in the evening.

However, how do we know this is due to the Sun and not due to some other weather phenomenon? As a check, the right graph is the average temperature for every hour on March 19 over thirty years (1981-2010). This removes any influence of weather phenomena and only reveals the diurnal variation due

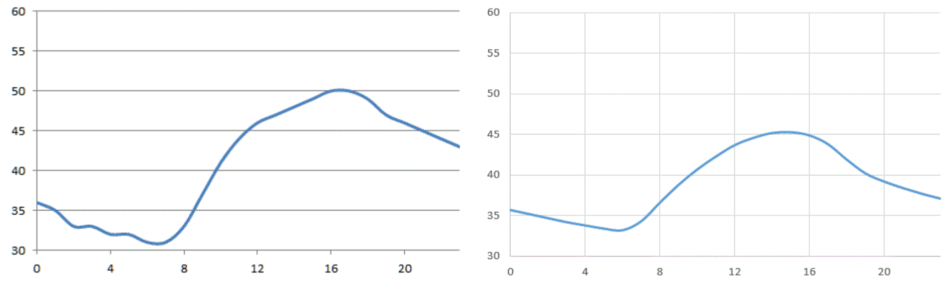

Figure 8.1: Hourly (CDT) temperatures (°F) for Chicago, IL. Left: observed temperatures on March 19, 2019 (NOAA). Right: 30-year average hourly temperatures for March 19 (NOAA).

to the Sun since its impact is the only consistent factor present in all thirty years.

Notice that we still have the coolest temperature in the morning (6 am) but the maximum temperature occurs at 3 pm, not 5 pm. Why is the peak at 3 pm, while the Sun is still out? The problem is even greater because this is Daylight Savings Time, so the local time is only 2 pm at maximum temperature.

To answer this, you have to realize that the temperature change depends on *two* things: the incoming radiation received from the Sun *and* the outgoing radiation emitted by Earth. If the incoming is greater than the outgoing then the temperature goes up and has not yet reached its maximum. If the reverse is true, and the outgoing is greater than the incoming, then the temperature goes down and has not yet reached its minimum. The maximum and minimum temperature is only reached when the incoming is *equal* to the outgoing.

At night, we have no incoming solar radiation, so it is easy to see that the outgoing radiation is greater and so the temperature goes down during the night.[i] However, during the day, the incoming radiation varies, reaching a peak at noon. As we know from chapter 7, the irradiance (the rate the radiation is emitted) depends on the temperature. The greater the temperature, the greater the irradiance. Consequently, as the ground warms up, it emits more radiation.

---

[i] Keep in mind that the atmosphere emits radiation too, so at night there is no danger of it continuing to cool down to absolute zero.

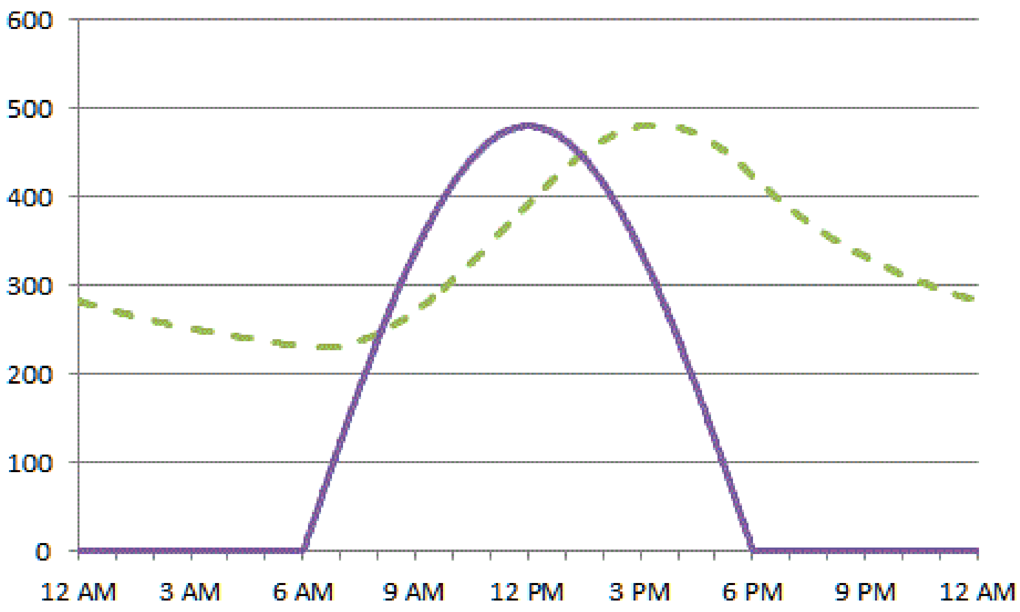

Figure 8.2: Typical hourly irradiances received (solid purple) and emitted (dashed green).

This is illustrated in Figure 8.2, which was obtained via a process described in Appendix C. The solid purple line represents the incoming solar radiation, which is zero at night (between 6 pm and 6 am) when the Sun is down and reaches a maximum at noon when the Sun is highest in the sky. The dashed green line represents the outgoing terrestrial radiation. It is never zero because the ground always has a temperature above absolute zero. However, it decreases during the night as the temperature cools.

Notice that the dashed green line (outgoing) is above the solid purple line (incoming) from about 2 pm to about 7:30 am. The temperature decreases during that time because the outgoing irradiance is greater than the incoming. On the other hand, between 7:30 am and 2 pm, the solid purple line (incoming) is above the dashed green line (outgoing) and the temperature increases.

If the temperature was constant, the crossover would be symmetric before and after noon. The crossover in the morning (from more outgoing to more incoming) and the crossover in the evening (from more incoming to more outgoing) would both occur the same time from noon, like four and a half hours before noon (7:30 am) and four and half hours after noon (5:30 pm). However, since the ground warms up during the day, it takes less time from maximum incoming (noon) to more outgoing (two hours; 2 pm) than it takes to reach maximum incoming (noon) from the crossover to more incoming (four and a half hours; 7:30 am).

Checkpoint 8.1: According to Figure 8.2, when should the minimum temperature at that location occur?

8.2 Equilibrium

in previous section we saw how we can use radiation to explain why it gets warmer during the day and cools down at night, with a maximum temperature in early afternoon and minimum temperature around sunrise.

We can use the same process to predict the average temperature globally. To do so, we need to identify how much solar radiation is absorbed by Earth then assume that is balanced by the radiation emitted by Earth.

WHY WOULD THEY BE BALANCED?

Because if Earth absorbed more than it emitted, it would warm up and, as the temperature increases, the emitted radiation would likewise increase (see section 7.3) until the emitted radiation equaled the absorbed radiation. Conversely, if Earth absorbed less, it would cool down and, as the temperature decreases, the emitted radiation would likewise decrease until, again, the emitted radiation equaled the absorbed radiation.

This radiative balance is called **equilibrium**.

So the first step is to identify how much solar radiation Earth absorbs. We already know from section 7.4 that the solar irradiance at Earth's position is equal to 1361 $\mathrm{W/m^2}$ (a value known as the solar constant).

Recall that the 1361 $\mathrm{W/m^2}$ number is the *concentration* of the power if you are facing the Sun. In other words, if you take a board of area 1 $\mathrm{m^2}$ and turn it to face the Sun, 1361 J of energy would hit the board every second. The bigger the board, the greater the energy it intercepts. As illustrated in Figure 8.3, Earth is like a board with an area equal to its cross-sectional

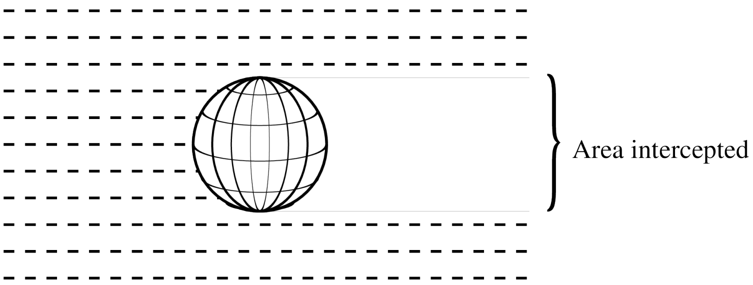

Figure 8.3: The portion of the incoming solar radiation that is intercepted by a sphere is equal to the cross-sectional area of the sphere (area of a circle).

area. Consequently, Earth intercepts a total power of 1361 W/m^2 times its cross-sectional area (πR^2_{earth}).

When we multiply 1361 W/m^2 by Earth's cross-sectional area (1.27×10^{14} m^2), we get that Earth intercepts the solar radiation at a rate of 1.73×10^{17} W.

Is this how much Earth absorbs?

No. There is a difference between how much it *intercepts* and how much it *absorbs*. Not all of the energy intercepted is absorbed. Some of it is reflected back out to space without being absorbed. Indeed, the fact that we can see Earth from space tells us that at least some of the light from the Sun must reflect off Earth.

On average, about 30.6% (0.306) of the incident energy is reflected (NASA). This fraction is known as the **albedo**, and will be indicated by the letter a.

Do we need to multiply the 1.74×10^{17} W by 0.306?

No, as that would be the fraction that is reflected. Instead we must multiply the intercepted energy by 0.694, which is $(1 - a)$.

After multiplying by 0.694, I get a absorption rate of 1.20×10^{17} W. In radiative equilibrium, Earth must likewise emit at a rate of 1.20×10^{17} W (i.e., Earth's luminosity).

Checkpoint 8.2: What fraction of the energy incident upon the atmosphere is actually absorbed by the atmosphere and Earth on average?

From section 7.3, we know that an object's luminosity depends on its size and temperature. Since we know what Earth's luminosity must be (to balance the absorbed solar radiation), we can use equation 7.2 ($P = A\sigma T^4$) and work backward to determine its temperature.

$$1.20 \times 10^{17}\ \mathrm{W} = (5.10 \times 10^{14}\ \mathrm{m}^2)(5.67 \times 10^{-8}\ \mathrm{W \cdot m^{-2} \cdot K^{-4}})T^4$$

This can be solved for T to get a temperature of 254 K.

Why are the two area values different?

One corresponds to the cross-sectional area (1.27×10^{14} m^2) and the other corresponds to the surface area (5.10×10^{14} m^2). When determining the solar radiation intercepted by Earth, we use the cross-sectional area (πR^2). When determining the energy emitted by Earth, use its surface area ($4\pi R^2$). For a sphere, the surface area is four times bigger than the cross-sectional area.

Checkpoint 8.3: (a) As seen from the Sun, what is the cross-section of Earth? (b) How does that compare to the total surface area of Earth?

Why is the temperature cooler than the average surface temperature of 288 K?

The radiative equilibrium temperature doesn't correspond to the *surface* temperature. As discussed in section 7.2, when reviewing infrared satellite pictures, the radiation emitted by the Earth/atmosphere system is not from Earth's surface. Most of it is emitted by clouds and other parts of the atmosphere. The temperature of 254 K, then, represents that average temperature of those parts of the Earth/atmosphere system.

The radiative equilibrium temperature 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) there is no other radiation emitted by the object. In our case, the blackbody temperature 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 object absorbs all visible light incident upon it. Such objects are called **blackbody** objects.

Checkpoint 8.4: Why is Earth's blackbody temperature less than its surface temperature?

Can we do the same with the Sun?

Yes (see problem 7.2). However, as we found with 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 the effective temperature of the Sun is 5772 K (see problem 7.2). 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 receive from the Sun is emitted from the Sun's surface.

> ✍ The inner sun is much hotter than the surface but its radiation is absorbed before reaching the surface.

Checkpoint 8.5: Which is warmer: Earth's blackbody temperature or its surface temperature?

8.3 Energy budget

Can we use radiation balance to predict the temperature at Earth's surface?

Yes, we can, but the problem is more complicated than just figuring out the incoming solar radiation and the outgoing emitted radiation from Earth. We also need to figure out what other radiation sources are around. All objects emit radiation. So the ground receives radiation not just from the Sun but also from the atmosphere.

> ✍ The incoming solar irradiance is sometimes called **insolation**.[ii] Thus, we could say the insolation outside of Earth's atmosphere is 1361 W/m^2.

We already mentioned that the Earth/atmosphere system absorbs 1.21×10^9 W, which is the solar constant (1361 W/m^2) times Earth's cross-section minus the fraction reflected (albedo). Not all of that gets absorbed by the ground. Some is absorbed by the atmosphere. In fact, the atmosphere absorbs about 33% of the incoming solar radiation.

Due to the absorption by the atmosphere, less solar radiation is incident upon the ground. However, the atmosphere, also emits radiation, and that radiation is also incident upon the ground.

Things get even more complicated. Most of the radiation emitted by the ground is absorbed by the atmosphere, so there is a feedback loop between the atmosphere and the ground (explored further in the next section).

In addition, we've only been considering radiation as an energy transfer mechanism. Two additional mechanisms that transfer energy from the ground are **convection** (i.e., the replacement of warm air right next to the ground with cooler air from above) and **latent cooling** (i.e., the cooling associated with the evaporation of water).

The various transfer mechanisms are all part of the global energy budget. An illustration for the mechanisms is provided in Figure 8.4.

What do the values in Figure 8.4 correspond to?

Those are irradiances, in units of W/m^2.

Why is the incoming solar irradiance equal to 340 W/m^2, not 1361 W/m^2?

The value in the figure is the globally averaged value. To get that value, we need to consider that "spread" what Earth over the entire surface of the globe. So, as discussed in section 8.2, Earth intercepts an amount equal to the solar constant (1361 W/m^2) times Earth's cross-sectional area. To find

---

[ii] 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.

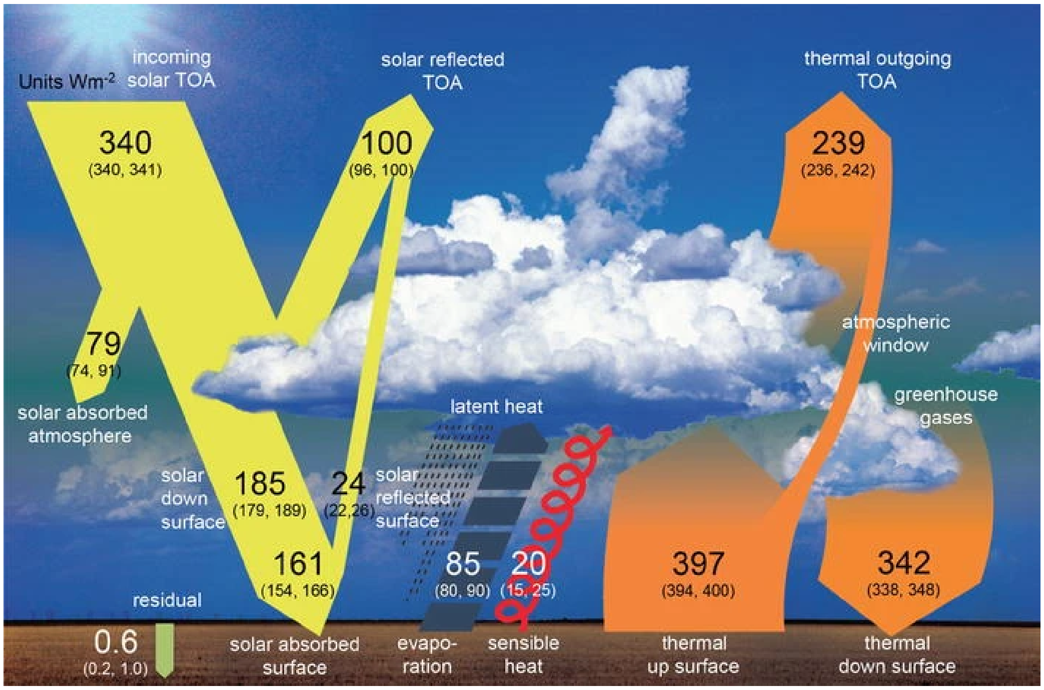

Figure 8.4: Earth energy budget in W/m^2 (from Wild et al 2013). Numbers in parentheses represent the uncertainty ranges.

the globally averaged irradiance, we then divide that value by Earth's surface area.

Since Earth's surface area is four times that of its cross-section, the result is equivalent to just dividing the solar constant by four. That gives the 340 value.

Notice how there is some uncertainty in the numbers. The values in Figure 8.4 are the authors' best guess as of 2000. Estimates can change. For example, according to the figure, the albedo is slightly under 0.3 (divide the solar reflected irradiance by the incoming solar irradiance) whereas NASA gives an albedo slightly over 0.3.

However, the general idea is consistent with what we've discussed so far. Notice how the atmosphere absorbs some of the incoming solar radiation as well as most of the outgoing radiation emitted by Earth. Notice also how the atmosphere emits radiation both downward to Earth and upward to space.

A careful tallying of the incoming and outgoing radiation reveals that there is slightly more energy absorbed by Earth than emitted. The difference, indicated by the downward-pointed green arrow with the label "residual", is consistent with global warming.

Checkpoint 8.6: How would a warmer atmosphere impact the amount of radiation emitted by the atmosphere and, in turn, the Earth's temperature?

8.4 Atmospheric windows

For an object to be in equilibrium (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 Earth/atmosphere system has 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 ground, and the atmosphere is colder than the ground.

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 (emitted by Earth) than solar radiation (emitted by the Sun).

As mentioned in section 7.5, we can consider the Sun the source of shorter wavelengths (ultraviolet, visible and near infrared) whereas Earth is the source of longer wavelengths (mid and far infrared). As you are probably aware, the atmosphere is mostly **transparent** to solar radiation, which is we can see through it (solar radiation being mostly visible wavelengths of light). What you may not be familiar with is that the atmosphere is mostly **opaque** to the infrared wavelengths that make up terrestrial radiation. When I say "opaque", I mean that the radiation is unable to propagate through the entire depth of the atmosphere without getting absorbed.[iii]

In other words, how much the atmosphere absorbs depends upon the frequency of the radiation.

[iii] For short distances, like a couple hundred meters or so, air is still relatively transparent to most wavelengths.

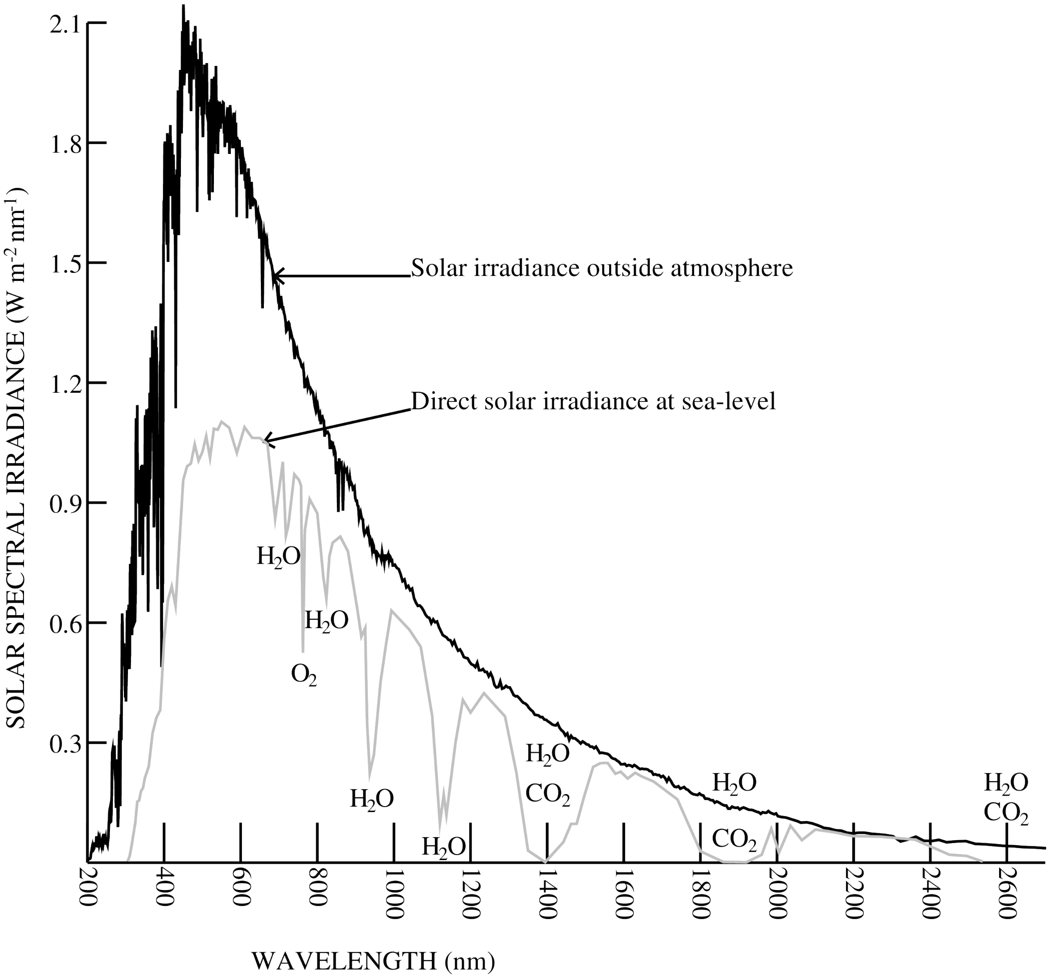

Figure 8.5: Solar spectral irradiance incident on the top of the atmosphere (black curve) and transmitted through the atmosphere to Earth's surface (gray curve), which principal gases responsible for atmospheric absorption noted. Repeated from Figure 7.6.

For example, consider Figure 7.6, which is repeated here as Figure 8.5.

As mentioned previously, the dark black line in Figure 8.5 represents the solar spectral irradiance received just outside Earth's atmosphere. Let's now look at the light gray line, which indicates the solar irradiance received at the surface of Earth (after passing through the atmosphere).

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 Earth's surface.[iv] This means that only 58% or so of the solar radiation reaches Earth's surface (divide 1.06 by 1.85).

Since about 30% of the solar radiation is reflected (due to Earth's albedo) and since some of the albedo is due to reflection off Earth's surface, you might expect that at least 70% of the solar radiation should reach Earth's surface, not 58%.

The difference is due to absorption by the atmosphere. As discussed in section 8.2, the atmosphere absorbs about 33% of the solar radiation that is not reflected. That is why the radiation received at 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 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 type of 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 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

---

[iv] Based on the average spectral irradiance between 500 and 600 nm.

[v] It is worth noting that this implies that most of Earth's albedo is due to atmospheric reflection, not reflection off the surface of Earth (see Figure 8.4).

[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). Note that ozone reacts with organic matter, which we are, so it is considered a pollutant when present at the bottom of Earth's atmosphere, where we reside.

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]

Between these locations, there are regions where almost all of the solar radiation gets through. Such regions are called **atmospheric windows**.[viii]

Checkpoint 8.7: What would be different if the atmospheric window was mainly in the ultraviolet instead of the visible light range?

8.5 The greenhouse effect

8.5.1 Qualitative treatment

As one can see by figure 8.5, 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 7.3, terrestrial radiation is entirely infrared and longer wavelengths. And, though it is not shown in figure 8.5, 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

---

[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 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 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 7.2.

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 8.6.

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 the Sun. 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 and atmosphere system is the same. It isn't as complicated as what was discussed in section 8.2 but it is more complicated than the zero-layer model because it treats the atmosphere as a separate source and absorber of radiation. This leads to several important differences with the zero-layer model.

First, let's look at the incident radiation from the Sun. In the single-layer model, part of that incident radiation is reflected off the atmosphere and part is reflected off Earth.[xi] Similarly, some of the incident radiation is absorbed by the atmosphere and some is absorbed by Earth.

Comparing the two models, 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 ground not only gets radiation from the Sun but also gets radiation from the atmosphere.

HOW DOES THE TOTAL RADIATION RECEIVED FROM BOTH THE SUN AND THE ATMOSPHERE COMPARE TO WHAT THE GROUND WOULD RECEIVE IF THERE WAS NO ATMOSPHERE?

How much the ground gets depends on how much the atmosphere absorbs of the terrestrial radiation vs. how much it absorbs of the solar radiation. The more terrestrial radiation it absorbs, the warmer it will get and, subsequently,

---

[xi] There are additional cascading reflections off the atmosphere and Earth again but these aren't shown since they are considered to be small.

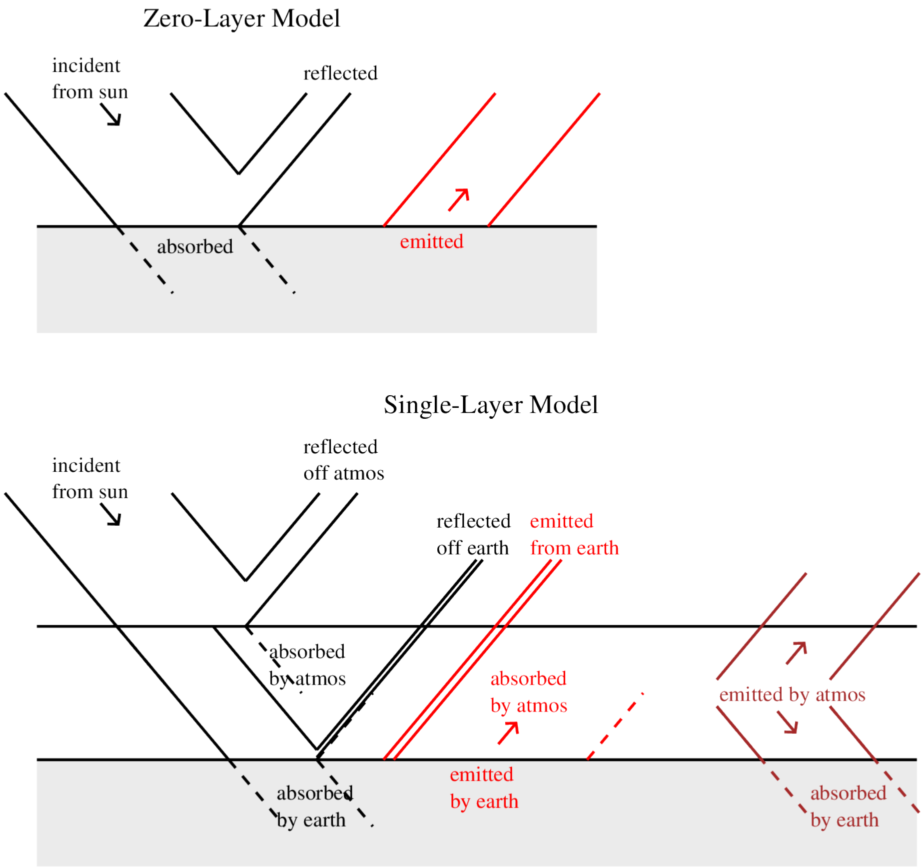

Figure 8.6: (top) Earth's radiation budget [top] considering the Earth/atmosphere as a single entity (with a single temperature) and [bottom] considering the atmosphere as being a single layer (with a single temperature) separate from Earth. In both cases, the incident radiation (from the Sun) is represented 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 more radiation it will emit back to the ground. On the other hand, the more solar radiation it absorbs, the less solar radiation that reaches the surface.

In our case, the atmosphere absorbs a much larger portion of the terrestrial radiation than the solar radiation and, as such, the ground 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.

> ✍ While the atmosphere, like a greenhouse, keeps things warmer than it would otherwise be, the mechanism is different. A greenhouse keeps the interior warmer due to a decrease in wind advection (much like your sweater) more so than an absorption of outgoing infrared radiation.

Checkpoint 8.8: According to figure 8.6, which provides more radiation to Earth's surface: solar radiation that passes through the atmosphere or radiation emitted by the atmosphere?

8.5.2 Quantitative treatment

We can get an estimate of what Earth's surface temperature should be, based on the balance of radiation discussed so far. It will be a rather crude model, but it will be sufficient to provide an estimate and also give you a sense of the physics responsible for the surface temperature being warmer than the blackbody temperature.

We already know that the solar irradiance averaged over the globe is equal to 340 W/m^2, or one-quarter of the solar constant value (see page 135). Since we will be using this value so often, I am going to write this as S_{avg} from now on.

For the irradiance emitted by the atmosphere, I will use F_{atmos}. For the irradiance emitted by the ground, I will use F_{sfc}.

In this single-layer model, there are three separate entities: the ground, the atmosphere, and space. For each, the equilibrium condition can be applied (assuming balance).

Energy balance - ground

We'll start by applying energy balance to the ground, meaning that the surface is emitting as much as it receives from the atmosphere and sun. The absorbed radiation from the Sun is <u>not</u> S_{avg}, as some of the incoming solar radiation is reflected and some is absorbed by the atmosphere prior to being incident upon the ground.

Since a represents the albedo, we'll say that aS_{avg} is reflected prior to passing through the atmosphere[xii] That means that $(1-a)S_{\text{avg}}$ passes through the atmosphere.

Of the radiation that isn't reflected, some is absorbed by the atmosphere. This is where things get interesting. The atmosphere absorbs a fraction of the solar radiation that passes through it, and it also absorbs a fraction of the terrestrial radiation that passes through it. The fraction absorbed is caled the **absorptivity**, and is indicated by α.

> ✍ Do not confuse the absorptivity α with the albedo a. The absorptivity is represented by the Greek letter "alpha" and represents the fraction of the radiation that is absorbed. The albedo is represented by the Roman letter "a" and represents the fraction of radiation that is reflected.

As noted earlier, the atmosphere is a better absorber of infrared (terrestrial) radiation than visible (solar), so there will be two different absorptivities, one for shorter wavelengths (solar) and one for longer wavelengths (terrestrial). We'll indicate these two absorptivities as α_{SW} and α_{LW}.

Due to atmospheric absorption of the solar radiation, then, the amount actually absorbed by the ground is $(1-\alpha_{\text{SW}})(1-a)S_{\text{avg}}$.

The ground also absorbs radiation from the atmosphere, which we are indicating as F_{atmos}. We'll assume that that entire amount is absorbed by the ground (i.e., no reflection).

Thus, the total amount of irradiance absorbed by the ground is $(1-\alpha_{\text{SW}})(1-a)S_{\text{avg}} + F_{\text{atmos}}$.

For the ground to be in radiative equilibrium, this must be balanced by the amount the ground emits, which we are indicating as F_{sfc}.

[xii] This, of course, is not entirely true. Certainly, some is reflected from the ground, as well as some reflected within the atmosphere. However, the general approach is still sound.

Setting the input equal to the output, we have the following.

$$(1 - \alpha_{\mathrm{SW}})(1 - a)S_{\mathrm{avg}} + F_{\mathrm{atmos}} = F_{\mathrm{sfc}} \tag{8.1}$$

This equation has two unknowns, F_{atmos} and F_{sfc}, assuming we know the incoming solar radiation, the albedo and the atmospheric absorptivity of solar radiation. We are particularly interested in F_{sfc} because that will allow us to determine Earth's surface temperature (via the Stefan-Boltzmann equation; see page 115). However, to get it we need to know F_{atmos} and we don't know that either.

Energy balance - Earth and atmosphere

To solve for the two unknowns in equation 8.1, we need another equation. To get the other equation, we can examine the radiation balance for the atmosphere or the entire system (Earth and atmosphere). It doesn't matter which one we look at (atmosphere or Earth/atmosphere system) – either one will give us another expression with the two unknowns.

Since it is mathematically a little simpler to look at the Earth/atmosphere system, we'll do that. For the system, the input is simply the insolation S_{avg}.

Leaving the system, we have the portion of the solar radiation that is reflected (due to the albedo), which is equal to aS_{avg}. We also have a large amount leaving due to emission from the atmosphere, which we are indicating as F_{atmos}. In addition, we have a small amount of radiation emitted from the ground that has managed to get through the atmosphere. Due to the atmospheric absorptivity of longwave radiation (α_{LW}), the amount that passes through the atmosphere is equal to $(1 - \alpha_{\mathrm{LW}})F_{\mathrm{Earth}}$.

Setting the input equal to the output (for the Earth/atmosphere system), we have the following.

$$S_{\mathrm{avg}} = aS_{\mathrm{avg}} + F_{\mathrm{atmos}} + (1 - \alpha_{\mathrm{LW}})F_{\mathrm{sfc}} \tag{8.2}$$

Surface temperature

Equation 8.1 and 8.2 can be combined to solve for the two unknowns, F_{atmos} and F_{sfc}. We are really just interested in F_{sfc}, because we can then determine the surface temperature from that.

Combining the two expressions, we get the followinng expression:

$$F_{\text{sfc}} = S_{\text{avg}} \frac{(2-\alpha_{\text{SW}})(1-a)}{(2-\alpha_{\text{LW}})} \tag{8.3}$$

The derivation is left as an exercise for the reader (see Show-me 8.1). To get the surface irradiance, then, we need values for S_{avg}, a, α_{SW} and α_{LW}. For our purposes, we'll use the values implied by the energy budget in Figure 8.4.

- The incoming solar irradiance is listed as 340 W/m^2.
- The albedo is the fraction reflected (100 W/m^2 in the figure) divided by the incoming (340 W/m^2). This gives an albedo of 0.294 (slightly under the NASA value of 0.306).
- The absorptivity of shortwave radiation is the fraction of the non-reflected irradiance (240 W/m^2) that is absorbed by the atmosphere (79 W/m^2 in the figure). This gives an α_{SW} equal to 0.329.
- The absorptivity of longwave radiation is the fraction of the terrestrial irradiance (397 W/m^2 in the figure) that is absorbed by the atmosphere. The figure doesn't show us that, but a slightly older diagram has values consistent with α_{LW} being 0.900.

Plugging in these values for S_{avg}, a, α_{SW} and α_{LW}, we get a surface irradiance equal to 364.64 W/m^2. This must equal σT^4, where T is the radiative temperature, and solving for T we get 283.2 K, which is significantly warmer than the effective temperature discussed before (254 K; see section 8.2).

As you can see in equation 8.3, if $\alpha_{\text{SW}} = \alpha_{\text{LW}}$ then $F_{\text{sfc}} = (1-a)S_{\text{avg}}$, which is equivalent to the zero-layer model and gives the 254 K temperature value. It is only because there is a *different* absorptivity of shortwave vs. longwave that leads to a surface temperature that is warmer than the blackbody temperature.

Checkpoint 8.9: Suppose the atmospheric window was in the infrared rather than visible, meaning that the atmosphere was a better absorber of infrared than visible. Would you expect the surface temperature to be warmer than 254 K, colder than 254 K, or equal to about 254 K? Explain your reasoning.

8.5.3 Multilayer models

Why does this approach give a surface temperature that is cooler than the U.S. Standard Atmosphere surface temperature (288 K)?

A significant problem with the single-layer model is that it treats the atmosphere as, well, a single layer. As one can see in Figure 8.4, the atmosphere emits much more radiation downward to Earth's surface than upward to space, for the same reason the surface emits more than the atmosphere – warmer objects emit more radiation.

A better model would have many layers, so that we could have the temperature vary throughout the atmosphere. We'd then have an equation for each layer. We could still solve the multiple equations but it would be more complicated.

> ✍ When considering thinner layers of the atmosphere, one needs to keep in mind that the absorptivity will be less for thinner layers. In addition, being thinner, they emit less radiation. It turns out that each layer would emit a fraction of the σT^4, with the fraction being equal to the absorptivity. The idea that the **emissivity** equals the absorptivity is known as **Kirchhoff's law of thermal radiation**.

The end result is that there would be more radiation downwoard emitted downward by the "real" atmosphere, leading to a higher surface temperature.

We've also neglected mechanisms that help cool the surface. For example, the warm Earth can be cooled by blowing cooler wind across it. This process is called **advection** or **convection** (see chapter 5 for more information). We also know that evaporation (from the oceans, for example) is a cooling process (see chapter 16). These two processes will somewhat counteract the warming that would result from a multi-layer model.

Checkpoint 8.10: What mechanisms, other than radiation, act to cool Earth's surface?

8.6 Global warming

Although the single-layer model is not perfect, it does provide a reasonable estimate of the average surface temperature. In addition, it allows us to explore what would happen if change some of the parameters. In this way, we can explore which is the most significant factor in **global warming**.

8.6.1 Increase in longwave absorptivity

Since the greenhouse effect results when $\alpha_{\mathrm{SW}} < \alpha_{\mathrm{LW}}$, it stands to reason that increasing the atmosphere's ability to absorb infrared radiation will lead to warmer surface temperatures. Indeed, raising α_{LW} from 0.90 (estimated above) to 0.95, equation 8.3 gives a temperature of 286.5 K (vs. 283.2 K before).

What would cause α_{LW} to increase?

Since the absorption of longwave radiation is due mainly to water vapor and carbon dioxide, an increase in those gases is expected to increase α_{LW} and, consequently, an increase in the surface temperature.

Will increasing those gases necessarily lead to an increase in α_{LW}? Isn't it possible that they already absorb 100% of the radiation at the frequencies they affect?

Even if they already absorb everything in that frequency, increasing them will still lead to an increase in surface temperature. The reason is that they also emit infrared.

But how can α_{LW} be greater than 1?

It can't, but that is the limitation of our single-layer model. A multi-layer model would have lots of layers. The thinner the layer, the lower the absorptivity, and so we don't really have to worry about the upper-limit for α_{LW}.

Which gas has a greater effect on the absorptivity of terrestrial radiation?

Water vapor has a greater effect, since there is more water vapor and it absorbs a greater fraction of the radiation. However, there is a limit to how much water vapor can be present before it condenses whereas there is no such limit on carbon dioxide. It is for this reason, and the fact that the burning of fossil fuels results in the emission of carbon dioxide, that carbon dioxide is considered to be the greater contributor to enhanced global warming.

Checkpoint 8.11: Based on equation 8.3,

$$F_{\text{sfc}} = S_{\text{avg}} \frac{(2 - \alpha_{\text{SW}})(1 - a)}{(2 - \alpha_{\text{LW}})}$$

what happens to Earth's surface temperature if the absorptivity of terrestrial radiation increases?

8.6.2 Decrease in shortwave absorptivity

WILL AN INCREASE IN GREENHOUSE GASES INCREASE THE ABSORPTIVITY OF SOLAR RADIATION AS WELL AS TERRESTRIAL RADIATION?

Not really. Although there is some infrared radiation associated with solar radiation (see Figure 7.6), most of the solar spectrum is in the visible and ultraviolet, which are mainly unaffected by the presence of water vapor and carbon dioxide.

WHAT ABOUT IF WE INCREASE THE AMOUNT OF OZONE?

It is interesting to note that **ozone** is a good absorber of ultraviolet, which leads to warming in the stratosphere.[xiii] So, one might wonder if increasing stratospheric ozone will affect the absorptivity of solar radiation and, in turn, surface temperatures.

The impact is probably negligible, however. Since stratospheric ozone already absorbs all of the ultraviolet radiation incident upon it, there may

---

[xiii] Ozone is an interesting gas. It reacts with organic matter so it is considered a pollutant yet it absorbs ultraviolet which is nice (since it prevents it from being transmitted to the ground). You may have heard that there is a problem with **ozone depletion** even while *too much* ozone is considered part of the air pollution problem. This apparent paradox is due to these two characteristics. It turns out that most of the ultraviolet absorption occurs high in the atmosphere (where it is too far away to interact directly with us). Ozone gas is necessary there. Ozone near the surface is small in comparison and so does little to absorb the ultraviolet yet by being near us it affects our lives. Thus, the desire is to maintain low levels of ozone near the surface and high levels in the **stratosphere** and **mesosphere**.

already be more than enough ozone to absorb the ultraviolet and so increasing stratospheric ozone won't change the absorptivity.

On the other hand, there has been concern about *decreases* in stratospheric ozone. We know that there has been dramatic decreases in stratospheric, such as the annual decreases of 50% or more observed over Antarctica as it comes out of its winter. This leads to more solar radiation making it to the surface. It is thus a reasonable question as to whether that may lead to warmer surface temperatures. However, the increase is in the ultraviolet region, which can have a significant impact on us (sunburn and skin cancer) but the energy in the ultraviolet region is still a small portion of the overall solar irradiance, so the impact on the surface temperature is likely to be small.

Still, let's suppose we lower α_{SW} from 0.329 (estimated above) to 0.279 (a decrease of 0.05). It turns out that equation 8.3 gives a temperature of 285.3 K (vs. 283.2 K before), just two-thirds of the warming resulting from a 0.05 increase in α_{LW}.

Checkpoint 8.12: Based on equation 8.3,

$$F_{\mathrm{sfc}} = S_{\mathrm{avg}} \frac{(2-\alpha_{\mathrm{SW}})(1-a)}{(2-\alpha_{\mathrm{LW}})}$$

what happens to Earth's surface temperature if the absorptivity of solar radiation increases?

8.6.3 Decrease in albedo

Recent observations seem to show that Earth's albedo is decreasing (possibly because of loss of snow/ice cover).[xiv] For example, the NASA Earth fact sheet gives 0.306 as the albedo whereas Figure 8.4 suggests it is 0.294, a difference of 0.012. What happens if we lower the albedo by another 0.012?

[xiv] Some people argue that global warming will actually lead to a higher albedo, not a lower albedo, due to the associated increase in water vapor (due to evaporation) leading to more clouds (and thus cooler temperatures). This may not be entirely valid, however. With global warming, both the atmosphere and the ground warms. Thus, an increase in water vapor will not necessarily lead to an increase in clouds (condensation).

One can see from equation 8.3 that a decrease in a, and the corresponding increase in $(1 - a)$, leads to a higher F_{Earth}. With a decrease in the albedo, a greater amount of radiation enters the Earth/atmospheric system. This implies that the surface temperature will rise.

We already know that equation 8.3 gives a temperature of 283.2 K. With an albedo of 0.282, we get a temperature of 284.4 K, a difference of 1.2 K (compared to 3.3 K when α_{LW} increased by 0.05).

Checkpoint 8.13: Based on equation 8.3,

$$F_{\text{sfc}} = S_{\text{avg}} \frac{(2 - \alpha_{\text{SW}})(1 - a)}{(2 - \alpha_{\text{LW}})}$$

what happens to Earth's surface temperature if the albedo increases?

8.6.4 Increase in solar irradiance

All this time, we have assumed that the solar irradiance is constant at 1361 W/m^2 just outside the atmosphere. In reality, it varies a little bit from day to day, with about a 12-year cycle, as seen in Figure 7.5. The day-to-day variation is typically associated with changes in solar activity.

This variation from day to day is very small, however. Indeed, a change of 1 W/m^2 out of 1361 W/m^2 is a change of about 0.07%. With such a change in solar radiation, the equilibrium temperature will change by only 0.05 K or 0.02% (i.e., about one-quarter of the relative change in solar irradiance).

In addition, there does not seem to be a trend upward or downward, meaning that long-term trends do not imply that changes in solar irradiance are causing global temperature trends.

> By the way, one way to change the solar irradiance is to simply change the distance between the Earth and the Sun. Although we say that the Earth's orbit has a radius of 1.496×10^8 m, it actually varies from 1.4709×10^8 m to 1.5210×10^8 m. This is a change of 1.7%. With such a change in distance, the solar irradiance would change by 3.4% (double the change in distance, since the irradiance depends on the square of the distance). The equilibrium temperature will then change by only 2.4 K or 0.8% (about one-quarter the relative change in solar irradiance).

Checkpoint 8.14: Based on equation 8.3,

$$F_{\text{sfc}} = S_{\text{avg}} \frac{(2 - \alpha_{\text{SW}})(1 - a)}{(2 - \alpha_{\text{LW}})}$$

what happens to Earth's surface temperature if the insolation increases?

Questions

Question 8.1: (a) Should the temperature continue to rise while sunlight is present during the day since there is radiation being absorbed during that entire time?
(b) Why would you expect the minimum daily temperature to occur a little bit after sunrise, not right at sunrise or just before sunrise?

Question 8.2: The average insolation over the entire globe is equal to one-quarter of 1361 W/m^2*. Why one-quarter?*

Question 8.3: If we increase the amount of carbon dioxide in the atmosphere, how would that affect the single-layer model in Figure 8.6?

Problems

Problem 8.1: The solar constant for Venus, being closer than Earth to the Sun, is 2613.9 W/m^2 *but, being covered in clouds, Venus 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[xv] *(i.e., 737 K)?*

Problem 8.2: (a) Suppose the air had a temperature of 273 K. What is the irradiance emitted by the air out to space? (see Stefan-Boltzmann law, equation 7.2, but note that this question is asking for the irradiance, not the

---

[xv] See NASA's Fact Sheet on Venus.

luminosity)
(b) Why is your answer in (a) greater than the average absorbed insolation by the Earth and atmosphere system, $(1 - a)S_{\text{avg}}$*?*

Problem 8.3: For the following questions, predict whether the surface temperature will increase, decrease or remain the same, and explain why the change in parameter would make the temperature change in that way.
(a) What would happen to the surface temperature if the atmosphere absorbed a greater fraction of terrestrial radiation (i.e.,α_{LW} *was increased)?*
(b) What would happen to the surface temperature if α_{SW} *increased?*
(c) What would the surface temperature be if there was no atmosphere at all (i.e., $\alpha_{\text{SW}} = \alpha_{\text{LW}} = 0$*)?*
(d) What would happen to the surface temperature if the atmospheric absorptivity was independent of wavelength (i.e., $\alpha_{\text{SW}} = \alpha_{\text{LW}}$*)?*

Problem 8.4: For the following questions, use equation 8.3 along with the values described on page 145 to get the terrestrial irradiance and then use equation 7.2 to get the temperature and confirm that you get 283.2 K, as indicated on page 145. It might help to put them into an Excel spreadsheet so you can change one parameter at a time. Once you have confirmed you get the same result, determine the new surface temperature under the following conditions.
(a) α_{LW} *is increased from 0.9 to 1.0 (with everything else the same as before)?*
(b) a was decreased from 0.294 to 0.290 (with everything else the same as before)?
(c) S_{avg} *was increased by 0.1% (with everything else the same as before)?*
(d) Based on your results, which of the following would have the greatest impact on the surface temperature: variations in the amount of carbon dioxide, variations in the amount of snow cover, or variations in solar output? Explain.

Derivations

Show-me 8.1: Derive equation 8.3 from equations 8.1 and 8.2.

Part C

Newton's Second Law

9. Vertical Balance

Introduction

Way back on page 9, a simple 3-step forecast method was described. The three steps were: (1) predict the temperature, (2) predict the pressure based upon the new temperature values, and (3) predict the wind based upon the new pressure values.

To do the first step, we utilize the physics discussed in part B. In this chapter we examine the physics behind the second step. The physics behind the third step is discussed in the last two chapters of this part of the text (chapters 12 and 13). This three-step process is then repeated, over and over, for very short forecasts, until the desired time period is reached.

9.1 Hydrostatic balance

In chapter 4, we discussed how the pressure decreases roughly exponential with height. However, we didn't explain *why* it does this. We'll do that in section 9.2. In this section, we'll tackle a simpler problem: why the pressure decreases with height at all.

It has to do with **vertical balance**. Air is pushed from high pressure to low pressure. So, vertical balance means that the pressure *must* decrease with height to produce an upward force (from high to low pressure) that exactly balances the downward force due to the gravitational attraction with Earth.

When a fluid, of which air is an example, is in vertical balance, we call it **hydrostatic balance**.

> ✍ The word "hydrostatic" comes from "hydro-", which refers to water, and "-static", which refers to something being stationary. In this case, the fluid is air, not water, but the "hydro" prefix is used for any fluid.

DOES THE AIR HAVE TO BE STATIONARY TO BE IN HYDROSTATIC BALANCE?

No. 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.

IF THE AIR IS NOT ACCELERATING VERTICALLY, HOW CAN THERE BE VERTICAL MOTIONS?

Keep in mind that hydrostatic balance does not require that the air be stationary. 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.

WHY DOES THE AIR HAVE TO BE IN HYDROSTATIC BALANCE?

It doesn't. There are, of course, times when the air accelerates upward or downward, such as in the middle of a thunderstorm. However, 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, as we will see, 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.

In a sense, if the air is not in hydrostatic balance, the resulting distribution would push the air back toward hydrostatic balance. For example, gravity pulls the air downward. If there was no upward force to counter that, the air would pile up at the bottom. Such a "piling up" would cause a large vertical pressure gradient, which would push the air back upward.

It turns out, then, that in hydrostatic balance, not only must the pressure decrease with height but it must decrease in a *particular* way, namely in a way that the upward force (due to the vertical pressure gradient) exactly counters the gravitational force pulling down. It is left as an exercise (see Problem 9.2) to show that these two forces are observed to be comparable in magnitude in general.

Mathematically, we can express this equivalence as follows:

$$\frac{\partial P}{\partial z} = -\rho g \tag{9.1}$$

The left side, as discussed previously, is the vertical gradient in the pressure and represents the upward force (per volume) due to the pressure gradient.

The right side represents the gravitational force (per volume)[i], where the variable g represents Earth's gravitational field and is approximately 9.8 N/kg near Earth's surface.[ii] The right side is negative to match the left side (since pressure decreases with height, the vertical pressure gradient is negative).

> ✍ We often treat the value of g as being constant with height, particularly within the troposphere, because the height of the troposphere (10-15 km) is small compared to Earth's radius (6371 km). The density, however, typically decreases significantly with height, which means the pressure gradient also decreases with height.

Note that the left side, like the right side, is equivalent to a force *per volume*. However, be aware that both are often referred to as forces, not forces per volume. For example, it is common to refer to $\partial P/\partial z$ as the vertical **pressure gradient force**, even though technically it is a force **per volume**.

How is the pressure gradient equivalent to a force per volume?

To see how the pressure gradient is also a force per volume, consider a tiny cube of air with sides of infinitesimal lengths δx, δy (not shown) and δz, as shown in the illustration.

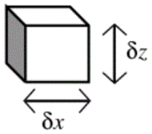

There are four vertically oriented faces (two of area $\delta x \times \delta z$ and two of area $\delta y \times \delta z$). There are two horizontally oriented faces (of area $\delta x \times \delta y$). When there is a vertical pressure gradient, the force per volume pushing up on the bottom ($P_{\text{bottom}}\delta x\delta y/\delta x\delta y\delta z$) 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).

Regardless, don't lose sight of the fact that when the atmosphere is in hydrostatic balance, the two forces balance, which necessarily means the pressure must decrease with height in a particular way such that the vertical pressure gradient force balances the gravitational force.

---

[i] The gravitational force on an object is equal to the product of the gravitational field g and the object's mass m. Since density is the mass per volume, we get the gravitational force per volume by multiplying the gravitational field g by the density ρ. In this case, we multiply g by the density of the air, since we want to know the gravitational force per volume on the air.

[ii] This value is actually relative to the rotating coordinate frame as Earth spins on its axis. It is within 0.03 N/kg of 9.8 N/kg as long as we are not higher than about 20 km above sea-level (which is about the maximum height of the troposphere).

> ✍ A common misconception is that the pressure drops because gravity decreases. As you can see from the expression, simply having gravity, whether it decreases with height, increases with height, or is constant with height, means that the pressure must decrease with height.

Checkpoint 9.1: If the atmosphere is in hydrostatic balance, does that mean that the pressure necessarily *has to decrease with height?*

9.2 Pressure decrease with height

Hydrostatic balance (equation 9.1) can be used to show that the pressure decreases more rapidly with height when the pressure is large (near the surface) and decreases less rapidly with height when the pressure is less (higher up). More specifically, hydrostatic balance implies that the pressure must decrease *exponentially* with height.

Taking the hydrostatic relationship (equation 9.1)

$$\frac{\partial P}{\partial z} = -\rho g$$

and replacing 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., $dP \propto -P$). Only the exponential function has this property, where the derivative of the exponential function is proportional to itself (i.e., $de^{-ax}/dx = -ae^{-ax}$). This implies, then, that the pressure decreases with height exponentially.

The actual relationship is as follows:

$$P_{\text{top}} = P_{\text{bottom}} \exp\left(-\frac{g}{RT_{\text{avg}}}\Delta z\right) \tag{9.2}$$

where Δz is the depth of the layer, with pressure P_{bottom} at the bottom and pressure P_{top} at the top, and T_{avg} is the average temperature of the layer (note that $\exp x$ is equivalent to e^x). A version of this is used in Project 9.1 to figure out how pressure decreases with height in the U.S. Standard Atmosphere.

How is this relationship obtained?

When the ideal gas law is used to replace P in equation 9.1 with ρRT, we get (after some rearrangement of terms):

$$\frac{dP}{P} = -\frac{g}{RT}dz$$

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 using the average temperature, T_{avg}, of the layer. In that case, we can consider a thicker layer, of depth Δz:

$$\Delta\ln P = -\frac{g}{RT_{\text{avg}}}\Delta z$$

We now set the top and bottom of the layer such that the top is at height z_{top} and the bottom is at height z_{bottom}. The equation then becomes

$$\ln P_{\text{top}} - \ln P_{\text{bottom}} = -\frac{g}{RT_{\text{avg}}}\Delta z$$

Since the difference in two logs is equal to the log of the ratio[iii], we have

$$\ln\left(\frac{P_{\text{top}}}{P_{\text{bottom}}}\right) = -\frac{g}{RT_{\text{avg}}}\Delta z \tag{9.3}$$

Take the exponential of both sides[iv] to get equation 9.2.

[iii] See footnote on page 100.

[iv] See the *other* footnote on page 100.

> ✍ Technically, the value of g decreases very slightly as we go up in the atmosphere and so an assumption of constant g would cause an error in the forecast, albeit very slight. To get around this, we can use the **geopotential** height as the vertical coordinate. The geopotential height is the actual height times g divided by the sea-level value of g.

Checkpoint 9.2: Why does equation 9.2 use T_{avg} 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 4.1:

$$P = P_0 e^{-z/H}$$

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**). That expression differs from the one we've derived here from hydrostatic balance in two ways. First, the hydrostatic balance expression is for any two points in the atmosphere. The bottom need not be at the surface. Second, the hydrostatic balance expression has RT_{avg}/g instead of H. What this means is that in a hydrostatic atmosphere, the scale height is proportional to the temperature as follows:

$$H = \frac{RT_{\mathrm{avg}}}{g} \tag{9.4}$$

Since the temperature varies with height, this means that the scale height varies with height also.

Checkpoint 9.3: According to equation 9.4, if the temperature of a layer increases, what happens to the scale height?

9.3 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.

9.3.1 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:

$$M = \frac{A}{g} P_0 \tag{9.5}$$

where A is the surface area, g is the gravitational field 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 Earth (recall that the surface area of a sphere is $4\pi R^2$).

Checkpoint 9.4: According to equation 9.5, how is the surface pressure related to the mass of the atmosphere?

HOW IS THIS RELATIONSHIP OBTAINED?

To obtain equation 9.5, we start with a mathematical representation of the total mass of the atmosphere:

$$M = \int dm$$

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.

Since mass is equal to the product of density and volume, we can replace dm by $\rho(Adz)$, where A is the area of the slice (so that Adz is like an infinitesimal volume):

$$M = \int \rho A dz$$

The problem with this expression is that it has ρ in it and ρ varies with height. To solve this problem, we use the hydrostatic relationship (equation 9.1) to replace ρ by $-(1/g)dP/dz$:

$$M = -\int \frac{A}{g} dP$$

The area of each slice A is the surface area of 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:

$$M = -\frac{A}{g}\int dP$$

The sum of all the dP's just gives ΔP, which gives

$$M = -\frac{A}{g}\Delta P$$

If we are considering the entire mass of the atmosphere, then ΔP is the difference in pressure between the bottom, where the pressure is P_0 and the top, where the pressure is zero. That means that $\Delta P = -P_0$ if we consider the entire atmosphere, leading to equation 9.5.

WHAT IF WE WANTED THE MASS ABOVE A PARTICULAR HEIGHT IN THE ATMOSPHERE?

As long as you know the pressure at that height, simply use the pressure at that height instead of the pressure at the surface. In other words, the atmospheric pressure that you measure at any altitude is the weight (Mg) per area (A) of all the air molecules above you: $P = Mg/A$.

WHAT IF WE WANTED THE MASS BELOW A PARTICULAR HEIGHT IN THE ATMOSPHERE?

Remember that the actual expression has a ΔP in it. As long as you know the pressure difference between two levels, you can use ΔP with the expression to get the mass in that layer.

Indeed, you can rearrange the expression to show that the pressure difference between any two levels in the atmosphere equals the weight per area of the air between the two levels:

$$\Delta P = \frac{Mg}{A}. \tag{9.6}$$

Checkpoint 9.5: According to equation 9.6, 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?

9.3.2 Warm advection and pressure falls

In chapter 5, 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 9.2, which shows how the pressure decreases with height.

$$P_{\text{top}} = P_{\text{bottom}} \exp\left(-\frac{g}{RT_{\text{avg}}}\Delta z\right)$$

Let's rewrite this in the following way. Let's make P_{top} equal to the pressure at the tropopause and P_{bottom} equal to the pressure at the surface. Solving for the surface pressure, we get:

$$P_0 = P_{\text{trop}} \exp\left(\frac{g}{RT_{\text{avg}}} z_{\text{trop}}\right) \tag{9.7}$$

Let's further assume that the height of the tropopause is fixed[v] 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 T_{avg} of the troposphere changes. If that mean temperature goes up then the surface pressure must go down (since T_{avg} is in the denominator).

Let's see what this means in terms of the polar front model that was discussed in section 1.2.3. It was mentioned that the center of storms tend to be areas

---

[v] This height isn't fixed. However, the stratosphere can be thought of as a "lid" of sorts (see chapter 11) and so the tropopause won't change significant over the time scale we are looking at.

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 T_{avg} east of the center and decreases it west of the center. From equation 9.7, the changing T_{avg} values result in the surface pressure decreasing east of the center and increasing west of the center, acting to move the system as a whole toward the east, consistent with what was discussed in section 1.2.3.

Checkpoint 9.6: If warm advection is occurring at a location, what is your short-term forecast for the surface pressure?

9.3.3 Thickness

In the hydrostatic relationship, Δz represents the depth of the layer for which the average temperature is T_{avg}. For this reason, Δz is known as the **thickness**[vi]. As you can see from equation 9.3, the height difference (or thickness) between two pressure levels depends upon the average temperature in the layer. The warmer the temperature, the larger the height difference (i.e., the thicker the layer).

The reverse is also true. The thicker the layer, the warmer the average temperature within the layer. In other words, the thickness can be used to get a sense of the average temperature.

In fact, it is common for meteorologists to plot the 1000-500 mb thickness (the difference in height between where the pressure is 1000 mb and where the pressure is 500 mb). The thicker this layer, the warmer the average temperature between 1000 mb and 500 mb.

Checkpoint 9.7: 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?

[vi] 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.

9.3.4 Sea-level pressure

We have seen that surface maps (see, for example, Figure 4.3 on page 57) show the sea-level pressure, not the actual pressure measured on the surface.[vii] Since they measure the surface pressure, not the sea-level pressure (which would require them to don't dig a hole down to sea-level and drop down a barometer), you might wonder how they figure out the sea-level pressure.

The answer is that they use the hydrostatic balance relationship. Basically, the process is the reverse of that described in section 9.3.3 for thickness. A "representative" temperature[viii] is used, as with thickness, but rather than using a fixed *pressure* difference and finding the corresponding difference in height we instead use a fixed *height* difference (the height of the surface above sea-level) and, from that, find the pressure difference.

> ✍ In chapter 1, it was mentioned that the METAR information includes an altimeter reading. Basically, the altimeter reading is the sea-level pressure used as a reference for aircraft who are using a standard temperature profile to determine their altitude. Consequently, it will be very similar to the actual sea-level pressure, with the difference depending on how close the actual temperature profile is to the standard.

Checkpoint 9.8: On a surface map, is the plotted pressure typically greater than, equal to, or smaller than the observed surface pressure?

9.4 Vertical gradient in density

As mentioned in section 4.1.3, the **density** of the air typically decreases with height. After all, we expect that the air is thinner the higher one goes up a

---

[vii] 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 (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.

[viii] The representative temperature is based somewhat on the air temperature at the surface.

mountain.

However, whereas the pressure *must* decrease with height (in order for the PGF to balance gravity), the density can *increase* with height (which produces optical effects called mirages; see Problem 9.4). Special circumstances are needed, though. In particular, the temperature has to decrease *drastically* with height.

How drastically?

Using the hydrostatic balance, we can show that temperature must decrease with height at a rate greater than g/R:

$$\frac{\partial \rho}{\partial z} > 0 \quad \text{if} \quad \frac{\partial T}{\partial z} < -\frac{g}{R}. \tag{9.8}$$

Since g/R is approximately 34 K/km, this means that the vertical temperature gradient must be less than -34 K/km.

The lapse rate is the negative of the vertical temperature gradient, so this means the density *increases* with height if the lapse rate is greater than 34 K/km. In chapter 4, we found that a typical lapse rate in the troposphere is about 6.5 K/km, so normally the density *decreases* with height.

A lapse rate of 34 K/km is rare, although it can occur in a shallow layer during the summer when the air just above a dark surface gets very hot compared to the air further up. This makes the density increase with height, and leads to the mirror-like surface that makes it look like there is water on the ground.[ix]

How is this relationship obtained?

We can get this relationship by combining the hydrostatic relationship (equation 9.1) with the ideal gas law.

Before doing so, I want to point out that the ideal gas law tells us that if P decreases (which it must do to maintain hydrostatic balance) then so must the product ρRT (since $P = \rho RT$). The gas constant R doesn't change, so that means that the product ρT must decrease with height.

[ix] This is because the speed of light depends on the density of the air. As light travels into a region where its speed is different, it bends. This process is called **refraction**, and is responsible for reflecting the light from the sky, producing the mirror-like mirage.

Normally, both ρ and T decrease with height. However, if T decreases very quickly, then ρ could *increase* with height, and the product ρT would *still* decrease with height.

To find out just how great the lapse rate must be for this to happen, we use the hydrostatic relationship (equation 9.1). Basically, we need to convert the hydrostatic relationship from a pressure-height relationship to a density-height relationship.

To do this, we start with the ideal gas law(equation 2.1).

$$P = \rho RT$$

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

$$\frac{\partial \rho}{\partial z} = -\frac{\rho}{T}\left(\frac{\partial T}{\partial z} + \frac{g}{R}\right) \tag{9.9}$$

According to this expression, the density will be constant with height (i.e., $\partial\rho/\partial z = 0$) if the stuff in parentheses is zero. That happens if the vertical gradient of temperature equals $-g/R$. If the vertical gradient of temperature is more negative than that, the density increases with height. Since the lapse rate is the negative of the vertical temperature gradient, this means that the density increases with height when the lapse rate is greater than g/R.

Checkpoint 9.9: If the atmosphere is in hydrostatic balance, does that mean that the density necessarily *has to decrease with height?*

Project

Project 9.1: Create a graph of the pressure profile of the U.S. Standard Atmosphere from sea-level up to 86 km, using equation 9.2 to calculate the pressure at each level:

$$P_{\text{top}} = P_{\text{bottom}} \exp\left(-\frac{g}{RT_{\text{avg}}}\Delta z\right)$$

where T_{avg} *is the average temperature of the layer between* z_{bottom} *and* z_{top}*,* R *is the gas constant, and* g *is the gravitational field (see appendix B.4). Note that this project is similar to Project 4.2 except that we are calculating the pressures differently. See the footnote*[x] *for a suggested method. Note that we are assuming* g *is constant, so use the geopotential heights, not the geometric heights, as mentioned in the note on page 160.*

Project 9.2: (a) How do your values compare to the ones calculated in the Project 4.2, which assumed a strictly exponential function with scale height H *equal to 8500 m? Describe any differences.*
(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 9.3: (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 average temperature of the troposphere was warmer, what would happen to your answer in (a)? Assume the tropopause is at the same height. Try it out.
(c) Based on your answer to (b), where would the height of the 500-mb level be higher: where the troposphere is warm or where it is cold? Explain, using your answer in (b).

Questions

Question 9.1: In order for the pressure gradient force to balance the gravitational force, the pressure gradient force must be directed upward.
(a) Does this necessarily mean that under hydrostatic balance the pressure decreases with height? Why or why not?

---

[x] 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_{top}, using the average temperature of the layer just below it (between z_{bottom} and z_{top}), the appropriate values of g and R, the difference in height between the two layers, Δz, and the pressure at the previous layer P_{bottom}. At the surface, set the pressure to be 1013.25 mb.

(b) Does this necessarily mean that under hydrostatic balance the pressure gradient decreases with height? Why or why not?

Question 9.2: 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?

Question 9.3: (a) In the northern hemisphere, where should the 1000-500 mb thickness be greater: toward the north or toward the south? Why?
(b) Is your answer to (a) consistent with the current 1000-500 mb thickness over the United States? How or how not?
(c) In general, the 500 mb heights are similar to the 1000-500 mb thicknesses. Based on that, which typically varies more: the 500-mb heights or the 1000-mb heights?

Question 9.4: 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). Provide your reasoning.
(b) Based upon the relationship between warm advection and pressure falls discussed in section 9.3.2 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? Provide your reasoning.
(c) Is your answer in (b) consistent with the observed 3-hour pressure change relative to the observed sea-level pressure? How or how not?

Problems

Problem 9.1: Based on the information obtained in Project 9.1, 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 4.3 on page 57). Which is greater?

Problem 9.2: (a) Estimate the gravitational force per volume, ρg, 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 9.1)? Should they be equal or should one be greater? If one should be greater, which one?

Problem 9.3: (a) In Problem 9.1, 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 9.4: (a) For dry air, what is the value of g/R? Convert your answer to units of °C/km.
(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 lapse rate (negative of vertical temperature gradient) for air near Earth's 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 air near Earth's surface as shown in the skew-T shown in Figure 4.5 (page 64). Note: This is a little tricky since height is not provided on a skew-T. However, you can use the information from Project 9.1 or Problem 9.1 to relate vertical pressure differences with vertical height differences.
(f) Compare the answers from (a) and (e). Does density decrease with height at Flagstaff as observed in Figure 4.5? Explain how you know.

Problem 9.5: If the atmosphere, with total mass of 5.136×10^{18} kg, had a constant density equal to its density at the surface, how thick would it be?

Problem 9.6: Use equation 9.5 to obtain the mass of the atmosphere. Why is it not equal to 5.136×10^{18} kg?

Problem 9.7: Venus is about the same size as Earth (its radius is about 95% of Earth's and its surface gravitational field is about 90% of Earth's) but its atmosphere has a surface pressure that is about 92 times that of Earth's. Based on this, what is the mass of Venus' atmosphere? Compare your answer with the value provided on NASA's Venus fact sheet.

Derivations

Show-me 9.1: 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, $\partial P/\partial z$ *(see section 9.1).*

Show-me 9.2: Show how equation 9.2 can be obtained from equation 9.3:

$$\ln\left(\frac{P(z_2)}{P(z_1)}\right) = -\frac{g}{RT_{\text{avg}}}\Delta z$$

Show-me 9.3: Show how equation 9.8 is obtained from equation 9.9.

Show-me 9.4: To derive equation refmass of atmos on page 161, we had to assume that ΔP *equals* $P(0)$ *if* M *is the entire mass of the atmosphere (see last step of the derivation). Why does* ΔP *equal* $P(0)$ *when* M *is the entire mass of the atmosphere?*

10. Lapse rates

Introduction

Once we have the pressures (see previous chapter), we can then figure out the wind (see chapters 12 and 13). Before we do that, however, I'd like to stop and think about how we got the pressures and what that means about vertical motions in the atmosphere.

To get the pressures, we needed to assume hydrostatic balance, which means there is no vertical forcing and thus no vertical accelerations. At the time, it was mentioned that the air can still *move* vertically even though it doesn't *accelerate* vertically but how do we get vertical motion without some kind of vertical forcing?

This is important because vertical motion is an important thing to predict – rising air cools (because of adiabatic expansion), which can cause clouds and precipitation if the air cools to the dew point.

We'll address this deficiency in this chapter and the next. It will be addressed, not by getting rid of hydrostatic balance but rather by identifying the situations when individual parcels are out of balance from the surrounding, balanced environment. By treating the parcels as separate from the environment, we can continue with our assumption of an environment in hydrostatic balance while allowing for individual parcels to accelerate up or down.

We'll approach this in two parts. In this chapter, we look at how a parcel's temperature changes as the parcel is raised or lowered in the atmosphere. In the next chapter, we'll compare that with the surrounding environment and explore what that means for the stability of the parcel.

10.1 Parcel vs. environment lapse rate

To predict a parcel's temperature as it rises or falls, we assume adiabatic expansion as it rises (and thus cools) and adiabatic compression as it falls (and thus warms), as we identified in chapter 6. In section 10.2, we'll identify the lapse rate experienced by the parcel as it rises or falls. The **lapse rate**, you may recall, is the rate at which the temperature decreases with height (see section 4.2.1).

Before we get to that expression, though, I want to emphasize the difference between the *parcel's* lapse rate and the *environmental* lapse rate.

The *environmental* lapse rate is the rate at which the *actual observed* temperature decreases with height. For example, the U.S. Standard Atmosphere describes an average atmosphere where the temperature in the troposphere is colder the higher up one is, with a lapse rate of 6.5°C/km. That means that if you had a really tall ladder, the temperature of the air at one height would be 6.5°C colder than the air 1 km below it.

> ✍ 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 this chapter, we are *not* looking at the environmental lapse rate. Instead, we are looking at an *individual parcel's* lapse rate. In other words, we are looking at how an individual parcel's temperature changes as it ascends or descends in the atmosphere. We'll call this the **parcel lapse rate**, to distinguish it from the **environmental lapse rate**.

> ✍ Recall that 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. Rather, it cools or warms because it is expanding against the air around it or being compressed by the air around it.

The variable abbreviation for the lapse rate is Γ. To distinguish between parcel lapse rate and the environmental lapse rate, I'll use subscripts, so that Γ_{env}, Γ_{d} and Γ_{m} will refer to the environmental lapse rate, a dry parcel's lapse rate, and a moist parcel's lapse rate, respectively.

The parcel lapse rate depends on whether the parcel is saturated. If it is not saturated, it will cool or warm because of expansion or compression. However, as we'll see in chapter 16, if the air is saturated then any cooling will result in condensation, which releases latent heat. Consequently, an ascending saturated parcel will cool less slowly than an unsaturated parcel. There is no significant[i] difference between a saturated and unsaturated parcel upon descending, since no condensation occurs when a parcel warms.

Checkpoint 10.1: If a parcel of air rises in the atmosphere, does its lapse rate (the variation in temperature experienced by the air parcel as is ascends or descends) necessarily equal the environmental lapse rate (the lapse rate of the environment through which the parcel passes)?

10.2 Adiabatic lapse rate

A parcel's *pressure* always equals the pressure of the surrounding environment but the parcel's *temperature* does not need to equal the temperature of the surrounding environment. As a parcel rises in the atmosphere, it expands so that its pressure matches that of the environment and, as mentioned in the previous section, it cools. In this section, we derive an expression that gives how an individual parcel's temperature decreases as it rises in the atmosphere (the parcel lapse rate).

The parcel's temperature, after adiabatic expansion (or compression) is given by adiabatic temperature equation (equation 6.3). However, rather than start with that expression, it is easier to start with the expression from which we derived the adiabatic temperature equation, namely the first law of thermodynamics assuming purely adiabatic changes (equation 6.2):

$$dT = \frac{1}{c_\mathrm{p}\rho} dP$$

---

[i] A saturated parcel has a slightly different specific heat because the specific heat of water vapor is different than the specific heat of dry air. However, only a small part of saturated air is water vapor, so the overall impact on the parcel's specific heat is small.

The reason why it is easier to start with this expression is because it is in terms of the derivative of pressure, and in that sense it is similar to the hydrostatic balance relationship (equation 9.1):

$$\frac{\partial P}{\partial z} = -\rho g$$

Notice how the first relationship relates temperature and pressure, and the second relates pressure and height. We can thus combine the two to get rid of the dependence on pressure. Basically, we use the hydrostatic relationship to replace dP in the adiabatic expansion expression with $-\rho g dz$. This gives the following relationship between temperature and height:

$$\frac{\partial T}{\partial z} = -\frac{g}{c_{\mathrm{p}}}$$

The particulars of the derivation is left as an exercise for the reader (see Show-me 10.1).

Since the **lapse rate** (Γ) is defined as the negative of the vertical temperature gradient ($-\partial T/\partial z$), the parcel's lapse rate under adiabatic conditions is:

$$\Gamma_{\mathrm{d}} = \frac{g}{c_{\mathrm{p}}} \tag{10.1}$$

This is called the **adiabatic lapse rate** and is indicated as Γ_{d}.

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.[ii] Thus, the "d" stands for "dry".

WHAT DOES THE ADIABATIC LAPSE RATE DEPEND UPON?

According to the expression, it only depends on g and c_{p}. 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}:

$$\Gamma_d = \frac{g}{c_{\mathrm{p}}} = \frac{9.8\ \mathrm{N/kg}}{1004.67\ \mathrm{J \cdot kg^{-1}K^{-1}}} \sim 10\ \mathrm{K/km}$$

[ii] As we'll see in chapter 16, if the air is saturated then any cooling will result in condensation, which releases latent heat. Consequently, an ascending saturated parcel will cool less slowly than an unsaturated parcel. There is no difference between a saturated and unsaturated parcel upon descending, since no condensation occurs when a parcel warms.

Thus, the adiabatic lapse rate (see equation 10.1) is about $+10\ \mathrm{K} \cdot \mathrm{km}^{-1}$. 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.

Checkpoint 10.2: In the U.S. Standard Atmosphere, the temperature 1 km above the surface is 8.5° C. If that 8.5° C air was brought down to the surface, what temperature would it then have? Would that air be warmer, or cooler, than the air that was already at the surface (15° C in Standard Atmosphere)?

10.3 Skew-T log-P diagrams

In section 4.5, 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 10.1 (light red, if colors are viewable).

These lines are called **dry adiabats** and 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 10.1: From the skew-T log-P diagram in figure 10.1, determine the temperature an air parcel would have if it initially had a temperature of −14°C at 540 mb and then was brought down to 800 mb.

Answer 10.1: First we identify the initial temperature and pressure of the parcel. This is indicated by the small black circle in Figure 10.1, which is between the −10°C and −20°C temperature lines (slanted solid blue lines going from lower left to upper right) and between the 500 mb and 600 mb pressure lines (horizontal solid blue lines). This temperature and pressure happens to coincide with the observed temperature at that pressure level, but it need not do so for the purpose of this example.

Next, starting from this location, we follow the direction of the dry adiabatic lapse rate lines (slanted red dashed curves going from upper left to lower

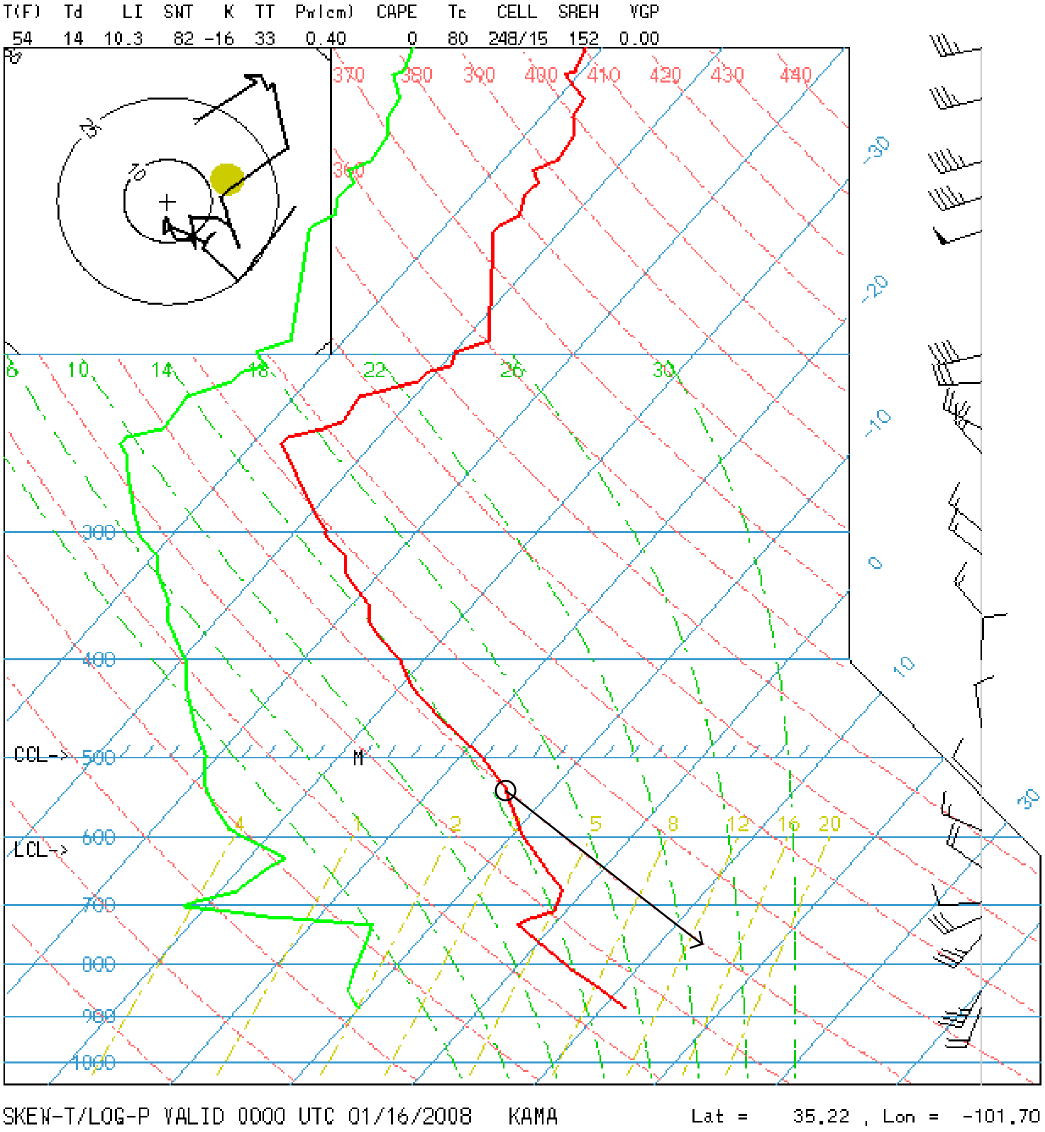

Figure 10.1: A skew-T log-P graph for Amarillo, Texas. Data is from 0000 UTC (7 PM local the previous day), 16 Jan, 2008. Horizontal solid lines are pressure (in mb). Skewed, solid lines are temperature (in °C). Small black circle and associated black arrow are discussed in the text.

right) toward the 800 mb pressure line. This is indicated by the solid black arrow. When we reach the 800 mb pressure line, we stop and look at what temperature that point corresponds to.

In this case, that point corresponds to a temperature arpimd 17°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.

Checkpoint 10.3: Which of the following is represented by the dry adiabats on the skew-T graph?
(a) The temperature of dry parcel would have as it moves vertically in the atmosphere,
(b) The temperature of the actual atmosphere,
(c) The temperature of the U.S. Standard Atmosphere.

10.4 Potential temperature

On the Skew-T diagram, the pressure lines are drawn every 100 mb and the temperature lines are drawn every 10°C. What about the dry adiabats?

The values of each dry adiabat is listed at the top of the Skew-T diagram. You'll notice that they have values like 360, 370, etc.

The units aren't given but it turns out these numbers are in kelvin (drawn every 10 K). They represent rather warm temperatures because the values correspond to the temperature a parcel located on that dry adiabat would have if it were compressed (or expanded) to 1000 mb.

Let's examine one of them as an example. The example I'll use is the dry adiabat that coincides with the black arrow that was drawn in the previous section. By incrementing 10 K multiple times from the labeled dry adiabats, one can see that the black arrow lies along the dry adiabat that has a value of 310 K.

To understand what that 310 K value means, continue down that dry adiabat to where it intersects the 1000 mb pressure line. You should find that the dry adiabat intersects the 1000 mb pressure line at a temperature of 37°C, which is equal to 310 K. In other words, the 310 K dry adiabat is labeled according to the temperature a parcel would have if it started on that dry adiabat and then was compressed (or expanded) to 1000 mb.

We call that reference[iii] temperature value the **potential temperature** and we indicate it mathematically by the Greek letter θ. To determine a parcel's potential temperature, then, we can either see which dry adiabat the parcel lies along or we can use the adiabatic temperature equation and calculate what the temperature would be if compressed (or expanded) to 1000 mb:

$$\theta = T \left(\frac{1000\text{ mb}}{P} \right)^{R/c_\text{p}} \tag{10.2}$$

The variables P and T represent the pressure and temperature, respectively, of the parcel before the air is compressed/expanded to 1000 mb.[iv]. I'll call this the **potential temperature equation**. It is really just the adiabatic temperature equation with 1000 mb as the new pressure.

Any parcel that has a temperature and pressure that puts it along the 310 K dry adiabatic curve has the *same* potential temperature value (310 K in this case). In other words, as a dry parcel rises or falls, its potential temperature (or reference temperature) must remain constant. The parcel's *temperature* changes (due to adiabatic warming or cooling) but its *potential* temperature does not, meaning 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 adiabats" we could instead call them the "potential temperature curves."

Checkpoint 10.4: 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?

[iii] Technically, the reference pressure could be any value, not just 1000 mb but in almost every instance people use 1000 mb as the reference pressure and I will use 1000 mb as the reference pressure exclusively.

[iv] By convention, potential temperatures are given in kelvin.

10.5 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, particularly near the ground. Thus, it is useful to examine what happens in such cases.

For example, suppose we wanted to mix the air in the lowest one kilometer of the atmosphere. Suppose further that the temperature at the surface is 10°C and the lapse rate is 5°C/km, such that the temperature is 5°C at an altitude of one kilometer.

If they were at the same height and pressure, we could just take the average of the two temperatures and that would be the temperature of the mixed air (assuming the same number of moles in each portion; see Show-me 10.2). In other words, we'd just take the average of 10°C and 5°C and get 7.5°C. However, we are mixing air at the surface with air one kilometer above. To mix those two parcels, we need to bring the air aloft down to the ground and visa-versa, and that results in adiabatic warming of the air that was aloft (as it is brought down) and adiabatic cooling of the air from the surface that is raised aloft.

To illustrate, envision that we split the one-kilometer layer in half, creating two half-kilometer layers. Let's assume the lower layer is entirely at 10°C at a height of zero, and that the upper layer is entirely at 5°C at a height of one kilometer.

If we take the top layer and move it to the surface, it will warm at the dry adiabatic rate of 10°C/km, which means it will warm to 15°C. When it mixes with the 10°C air there, the result is 12.5°C (where 12.5 is the average of 10 and 15). It is at the same pressure, after all, so we can take the straight average.

Meanwhile, if we take the bottom layer and move it to a height of one kilometer, it will cool at the dry adiabatic rate of 10°C/km, which means it will cool to 0°C. When it mixes with the 5°C air there, the result is 2.5°C (where 2.5 is the average of 0 and 5).

So, the end result is a one-kilometer layer that is 12.5°C on the bottom and 2.5°C on top. The bottom has *warmed* (from 10°C to 12.5°C) and the top has cooled (from 5°C to 2.5°C).[v]

---

[v] In fact, citrus growers take advantage of this when the temperature dips below freez-

Notice, however, that the *new* lapse rate is 10°C/km, which is the adiabatic lapse rate. Consequently, the air at the top of the layer would have the same temperature as the air at the bottom, if the top air was brought down to the bottom.

In other words, the well-mixed layer has the same potential temperature throughout the layer. Such a well-mixed layer is shown near the surface in Figure 10.1.

> ✍ As you'll see in section 11.4, we refer to a layer of constant potential temperature as having **neutral** stability.

In fact, it turns out to be relatively easy to determine the temperature after mixing if we focus on the potential temperature. When two parcels from different heights mix together, the final **potential** temperature will be an average of the initial **potential** temperatures.

For example, in the scenario considered earlier, the temperature at the bottom was initially 10°C. Let's suppose the pressure there is 1000 mb. That means its potential temperature at the surface is also 10°C. Meanwhile, the temperature at the top was initially 5°C. Since it was at a height of one kilometer, its potential temperature would be 15°C (since it would warm 10°C/km when brought down to the surface).

The average of those two potential temperatures is a potential temperature of 12.5°C, which agrees with what we got before, with air colder on top than on the bottom. Mixing the air near the surface with the air above tends to result in the surface air warming and the air above it cooling because the potential temperature usually *increases* with height (as opposed to the temperature typically *decreases* with height).

After mixing, the environmental lapse rate is equal to the adiabatic lapse rate (g/c_{p}, which is about 10 K/km), which is greater than the typical tropospheric lapse rate of 6.5 K/km (as represented by the U.S. standard atmosphere). There are two reasons for this difference. One is that a typical atmosphere only partially mixed, not totally well-mixed. We also have moisture in the atmosphere and the condensation of water vapor gives off heat, which affects the lapse rate (see section 16.2).

---

ing. 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.

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?

A house differs from the troposphere in two ways. First, in the troposphere the source of heat is the ground (due to it absorbing solar radiation) whereas in a house it is throughout the house (due to the heating system) or the roof (due to it absorbing solar radiation). Second, the air in a house is typically not well-mixed.

Doesn't hot air rise?

As will be discussed in chapter 11, 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 (the top of a mountain being colder than the valley).

Checkpoint 10.5: What is it about the layer of air near the surface as shown in figure 10.1 that implies the layer is well-mixed?

Project

Project 10.1: Consider a 5° C dry air parcel at 800 mb.
(a) Using a blank skew-T log-P diagram, determine the temperature the air parcel would have if it were expanded to 500 mb. Make sure you identify the location on the skew-T that corresponds to the original temperature and pressure, and also draw a curve that parallels the dry adiabat curves (the straight lines on the skew-T log-P diagram that have a slope opposite those of the temperature) to the point of the new temperature and pressure.
(b) Check your answer in (a) by using equation 6.3 (see page 92).

Problems

Problem 10.1: In project 9.1, you determined the air pressure every kilometer of the **U.S. standard atmosphere** *assuming hydrostatic balance. From that project, it was found that at an altitude of two kilometers, where the*

temperature is 275.15 K, the pressure is 795.15 mb. If that air (with a temperature of 275.15 K) is brought down to the surface, the temperature warms.
(a) Determine the parcel's new temperature by using the adiabatic temperature equation (equation 6.3), given that the parcel's new pressure is 1013.25 mb.
(b) Determine the parcel's new temperature by using the dry adiabatic lapse rate (equation 10.1), given that the parcel's new height is zero.
(c) Do both methods predict the same temperature? Should they?
(d) How does the new temperature compare with the temperature the parcel had prior to being brought down to the surface?
(e) How does the new temperature compare with the temperature of the air typically found at the surface (according to the U.S. Standard Atmosphere)?

Problem 10.2: In the U.S. Standard Atmosphere, the temperature 1 km above the surface is 8.5° C, which is 6.5° C colder than the temperature at the surface (15° C). 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 10.3: (a) From the skew-T log-P diagram in figure 10.1 (on page 178), 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.

Problem 10.4: If air at 800 mb has a potential temperature 300 K, what is its potential temperature if brought up to 500 mb?

Derivations

Show-me 10.1: Derive equation 10.1 by combining the hydrostatic balance relationship (equation 9.1) with the first law of thermodynamics (equation 6.5) for adiabatic processes ($dQ = 0$)*.*

Show-me 10.2: Using the $PV = nR^*T$ *formulation of the ideal gas law, show that when two parcels of identical pressures and number of moles are mixed, the resulting temperature equals the average of the two original temperatures. Apply to each individual parcel and the mixed parcel, using the fact that* n *of each parcel is the same,* $2n$ *is the number of moles of the mixed parcel, and the mixed parcel has the same pressure as the two initial parcels but a volume equal to the sum of the initial two parcel volumes.*

11. Stability

Introduction

As mentioned in the previous chapter, we are assuming that the atmosphere is in hydrostatic balance. That means that there are no vertical accelerations. This doesn't mean there can't be vertical motions, but without vertical acceleration it is pretty difficult to get vertical motions (especially considering that the air at the surface can't be moving vertically).

Despite this constraint, we can still get vertical accelerations. The way around this is to have individual *parcels* that are not in equilibrium with the surrounding air. In other words, we'll assume the surrounding air is in hydrostatic balance but an individual parcel *in* the atmosphere may not share the same characteristics (i.e., may be warmer or colder) and thus may *not* be in hydrostatic balance with the rest of the atmosphere.

We'll approach this in two parts. First, I'll review how an individual parcel can be accelerated even though the environment around it is in hydrostatic balance. We'll then compare the parcel's lapse rate, as it is raised or lowered in the atmosphere, with the lapse rate of the surrounding environment and explore what that means for the acceleration of the parcel as it moves up or down in the atmosphere.

11.1 Buoyancy

As you probably already know, hot air rises and cold air sinks.

Let's look at this more closely.

As discussed in section 9.1, there are two forces acting on an air parcel: the gravitational force pulling downward and the pressure gradient force pushing upward.

> ✍ As we did in section 9.1, we'll deal with 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 just force, I will still tend to use F as the variable abbreviation since it has many of the same properties as force. Just remember that the units are different.

During hydrostatic balance, it doesn't matter if the air is hot or cold. Either way, the air won't accelerate vertically because the two vertical forces (gravity pulling down and pressure gradient pushing up) balance.

However, what happens if the air is warmer or colder than the *surrounding*, vertically-balanced air?

In that case, the forces on the *parcel* won't be balanced.

To see why, let's first examine the forces on the *surrounding*, vertically-balanced air. The gravitational force per volume 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, equation 9.1 applies:

$$\frac{\partial P}{\partial z} = -\rho g$$

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 volume 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 volume is greater and thus the parcel accelerates downward (in the direction of the gravitational force). In other words, cold air sinks, but only if it is colder than the *surrounding* air.

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). In other words, warm air rises, but only if it is warmer than the *surrounding* air.

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.

Checkpoint 11.1: Does warm air always rise?

Mathematically, we can write the gravitational force per volume on the parcel as $\rho' g$ (where ρ' is the density of the parcel). That means the force imbalance on the air parcel is $\partial P/\partial z - \rho' g$. Since the surrounding air is in balance, we can replace $\partial P/\partial z$ with ρg (where ρ is the density of the surrounding air). Making this replacement we have that the force imbalance on our parcel is as follows:

$$\vec{F}_{\text{buoyancy}} = g(\rho - \rho')\hat{k} \tag{11.1}$$

where the force imbalance $\vec{F}_{\text{buoyancy}}$ is known as the **buoyancy** force and $\hat{k}$ indicates the upward direction. 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.

Checkpoint 11.2: Why is it said that equation 11.1 implies that an air parcel that is less dense than its surroundings will experience an positive (upward) buoyant force.

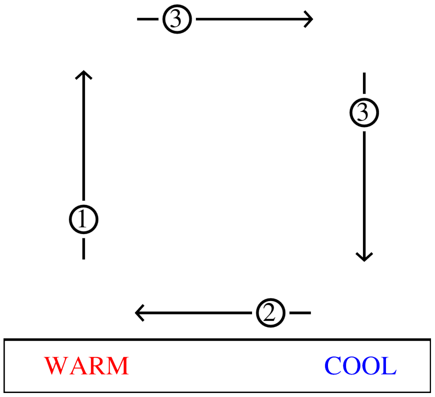

Figure 11.1: A representation of a vertical slice of the atmosphere, showing the circulation set up when there is a temperature difference at the surface.

11.2 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).

We can understand why it switches by using the ideas discussed so far. If the environment is in hydrostatic balance, then a parcel within that environment will rise if it becomes warmer than the environment.

Consider the surface near the shore. During the day, the land heats up more than the sea.[i] If the air was in hydrostatic balance before the heating, the air above the land will be buoyant and will rise. This is illustrated by arrow ① in Figure 11.1.

---

[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.

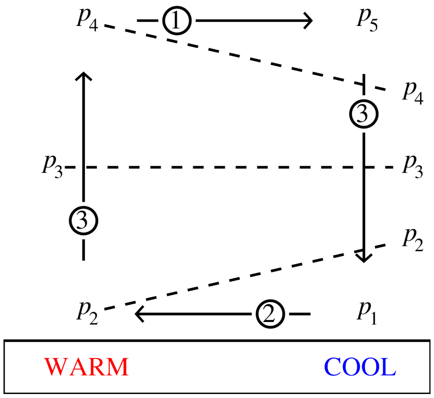

Figure 11.2: As in Figure 11.1, with the addition of isobars. The surface pressure is p_1 at the surface in the cool region and p_2 at the surface in the warm region. Pressure decreases with height, so that $p_1 > p_2 > p_3 > p_4 > p_5$.

To fill in the "gap," so to speak, the air near the surface in the surrounding areas then moves toward the warm region. This is indicated by arrow ② in Figure 11.1. When the wind is blowing from the cooler sea to the land, this is called the **sea breeze**. The reverse situation, called the **land breeze** occurs at night, when the land is cooler than the water (since land cools down more quickly than the water).

A circulation is then set up, as indicated by the addition of the other arrows in Figure 11.1.

Couldn't both regions be in hydrostatic balance at the same time?

Yes. After all, just because the air over the land warms up does not necessarily mean it is no longer in hydrostatic balance. However, it turns out the circulation will be set up regardless.

To see why, consider the situation illustrated in Figure 11.2. This figure is the same as Figure 11.1 but I added some dashed lines to indicate the air pressure.

If both regions are in hydrostatic balance then the vertical pressure gradient dP/dz equals $-\rho g$ for both regions. However, the warmer region will have a lower density, and thus will have a smaller vertical pressure gradient. This is why the warm region has fewer dashed lines above it than the cool region does. In other words, the thickness is greater in the warm air.

Because of this difference, the air pressure is greater at the top of warm column than at the top of the cool column. That difference in pressure pushes the air toward the cool region, as indicated by arrow ① in Figure 11.2.

Mind you, this is at the top of the column. With more air being placed above the cool region, the surface pressure rises, since by hydrostatic balance the surface pressure indicates how much air is above it.

At the surface, then, there is greater pressure above the cool air, and this pushes the air from the cool region to the warm region. This is indicated by arrow ② in Figure 11.2. During the day, arrow ② in Figure 11.2 indicates the sea breeze, as the air flows from the cooler water to the warmer land.

The same circulation is then set up as we described before, as indicated by the addition of the other arrows in Figure 11.2.

Checkpoint 11.3: 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?

11.3 Equilibrium level

We can now predict 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. We also can predict what happens to that parcel's temperature as it rises or falls in the atmosphere (from chapter 10).

To determine if the parcel will *continue* to rise or fall, we need to compare the parcel's "new" temperature with the temperature of the surroundings. As long as the parcel's temperature remains warmer than the surroundings, it will continue to rise. Conversely, as long as the parcel's temperature remains colder than the surroundings, it will continue to fall.

Basically, the parcel will continue to rise or fall until it reaches an environment where its temperature equals the temperature of the environment. That level is called the **equilibrium level**.

To find this level, we need to know two things. First we need to know how the parcel's temperature changes as it rises or falls. We've already addressed that previously and can either use equation 10.1 or a skew-T log-P chart.

The second thing we need to know is the temperature of the environment. Since that information is provided on a skew-T log-P chart, 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 11.3.

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).

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.[ii] 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 −21°C parcel "introduced" at 400 mb. This parcel is indicated by the thick gold circle on the diagram.

[ii] 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.

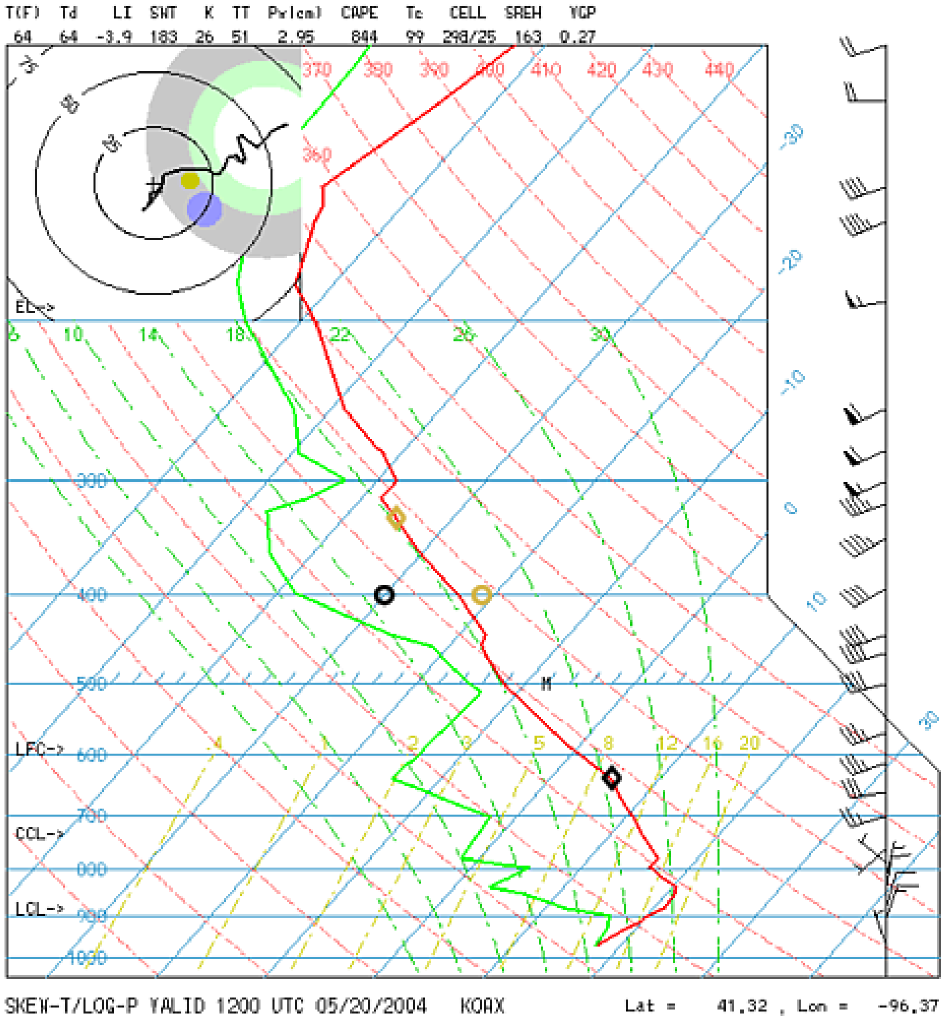

Figure 11.3: A skew-T log-P graph for Omaha, Nebraska. Data is from 1200 UTC (7 AM Central Daylight Time), 20 May, 2004. The gold and black circles and diamonds are used to indicate parcels discussed in this chapter.

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 330 mb, which I've indicated with the thick gold diamond.

HOW DO WE GET THE AIR TO BE WARMER OR COOLER THAN THE SURROUNDINGS?

For air near the surface, the air can warm or cool simply by being next to the ground, which warms up during the day as it absorbs solar radiation. Keep in mind, though, that a warm parcel doesn't necessarily rise – 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 or cooler than the environment is to cause it to ascend or descend. That requires some other force. For example, if the wind sends the air into a mountain, the ground may force the air to move upward. Or, when two air masses meet (as with a front[iii]), some air may be pushed up or down at the boundary.footnoteConsider, for example, how squeezing some clay in your hand would make some ooze out through the gaps in your fingers.

Checkpoint 11.4: Suppose a $-20°C$ *parcel is introduced at 500 mb over Omaha, Nebraska, at 1200 UTC 20 May 2004. Where is its equilibrium level?*

11.4 Stable, neutral and unstable layers

In this section, we examine **stability**.

[iii] The two air masses could have different temperatures, which may introduce air of a different temperature into a region.

To understand stability, consider a marble in the bottom of a bowl. When the marble is moved to one side, it rolls back to the bottom. This is considered to be a **stable** situation because the marble returns to its equilibrium 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 to be an **unstable** situation because the marble continues to move away from its equilibrium position when displaced slightly from it.

In the atmosphere, we deal with air parcels, not marbles. Like the marbles, though, we'll assume the air at all levels is in equilibrium, meaning that the air is the same temperature as its surroundings.

The question, like with the marbles, is whether air parcels will return to its position if given a push up or down. Unlike the marbles, though, what forces the air parcel back toward its equilibrium position or away from its equilibrium 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 equilibrium 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 equilibrium 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.

Checkpoint 11.5: (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?

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 Γ that is less than g/c_p or about 10 K/km:

For example, the sounding for Omaha, Nebraska, as shown in figure 11.3 (page 192), 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.

Checkpoint 11.6: What is it about the air between 620 mb and 570 mb Omaha, Nebraska, in figure 11.3 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.

> ✍ Contrary to what people say about warm air rising, warm air only rises if it is warmer than the air on either side. The air can be warmer than the air above it (or colder than the air below it) and still be in a perfectly stable situation. Even if 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 11.3.

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).

Checkpoint 11.7: (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.

Questions

Question 11.1: Draw a figure, like that in Figure 11.2, which illustrates the surface and upper-air pressures over water and land at night. Explain how your figure leads to the wind coming off the land at night.

Question 11.2: 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.*

Question 11.3: 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?

Problems

Problem 11.1: Suppose the air at 700 mb is $-10°C$*. If we warm an air parcel there to* $-9°C$ *(i.e., 1°C above its surroundings), what is the magnitude of 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 the pressure of the air around it.*

Problem 11.2: (a) What is the potential temperature at the surface for Amarillo Texas, as shown in Figure 10.1 on page 178?
(b) What is the potential temperature at 800 mb for Amarillo Texas, as shown in Figure 10.1?
(c) What is it about the air between the surface and 800 mb for Amarillo Texas that indicates the layer is neutral?

Problem 11.3: (a) From the skew-T log-P diagram in Figure 10.1 (on page 178), is the air between 700 mb and 720 mb stable or unstable? How can you tell? (hint: compare with adiabatic lapse rate lines)
(b) Is there any level above 700 mb where the air is unstable? If so, where? How about neutral?

Problem 11.4: Consider the air above the surface as indicated by the skew-T log-P diagram in Figure 10.1. Suppose we take a parcel of air at 850 mb and heat it to 14° C *(which is warmer than the air that is surrounding it at that level).*
(a) Use the adiabatic lapse rate lines on the skew-T to determine how high it will rise. Provide the pressure level and the temperature the parcel will have at that level.
(b) Using equation 10.2 (potential temperature 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.

12. Wind Direction

Introduction

From hydrostatic equilibrium, we can deduce the pressures based on the temperature (see chapter 9). Once we have the pressures, we can then figure out the wind because, as you probably already recognize, the wind is intimately tied to the air pressure.

However, the process is not as simple as you might think. Basically, the air is forced from high pressure to low pressure. However, when we look at weather maps, we find that the wind tends to be directed *around* the low pressure center rather than *toward* the low pressure center. In other words, the wind is directed *parallel* to the isobars rather than *across* the isobars. Similarly, when we have a high pressure area, the wind is likewise directed *around* the high pressure center rather than *away* from it. Even more curious, the direction of the wind is opposite in opposite hemispheres (northern vs. southern).

All of this leads to the following two questions:

- Why doesn't the wind blow from high to low pressure?
- Why is the direction opposite in the southern hemisphere vs. the northern?

In this chapter, we'll answer both in a conceptual way that will allow us to predict the direction of the wind.

> ✍ Air moving horizontally is referred to as the **wind**. Air moving vertically is usually referred to as **updrafts** or **downdrafts**.

12.1 A conceptual approach

As mentioned in the introduction, we find that the wind is generally directed *around* the center of lowest pressure, not *toward* the lowest pressure. In addition, the direction is counter-clockwise in the northern hemisphere and clockwise in the southern hemisphere (with the directions reversed around high pressure centers).

Although one must be fairly quantitative to predict the wind *speed* (see chapter 13), a purely conceptual approach is sufficient to explain why the *direction* is the way it is. In this chapter, we are focusing on the wind direction, so a conceptual approach is sufficient.

12.1.1 Conservation of angular momentum

To understand why the air spins around a low pressure center, consider figure skaters. They can spin very quickly when they bring their arms toward their body. Why do they spin?

First, we need to recognize that the skaters don't start at rest. In other words, if we look carefully, the skaters are *already* spinning before they bring their arms in. By bringing their arms in toward their body, the skaters don't *start* to spin but rather start to spin *faster*. The 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 has 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.[i]

An object's angular momentum about an axis is the product of the object's **rotational inertia**[ii] about the axis and the object's **angular velocity**[iii] 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

---

[i] By definition, torque is equal to the product of F_{T} and r, where F_{T} is the component of the applied force about the rotation axis and r is how far the force is applied from the axis of rotation.

[ii] Or,**moment of inertia**.

[iii] Or, **rotation rate**.

axis of rotation.[iv] The further the mass from the axis, the greater the rotational inertia (meaning more **torque** is 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.

Checkpoint 12.1: When skaters bring their arms toward their body, they spin faster. What happens to their angular momentum during this time?

12.1.2 Frames of reference

How does this apply to the case of wind around a low pressure?

The air, like the skater, is already rotating because it is on Earth, which is rotating. Then, when the air is brought in toward the low pressure center, the air spins faster, just like how skaters spin faster when they bring their arms in.

The problem with the air is that we are observing the air on a rotating frame of reference (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 their arms, their rotational inertia (about the rotating axis of the turntable) will increase and their 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.

[iv] The rotational inertia is defined as $m(r^2)_{\text{avg}}$, where $(r^2)_{\text{avg}}$ is the mass-weighted average square radius of the object.

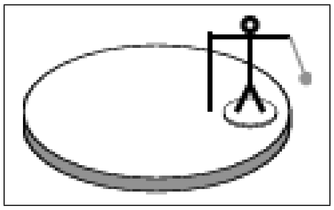

Now consider the situation illustrated in the figure. The person and the rotating turntable are placed on an even larger rotating platform (like the rotating platforms found at many playgrounds). What happens if the larger platform is rotating?

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?

If the person lets go of the platform and the turntable is free to rotate, then the person will rotate *faster* in the same direction they were rotating before pulling their arms in toward the body. Since the person is *already* rotating with the large platform, they 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 begins to rotate rather than remain stationary 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 the ground is not accelerating. Thus, someone on the ground sees the situation correctly (the person on the turntable rotates faster because that person 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 Earth. When looking down from above the northern hemisphere, 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).

What about around a high pressure center?

A similar process is at work for wind around a center of high pressure. Since stationary air relative to Earth is actually rotating cyclonically (with Earth), the air will spin 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 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.

> ✍ 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 the name: **cyclones**). Conversely, meteorologists use the word **anticyclonic** to describe the reverse direction – the direction that the air spins around a center of high pressure.

What happens in the southern hemisphere?

In the southern hemisphere, the direction of rotation will be clockwise because Earth spins clockwise as seen by someone looking down from above the southern hemisphere.

> ✍ The direction is actually the same in both hemisphere – with Earth's rotation – 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) by 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.

Checkpoint 12.2: In which direction does the air in a hurricane spin in the northern hemisphere: the same direction as the rotation of Earth, or opposite that direction? What about in the southern hemisphere?

12.2 Jet stream

A common feature of upper-level maps is the presence of a band of fast westerly winds called the **jet stream**. In Figure 12.1 this is seen as the shaded region of wind speeds greater than 70 knots, a region that dips southward along the west coast of the United States, cuts along southwest states and then rises northward along the eastern states.

We can now explain the physics behind this feature.

This band of very high wind speeds is a consequence of the temperature of the troposphere. More specifically, it is a consequence of how the average temperature of the troposphere (from surface to tropopause) is colder near the poles and warmer near the tropics.

To understand why, let's first consider the sea breeze, which is also set up by a temperature difference, but on a smaller scale. As mentioned in section 11.2, there are two ways we can describe the mechanism behind the sea breeze. One is to recognize that warmer air is less dense and so, if warm air lies (horizontally) next to cold air, the warm air will rise and the cold air will sink. The other way to describe the mechanism is to examine the vertical and horizontal pressure gradients. Recall from section 9.3.3 that

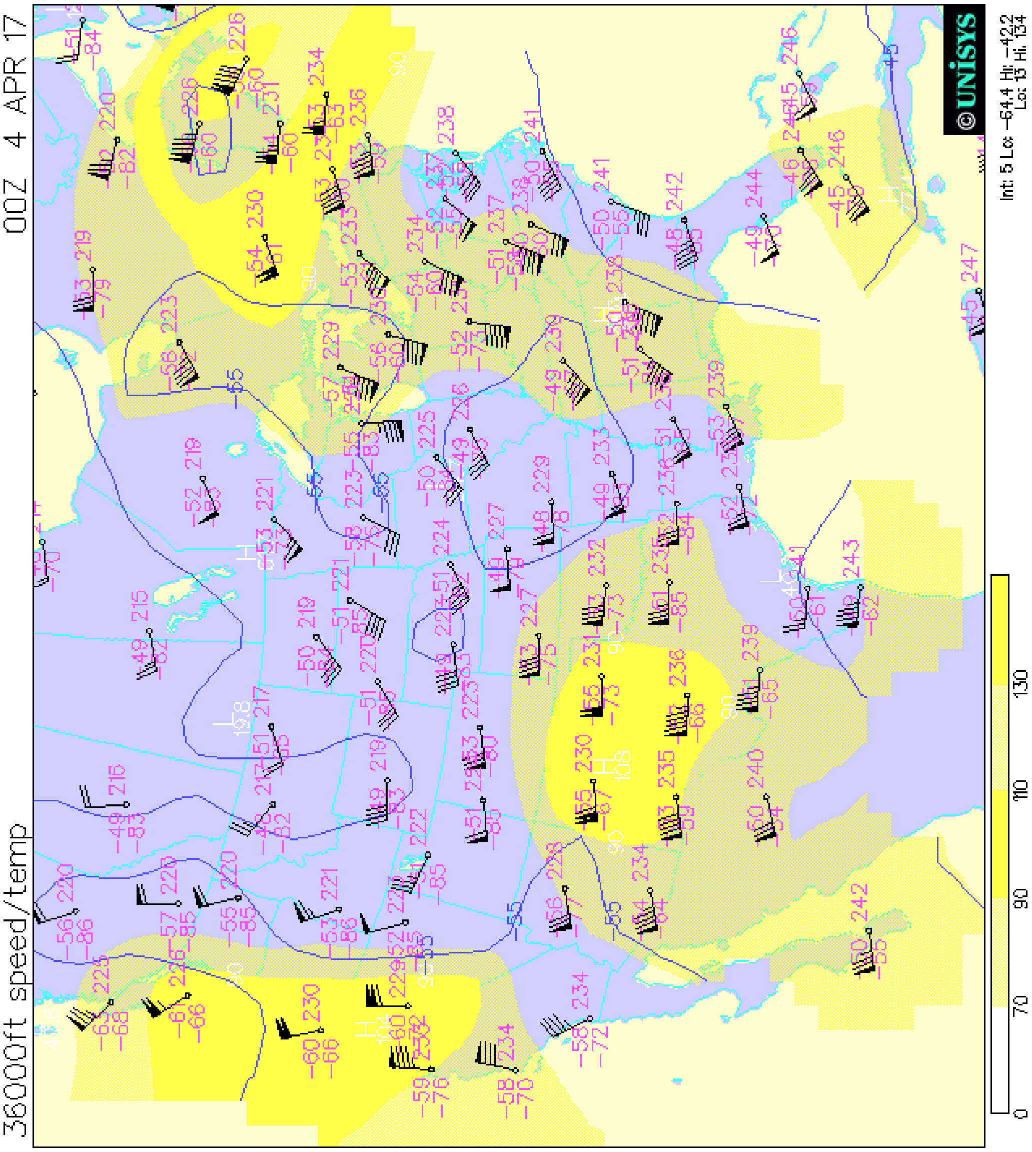

Figure 12.1: A plot of data at 36,000 ft for 0000 UTC on April 4, 2017, including wind flags and isotherms. Wind speeds over 70 knots are shaded.

warm layers are "thicker," meaning that the isobars are spread farther apart (in the vertical), because the vertical pressure gradient depends upon the air temperature (air pressure decreasing more slowly with height in warmer air). This is illustrated in Figure 11.2 with respect to the sea breeze, where each column has the same pressure at the surface (since the same mass of air is above each region) but, due to the different vertical pressure gradients, there exists a horizontal pressure gradient at the top of the layer, with higher pressures over the colder air.

For the sea-breeze, the effect is only a couple thousand feet deep. However, the same thing happens in the troposphere. Since the troposphere is so much deeper, the effect is amplified, such that there is a significant horizontal pressure gradient at the top of the troposphere, with higher pressures over the colder air. This can be seen in Figure 12.1, with pressures of around 220 mb to the north and pressures of around 240 mb to the south.

As discussed in the previous section, Earth's rotation comes into play at this scale and the wind that gets set up flows parallel to the isobars, not perpendicular to the isobars (as with the sea breeze). In the northern hemisphere the wind flows such that low pressure is on the left, which leads to westerly winds at the top of the troposphere.

> ✍ Notice that there doesn't need to be a temperature difference at the *top* of the air columns. Rather, what is needed is that the average temperature of throughout each air *column* is different. For the jet stream, that means that the entire layer from the surface to the tropopause has, on average, a temperature that is warmer on one side of the jet stream and colder on the other side.

Checkpoint 12.3: Which way does the jet stream typically flow: toward the east or toward the west?

12.3 Size of effect

WHY DOES THE SEA BREEZE BLOW ACROSS THE ISOBARS (TOWARD LAND AND LOWER PRESSURE) WHEREAS THE JET STREAM BLOWS PARALLEL TO THE ISOBARS?

Whether the wind blows parallel to the isobars or across the isobars depends on the size of the system (or the time involved). The system has to be large enough that the rotation rate of Earth becomes a significant factor.

So this isn't 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 Earth's surface. 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 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 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 Earth around its axis. So, we first need to determine the angular velocity of Earth around its axis (see next section).

Checkpoint 12.4: Does the water always go down the drain in a particular direction north of the equator?

12.3.1 Angular velocity of Earth

An object's **angular velocity** is the same as its **rotation rate** (or rotational speed).[v] To indicate the angular velocity, one provides the angle the object 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).[vi] 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 Earth with the velocity of a point on Earth's surface. Those equations are derived assuming radians as the angle unit. Thus, to use those equations, you need to use radians as the angle unit.

> ✍ A lower-case omega (ω) will be used to indicate the angular velocity of an object, except in the case of Earth, where a capital omega (Ω_e) will be used to emphasize its importance.

So what is the angular velocity of Earth?

Earth rotates once every 24 hours.

Since angular velocity is the angle the object rotates divided by the time to do so, Earth's angular velocity is one revolution per day or, using unit abbreviations, 1 rev/d.

> ✍ Notice that the rotation rate of Earth is slower than the hour hand of a clock (2 revolutions per day).

What about in units of revolutions per second?

In units of revolutions per second, that would be

$$\frac{1 \text{ rev}}{1 \text{ d}} = \frac{1 \text{ rev}}{(24 \text{ h}) \times (3600 \text{ s/h})}$$

---

[v] 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.

[vi] This value of π has been rounded. More precisely, it is the value of a circle's circumference to its diameter.

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 Earth's angular velocity in units of radians. There are 2π radians in one revolution.

In units of rad/s, Earth's angular velocity would be

$$(1.157 \times 10^{-5})\frac{1 \text{ rev}}{1 \text{ s}} = (1.157 \times 10^{-5})\frac{2\pi \text{ rad}}{1 \text{ s}}$$

or 7.272×10^{-5} rad/s.

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 (length of arc divided by length of radius). So, when used with other units we typically don't include the radian name.

Second, I will replace the "7" with a "9", so that the angular velocity is given as follows:

$$\Omega_{\text{e}} = 7.292 \times 10^{-5} \text{ s}^{-1}$$

This is the value that you'll find in appendix B.4.

Why is the value equal to 7.292×10^{-5} s^{-1} instead of 7.272×10^{-5} s^{-1}?

This small difference can be accounted for by recognizing that 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 Earth spinning on its axis.

Earth, however, does not stay in one place as it spins. Earth is also orbiting the Sun. We must take that motion into account. The time it takes for Earth to rotate once relative to the fixed stars is called the **sidereal day** (pronounced: sigh-DEE-ree-al). In problem 12.1, you will calculate the actual angular speed of Earth and show that it 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 Earth shares the same rotation rate (see problem 12.2).

Checkpoint 12.5: The rotation rate of 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?*

12.3.2 Angular velocity of the air

The purpose of identifying the angular velocity of 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 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 Earth, which is cyclonic (and, as we've seen above, is close to one 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 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 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 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 Earth).

Only then would motion toward a center result in greater anticyclonic motion, both relative to the stars and to Earth.

How fast is that?

As we've already seen, that is equal to one 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 one revolution per day, that would mean the air around its edge would have to travel 6280 km each day. In meters per second, that would be

$$\frac{6.28 \times 10^6 \text{ m}}{(24 \text{ h}) \times (3600 \text{ s/h})} = 72.92 \text{ m/s}$$

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 likely 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

$$\frac{1.57 \text{ m}}{(24 \text{ h}) \times (3600 \text{ s/h})} = 18 \times 10^{-6} \text{ m/s}$$

This speed is so slow that imperceptible motions are sufficient to create rotation rates around the drain equal in magnitude to 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.

Checkpoint 12.6: (a) If the air is calm, what is its angular velocity relative to the stars? (b) What is its angular velocity relative to Earth?

Project

Project 12.1: Observe a surface map with isobars and wind (plotted as wind barbs). Roughly speaking, are the winds directed (a) away from high pressure regions and toward low pressure regions (across the isobars), (b) clockwise around high pressure regions and counter-clockwise around low pressure regions (parallel to the isobars), or (c) a combination of (a) and (b)?

Project 12.2: Observe a map of 500-mb heights with wind (plotted as wind barbs). Roughly speaking, are the winds directed (a) away from high heights and toward low heights (across the contours), (b) clockwise around high heights and counter-clockwise around low heights (parallel to the contours), or (c) a combination of (a) and (b)?

Questions

Question 12.1: Describe an observation you can make (here on Earth) other than the direction of winds around low pressure that supports the contention that Earth rotates counter-clockwise when looking down above the northern hemisphere. For example, consider the apparent movement of the stars and sun.

Question 12.2: (a) Why do we say that 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 map showing the winds in the southern hemisphere. Was your prediction correct?

Question 12.3: Suppose the wind at 300 mb (near the tropopause) is blowing from south to north.
(a) Does that mean the temperature at that level (300 mb) but west of the location is different than the temperature at that level east of the location? If so, which is colder and why?
(b) Does that mean the temperature at that level (300 mb) but south of the location is different than the temperature at that level north of the location? If so, which is colder and why?
(c) Does that mean the average temperature of the air column from the surface

to 300 mb but west of the location is different than the temperature at the air column east of the location? If so, which is colder and why?
(d) Does that mean the average temperature of the air column from the surface to 300 mb but south of the location is different than the temperature at the air column north of the location? If so, which is colder and why?

Problems

Problem 12.1: (a) Calculate the number of solar days in one full orbit around the Sun, taking into account that there is an extra day (beyond the normal 365 days) every leap year except for when the year is divisible by 100 and not also divisible by 400.[vii]
(b) The number found in (a) is the time it takes to go completely around the Sun. How many times does Earth rotate during that time? Hint: Keep in mind that Earth is also orbiting the Sun.
(c) The value in (a) is the time it takes for 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 12.2: Assuming Earth rotates once every 24 hours, determine the *<u>angular</u> speed (in radians per second) of a point on Earth's surface (relative to Earth's axis) on the equator, at 45°N, and at 85°N.*

[vii] If a leap year occurred every fourth year, without any exceptions, that would mean there would be 365.25 solar days in one full orbit. Given the exception every 100 years (except for years divisible by 400), your answer will be a little off from 365.25 solar days.

13. Wind Speed

Introduction

In chapter 12, we explained in a rather conceptual way why the wind flows in the direction it does. In this chapter, we'll use a more quantitative approach to predict the wind *speed*. To do so, we'll analyze the wind in terms of components, so I'll first discuss that before going into how to predict the wind speed.

13.1 Wind components

In your introductory physics class, you may remember learning that we can describe two-dimensional vectors in terms of components or in terms of magnitude and direction. So far, we've described the wind in terms of wind speed and direction. However, we can instead 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 eastward (the zero degree direction) and another northward (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[i].

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 east of the origin. A negative value corresponds to a position west of the origin.

Just as the three-dimensional position is given in terms of three components (x, y and z), the three-dimensional velocity is also given in terms of three components, which we'll indicate as u, v and w. These components correspond to dx/dt, dy/dt and dz/dt, respectively.

> ✍ As with position, a positive value of u corresponds to an eastward velocity and a positive value of v corresponds to a northward velocity. Consequently, for a north wind (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 $u\hat{\imath} + v\hat{\jmath}$.

> ✍ 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.

---

[i] 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 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 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.

Checkpoint 13.1: 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? What about* u*?*

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 eastward and increasing counter-clockwise, a wind velocity of magnitude V and direction θ, $V\cos\theta$ would be equivalent to the component directed eastward and $V\sin\theta$ would be equivalent to the component directed northward:

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

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

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

Checkpoint 13.2: For an angle of 60°, and without a calculator, which is greater: the cosine of 60° or the sine?

Describing winds in terms of components is not only useful for writing expressions for the wind speed (see next section) but it turns out that numerical models (see section 1.3) must use equations in component form, since they represent the atmosphere in terms of a grid. This is why numerical models

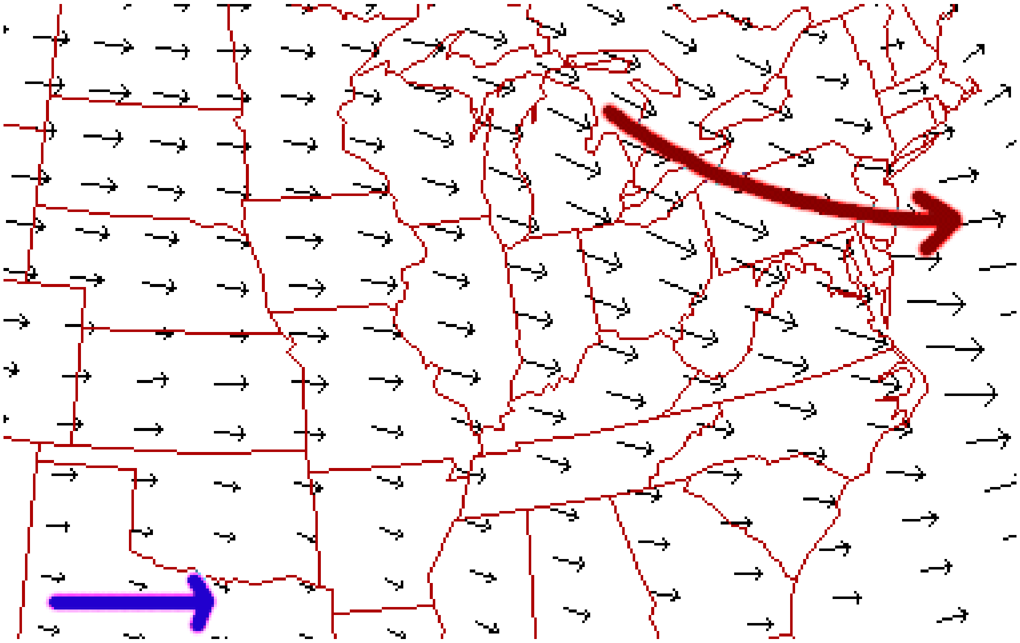

Figure 13.1: A vector wind plot. (source: University of Illinois WW2010 ProjectDataStreme)

produce wind maps using **wind vectors** plotted at regular intervals along a grid, as shown shown in Figure 13.1.

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). 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.

Checkpoint 13.3: In Figure 13.1, toward which direction is the wind blowing over Kansas?

13.2 Geostrophic balance

Now that we are familiar with the practice of describing wind in terms of components, let's 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**, which describes what happens when there is no friction and the air is not accelerating horizontally (meaning it is not only staying the same speed but is also not turning). This is close to the same condition explored earlier, where the wind flows parallel to the isobars at upper levels.[ii]

The geostrophic balance relationship states that when there is no friction and the isobars are straight then the horizontal wind flows parallel to the isobars with a speed proportional to the pressure gradient as follows:

$$V_{\mathrm{g}} = \frac{1}{f\rho}\nabla P \tag{13.1}$$

where f is a parameter known as the **Coriolis parameter**, which we will discuss later. You should already be familiar with the other variables, ρ and ∇P, which are the density and pressure gradient, respectively. 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 13.1 come from?
- How should equation 13.1 be modified if there is friction or 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 horizontal 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 account that Earth is spinning via "extra" force

---

[ii] If the air is accelerating (either speeding up, slowing down or curving), we can get a qualitative sense of what the wind speed is likely to be by comparing how it would be different than the geostrophic solution.

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 (see appendix D for derivation):

$$\frac{dv}{dt} = -\frac{1}{\rho}\frac{\partial P}{\partial y} - F_y - (2\Omega_e u \sin\phi) - \frac{u^2 \tan\phi}{a} - \frac{vw}{a}$$

This looks complicated but let's identify the terms.

First, notice that this expression is just for the rate of change of v, which is the northward or y component of the wind. As discussed in the previous section, u represents the eastward or x component of the wind. The expression for du/dt looks similar and the differences are not crucial to our discussion at the moment.

Consequently, the left-hand side (dv/dt) represents the acceleration of the northward component of the wind. 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. However, if we ignore that term then our expression will not provide an accurate prediction of the wind near Earth's surface. In particular, it will over-estimate the wind speed near Earth's surface.

The remaining terms on the right-hand side are all a consequence of the spinning Earth. As discussed in chapter 12, 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 Earth, in the denominator). As mentioned above, we will also ignore friction. This results in the following simplification:

$$\frac{dv}{dt} = -\frac{1}{\rho}\frac{\partial P}{\partial y} - (2\Omega_{\mathrm{e}} u \sin\phi)$$

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 of 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 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:

$$(2\Omega_{\mathrm{e}} u \sin\phi) = -\frac{1}{\rho}\frac{\partial P}{\partial y} \tag{13.2}$$

This assumption is called **geostrophic balance**.

Our final steps are to replace $2\Omega \sin\phi$ with f, the Coriolis parameter,

$$f = 2\Omega_{\mathrm{e}} \sin\phi. \tag{13.3}$$

and solve for the wind speed u, which is the eastward component of the wind. There will also be an expression for v, which is the northward component of the wind. Combining the two expressions gives the geostrophic balance expression (equation 13.1), which is written in vector form rather than just for one component.

> ✍ I think f is used to indicate the Coriolis parameter because it has units of frequency. Keep in mind that it is not a force, a force per mass or a force per volume. We have been using F for those.

Equation 13.2 has a negative sign but equation 13.1 does not. Why?

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.[iii]

[iii] 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 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.

WHAT HAPPENS IF THE WIND IS MOVING AROUND A CENTER (LIKE A LOW PRESSURE CENTER)?

We've assumed no acceleration ($dv/dt = 0$). 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).

Checkpoint 13.4: Does the Coriolis parameter depend upon latitude? If so, is it greater near the poles or near the equator?

13.3 Determining the geostrophic wind speed

To obtain the geostrophic wind speed, we use equation 13.1:

$$V_{\mathrm{g}} = \frac{1}{f\rho}\nabla P$$

This means we must 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 13.1: 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 13.1: First we calculate the Coriolis parameter from equation 13.3:

$$\begin{aligned} f &= 2\Omega_{\mathrm{e}} \sin\phi \\ &= 2(7.292 \times 10^{-5}\ \mathrm{s}^{-1}) \sin 40^\circ \end{aligned}$$

which gives a value of $9.4 \times 10^{-5}\ \mathrm{s}^{-1}$.

Assuming a density of $1.22\ \mathrm{kg/m}^3$, we apply equation 13.1 to get

$$\begin{aligned} V_{\mathrm{g}} &= \frac{1}{f\rho}\nabla P \\ &= \frac{1}{(9.4 \times 10^{-5}\ \mathrm{s}^{-1})(1.22\ \mathrm{kg/m}^3)} \frac{(4 \times 10^2\ \mathrm{Pa})}{(300 \times 10^3\ \mathrm{m})} \end{aligned}$$

which gives a value of 11.7 m/s.

> ✍ For mid-latitudes (30 to 50 degrees), the Coriolis parameter varies from about $7 \times 10^{-5}\ \mathrm{s}^{-1}$ to about $11 \times 10^{-5}\ \mathrm{s}^{-1}$. At 43 degrees latitude, the Coriolis parameter has a value close to $10^{-4}\ \mathrm{s}^{-1}$. In problem 13.2, you will determine the value of the Coriolis parameter at your location (using equation 13.3) but it should be around $10^{-4}\ \mathrm{s}^{-1}$.

Checkpoint 13.5: If the pressure gradient increases, what happens to the geostrophic wind speed?

13.4 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 13.1).

Unfortunately, it is difficult to calculate the geostrophic wind from upper-level maps using equation 13.1:

$$V_{\mathrm{g}} = \frac{1}{f\rho}\nabla P$$

This expression requires us to first calculate pressure gradient ∇P. However, upper-level maps are typically drawn for a constant pressure level, like the 300 mb map shown in Figure 13.2.footnoteThe 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 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 9.1. Before computers, this would be a long, involved process. The pressure is the *same* everywhere on a constant pressure map, so we can't determine the pressure gradient from the map.

What do the lines on a constant-pressure map correspond to?

On constant-pressure maps, they plot the height where the particular pressure is observed. The lines are contours (i.e., constant height) and indicate the height of the pressure level.

> ✍ Lines of constant height are called **contours**. Lines of constant pressure are called **isobars**.

How does a constant pressure map differ from a constant height map?

Assuming the same approximate level is being plotted, the overall pattern shown by contours on a constant pressure map will be the same as that shown by isobars on a constant height map. In other words, the "lows" are in the same location and the "highs" are in the same location (see footnote[iv] for why).

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 13.1 in terms of the height gradient $\nabla_p Z$ instead of the pressure gradient ∇P:

$$V_{\mathrm{g}} = \frac{g}{f}\nabla_p Z \tag{13.4}$$

[iv] This is because pressure decreases with height. For example, consider a situation where 300 mb is located at height H over location A and at height $H + \Delta H$ over location B. In other words, 300 mb occurs at a lower height over location A. Since pressure decreases with height, then the pressure must be less than 300 mb at height $H + \Delta H$ over location A. In other words, at the same height, the pressure is lower over location A.

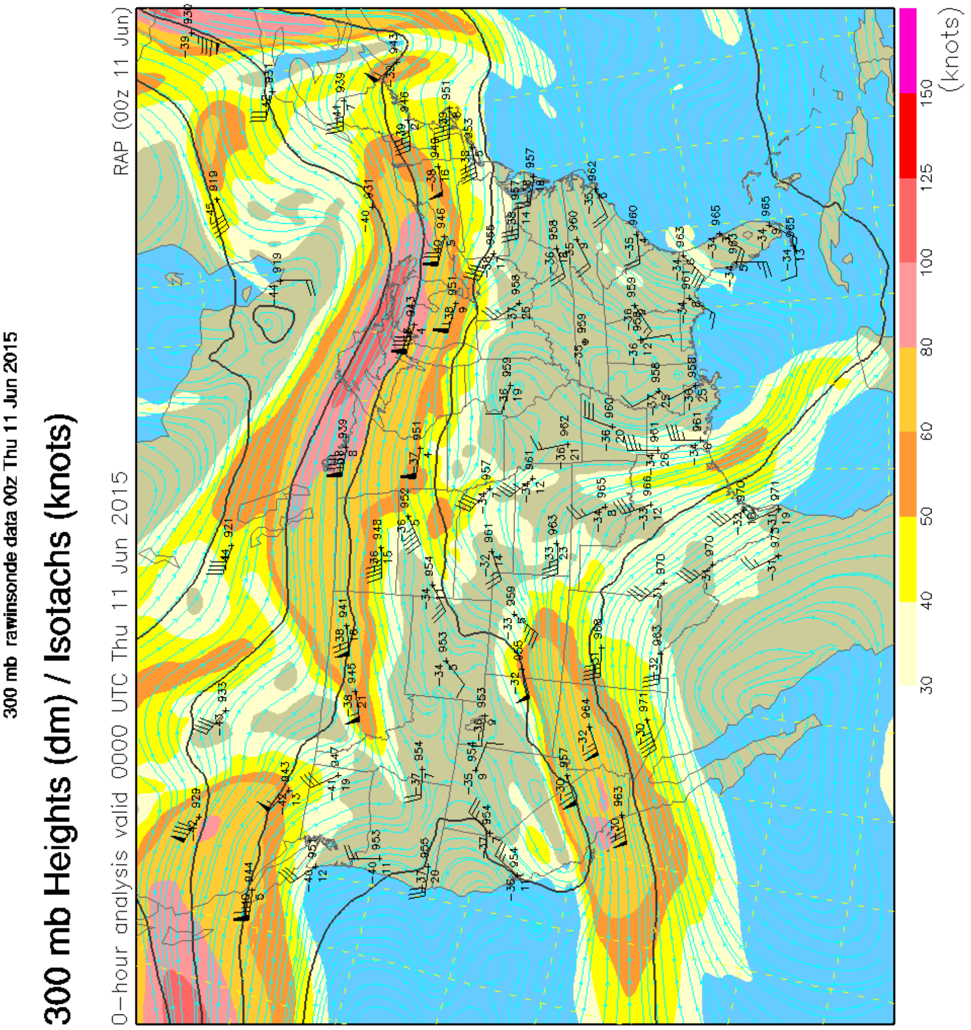

Figure 13.2: A map of 300-mb heights (solid lines), streamlines (light blue) and wind speeds (shaded). Source: UCAR.

Where does equation 13.4 come from?

To obtain equation 13.4, I replaced the pressure gradient ∇P in equation 13.1

$$V_{\mathrm{g}} = \frac{1}{f\rho}\nabla P$$

with $\rho g \nabla_p Z$.

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 9.1.

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[v], we know that

$$\left.\frac{\partial P}{\partial x}\right|_z = \left.\frac{\partial P}{\partial z}\right|_x \left.\frac{\partial z}{\partial x}\right|_P \tag{13.5}$$

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 13.4 instead of equation 13.1.

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.

---

[v] If you don't recognize this from your calculus class, ask your instructor to derive it.

Checkpoint 13.6: 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)?

13.5 Gradient winds

WHAT HAPPENS IF THE AIR ENCOUNTERS A REGION WITH A DIFFERENT 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 13.1 doesn't apply.

Recall that geostrophic balance means that there is a balance between the pressure gradient force and the Coriolis force. Consequently, the winds aren't accelerating.

IF THE PRESSURE GRADIENT FORCE DOESN'T BALANCE THE CORIOLIS FORCE DOES THAT MEAN THE AIR IS SPEEDING UP OR SLOWING DOWN?

Not necessarily. If there are other forces present, we can still have a balance between the forces but it just won't be geostrophic balance.[vi] Such a situation exists when there is friction, which we'll examine in section 13.6. In that case, there will still be a balance except that three forces will be involved: friction, pressure gradient and Coriolis.

If there are no other forces acting and the Coriolis force does not exactly balance the pressure gradient force, then the air must accelerate. Just because the air accelerates, though, does not mean it slows down or speeds up. It could just change direction.

For example, consider the case of air circulating around a low or high pressure center. The net force is not zero. However, the wind speed doesn't change. Rather, the wind direction does.

If we have a low or high pressure center, the winds will go in a circle around the pressure center. Even if the wind speed is constant, we know that circular

---

[vi] If the winds aren't geostrophic then we say that the winds **ageostrophic**.

motion involves an acceleration (since the direction is changing) and so this would not be a case of geostrophic balance. Because of that, we do not expect the wind speed to equal the geostrophic wind speed when we examine the circulation around a low pressure center (see, for example, project 13.1).

In this situation, the wind is called the **gradient wind** and will not have the same speed as it would have in the geostrophic case.

13.5.1 Flow around a low pressure center

Let's first consider the scenario when air moves around a center of low pressure. Since the air is moving in a circle, it must be accelerating toward the center of the circle. As such, the net force must equal the centripetal force given by equation 13.6:

$$f_{\mathrm{c}} = \frac{v^2}{r} \tag{13.6}$$

where v is the speed of the object and r is the radius of the circular motion.[vii] Consequently, even with no other forces acting, the pressure gradient force and the Coriolis force do not balance.

For air flow around a low pressure area, the pressure gradient force is directed inward. The Coriolis force must be directed outward. The difference must equal the centripetal force.

Let's suppose that we are east of the low center in the northern hemisphere. The pressure gradient is equal to $(1/\rho)\partial P/\partial x$ and is directed toward the west. The Coriolis force (assuming the winds are northward) is fv and is directed eastward. Take the difference and setting it equal to the centripetal force (per mass), we get

$$\frac{v^2}{r} = \frac{1}{\rho}\frac{\partial P}{\partial x} - fv.$$

Before solving for v, let's divide through by f,

$$\frac{v^2}{fr} = \frac{1}{f\rho}\frac{\partial P}{\partial x} - v$$

[vii] Equation 13.6 can be obtained by recognizing that $\vec{v}$ is a vector whose direction changes at a rate of 2π radians every period T. If we rewrite $\vec{v}$ as $v\hat{v}$, where v is the magnitude and $\hat{v}$ is the direction, then from the definition of acceleration we have $\vec{a} = d\vec{v}/dt = v(d\hat{v}/dt) = 2\pi v/T$. Since v is equal to the circumference $(2\pi r)$ divided by the period T, we can replace $2\pi/T$ by v/r.

and then use the value of the geostrophic wind (equation 13.1) to replace the first term on the right by v_{g}. We then end up with

$$\frac{v^2}{fr} = v_{\mathrm{g}} - v$$

or

$$v_{\mathrm{low}} = v_{\mathrm{g}} - \frac{v^2}{fr}. \tag{13.7}$$

This means that the wind speed around a low is smaller than what would be predicted by geostrophic balance.

> ✍ The wind speed will be lower than geostrophic when upward motion is advecting lighter winds from lower in the atmosphere. That this occurs east of upper-level troughs (of low pressure) is consistent with the analysis above.

13.5.2 Flow around a high pressure center

Let's now suppose that we are east of a *high* center in the northern hemisphere. The net force should once again be directed toward the west. However, the pressure gradient is directed toward the east this time. And, the Coriolis force is directed toward the west. Again taking the difference and setting it equal to the centripetal force (per mass), we get

$$\frac{v^2}{r} = fv - \frac{1}{\rho}\frac{\partial P}{\partial x}.$$

Dividing through by f and using the value of the geostrophic wind (equation 13.1) to replace the second term on the right by v_{g}, we end up with

$$\frac{v^2}{fr} = v - v_{\mathrm{g}}$$

or

$$v_{\mathrm{high}} = v_{\mathrm{g}} + \frac{v^2}{fr}. \tag{13.8}$$

This means that the wind speed around a high is larger than what would be predicted by geostrophic balance.

Checkpoint 13.7: For gradient flow around a high pressure system, which is larger: the Coriolis force or the pressure gradient force?

13.6 Frictional force (viscosity)

Don't be surprised if the observations are still inconsistent with the analysis presented so far. It turns out that there may be an additional force acting on the air that we need to take into account. This force is due to the air's **viscosity**.

Viscosity describes how a fluid interacts with itself. A lightly viscous fluid flows more easily than a highly viscous fluid. Peanut butter has a higher viscosity than molasses, which has a higher viscosity than oil, which has a higher viscosity than air.

Viscosity by itself doesn't necessarily slow down a parcel. Rather, it depends on what the air around the parcel is doing. If the parcel is surrounded by slower air, the parcel will be slowed down and the wind speed will likely be less than geostrophic. On the other hand, if the parcel is surrounded by faster air, the parcel will be sped up and the wind speed will likely be greater than geostrophic.

For example, let's consider the wind at the surface. This wind is typically measured at a height of 2 meters above the ground. At the ground, the wind speed can be considered to be zero[viii]. Consequently, the wind at 2 meters will be slowed down due to its proximity to the stationary air just below it. The greater the viscosity, the greater the effect.

The wind speed gradient is called the **wind shear**.

The vertical wind shear can act to slow down the air (as just discussed) or it can act to speed it up. For example, the wind above 2 meters typically is faster than the wind at 2 meters. Consequently, the wind shear above 2 meters is acting to speed of the wind.

So what happens? If the wind below is less and the wind above

---

[viii] This is called the "no slip" assumption.

IS GREATER, DOES THAT ACT TO SPEED UP THE WIND, SLOW DOWN THE WIND, OR NEITHER?

It depends on which wind shear is greater. The effect depends on the vertical **gradient** of the wind shear.

Near the ground, the wind shear typically decreases with height. This means that for an air parcel near the ground, the air below it is moving so slow that it is acting to slow down the air parcel more than the air above it is acting to speed it up. Consequently, viscosity acts to slow down the air near the ground. This effect is called **surface friction** because, like friction, it acts to slow down the air.

Whereas in geostrophic balance the wind is directed parallel to the isobars, with friction the wind tends to be directed more toward the region of lower pressure. This is because the speed is slower than geostrophic and so the Coriolis effect is less.

> ✍ There is also an viscosity effect due horizontal gradients in the horizontal wind shear. However, horizontal gradients of the horizontal wind shear are small and so we have neglected it.

Checkpoint 13.8: If friction is present, which way does the air blow: somewhat toward low pressure, somewhat toward high pressure, or parallel to the isobars?

Project

Project 13.1: (a) Observe a map of 500-mb heights with wind (plotted as wind barbs; select "conterminous US" as the region to plot). Choose a wind observation on the map where the winds are straight and uniform. What is the speed (in m/s) at that location?
(b) Measure the height gradient at that location.
(c) Using equation 13.4, determine the geostrophic wind speed corresponding to the height gradient calculated in part (b). How does that correspond to the observed wind speed?

Questions

Question 13.1: Should the height pattern on a 500-mb map be the same as the pressure pattern on an 18,000-ft map? In other words, are high heights in the same general region as high pressures? Why or why not?

Question 13.2: 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.

Problems

Problem 13.1: (a) Determine the speed of a point on Earth's surface (relative to Earth's axis) for a point on each of the following latitudes: on the equator, at 45° N, and at 88° N.
(b) One assumption made in the geostrophic wind equation is that the wind speed is negligible compared to Earth's speed at that location. Is this an valid assumption for the three locations identified in part (a)? If so, which ones and why? If not, why not? You'll have to compare the values in (a) with a typical wind speed.

Problem 13.2: Calculate the value of the Coriolis parameter, f, for your location.

Problem 13.3: (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 problem 5.1).
(b) Indicate the orientation of the wind vector in degrees according to the (i) meteorological system described on page 72, where 0° indicates wind from the north, with the angle increasing clockwise around the circle, and the (ii) mathematical system described in this chapter.
(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.

Derivations

Show-me 13.1: In section 13.2, the geostrophic wind relationship (equation 13.1),

$$V_{\mathrm{g}} = \frac{1}{f\rho}\nabla P$$

was derived from the dv/dt *expression (equation D.2) in Appendix D. In that appendix, a similar expression is provided for the* du/dt *component (see equation D.1):*

$$\frac{du}{dt} = -\frac{1}{\rho}\frac{\partial P}{\partial x} - F_x - 2\Omega w\cos\phi + 2\Omega v\sin\phi - \frac{uw}{a} + \frac{uv\tan\phi}{a}$$

Show how the geostrophic relationship (equation 13.1) can be derived from the du/dt *expression. Be sure to specify the assumptions that go into each step. In other words, if you drop a term from the expression because it is negligible, explain why that term would be negligible.*

Part D

Moisture

14. Humidity

Introduction

Using what we know so far, we can make a pretty good numerical model, which can reliably simulate the genesis and decay of storms and the motion of fronts. However, we are missing one key ingredient: moisture. Without moisture, the storms and fronts won't generate any clouds or precipitation.

To address the deficiency we examine moisture in this part of the book. Before getting to how moisture is involved in the development of clouds and precipitation, we'll first examine how we quantify the amount of moisture that is present. Then, in chapter 15 we examine how clouds and precipitation are formed. Finally, in chapter 16 we'll examine how the condensation of water vapor helps strengthen storms by releasing latent heat.

But first we need to quantify the amount of moisture in the air.

14.1 Water and its states

Whereas clouds and precipitation consist of solid or liquid water (see chapter 15, the moisture we are trying to quantify in this chapter is 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 **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.

Checkpoint 14.1: Can one "see" water vapor?

14.2 Dew point

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 measured by cooling the air. 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 more water vapor that is present, the less you have to cool the air to get condensation, and the higher the temperature at which condensation occurs and the stickier it feels.

The condensation temperature 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.[iii]

> ✍ In chapter 16, 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.

The dew point only depends on the amount of water vapor that is present. So, as water vapor condenses, the dew point decreases (since condensation removes water vapor and changes it into liquid water).

Consequently, if the temperature continues to cool, more and more water vapor condenses, making the dew point decrease as well. Basically, once the

---

[iii] This typically happens during the night as the air cools due to radiation (see chapter 7) and the dew is seen in the morning on plants and other surfaces.

temperature reaches the dew point, any further cooling of the temperature will necessarily lower the dew point along with the temperature.

> Note that, if the temperature warms, the dew point doesn't then necessarily rise as well. The dew point only rises if the water vapor increases, and that requires evaporation into the air. If cloud droplets are present then those droplets would evaporate, making the dew point rise, but without that source of liquid water, the dew point wouldn't rise.

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 in figure 14.1, where the observed dew point is indicated by the green squiggly line to the left of the red squiggly line that represents the observed temperature. 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.

Checkpoint 14.2: According to the skew-T log-P diagram on page 243, what is the dew point at 900 mb?

14.3 Fractional content

While the dew point is a straightforward way of indicating how humid the air is, it doesn't directly tell us the *amount* of water vapor that is present. After all, dew point is measured in degrees, which is unit of temperature, not an *amount* of water vapor. For this reason, we tend to use other measures of the water vapor content when needing to do mathematical calculations involving water content (like predicting whether it will rain).

There are two general approaches. One approach, which will be discussed in the next section, is to indicate the vapor content in terms of the pressure it exerts.

In this section, we examine how to indicate the vapor content in terms of the fraction of the air that is water vapor.

Before continuing, we should recognize that even under the most humid conditions, water vapor makes up a tiny fraction of the total air, maybe 2-3% at most. If you think about it, that makes sense, since there is more than

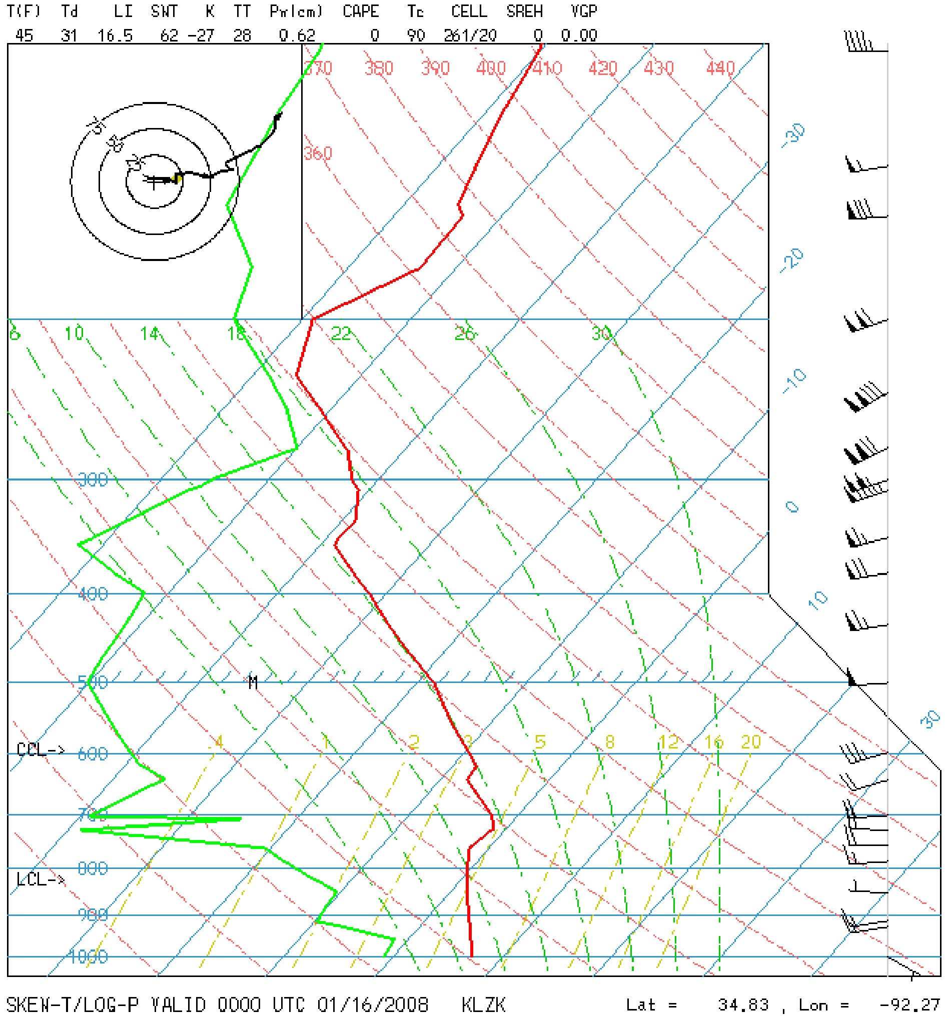

Figure 14.1: A skew-T log-P graph for Little Rock, Arkansas. Data is from 0000 UTC, 16 Jan, 2008 (local time would be 6 PM the previous day). Horizontal solid lines are pressure (in mb). Skewed, solid lines are temperature (in °C).

enough oxygen even in the most humid conditions. That is because oxygen still makes up roughly 20% of the air, with nitrogen making up most of the rest.

14.3.1 Specific humidity

The **specific humidity**, indicated by q, is the fraction of the air that is water vapor. Mathematically, we can write it as follows:

$$q = \frac{m_{\mathrm{v}}}{m_{\mathrm{T}}}. \tag{14.1}$$

where m_{v} and m_{T} represent the mass of the vapor and the total air content, respectively.

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 it remains the same as a parcel expands or contracts (so the ratio of water vapor to the total air remains the same) as long as the water vapor content doesn't change (meaning no moisture condenses or evaporates). 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 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.

✍ We usually don't measure the specific humidity directly. Rather, we obtain it the dew point value or the vapor pressure value (see, for example, Show-me 14.1).

Checkpoint 14.3: Why is the specific humidity typically written with units of g/kg?

14.3.2 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 (represented by r) is the ratio of water vapor to the *dry air only*:[iv]

$$r = \frac{m_{\mathrm{v}}}{m_{\mathrm{d}}} \tag{14.2}$$

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:

$$q = \frac{r}{1+r} \tag{14.3}$$

Typically, $r \ll 1$ so the denominator is approximately 1 and $q \sim r$.

WHERE DID EQUATION 14.3 COME FROM?

Equation 14.3 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

$$\begin{aligned} \frac{1}{q} &= \frac{m_{\mathrm{v}} + m_{\mathrm{d}}}{m_{\mathrm{v}}} \\ &= 1 + \frac{1}{r} \\ &= \frac{r+1}{r} \end{aligned} \tag{14.4}$$

Simply invert to get equation 14.3.

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

Why would one use mixing ratio over the specific humidity?

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. However, their values are very similar and both have the property of remaining the same for a parcel undergoing expansion and compression if no water changes phase, so in practice there really isn't any advantage of one over the other.

Checkpoint 14.4: 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 14.1 on page 243). 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.

Checkpoint 14.5: Answer the following questions based on the skew-T log-P chart on page 243.

(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 dew point was equal to the observed temperature value?

14.4 Vapor pressure

In this section, we examine the approach where the vapor content is expressed in terms of the pressure it exerts. The greater the amount of water vapor, the greater the pressure that it exerts. As we will see, this has the advantage of allowing us to use the ideal gas law with water vapor.

Does the water vapor exert the same pressure as the air?

No. The water vapor pressure (which we typically just call the **vapor pressure**) is the pressure of the water molecules alone, and since the water vapor is a small fraction of the air, the vapor pressure is likewise a small fraction of the total air pressure. The vapor pressure would be equal to 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.

Indeed, whereas the air pressure near Earth's surface is about 1000 mb, the pressure of the water vapor is only 25 mb or so at most, even when it is really humid, consistent with the fact that water vapor hardly ever makes up more than 3-4% of the air (see previous section).

Checkpoint 14.6: Which is typically greater: the vapor pressure or the pressure exerted by oxygen?

14.4.1 vapor pressure and dew point

Since the vapor pressure, like the dew point, represents the water vapor content, it should be no surprise that the higher the dew point, the greater the vapor pressure and the stickier it feels. Indeed, there is a one-to-one correspondence between the dew point and the vapor pressure, and we can obtain one when we know the other.

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.

✍ While a table is sufficient for our purposes, an equation is best for a numerical model. The simplest is the **Tetens equation**. [v] The Tetens equation 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.

Checkpoint 14.7: According to a vapor pressure table, what vapor pressure corresponds to a dew point of 20° C?

14.4.2 Ideal gas law

In chapter 2, we introduced the ideal gas law, which can be written as follows (equation 2.1):

$$P = \rho R T$$

where P, ρ and T are the pressure, density and temperature, respectively. The R represents the **gas constant**. As was mentioned in chapter 3, different gases have different values of R because the value depends on the composition of the gas. For dry air (air without any water vapor in it), R has a value of $287.06\ \mathrm{J \cdot kg^{-1} K^{-1}}$. More specifically, the gas constant is the universal gas constant divided by the molar mass (equation 3.1):

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

where R^* is equal to $8.3145\ \mathrm{J \cdot mol^{-1} K^{-1}}$ and μ is equal to 0.02896443 kg/mol and 0.01801528 kg/mol for dry air and water vapor, respectively.

Now that we are considering the presence of water vapor, the value of R will be different (since μ will be different). After all, to use the ideal gas law, we need the value of R for *all* of the air, not just the *dry* air.

It turns out that the gas constant for air will be close to the value for dry air alone, since water vapor makes up only a small portion of the air. However, 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. What we need is to derive an expression that gives us the value of R as a function of the vapor pressure.

To do this, we first 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 ideal gas.

---

[v] Named after Otto Tetens, a German meteorologist who published the equation in 1930. Sometimes referred to as Tetens' formula.

Since the water vapor pressure is different than the air pressure, we'll use different variable abbreviations for the pressure of each. In particular, whereas we used P_d for the pressure of the dry air, we'll use e to represent the vapor pressure.[vi]

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

Checkpoint 14.8: What letter do we use for the variable abbreviation of vapor pressure?

Using P_d and e for the pressure of the dry air and the water vapor, respectively, let's write the ideal gas law three times: once for dry air, once for water vapor, and once for the total amount of air:

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

where R_d and R_v are the gas constants for dry air and water vapor, respectively, and P and R are the pressure and gas constant, respectively, for the total air, including water vapor.

Why doesn't T have subscripts to distinguish between the temperature in the three cases?

We will assume that the water vapor and the rest of the air are in thermal equilibrium with each other, so they have the same temperature.

Does the vapor pressure depend upon the temperature then?

No. I admit that the ideal gas law seems to suggest it does, but it doesn't. The vapor pressure only changes if you add or subtract water vapor. If you raise or lower the temperature without adding or subtracting water vapor then the vapor pressure will remain the same.

The reason is that the atmosphere isn't a closed container. If the temperature decreases, then the air compresses (i.e., the density goes up) to maintain

---

[vi] 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.

the same pressure as the surroundings. So, as long as the pressure of the surroundings don't change then the vapor pressure won't change either.

Still, the ideal gas law still applies. Just remember that the density is not constant.

In any event, our task is to combine the three ideal gas law expressions, to get an expression for R, the gas constant for the total air. The challenge of solving for R is a relatively straightforward one, mathematically. Basically, we have three independent expressions. As long as we have no more than three unknowns (one of which is R), we can solve for any of the unknowns.

We'll assume we can measure the total pressure P, the vapor pressure e and the temperature T. We also know the values of R_d and R_v (since we know the molar masses for dry air and water vapor, respectively). We are thus left with five unknowns: P_d, R and the three densities.

These are too many unknowns. We need two more relationships, involving the same variables.

The first additional relationship we'll use 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:

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

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:

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

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:

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

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

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

By the way, the ratio R_d/R_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:

$$\epsilon = \frac{R_\text{d}}{R_\text{v}} \approx 0.622 \tag{14.5}$$

If you want, you can replace every occurrence of R_d/R_v with ϵ. Thus, the equation for the density ρ can be written as

$$\rho = \frac{P - e + \epsilon e}{R_\text{d}T} \tag{14.6}$$

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

$$R = R_\mathrm{d} \frac{P}{P - e + \epsilon e} \tag{14.7}$$

> ✍ 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.

Checkpoint 14.9: 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?

14.5 Relative humidity and saturation

While the quantities discussed so far are used by meteorologists, non-meteorologists tend to use the **relative humidity**.

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 equaled the temperature. Thus, 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.

In other words, the saturation vapor pressure depends upon the air temperature. The warmer the air temperature, the greater the saturation vapor pressure.

To find the saturation vapor pressure, use the vapor pressure table, but find the vapor pressure that corresponds to the air temperature, not the dew point.

> ✍ This assumes a flat water surface. Curved water surfaces require a higher vapor pressure for saturation than that indicated in the table.

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

$$\mathrm{RH} = \frac{e}{e_\mathrm{s}} \tag{14.8}$$

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.

Checkpoint 14.10: (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?

14.6 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 reason is that the temperature, 100°C, corresponds to the boiling point of water and the pressure, 101.325 kPa (1013.25 mb), corresponds to the standard sea-level pressure.

This isn't just a coincidence. To understand what is going on, 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 (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 the saturation vapor pressure equals 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.

Checkpoint 14.11: (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

Project 14.1: Download the saturation vapor pressures (in kPa) and import the values into a spreadsheet program like Excel. Create a graph of the saturation vapor pressure over water with respect to temperature (for the tem-

peratures for which values are available). The curve that results is called the **saturation curve***.*

Questions

Question 14.1: 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.

Question 14.2: Which is more likely, an air pressure of 1000 mb or a vapor pressure of 1000 mb, or are they equally likely? Why?

Question 14.3: 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 condensation and form a cloud, the dew point must be greater than the temperature. However, the dew point curve on the skew-T never gets to the right of the temperature curve, even when there are clouds. Explain.

Problems

Problem 14.1: Suppose the air pressure was 1000 mb but half of the molecules present in the air were water molecules. What would the vapor pressure be?

Problem 14.2: (a) Suppose a bubble forms inside a liquid via evaporation. Using the information from project 14.1, for what temperature will the pressure inside the bubbles (the saturation vapor pressure) be equal to standard sea-level pressure (see appendix B.2)?
(b) Based on your answer to part (a), what is the boiling point in a place like New York City, where the elevation is near sea-level?
(c) Given the answers in (a) and (b), what should be the boiling point in a place like Denver, where the surface pressure is about 850 mb? Explain.
(d) One way to cook a "hard boiled" egg is to place the egg in water, bring the water to boiling and then turn off the burner and let the egg sit in the water

for several minutes. At which location (Denver or New York City) would you need to let the egg sit in the water longer? Explain.

Problem 14.3: (a) Give an estimate for the temperature of the water vapor is the room you are in.
(b) Suppose an endless source of liquid water was present in the room. What would be the vapor pressure (the pressure of the water vapor in the room) once equilibrium is reached (meaning as much vapor evaporating from the liquid as condensing into the liquid)?

Problem 14.4: Using the information downloaded in Project 14.1, estimate the dew point when the vapor pressure is 5 mb and the air temperature is 20°C. Explain your reasoning.

Problem 14.5: If the dew point equals the temperature, what is the relative humidity?

Problem 14.6: (a) Identify the observed dew point and air temperature at the surface according to the skew-T in Figure 14.1 (on page 243).
(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.
(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 (divide one by the other). 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 14.7: 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.

Problem 14.8: As one can show (see Show-me 14.1), 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 14.9: Suppose we want to calculate the density of the air on a day when the vapor pressure is 6 mb, the total air pressure is 1013 mb and the *air temperature is 15°C. We'd like to use the ideal gas law but we only know the gas constants for pure water vapor (461.52 J/kg/K) and for dry air (287.06 J/kg/K), not a mixture of the two (as we have in this case). There are two ways we can do this, one method outlined by steps (a), (b) and (c) below, and the other method outlined by steps (d) and (e) below. When doing this do not round intermediate steps, and report your answers to the thousandth place.*

(a) Use the ideal gas law to calculate the density of just the water vapor portion of the air. [vii]

(b) Use the ideal gas law to calculate the density of just the non-water vapor portion of the air (the dry air part). [viii]

(c) Based upon your answers to parts (a) and (b), what is the density of the air as a whole (add the two values obtained so far)?

(d) Calculate the value of the gas constant for the total air via equation 14.7.[ix]

(e) Use the ideal gas law to calculate the density of the air as a whole, using the gas constant found in part (d) along with the temperature and pressure of the air as a whole.

(f) Compare the values obtained from the two methods (parts c and e). Do they agree? Should they agree?

Derivations

Show-me 14.1: 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.*

[vii] Since you are applying the gas law to water vapor only, remember that the values of pressure, temperature and R must be appropriate for water vapor, not dry air.

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

[ix] Alternatively, you could get the weighted average molar mass and then divide R^* by that. Notice that the average molar mass is less than that for dry air, since water vapor molecules are lighter than diatomic nitrogen and oxygen. That means the gas constant will be larger than what it is for dry air.

15. Clouds and Precipitation

Introduction

Clouds and precipitation are both made up of liquid or solid water. The major difference is that clouds consist of water suspended[i] in the air whereas **precipitation** is the water that has fallen to the ground. In this chapter we examine how each forms.

15.1 Cloud Formation

15.1.1 Cloud condensation nuclei

As we know from chapter 14, water vapor condenses when the temperature drops to the dew point. At Earth's surface, the water condenses onto plants and other objects, forming dew. What happens above 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 particles are called **Cloud Condensation Nuclei** (or **CCN**, for short).

CAN'T THE WATER VAPOR JUST CONDENSE ON ITS OWN WITHOUT THE HELP OF CLOUD CONDENSATION NUCLEI?

Actually, it is very hard to get water to condense on its own (for reasons that will be discussed in section 15.2). Not that you can tell, though, since there is usually no shortage of such CCN (the number and size will vary depending on location).

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

---

[i] In section 15.4 we'll examine how cloud droplets are suspended in the air.

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.[ii]

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 and can can grow big enough to impact visibility even if the relative humidity is less than 100%. This is the mechanism behind **smog** and **haze**.

Checkpoint 15.1: (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?

15.1.2 Drop sizes

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 15.4 and Show-me 15.2, 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.

[ii] You may wonder how they measure the number of CCN in the air. One way is via an **Aitken** (condensation) **nucleus counter**. Saturated air is expanded rapidly so that it becomes supersaturated by several hundred percent with respect to water. By expanding the air very rapidly, the air is cooled and the saturation vapor pressure decreases a lot. The vapor pressure is much higher than the new saturation vapor pressure and even the tiniest aerosol particles can grow. In this way, all of the nuclei become droplets and are easier to spot (and thus count). If the air is not expanded very quickly, the will still cool adiabatically but it won't reach such high supersaturations because the vapor will end up condensing only on the larger particles.

Checkpoint 15.2: 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)?

15.1.3 Lifting condensation level (LCL)

Since a cloud forms as soon as the air temperature cools to the dew point, we can use that to predict the height that a cloud will form if we raise the air. As you may recall from chapter 11, a dry parcel cools at the dry adiabatic rate of g/c_p, which is approximately 10°C/km, and we can use that to predict when the temperature reaches the dew point. Once the temperature reaches the dew point, condensation occurs and a cloud forms. The level at which this occurs is thus called the **lifting condensation level** or **LCL**, for short.

Does it still cool at the dry adiabatic rate if water vapor is present?

As long as there is no condensation, the parcel will still cool at the adiabatic rate, and condensation won't happen until the parcel temperature reaches the dew point.

Does the dew point remain the same as the parcel rises?

Unfortunately, no. The dew point depends on the vapor pressure. As the parcel rises, it expands and the vapor pressure decreases along with the total pressure. That means the dew point likewise decreases.

Then how do we know when saturation is reached?

Rather than focusing on the temperature and the dew point, it is easier to focus instead on the *mixing ratio*. As discussed in section 14.3.2, the mixing ratio represents the ratio of water vapor to dry air that is present. That ratio remains the same upon compression and expansion as long as no water vapor is added or removed. Consequently, until the LCL is reached, the mixing ratio of the parcel will remain the same.

For example, consider the skew-T log-P chart for Amarillo in figure 15.1 (repeated from page 178). The mixing ratio value is given by the dashed gold lines that go from 600 mb down to the bottom of the skew-T chart. The labels are just above the 600 mb line, and go from 0.4 g/kg to 20 g/kg.

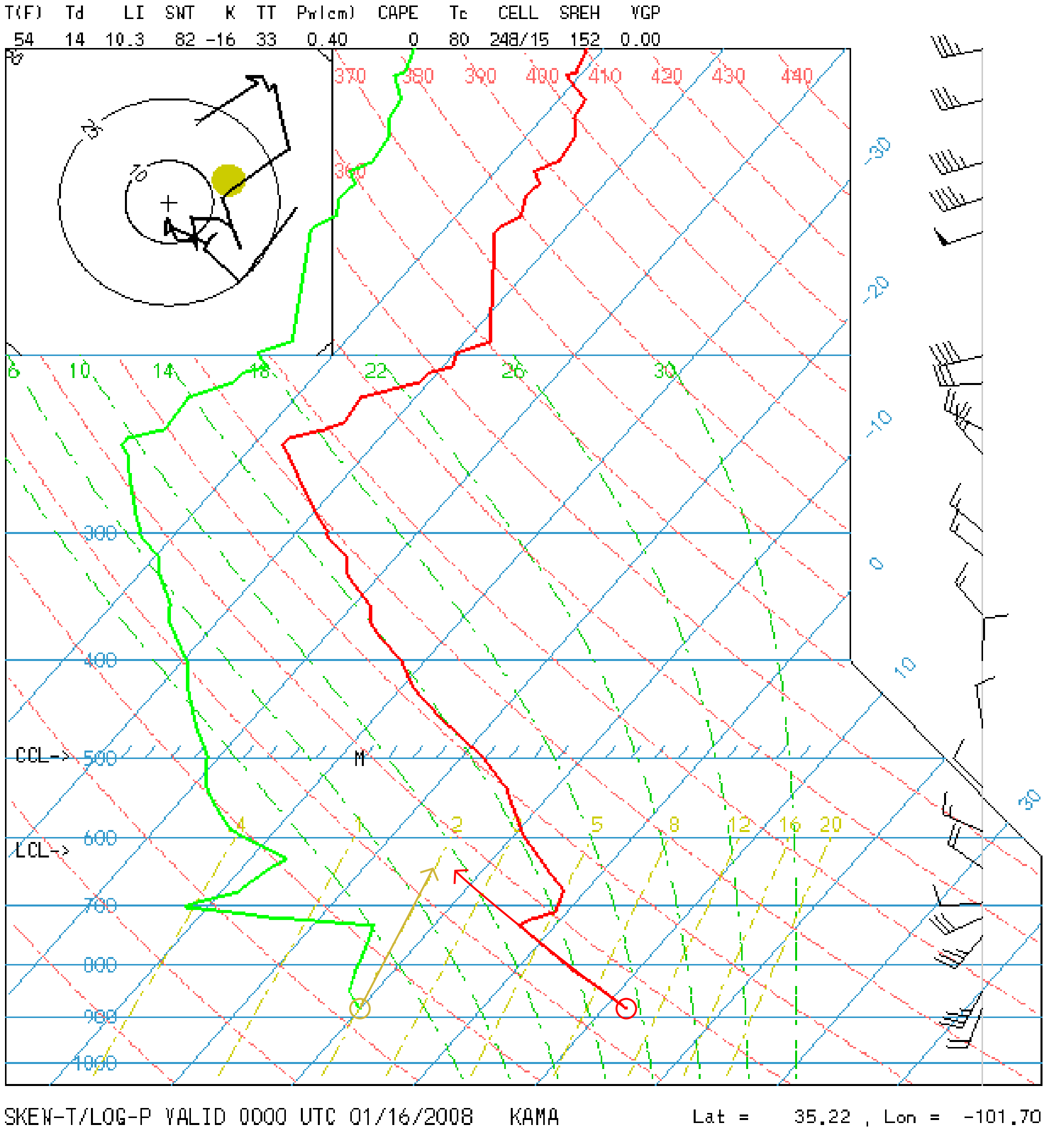

Figure 15.1: A skew-T log-P graph for Amarillo, Texas, repeated from Figure 10.1. Data is from 0000 UTC (7 PM local the previous day), 16 Jan, 2008. Horizontal solid lines are pressure (in mb). Skewed, solid lines are temperature (in °C). The small red and gold circles and the associated red and gold arrows are discussed in the text.

According to the diagram, the air at the surface has a dew point of −10°C, a mixing ratio of 2 g/kg and a temperature of around 12°C. The small gold circle points us to dew point temperature (slanted solid blue lines) and mixing ratio (slanted dashed gold lines). The small red circle points us to the temperature (slanted solid blue lines).

If we then raise the air, it expands and cools, but its mixing ratio remains the same (at 2 g/kg; follow the gold arrow). As you can see, the dew point doesn't stay at −10°C. Instead, the dew point cools slightly with height.

Meanwhile, the temperature cools at the dry adiabatic rate, as indicated by the red curves that slope to the left with height (follow the red arrow).

Saturation is reached where the two arrows intersect, which is about 620 mb. Any further lifting will result condensation. Via this process, then, we can estimate where the bottom of the cloud is, which is the location of the lifting condensation level (620 mb in this case).

✍ | According to the skew-T diagram, their estimate of the LCL (see left side of skew-T chart) is just above my estimated location.[iii]

Checkpoint 15.3: As an air parcel is raised, but before the LCL, what happens to the parcel's mixing ratio and potential temperature?

15.1.4 Saturation due to mixing

To form a cloud, we must cool the air until the air becomes saturated. One way to do this is to raise the air to the LCL. Another way is to mix the air with air that is colder.

To predict whether mixing will result in a cloud or not, we need to consider two effects. Mixing not only changes the parcel's temperature but it also changes its moisture content.

[iii] This particular skew-T log-P chart calculates the LCL using values based on uniform mixing within the lowest 110 mb, as this is somewhat more realistic (since the real atmosphere gets mixed somewhat during lifting). Since the potential temperature is pretty uniform throughout that layer (in figure 10.1), 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.

For example, consider two parcels: one warm and one cool. If both parcels are unsaturated, the cooler air will be drier and, consequently, mixing it with a warm, moist air parcel, while technically leading to cooling it, will also dry it out.

So can mixing a parcel with a colder one actually result in saturation?

Yes.

This is illustrated quite easily by exhaling on a cold day (see problem 15.2). As most people know, the result is a cloud (i.e., saturated air), even though both your breath and the outside air are not saturated prior to the mixing. A similar process occurs when jet exhaust mixes with the cold air outside the plane. The jet exhaust has a lot of moisture in it (due to the products of combustion) yet is not saturated.[iv] The hot, unsaturated, but very moist, exhaust mixes with the cold, unsaturated air outside. The result is air that is **supersaturated**. The water vapor condenses and produces what is called a **condensation trail** or **contrail**.

The best way to explain why this is possible is to go through an example.

Let's consider two parcels, one at 10°C and another at 20°C, with equal number of moles and at the same pressure. As mentioned in section 10.5 (and shown in Show-me 10.2), the temperature of the mixture is 15°C, the average of the two initial temperatures.

Now let's suppose that each parcel is saturated at 100% **relative humidity**. That means that the vapor pressures were 12.3 mb and 23.4 mb, respectively, prior to mixing (values obtained from the saturation vapor pressure data table – see project 14.1).

If the combined parcel had a vapor pressure equal to the average of these two (see Show-me 15.3), the vapor pressure would be 17.8 mb.

However, if you check the saturation vapor pressure curve, you find that 17.8 mb is about 0.8 mb *greater* than what is needed for 100% relative humidity. In other words, the relative humidity is now *greater* than what it was in *either* parcel prior to mixing. That means the result is a supersaturated parcel even though the initial two air parcels were not supersaturated.

[iv] If you look at the exhaust immediately after leaving the engine, whether it is the exhaust for a car, a power plant or an airplane, you will see that the exhaust is clear. It is not until the exhaust mixes with the outside air that a cloud is formed.

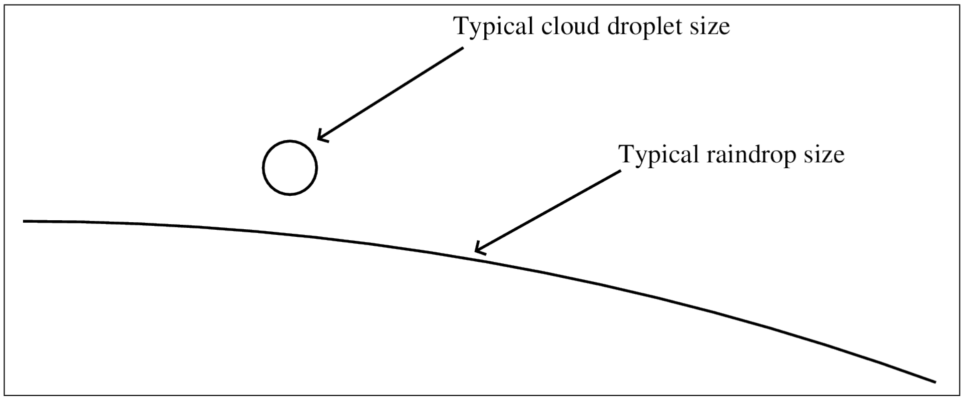

Figure 15.2: The relative sizes of a cloud droplet and a raindrop. Only a small portion of the raindrop is drawn, since it is so much larger than the cloud droplet.

✍ It is because the saturation mixing ratio curve is not linear that the relative humidity of the combined parcel can be greater than a straight average of the initial relative humidities.

Checkpoint 15.4: In the example provided above, what is the relative humidity of the initial parcels and the combined parcel?

15.2 Rain and snow formation

15.2.1 Rain

Raindrops are much bigger than cloud droplets. For example, raindrops range in size between 1 and 10 mm in diameter compared to 0.001 and 0.1 mm for cloud droplets. The dramatic size difference is illustrated in Figure 15.2.

While you might think that cloud droplets grow into raindrops via condensation of water vapor but that is not the case. It turns out that even if you condensed all of the water out of the air, the cloud droplets wouldn't grow to the size of raindrops. There are just too many CCN's in the atmosphere – distributing the moisture across all those CCN's means the droplets are too small (as shown in Show-me 15.4).

So, how do we get raindrops?

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 **collision and coalescence** and the other is called **diffusion**.

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

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). This has to do with the curvature of the droplet surface. The surface of the smaller cloud droplets are curved more than the surface of the larger droplets, a difference that is apparent in Figure 15.2, which illustrates the relative size of raindrops vs. cloud droplets.

The greater curvature means that there are more opportunities for a water molecule to evaporate compared to the volume of water present. This means the saturation vapor pressure is greater for smaller droplets.[v]

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

> ✍ As we will see in section 15.2.2, 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.

---

[v] This is also why the condensation of water into droplets requires the presence of cloud condensation nuclei and won't condense into droplets on its own (see section 15.1.1). It is also the reason why boiling requires imperfections in the container (called nucleation sites) or dissolved gases in order for the bubbles to form. Without them, bubbles won't form and the liquid can be superheated above the boiling point, at which point one risks the water "exploding" when a bubble finally does form.

Checkpoint 15.5: (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.

15.2.2 Snow

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**[vi].

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.

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?

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

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, and with only one ice nucleus per liter, they can get pretty large. The only thing stopping their growth is their weight. After a certain size, their fall speed (see section 15.4) becomes large enough that they fall out of the sky.

This **ice crystal process** is much more efficient than the **collision and coalescence** process described in section 15.2.1 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.[vii]

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

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

What if the temperature throughout the cloud is greater than freezing?

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

If the temperature is warmer than freezing throughout, the precipitation process is through collision and coalescence, since no ice crystals will exist at temperatures above freezing.

Actually, this is probably *still* the process as long as the entire cloud is warmer than −15°C, even though −15°C is below freezing.[viii] I say this because without freezing or contact nuclei, water droplets will not spontaneously freeze when the temperature is less than 0°C. Such colder-than-freezing liquid water droplets are called **supercooled** droplets.

In fact, for most cloud droplets, spontaneous freezing (called *homogeneous* or *spontaneous nucleation*) won't occur until the temperature falls to −35°C or below, and for raindrops, this won't occur until the temperature falls to −20°C or so.

Because of this, if the temperature is warmer than −15°C, the cloud likely consists of supercooled water droplets and only consists of ice crystals if the temperature is colder than −20°C.[ix]

Checkpoint 15.6: The temperature profile in the skew-T graph shown on page 243 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?

15.3 Other precipitation types

We've discussed how rain and snow forms. Snow forms when the cloud is made up of ice crystals, and the ice crystals grow or combine to form snow flakes, which then fall to the ground without melting. This can only occur if the falling snow flakes don't encounter any air where the temperature is above freezing.

[viii] A cloud that is liquid water throughout is called a **warm cloud** or **water cloud**.

[ix] Clouds that are entirely ice are called **ice-crystal clouds**. Clouds with a mix are called, unsurprisingly, **mixed clouds**.

If the falling snow flakes encounter air where the temperature is above freezing, the snow flakes will melt, forming liquid water droplets. Only liquid water can exist above freezing. The melting point of water is 0°C. That means that if the temperature is above 0°C then water is in the liquid form (or soon will be). This type of precipitation is called **rain**. Rain is made up of liquid water raindrops.

> ✍ If the entire atmosphere is above freezing except for a small layer near the ground, the snow flakes may only partially melt prior to hitting the ground.

In this section, we'll explore how other types of precipitation forms.

15.3.1 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 above the ground is below freezing, there can be a layer of air higher up that is 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?

As mentioned on page 269, 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.[x]

So water doesn't automatically freeze when less than 0°C?

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 particles in the air can also freeze water droplets on contact. Such particles 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 14.6).

So freezing can occur below 0°C. What about melting?

Ice won't melt if below 0°C. 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.

Checkpoint 15.7: 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?

---

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

15.3.2 Sleet

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

As mentioned above, pure water doesn't necessarily freeze when the temperature falls below 0°C. However, ice crystals are built upon **freezing nuclei** and so any melted snow flake or ice crystal already has an impurity that will allow it to freeze when the temperature falls below 0°C. A partially melted snow flake or ice crystal is even more likely to freeze if it falls below 0°C.

This causes the droplets to freeze into little spheres of ice before they hit the ground if they encounter a larger enough sub-freezing layer. 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.

WHAT IS THE DIFFERENCE BETWEEN SNOW AND SLEET?

Although both are the solid form of water, 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.

Checkpoint 15.8: What is the difference between freezing rain and sleet?

15.3.3 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.

Checkpoint 15.9: What is the difference between sleet and hail?

15.4 Fall speeds

15.4.1 Cloud droplets

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 the **fall speed** (or **terminal velocity**) .

Because of the low fall speed 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.

To find how much upward motion is needed, we first have to find the fall speed. At the fall speed, the gravitational force balances the air drag. One can show (see Show-me 15.1) that the fall speed v depends on the square of the droplet radius r as follows:

$$v = \frac{2}{9}\frac{r^2 \rho_{\text{water}} g}{\eta} \tag{15.1}$$

where g is 9.8 N/kg and η is the viscosity of air (1.84×10^{-5} Pa·s at 20°C according to the Gas Viscosity Calculator).

As stated before, typical cloud droplets have a diameter between 1 and 100 μm. Plugging in an average diameter of 10 μm (a radius of 5 μm) and a water density of 1000 kg/m^3, I get a fall speed of about 3 mm/s.[xi]

An upward speed of 3 mm/s is easily within the range of possible upward motions in the atmosphere, as seen in a plot of vertical velocity at 700 mb (colors indicate vertical motion; units are in mb/s but a mb/s is roughly equal to a cm/s).

[xi] For the range of diameters from 1 μm to 100 μm, I get a fall speed range from 0.03 mm/s to 30 cm/s. It is a large range because the radius is squared in the expression.

Checkpoint 15.10: When something is falling at its fall speed, what is the net force exerted on it: upward, downward or zero?

15.4.2 Precipitation

Because a raindrop is larger, there is a larger gravitational force on the raindrop. This allows the raindrop to attain a larger fall speed.[xii] For example, the fall speed of raindrops is about 5 to 10 m/s whereas the fall speed of cloud droplets is around 3 mm/s (as discussed in the previous section).

The fall speed for a raindrop can be obtained in much the same way as with cloud droplets. Namely, at the fall speed, the gravitational force balances the air drag. However, for raindrops the air drag tends to be proportional the square of the speed rather than being linear with speed (as with cloud droplets). Consequently, the fall speed for raindrops v depends on the square root of the droplet radius r as follows (see Show-me 15.1):

$$v = \sqrt{\frac{8rg}{3C}\left(\frac{\rho_{\text{water}}}{\rho_{\text{air}}}\right)} \tag{15.2}$$

where g is 9.8 N/kg and C is the drag coefficient (typically between 0.5 and 1 for sea-level temperature and pressure).

As stated before, typical rain drops have a diameter between 1 and 10 mm. Plugging in an average diameter of 5 mm (a radius of 2.5 mm), an average drag coefficient of 0.75, and densities of 1.2 kg/m^3 and 1000 kg/m^3 for air and water, respectively, I get a fall speed around 8.5 m/s.[xiii]

As shown on a plot[xiv] of vertical velocity at 700 mb, it is unlikely[xv] that

---

[xii] 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 (proportional to the cube of the radius vs. the square of the radius or linear with radius).

[xiii] For the range of diameters from 1 mm to 10 mm, I get a fall speed range from 3.8 m/s to 12 m/s.

[xiv] Colors indicate vertical motion; units are in mb/s but a mb/s is roughly equal to a cm/s.

[xv] Upward velocities of order 10 m/s or greater are sometimes seen in strong thunderstorms. Indeed, such updrafts are necessary for the formation of hail.

the vertical velocity will be 10 m/s and thus raindrop-size drops will not be suspended and instead will fall to the ground.

✍| Equations 15.1 and 15.2 give the same fall speed for a radius of 128 μm.

Checkpoint 15.11: If the air is experiencing upward motion at 50 cm/s, would rain drops likely fall, rise or neither? What about cloud droplets? Why?

Project

Project 15.1: Obtain a skew-T log-P chart for a location where the dew point and temperature differ at the surface 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?

Questions

Question 15.1: Describe the process whereupon sleet, snow, and hail are produced.

Question 15.2: Why are CCN necessary to produce clouds? In other words, what is it about CCN that allows clouds to form that would not form otherwise?

Question 15.3: Cirrus clouds are very high in the troposphere. What particles are cirrus clouds made of and how did these particles form? Consider the environment of cirrus clouds for your answer.

Question 15.4: Describe the process whereupon rain is produced in (a) a water cloud (liquid water throughout) and (b) a mixed cloud (some ice cystals mixed in with liquid water droplets).

Question 15.5: (a) From a table of saturation vapor pressures, there is a range of temperatures (between $-15^{\circ}C$ *and* $0^{\circ}C$*) where the saturation vapor pressure over water and the saturation vapor pressure over ice are both provided. Over this temperature range, 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?*
(b) At the temperature indicated in part (a), 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.
(c) At the temperature indicated in part (a) and the vapor pressure indicated in part (b), would a water droplet grow (due to condensation) or would it diminish (due to evaporation)?
(d) At the temperature indicated in part (a) and the vapor pressure indicated in part (b), would an ice crystal grow (due to deposition) or would it diminish (due to sublimation)? Provide the reasoning behind your answer.

Question 15.6: In section 15.1.4, it was shown that a cloud can produced by mixing two air parcels even if neither parcel, by itself, is saturated. Which of the following are utilized in the example?
(a) The average of the two temperatures (breath and surrounding air) is the average of the individual temperatures (of breath and of surrounding air).
(b) The average of the two vapor pressures (breath and surrounding air) is the average of the individual vapor pressures (of breath and of surrounding air).
(c) The average of the two relative humidities (breath and surrounding air) is the average of the individual relative humidities (of breath and of surrounding air).

Problems

Problem 15.1: A typical cloud droplet has a radius of 10 μ*m and a typical rain drop has a radius of 1 mm. Based on this, how many typical cloud droplets are required to make a typical rain droplet? Note: compare volumes.*

Problem 15.2: Following the process discussed in section 15.1.4, show that you can create a cloud by exhaling on a cold day when the temperature is

0°Feven if the relative humidity is 0%. You'll need to estimate the temperature of your breath (it need not be exact) and the vapor pressure of your breath (you can assume it is 90% of the saturation value so that it is almost, but not quite, saturated). Be explicit about how your calculations show that a cloud can result.

Problem 15.3: Estimate a height of a typical cloud above the ground. How long does it take for (a) a typical cloud drop and (b) a typical rain drop to fall from that height assuming no vertical motion of the atmosphere?

Problem 15.4: In section 15.1.1, 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:*

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

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.*

(a) 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.*

(a) Repeat the calculation for when the mixing ratio decreases by 8 g/kg (eight times the value in part a). This would correspond to condensing all of the water vapor present when the temperature of air at 1000 mb is 10°C and the relative humidity is 100% (as can be confirmed from a skew-T diagram at 1000 mb).

(c) Compare your answer in (b) to the size of a typical raindrop. What does your answer imply about the likelihood that we can get raindrop-sized droplets from condensing water onto CCN without any other mechanism (like collision and coalescence, or diffusion)?

Derivations

Show-me 15.1: Derive equations 15.1 and 15.2 by setting the gravitational force mg *equal to the air drag* F_{drag}*, where*

$$\begin{aligned} F_{\text{drag}} &= 6\pi\eta r v && \text{(for small drops)} \\ F_{\text{drag}} &= \tfrac{1}{2}\rho_{\text{air}} A C v^2 && \text{(for large drops)} \end{aligned}$$

Assume spherical drops (so the volume is $\frac{4}{3}\pi r^3$ *and the cross-sectional area* A *is* πr^2*).*

Show-me 15.2: Derive the relationship in Problem 15.4. 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 15.3: Show that the average vapor pressure of two parcels is equal to the vapor pressure that results when the two parcels of identical number of moles and identical pressures are mixed together. Use the $PV = nR^*T$ *formulation of the ideal gas law and the results of Show-me 10.2.*

16. Latent Heating and Cooling

Introduction

As mentioned in 14, condensation of water vapor helps strengthen storms by releasing latent heat.

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 liquid evaporates, it absorbs energy from the surroundings, cooling it off. Conversely, when a vapor condenses, it releases energy to the surroundings, warming it up.

In this chapter, we examine this process of **latent heating**.

16.1 Latent heat

As stated in the introduction, 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 (to the environment) 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.

Checkpoint 16.1: During evaporation, the bonds between the liquid molecules are broken. When a bond is broken, as during the transformation from liquid to vapor, is energy released to the environment (warming it up) or is it absorbed from the environment (cooling it down)? Based on your answer, does evaporation warm up the air does it cool down the air?

16.1.1 Cooling examples

Sweating

As long as the air is unsaturated, any liquid water that is present will evaporate. As mentioned above, evaporation leads 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).

Checkpoint 16.2: 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

An example of latent cooling is seen when you have one hand wet and one hand dry and you wave them both in the air. Even if the water is the same temperature as the air, the wet hand will feel cooler.[i]

This is because the water evaporates into the air, cooling your hand. The drier the air, the cooler your hand will feel.

In fact, one way of determining the moisture content of the air is to use a **sling psychrometer**, which is basically just two thermometers waved in the air, with one thermometer being dry and the other with its bulb wrapped in a wet cloth. The temperature of the wet thermometer is called the **wet-bulb temperature**[ii], for obvious reasons, and the cooler the wet-bulb temperature, relative to the air temperature (measured by the dry thermometer), the drier the air must be.

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).

Checkpoint 16.3: Which is higher: the dew point or the wet-bulb temperature?

16.1.2 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. More often, 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 6, rising motion causes the air to cool. Such cooling can lead to condensation

---

[i] A similar thing is seen when you use a blow drier to send air through a moist towel or piece of fabric. Even if the blow drier is producing hot air, the air on the other side of the fabric will be relatively cool because of the evaporation of the water from the fabric.

[ii] According to the *Glossary of Weather and Climate*, a dash is used for wet-bulb temperature and wet-bulb thermometer, although "wet bulb" can be used in place of either one.

(and clouds) and this introduces latent heating that we cannot ignore. This is examined in the next section.

16.2 Moist adiabatic lapse rate

As discussed in chapter 14, condensation occurs when the air cools to its dew point and, as we know from chapter 6, 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 10.1), we'll call the non-saturated rate the **dry adiabatic** lapse rate (indicated as Γ_{d}).

> ✍ 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 15). In that case, the process would no longer be adiabatic and would more appropriately referred to as non-adiabatic or diabatic.

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.

Checkpoint 16.4: Suppose we have air at 1000 mb and 0°C with a vapor *pressure of 6 mb. If the air is then lifted, it will cool. What would be the temperature of the air after being lifted 1 km: greater than* $-10°C$*, equal to* $-10°C$ *or less than* $-10°C$*?*

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, it is easier to just identify the lapse rate via the **skew-T log-P** chart.

On a skew-T log-P diagram (see, for example, page 178) 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 curves (i.e., dry adiabatic lapse rate) near the top of the chart.

These represent the moist adiabatic lapse rate curves.[iii] 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).

> ✍ For Amarillo, Texas, as shown in figure 10.1 on page 178, the observed temperature between 620 mb and 450 mb just happens to be roughly parallel to the moist adiabatic lapse rate curves.

Why do the moist adiabatic lapse rate curves match the dry adiabatic lapse rate curves at low pressures (i.e., in the upper part of the skew-T diagram)?

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

---

[iii] Meteorologists tend to call the dry and moist adiabatic lapse rate curves the dry and moist **adiabats**.

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).

✍ This is the mechanism behind the **rain** shadow effect, where dry, hot air is located on one side of a mountain chain. As the air moves up one side, it cools and the water condenses. That water precipitates out of air at the top of the mountain. As the air then descends on the other side of the mountain it warms adiabatically, ending up warmer than when it started because of the release of latent heat. This is explored in problem 16.5.

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 degrees 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.

Checkpoint 16.5: Suppose we have air at 1000 mb and 0°C with a vapor pressure equal to the saturation value at that temperature (i.e., 6 mb). The air is then lifted 1 km.

(a) What happens to the parcel's temperature as it rises: increase, decrease or remain the same?

(b) What happens to the parcel's potential temperature as it rises: increase, decrease or remain the same?

(c) What happens to the parcel's equivalent potential temperature as it rises: increase, decrease or remain the same?

16.3 Stability

The presence of water vapor and the potential for condensation affects the stability of the atmosphere.

Recall from chapter 11 that a parcel will rise if it is warmer than the surrounding air (and descend if it is cooler). To illustrate, consider the sounding from Omaha, Nebraska, on 1200 UTC 20 May 2004, as plotted on the skew-T

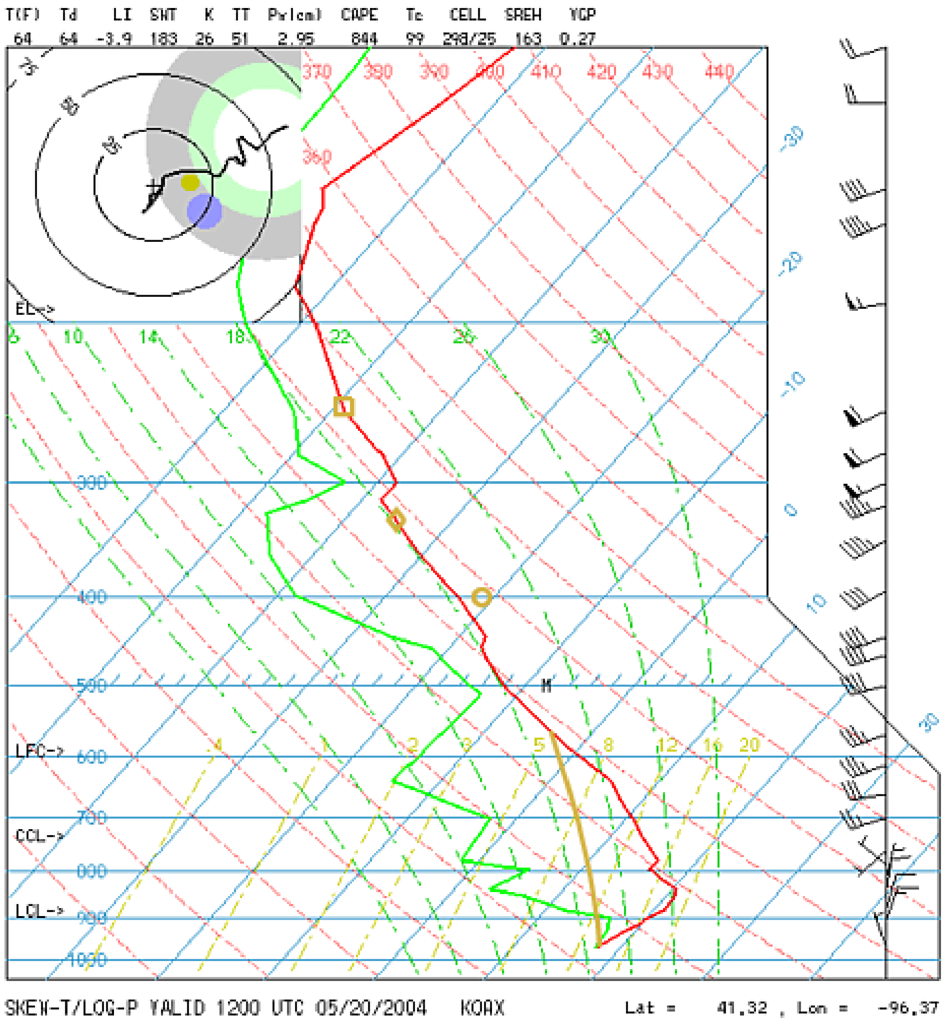

Figure 16.1: A skew-T log-P graph for Omaha, Nebraska, repeated from figure 11.3. Data is from 1200 UTC (7 AM Central Daylight Time), 20 May, 2004. The gold circle, diamond, square and curve are discussed in the text.

log-P chart in figure 16.1. We had considered this skew-T back in chapter 11, where we discussed how a −21°C parcel introduced at 400 mb (see thick gold circle on the diagram) would rise since it is warmer than the environment at that level.

Back in chapter 11, we assumed the parcel was unsaturated, which meant that it cooled at dry adiabatic rate, represented by the dashed red lines on the skew-T diagram, leading to a temperature around −35°C at around 330 mb, where it happens to be in equilibrium with the environment.

Let's now consider what happens if that same −21°C parcel is saturated.

If it is saturated then as it rises it cools according to the moist adiabatic lapse rate, not the dry adibatic rate, and its equivalent potential temperature remains the same. In this case, the equivalent potential temperature is about +19.5°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 250 mb, which indicates the equilibrium level for this particular parcel.

16.3.1 Level of free convection

During the previous discussion (on page 191), it was mentioned that we can get individual parcels to be warmer than the environment by forcing them to ascend. It turns out that if a parcel is totally dry, it is very hard to get the parcel to become warmer than the environment, since the dry adiabatic rate is 10°C/km and the atmospheric lapse rate is typically about 6.5°C/km. Consequently, ascending parcels will tend to cool so quickly that they will remain cooler than the environment.

If the parcel is saturated, however, it becomes much easier to cool slowly enough that at some point it becomes warmer than the environment.

For example, consider the air at the surface in Omaha, Nebraska, as shown in the skew-T of figure 16.1. 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 570 mb or so, the air we've pushed up remains colder than the environment. At 570 mb, though, it now equals the air that is there at 570 mb. And, any further lifting results in it being *warmer* than the environment air at those altitudes. If we release the air parcel at a pressure less than 570 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 −21°C temperature at 400 mb that was explored earlier (gold diamond in the skew-T chart) and eventually to the 250 mb equilibrium level (gold square in the skew-T chart).[iv]

Note that the LCL, LFC and EL abbreviations on the skew-T indicate the Lifting Condensation Level (LCL), the Level of Free Convection (LFC) and the Equilibrium Level (EL) of the raised surface air. Those levels, indicated on the skew-T are not too far off from what we determined above.

WHY ARE THEY NOT EXACTLY THE SAME AS WHAT WE DETERMINED?

There levels are a little different than the levels obtained by us because the group that produced the graphs first assume that the lowest levels are mixed before rising, as that is more realistic than just having the surface air rise without mixing. Mixing the air increases the surface temperature, though, since the air above the surface has a higher potential temperature (see section 10.5). The mixing ratio would remain roughly the same in this case, though, 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), consistent what is indicated on the skew-T diagram.

How the mixing impacts the LFC and EL depends on the situation. In the case shown in figure 16.1, the mixing ratio is the same as it was before the mixing and, as the parcel rises, that mixing ratio remains the same as the parcel cools. Given the way the mixing ratio line (dashed gold) is tilted, that means the parcel is warmer at the LCL than it would be otherwise. That means it doesn't have to rise as far to become warmer than the environment (lower LFC) and stays warmer longer with additional lifting (higher EL). The lower LFC and higher EL (in this case) are consistent with what is shown on the skew-T diagram.

Checkpoint 16.6: Not all soundings have a level of free convection. Given the sounding for Omaha, Nebraska, shown in figure 16.1, 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%).

[iv] This is just because I had chosen the initial parcel characteristics with the intention of matching this situation.

16.3.2 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?

As one might guess, latent heating will tend to make things more unstable.

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 (see section 15.1.3), 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).

Checkpoint 16.7: Does having water vapor in the air make the air more stable or less stable?

16.4 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 kilogram). According to this relationship, the more mass that changes state, the more energy that is released/absorbed.

> ✍ The value varies somewhat with temperature and thus it is technically more appropriate to 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.

16.4.1 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. The only difference is that it is released to the environment in one case (condensation) and it is absorbed from the environment in the other (evaporation).

Checkpoint 16.8: Why doesn't the latent heat of vaporization have units of heat?

16.4.2 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).

$$L_\mathrm{s} \sim 2.83 \times 10^6 \ \mathrm{J/kg}.$$

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 used that property to explain snow formation in section 15.2.2.

Checkpoint 16.9: 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?

16.4.3 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**.[v]

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.

$$L_\mathrm{f} \sim 3.35 \times 10^5 \text{ J/kg}.$$

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):

$$L_\mathrm{s} = L_\mathrm{v} + L_\mathrm{f}.$$

This should make sense, since it should take more energy to break the strong attractions of the solid (latent heat of sublimation) than the weaker attractions of the liquid (latent heat of evaporation), with the difference being the energy needed to go from the solid to the liquid bonds (latent heat of fusion). The difference is not great (only 335,000 J/kg), since the bonds between the water molecules in the solid phase (2.83 million J/kg) are only a little bit stronger than the bonds between the water molecules in the liquid phase (2.50 million J/kg).

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 maintain the same temperature during the change of state.

For example, consider a raindrop falling through a layer that is below freezing (see discussion on sleet in section 15.3). As the raindrop falls through the

---

[v] 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.

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 16.1 shows that if kept insulated, only a fraction of the drop will freeze and we wouldn't get sleet (see section 15.3.2).

Checkpoint 16.10: 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?

16.4.4 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 6.5)

$$\frac{dQ}{m_{\mathrm{T}}} = c_{\mathrm{p}} dT - \frac{dP}{\rho}$$

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.

> ✍ Since our focus is on the heating and cooling of air, we'll use the value of c_{p} for air, which is 1004.67 $\mathrm{J/(kg \cdot K)}$. The mass m corresponds to the mass of air that is heated or cooled. However, the energy can also go into heating or cooling water. In that case, we would have to use the value of c for water, which is 4160 $\mathrm{J/(kg \cdot K)}$. Sometimes the specific heat of water is given in units of calories and grams, which is 1 $\mathrm{cal/(g \cdot {}^{\circ}C)}$.

As mentioned in the previous section, if some water evaporates or condenses then the heat absorbed or released is 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.

Before moving on, it is important to recognize that we need to distinguish between the water that evaporates or condenses and the air that warms up as a result of the latent heating or cooling. I'm using m_v for the former and m_T for the latter.

With this in mind, we can plug the above relationship for dQ into the first law of thermodynamics to get the following:

$$\frac{L\ dm_\mathrm{v}}{m_\mathrm{T}} = c_\mathrm{p} dT - \frac{dP}{\rho}$$

or

$$dT = \frac{1}{c_\mathrm{p}} \frac{L\ dm_\mathrm{v}}{m_\mathrm{T}}. \tag{16.1}$$

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 16.1 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 14.3.1).[vi]

Consequently, equation 16.1 becomes

$$dT = \frac{L\ dq}{c_\mathrm{p}} \tag{16.2}$$

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.

[vi] Again, keep in mind that we are dealing with the air, since the air is warming/cooling. So, c_p is for air.

Checkpoint 16.11: In equation 16.2, 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?*

Project

Project 16.1: Obtain 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 project 15.1. 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).

Questions

Question 16.1: (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? Why would this be?

Question 16.2: (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.

Problems

Problem 16.1: 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 $L_s = L_v + L_f$.

Problem 16.2: Given the sounding for Omaha, Nebraska, shown in figure 16.1 (on page 286), 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 16.3: (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 16.4: What happens to the parcel's mixing ratio and potential temperature (increase, decrease or stay the same) as the parcel is raised above the lifting condensation level? Provide your reasoning.

Problem 16.5: A north-south oriented mountain range is experiencing westerly winds. In the valley on the west side, the pressure is 1000 mb and the temperature T *and dew point* T_d *are both* $20°$ *C.*
(a) Use a skew-T graph to estimate T *and* T_d *at the top of the mountain range where the pressure is 800 mb.*
(b) Where, if at all, do you expect rain? Assume any condensation associated with a decrease in dew point immediately falls out of the air as rain)
(c) Use a skew-T graph to estimate T *and* T_d *in the valley on the east side, after passing over the mountain.*
(d) Why is it both warmer and drier in the valley downwind of the mountain range than on the upstream side of the mountain range? This is an effect called the **rain shadow effect**.

Problem 16.6: 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: There are several ways to this. One way is to first get the amount of water that condenses out then figure out the temperature that results from condensing that amount of moisture. You could assume a reasonable temperature then use the skew-T *to get the mixing ratio at saturation, as an approximate value for the specific humidity, or you could use the equation in Show-me 14.1 and then equation 14.3. An alternate way is to follow the moist adiabat up on the skew-T, as though the parcel was being raised really high and cooling adiabatically, then follow dry adiabat down on the skew-T, as though the parcel was being lowered back to the surface and warming adiabatically.*

Derivations

Show-me 16.1: Suppose a blob of water has a temperature of -20°C when 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).

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 the Tetens equation
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	Earth's radius
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 Earth's surface)

t	time
T	period
T	temperature
T_1	a constant used in the Tetens equation
T_2	a constant used in the Tetens equation
T_a	temperature of the atmosphere
T_e	Earth's surface temperature

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

A.3 Metric prefixes

Prefixes are then used to represent factors of ten for each unit (see table A.1). For example, the “kilo” represents 10^3. Consequently, a kilometer is equivalent to a thousand meters and a kilogram is equivalent to a thousand grams.

Some notes to keep in mind:

- You can use the table, rather than memorize all the prefixes, but you

Abbreviation	**Name**	**Quantity**
Y	Yotta-	10^{24}
Z	Zetta-	10^{21}
E	Exa-	10^{18}
P	Peta-	10^{15}
T	Tera-	10^{12}
G	Giga-	10^{9}
M	Mega-	10^{6}
k	kilo-	10^{3}
h	hecto-	10^{2}
da	deka-	10^{1}
d	deci-	10^{-1}
c	centi-	10^{-2}
m	milli-	10^{-3}
μ	micro-	10^{-6}
n	nano-	10^{-9}
p	pico-	10^{-12}
f	femto-	10^{-15}
a	atto-	10^{-18}
z	zepto-	10^{-21}
y	yocto-	10^{-24}

Table A.1: Metric prefixes (http://physics.nist.gov/cuu/Units/prefixes.html)

should become familiar with the most commonly used prefixes, such as Mega-, kilo-, centi-, milli- and micro-.

- All prefixes are indicated by a single Roman letter except for micro-, which is represented by the Greek letter "mu"[i], and deca-, which is represented by two letters (da).
- Many people get centi and milli confused. If it helps, consider that the centi- prefix comes from the same root as "century" (100 years) and the milli- prefix comes from the same root as "mile" (1000 paces) and "millennium" (1000 years).

---

[i] Some people use the Roman letter u because it is easier to write.

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 $\cdot$ 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}
Planck's constant (h)	$6.62607015 \times 10^{-34}$ J·s
Boltzmann constant (k_B)	1.380649×10^{-23} J/K

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 $\mathrm{kg/m}^3$
specific heat of liquid water at 0°C	4218 J kg^{-1} K^{-1}
specific heat of ice at 0°C	2106 J kg^{-1} K^{-1}

B.4 Astronomical Properties

Earth's angular velocity	$7.2921151467 \times 10^{-5}\mathrm{s}^{-1}$
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 Earth radius	6.371×10^6 m
mean Sun radius	6.96×10^8 m
Earth mass	5.9736×10^{24} kg
Sun mass	1.9891×10^{30} kg
solar constant	1361 $\mathrm{W/m}^2$

C. Simulating the diurnal cycle

This appendix describes how to use Excel to simulate how the temperature evolves over the course of 24 hours, 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.

The Excel spreadsheet will utilize seven columns. The first row contains the titles of each column. Each row below the header contains the information for one minute during the day, starting at midnight and going to the following midnight.

C.1 Time

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”. This can be done by clicking on or highlighting the entire column and then choose Format/Cells. Click on the Number tab, and choose the Time category and the “1:30 PM” type.

Set the next cell (A3) in the column as equal to 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.

Copy the information in cell A3 to as many rows as necessary to get a record for every minute of the day up to midnight.

C.2 Temperature

Column B contains the temperatures at various times throughout the day, so title the second column “Temp” (in cell B1).

We'll start with a temperature of 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.

Set the first cell (B2) as 282 (Kelvin).

Eventually you will determine the change in temperature for each minute and put that information in column G. For the time being, that column is blank. However, in anticipation of that column being populated, set the next cell (B3) in the "Temp" column as the previous cell (B2) plus the change in temperature calculated for that minute (G2).

Copy the information in cell B3 to as many rows as was calculated for the time. It should automatically update the references to column G.

C.3 Solar irradiance

Title the third column "Fsolar" (in cell C1) to represent the radiation from the Sun absorbed at Earth's surface.

We'll assume that there is no radiation from the Sun during the night (from 6 pm to 6 am) and we'll assume the solar radiation during the day goes from zero at sunrise to a maximum of 494 W/m^2 at noon and then back to zero at sunset.

The value of 494 W/m^2 is obtained by recognizing the following:

1. The solar constant (1361 W/m^2) is the solar irradiance incident upon Earth outside of the atmosphere. By the time the radiation reaches Earth's surface, we need to subtract out the reflected amount (i.e., due to Earth's albedo) and the amount absorbed by the atmosphere.
2. This value is the maximum that occurs when the Sun is directly overhead. As mentioned on page 135, on average across the surface of the globe the value is one-quarter of this, or 340 W/m^2.
3. As mentioned on page 131, about 30.6% is reflected. It turns out that an additional 33% of the remaining is absorbed by the atmosphere. That leaves 157 W/m^2.
4. In order to get that as the average, the maximum during the day must be π times that average, which gives the 494 W/m^2 number.

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

$$494 \sin\left(\frac{(H-6+M/60)\pi}{12}\right)$$

where H and M stand for the hour and minute of cell A2. The Excel commands are "HOUR" and "MINUTE". So, when you type the command into cell C2, remember to use the "=" sign at the beginning, use an asterisk for any multiplication, use "HOUR(A2)" and "MINUTE(A2)" for the hour and minute, and use "PI()" to get the value of π.

Copy the information in cell C3 to as many rows as was calculated for the time. If you did this correctly, you should find that the values in this column are zero for all times between midnight to 6 am, reaches a maximum of 494 W/m^2 at noon, decreases back to zero at 6 pm and then remains zero the rest of the day.

C.4 Atmospheric irradiance

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 217 W/m^2. In reality, this wouldn't be constant but it helps simplify the calculations.

Title the fourth column "Fatmos" (D1) to represent the radiation from the atmosphere absorbed at Earth's surface. Since we are assuming this is constant at 217 W/m^2 throughout the day, put 217 into the first cell (D2) and then copy that to all of the other rows in this column.

C.5 Total absorbed irradiance

Title the fifth column "Fin" (E1) to represent the total radiation absorbed at Earth's surface. Set the first cell (E2) equal to the sum of the solar and atmospheric contributions (C2+D2) and then copy that to the rest of the rows.

C.6 Total emitted irradiance

Title the sixth column "Fout" (F1) to represent the radiation emitted by Earth's surface. This will depend on the temperature and can be calculated via the Stefan-Boltzmann law (equation 7.2):

$$F_{\text{out}} = \sigma T^4$$

where $\sigma = 5.67 \times 10^{-8}$ W·m^{-2}·K^{-4}. Use the temperature in column B for T. Copy to rest of rows in this column.

C.7 Change in temperature

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 C.1 (derived in next section):

$$\frac{dT}{dt} = \frac{F_{\text{net}}}{ch\rho}$$

to calculate the change in temperature. Set the time increment to be 1 minute (remember to convert to seconds), use 1004.67 $\text{J} \cdot \text{kg}^{-1}\text{K}^{-1}$ 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), 1.2 kg/m^3 for ρ and 270 m for h (this represents the layer of the atmosphere near Earth's surface). If you did everything correctly, you should get a temperature change of about -0.031 K in the cell (G2).

In reality, the lowest 270 m of the atmosphere doesn't absorb 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 270-m layer.

You'll find that, over the course of the entire 24-hour day, 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.

C.8 Derivation

The following equation was provided in the previous section:

$$\frac{dT}{dt} = \frac{F}{c \cdot h \cdot \rho} \tag{C.1}$$

It turns out this expression can be derived from the first law of thermodynamics. The expression is not a particularly useful expression outside of this exerise, since it requires assumptions that aren't particularly valid, but it provides a reasonable estimate of the temperature change due to radiation absorption.

First, we recognize that, for the ground, the energy received dQ is equal to $cmdT$, where c is the specific heat of the ground, m is the mass that absorbs the radiation, and dT is the resulting change in temperature.

Here we are interested in the time tendency (i.e., the rate at which the temperature changes), so we can divide both sides by dt to get:

$$\frac{dT}{dt} = \frac{1}{c \cdot m}\frac{dQ}{dt}$$

In this case, dQ/dt represents the energy per time contributed by the radiation imbalance. Since the irradiance F represents the concentration of energy provided per time, dQ/dt can be replaced by FA, where A is the area receiving/emitting the radiation:

$$\frac{dT}{dt} = \frac{F \cdot A}{c \cdot m}$$

To get rid of mass and area, multiply top and bottom by h, the depth that is absorbing the radiation,

$$\frac{dT}{dt} = \frac{F}{c \cdot h}\frac{A \cdot h}{m}$$

We then replace $m/(A \cdot h)$ by the density ρ (since $A \cdot h$ is equal to the volume) to get equation C.1.

The expression tells us that ground will warm up in a way that is proportional to the net irradiance it receives, F, and inversely proportional to the specific heat of the ground (larger specific heat means it can absorb more radiation and not warm up as much), the depth that absorbs the radiation (absorbing over a deeper layer means it won't warm up as much) and the density (a more dense layer won't warm up as much).

Like the ground, the air both absorbs and emits radiation. It isn't as efficient (see chapter 8) but it does both. And we can use the same expression for air, except with c_p for the specific heat.[i]

Example C.1: Earth's irradiance is 390 W/m^2 at its surface (see Example 7.1). Suppose all the radiation is absorbed by the lower 4 km of the atmosphere, where the average density is roughly 1 kg/m^3. Determine the change in temperature experienced by that column of air at night (no solar radiation) during one hour assuming the column totally absorbs the terrestrial radiation (it doesn't) and doesn't emit radiation (it does).

Answer C.1: Plugging in, we get:

$$\frac{dT}{dt} = \frac{F}{c_\text{p} h \rho} = \frac{(390\ \text{W/m}^2)}{(1004.67\ \text{J} \cdot \text{kg}^{-1}\text{K}^{-1})(4000\ \text{m})(1\ \text{kg/m}^3)} = 1.32 \times 10^{-4}\ \text{K/s}$$

Multiply by the time (3600 s) to get 0.35 K.

---

[i] The first law of thermodynamics, for gases, has a pressure term (see equation 6.5). However, 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).

D. Equations of Motion

In this appendix, I derive the three equations of motion, one equation for each component direction.

$$\begin{aligned}
\frac{du}{dt} &= -\frac{1}{\rho}\frac{\partial P}{\partial x} - F_x - 2\Omega w \cos\phi + 2\Omega v \sin\phi - \frac{uw}{a} + \frac{uv\tan\phi}{a} && \text{(D.1)} \\
\frac{dv}{dt} &= -\frac{1}{\rho}\frac{\partial P}{\partial y} - F_y - (2\Omega u \sin\phi) - \frac{u^2\tan\phi}{a} - \frac{vw}{a} && \text{(D.2)} \\
\frac{dw}{dt} &= -\frac{1}{\rho}\frac{\partial P}{\partial z} - g - F_z + (2\Omega u \cos\phi) + \frac{u^2+v^2}{a} && \text{(D.3)}
\end{aligned}$$

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 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 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 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, 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 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 Earth, since u is measured relative to 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 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:

$$\begin{aligned} dw/dt &= \partial(v_r\cos\phi)/\partial t \\ &= \Big((\Omega^2 a\cos\phi) + (2\Omega u) + (u^2)/(a\cos\phi)\Big)\cos\phi \\ &= (\Omega^2 a\cos^2\phi) + (2\Omega u\cos\phi) + (u^2/a) \qquad \text{(D.6)} \end{aligned}$$

And a similar expression for v:

$$\begin{aligned} dv/dt &= \partial(-v_r\sin\phi)/\partial t \\ &= -\Big((\Omega^2 a\cos\phi) + (2\Omega u) + (u^2)/(a\cos\phi)\Big)\sin\phi \\ &= -(\Omega^2 a\sin\phi\cos\phi) - (2\Omega u\sin\phi) - (u^2\tan\phi/a) \qquad \text{(D.7)} \end{aligned}$$

---

[i] The error comes about because of two things. First, Earth is not a perfect sphere. Consequently, the distance from the center of Earth to 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 Earth's surface. In both cases, the variation is small compared with the average radius of Earth.

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

$$\frac{du}{dt} = -\frac{uw}{a} + \frac{uv\tan\phi}{a} - 2\Omega w\cos\phi + 2\Omega v\sin\phi \tag{D.8}$$

D.3 Motion around line of longitude

When applied around a line of longitude, the center of the circle is the center of Earth. Consequently, the distance r, which represents the distance from the object to the axis is equal to a, the radius of 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 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 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 (Earth) is moving and so u_* includes the motion of Earth (equal to $\Omega r \cos\phi$, where Ω is the rotational velocity of 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.

$$\partial\hat{\imath}/\partial\phi = 0$$

$$\begin{aligned}
\partial\hat{\imath}/\partial\lambda &= -\hat{k}\cos\phi + \hat{\jmath}\sin\phi \\
\partial\hat{\jmath}/\partial\phi &= -\hat{k} \\
\partial\hat{\jmath}/\partial\lambda &= -\hat{\imath}\sin\phi \\
\partial\hat{k}/\partial\phi &= \hat{\jmath} \\
\partial\hat{k}/\partial\lambda &= \hat{\imath}\cos\phi
\end{aligned}$$

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:

$$\begin{aligned}
\frac{d\hat{\imath}}{dt} &= (u_*\tan\phi/r)\hat{\jmath} - (u_*/r)\hat{k} \\
\frac{d\hat{\jmath}}{dt} &= -(u_*\tan\phi/r)\hat{\imath} - (v/r)\hat{k} \\
\frac{d\hat{k}}{dt} &= (u_*/r)\hat{\imath} + (v/r)\hat{\jmath}
\end{aligned}$$

Plugging in to what we had before, the right side of equation D.14 can be written as

$$-u_*((u_* \tan\phi/r)\hat{j} - (u_*/r)\hat{k}) - v(-(u_* \tan\phi/r)\hat{i} - (v/r)\hat{k}) - w((u_*/r)\hat{i} + (v/r)\hat{j})$$

or, rearranging terms,

$$(u_* v \tan\phi/r - wu_*/r)\hat{i} + (-u_*^2 \tan\phi/r - wv/r)\hat{j} + (u_*^2/r + v^2/r)\hat{k}).$$

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 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 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 $-\hat{\jmath}$ 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 $-\hat{\jmath}$ 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 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 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 $\hat{\jmath}$ 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 $\mathrm{m/s^2}$ (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 $\hat{j}$ 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 $u_*\hat{i} + v\hat{j} + w\hat{k}$ (since v and w are measured relative to our shifted coordinates). The only difference is that $\partial\hat{j}/\partial\phi$ 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.

E. Molar masses

E.1 Introduction

In the chapter 3, the molar mass of dry air was determined from the molar masses of the elements that make up the dry air. In this appendix, I show how the molar mass of each element can be determined from the masses of the sub-atomic particles that make it up. This information is not needed to understand what is in the chapters and thus has been relegated to the appendix. 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.

E.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.

Checkpoint E.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?

E.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?

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.

---

[i] The electrons would "fall" into the nucleus if not for their kinetic energy.

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.

Checkpoint E.2: The most stable configuration of a nitrogen nucleus has seven protons and seven neutrons. For each of the following changes, is the atom still considered to be nitrogen?
(a) If the number of protons in the nucleus changes.
(b) If the number of neutrons in the nucleus changes.
(c) If the number of electrons surrounding the nucleus changes.

E.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 electrons if we do not need to be accurate within 0.1% (i.e., one out of a thousand).

Checkpoint E.3: What fraction of an oxygen-16 atom's total mass is made up of just the mass of the electrons?

[ii] Source: physics.nist.gov

E.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 37, a mole represents 6.0221415×10^{23} objects, a number 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 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 (where the average molar mass of oxygen was shifted from exactly 16 g/mol to 15.9994 g/mol).[iii]

Checkpoint E.4: Is the determination of a mole simply a matter of convenience or does it represent something that is set by the physical universe?

[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").

E.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. Instead of 12 g/mol, 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:

$$\begin{aligned} 6\mu_p + 6\mu_n + 6\mu_e &= 6(1.00727646 \text{ g/mol}) \\ &\quad +6(1.00866591 \text{ g/mol}) \\ &\quad +6(0.00054858 \text{ g/mol}) \\ &= 12.0989397 \text{ g/mol?} \end{aligned}$$

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]

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 is significant when compared with the masses. In comparison,

---

[iv] By the way, in a sense there is also a mass defect associated with the neutron. A 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.

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 attraction is greater than the electric repulsion. Technically, the energy we are talking about is the net effect of the two together.

Checkpoint E.5: During a chemical reaction, energy can be released or absorbed. For each chemical reaction listed below, 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?

E.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 (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:

$$E = mc^2 \tag{E.1}$$

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.

Checkpoint E.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?

E.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).

Checkpoint E.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?

E.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 E.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 E.1: What is the average molar mass of oxygen?

Answer E.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.

Checkpoint E.8: In the note I state that the relative abundances should add up to one. Why should that be so?

---

[v] By convention, a horizontal line on top of a variable abbreviation represents an average of some kind.

Table E.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

Questions

Question E.1: 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 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?

Problems

Problem E.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 2.4)?

Problem E.2: From the molar masses of the isotopes of nitrogen provided in table E.1, calculate the average molar mass of diatomic nitrogen. Compare your answer to that given in table 3.1. Should it be the same? If so, why? If not, which should be larger and why?

Problem E.3: 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 E.1, determine the mass defect for the nitrogen-14 nucleus. Note that one eV is equivalent to $1.602176534 \times 10^{-19}$ *J.*

Problem E.4: (a) Given the molar masses of an individual proton, neutron and electron, and the mass defect determined in problem E.3, calculate the molar mass of the nitrogen-14 atom.
(b) Compare your answer to that given in table E.1. Should it be the same? If so, why? If not, which should be larger and why?

F. Atmospheric Evolution

Introduction

In chapter 3, it was mentioned that nitrogen makes up the bulk of the atmosphere. You might wonder why this is.

To answer this, we need to examine how the atmosphere came to be. Whether 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)[i].

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 atmosphere altogether. Other gases return to the Earth via a process called deposition. The first process (escape) is discussed in section F.1, and the second process (deposition) is discussed in section F.2.

F.1 Escape

To escape from the atmosphere altogether, a molecule has to be going so fast that Earth's gravity is not strong enough to slow it down enough to bring it back toward Earth. This speed is known as the **escape velocity** and can be determined according to the following equation.

$$v_{\mathrm{e}} = \left[\frac{2Gm_{\mathrm{earth}}}{r}\right]^{1/2} \tag{F.1}$$

---

[i] Wallace, John M., and Peter V. Hobbs, 1977: Atmospheric Science, An Introductory Survey, Academic Press, 467 pp.

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 Earth (5.98×10^{24} kg) and r is the distance of the molecule from the center of 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 ($W = \Delta E_k$).

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 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 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 $F_\mathrm{avg}\Delta x$ (where Δx is the distance moved and F_avg is the average force exerted) we have the following relationship.

$$\Delta\left(\frac{1}{2}mv^2\right) = F_\mathrm{avg}\Delta x \tag{F.2}$$

The force is the gravitational force

$$F_\mathrm{g} = G\frac{m_\mathrm{earth}m}{r^2} \tag{F.3}$$

and this varies as the object moves away from Earth.

If the gravitational force was constant, we could just plug the expression for the gravitational force into equation F.2. However, the gravitational force on the object is not constant, in the sense that the force decreases as the object moves away from 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 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.

$$F_{\text{avg}} = \int \frac{F(x)dx}{\Delta x} \tag{F.4}$$

Since $F(x)$ is given by (F.3) above (with r instead of x) we actually have an expression that looks like it is the derivative of something:

$$\begin{aligned}
F_{\text{avg}} &= \int \frac{F(x)dx}{\Delta x} \\
&= \int G\frac{m_{\text{earth}}m}{\Delta x} \times \left(x^{-2}dx\right) \\
&= \int G\frac{m_{\text{earth}}m}{\Delta x} \times d\left(-x^{-1}\right) \\
&= G\frac{m_{\text{earth}}m}{\Delta x} \times \int d\left(-x^{-1}\right) \\
&= G\frac{m_{\text{earth}}m}{\Delta x} \times \Delta(-x^{-1})
\end{aligned}$$

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 integral 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 (F.2) we get

$$\frac{1}{2}mv^2 = Gm_{\text{earth}}m \times \frac{1}{r} \tag{F.5}$$

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.

Checkpoint F.1: (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 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 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 molar 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.

Checkpoint F.2: Is the atmosphere losing hydrogen to space? What about nitrogen?

F.2 Deposition

Deposition is the process of removing constituents from the atmosphere 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 (e.g., the gasoline in a car), oxygen is utilized with the hydrocarbon to produce water vapor and carbon dioxide. A similar process occurs during respiration except with carbohydrates. This is why carbon dioxide is exhaled when animals and plants undergo respiration.

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).

Checkpoint F.3: Does the burning of fossil fuels create carbon or does it just move it from one form to another?

Problems

Problem F.1: (a) Determine the escape velocity for an object at Earth's surface.

(b) Determine the escape velocity for an object at a point 500 km above Earth's surface.
(c) Compare the two values in (a) and (b). Is the difference significant? Should it be? Explain.

Derivations

Show-me F.1: Show that equation F.1

$$v_{\mathrm{e}} = \left[\frac{2Gm_{\mathrm{earth}}}{r}\right]^{1/2}$$

can be obtained from equations F.2, F.3 and F.4. In other words, repeat the derivation produced in section F.1 (stating the mathematics taken in each step) and add the algebra needed to go from equation F.5

$$\frac{1}{2}mv^2 = Gm_{\mathrm{earth}}m \times \frac{1}{r}$$

to equation F.1.

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** 22.9897693	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** 132.905452	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** *	Unununpentium 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** *