Weather and Atmosphere

What can we say about the atmosphere of the earth from an estimation/scaling perspective? We can establish the pressure and temperature profiles, assess the formation of clouds and associated energy, estimate the particle size within, and density of, a cloud, and discover the magnitude of an electric field that will allow air to break down into an arc of lightning.

Scale Height

How does the pressure in the atmosphere scale with height? At some height, z, the weight of a thin slice is $dW = gdm = \rho gdAdz$, so that the pressure at this slice is the weight of overburden divided by the area:

$$P = W/dA = \int_z^\infty \rho g dz'.$$

But ρ is a function of z, so we have to snoop this out. If we rearrange the ideal gas law: $PV = NkT$ to $n = P/kT$, and recognize that $\rho = nm_{\text{molecule}}$, where for air, $m_{\text{molecule}} = Am_{\text{amu}}$, where $A \approx 29$ is the average mass number for air molecules. We have that $\rho = P(z)Am_{\text{amu}}/kT$, so the pressure overhead becomes:

$$P(z) = \frac{Am_{\text{amu}}g}{kT}\int_z^\infty P(z')dz',$$

which means that $P(z)$ must have an exponential form. Perhaps this is more clear if we differentiate both sides, and note that because z is the lower limit of the integral, the derivative picks up a negative sign: $P'(z) = -\frac{1}{h}P(z)$, where $h = kT/Am_{\text{amu}}g$ represents the scale height, so that

$$P(z) = P_0 e^{-z/h};\ h = \frac{kT}{Am_{\text{amu}}g} \approx \frac{4\times 10^{-21}\,\text{J}}{29\cdot 1.66\times 10^{-27}\,\text{kg}\cdot 10\,\text{m/s}^2} \sim 8\,\text{km}.$$

We can also get a handle on the scale height by asking how high the atmosphere would be if it had uniform density to some hard cutoff. This stems from the property that the integral of $e^{-z/h}$ from zero to ∞ is just h. A surface pressure of $P_0 \approx 10^5$ Pa means an overburden of 10^4 kg/m^2, which at the sea-level density of 1.3 kg/m^3 translates to about 8 km.

Lapse Rate

We pretended above that temperature was constant with height in order to pull it out of the integral. We know this is not strictly true, even though we might expect the temperature to vary by something like 10% over the same height that the pressure sees a factor of two change. We term the rate at which temperature decreases with height the *lapse rate*.

We can get at the lapse rate by assuming that a parcel of air starting from sea level and rising under its own buoyancy will expand and cool in the process, but will not have a chance to exchange heat with surrounding air. Air is, after all, a poor conductor, poor radiator, and we exclude turbulent mixing from our consideration here. Thus the process is *adiabatic*, and follows the prescription that PV^γ is constant. Armed with the functional form of pressure, we can say:

$$P_0V_0^\gamma = PV^\gamma = P_0e^{-z/h}V^\gamma \rightarrow \left(\frac{V}{V_0}\right)^\gamma = e^{z/h},$$

leading to $V = V_0e^{z/h\gamma}$. We also know that $PV = NkT$, so that

$$PV = P_0e^{-z/h}V_0e^{z/h\gamma} = NkT = NkT_0e^{-\frac{z}{h}\left(1-\frac{1}{\gamma}\right)},$$

where we have used the fact that $P_0V_0 = NkT_0$. This leads to a scaling for temperature:

$$T = T_0 e^{-\frac{z}{h}\left(1-\frac{1}{\gamma}\right)}.$$

For air, $\gamma = \frac{7}{5}$, so the scale height for temperature is 3.5 times higher than it is for pressure, or 28 km. This means that the gradient at the ground is something like $\frac{\partial T}{\partial x}\big|_{z=0} \approx (280\,\mathrm{K})/(28\,\mathrm{km}) = 10$ K/km. This is known as the *dry lapse rate*, and applies to dry air.

When the air is moist, rising air will cool to the *dewpoint* of the air: the temperature at which the air becomes saturated with water, and water must condense. When this happens, the condensing water releases its heat of vaporization, thereby adding heat to the air, slowing the rate of cooling as it rises. A typical wet lapse rate is ~ 5 K/km, or about 3 °F/(1000 ft).

Atmospheric Stability

Sometimes the sky is calm, clouds form layers, everything is stable. Sometimes the sky is turbulent, full of puffy, dynamic clouds—sometimes shooting up to the tropopause in a towering cumulonimbus. What determines these two states?

It's all about the lapse rate and buoyancy. If a parcel of air happens to rise a little, cools adiabatically (whether in the wet or dry state), and finds that it is cooler than the surrounding air, it will be more dense than the local air and settle back down. We say that there is *negative feedback* in this system, *stable* against perturbations. If the ambient air has a temperature gradient that is less steep than the adiabatic lapse rate, stability reigns. If, on the other hand, the temperature gradient is steeper than the lapse rate, a parcel of air that happens to rise and cool will nonetheless find itself warmer than the surrounding air. It will be less dense and therefore buoyant, and rise more. So there is *positive feedback* in this case, and this leads to a runaway amplification of temperature perturbations. We call this *unstable* air, promoting turbulence, cloud formation, etc. Note that once condensation starts, the lapse rate of the (cloud) parcel decreases, so that it is even more unstable and will tend to build to the tropopause.

What sets the ambient gradient? Lots of things contribute: solar forcing on the surface, cold fronts, warm fronts, inversion layers (ultimate stability), radiation cooling of the surface. An example scenario: a sunny day deposits a lot of heat onto the ground, warming the air in contact with the ground. Convective airflow will mix this into the atmosphere, setting up a possibly steep temperature gradient. This in turn promotes instability, cloud formation, thunderstorms.

Moisture Content of Air

An evacuated flask with water in it will reach a vapor pressure equilibrium in which the rate of thermally-excited water molecules ejected from the liquid will equal the rate of previously-ejected water molecules by chance crashing back to the water surface. This *saturation vapor pressure* is the same whether or not air is present in the flask. In other words, the water molecules simply co-exist with air molecules in the saturated equilibrium state. It is helpful to think in terms of partial pressure. At 100°C, the saturated vapor pressure for water is 760 Torr. At 20°C, it is 17.5 Torr, corresponding to a mass density of 17 g/m^3.

We can estimate the rate of water molecules leaving the surface as the probability of a water molecule having enough energy to escape, times the number of molecules on the surface of the water, divided by the timescale it would take to replace the ejected molecule's place on the surface. Symbolically, $\Gamma_{\mathrm{out}} = pN/\tau$. In reverse order, we can approximate τ as the mean-free-time of a particle in the liquid, as this is the timescale for which a molecule will replace the ejected molecule's seat on the surface. We have $\tau = \lambda/v \sim 1/n_\ell \sigma v$, where λ is the mean free path, and v the molecule's thermal velocity. We have approximated $\lambda \sim 1/n\sigma$ in the usual manner, using an ℓ subscript to denote liquid state. The number of particles on the surface is simply the area of the surface divided by the cross-sectional area for a molecule, or $N \sim A/\sigma$. Finally, the probability of being ejected from the surface must follow a Boltzmann distribution: $p(E)dE = e^{-E/kT}dE$. If we integrate this from some energy threshold, $E_{\rm th}$, to infinity, and normalize for unit probability starting at $E_{\rm th} = 0$, we find that the probability factor is $e^{-E_{\rm th}/kT}$. Putting it together, we have $\Gamma_{\rm out} = An_\ell ve^{E_{\rm th}/kT}$.

The rate of incident water molecules can be approximated by the random-walk behavior of the molecules. In a volume above the water, $N = n_{\rm g}V = n_{\rm g}A\lambda$, water molecules remain for a time $\tau \sim \lambda/v$, during which half will move downward and encounter the water, while half will move upward. So the rate of influx would be $\Gamma_{\rm in} \sim \frac{1}{2}N/\tau \sim \frac{1}{2}n_{\rm g}Av$.

Equating the two timescales, and recognizing that the area and velocity factors are the same for both, we have that $n_{\rm g} \sim 2n_\ell e^{-E_{\rm th}/kT}$. Since the mass density is proportional to number density with the same mass multiplier (same molecule), we can replace the n values with ρ values. Matching the relationship to data suggests that we are off by about a factor of 150, so that we should use $\rho_{\rm g} \sim 300\rho_\ell e^{-H_{\rm vap}/kT}$, where we have replaced $E_{\rm th}$ with $H_{\rm vap}$, or the heat of vaporization for water. $H_{\rm vap} = 40.65$ kJ/mol; 2257 kJ/kg; 0.54 kcal/g; 0.42 eV/molecule. An online formula for the partial pressure of water vapor at saturation is $P \sim e^{20.386}e^{-5132/T}$, in Torr. If we identify $E_{\rm th} = 5132/k$, we find that $E_{\rm th} = 0.44$ eV/molecule—which is just our heat of vaporization. I have not tracked down the 10^2 discrepancy in the derivation above, but am at least gratified by the correct scaling.

At 273 K, we get 4.9 Torr, at 293, it's 17.6 Torr, and at 310 K, it's 46 Torr. Highly nonlinear with absolute temperature.

Evaporation Energy Budget

Now that we know something about the heat of vaporization of water, we can ask a few questions about weather.

First, how much of the sun's energy goes into driving the evaporation/rain cycle on the planet? We can estimate the annual amount of rain an average location on earth receives as being between 0.5 m and 1.0 m. If we pick something in the middle, we find a convenient number of 2 mm per day. Since all this rain was evaporated at some point, we must have 2 ℓ, or 2 kg of water evaporation per square meter per day. This requires about 4500 kJ of energy ($H_{\rm vap} = 2257$ kJ/kg) per day, turning into a 24-hour average rate of about 50 W. Since the global average solar energy absorbed (not immediately reflected) by the earth system is $0.7 \cdot 1370/4$ W/m^2, or 240 W/m^2, the water part constitutes just over 20% of the total energy budget! It's huge!

How much of a cloud's energy is in potential energy of lofted air vs. heat of vaporization? One kilogram of water takes 2257 kJ to vaporize, and 50 kJ to raise 5 km—the average height of water in a giant cumulonimbus). So only 2% of the energy is gravitational. Then in hydroelectric systems, we hang on to less than 1% *of this*, since only the last few hundred meters of "fall" back to the ocean is captured. So we might say that of the energy that hydroelectric systems *do* capture, it is only 0.02% of the solar energy that went into sending that water through the dam. 15% photovoltaics don't sound so bad in this context!

Cloud Properties

How heavy is that cloud, and what is the water drop size? These questions showcase a number of physical estimation techniques, including inertial drag, viscous drag, optical depth, diffraction, rainfall, and also the results from above.

Cloud Density

Let's start simply. How much rain might a 10 km high cumulonimbus cloud dump if it completely flushed itself of water? It would be downpour to behold, and would likely produce between an inch and ten inches of rain. A geometrical mean would say three inches, or 7.5 cm. We can flub this up to 10 cm while we're bring approximate. 10 cm out of 10 km is one part in 10^5, so the density of water in the cloud is $\rho_{\rm c} \sim \rho_{\rm w}/10^5 \approx 10$ g/m^3, or just under 1% the density of air.

We can also state from experience that when our jet airplane thrusts into a cloud, we do not notice a deceleration from increased drag. Since inertial drag force, $F_\mathrm{d} = \frac{1}{2}c_\mathrm{D}\rho A v^2$ is proportional to the density of the medium, a sharp step in total density would catch our attention. From this we might venture that the cloud over-density is no more than about 1% of the air density.

A cloud forming out of moist air at 20°C will condense the water as it cools, but the total density of water molecules remains the same across the condensation threshold. The saturated density at 20°C is about $(17.5/760)\times(18/29) = 1.4\%$ times that of air (17 g/m^3). The bottoms of most clouds are not at ground level where the temperature is 20°C. If the base is at 2 km, the dry adiabatic lapse rate applies over this range, so we expect the cloud base to be closer to 0°C. It is at this point that the air becomes saturated, so the density will be closer to the saturation density at this temperature. Scaling like $e^{-H_\mathrm{vap}/kT}$, we would expect a density ratio of $e^{-\frac{H}{k}\left(\frac{1}{T_1}-\frac{1}{T_2}\right)}$, which works out to about 30% in this case. So we expect a density around 6 g/m^3, or 0.5% air density.

Now that we have several estimates for the density of a cloud, all resulting in something like 1% air density, we can estimate the weight/mass of a cloud. A small puffball of a cloud 100 m on a side would have about 10^{-2} kg/m^3 $\times$ 10^6 m^3 $= 10^4$ kg of water—about the same as ten cars in weight. Yet it floats!

Droplet Size

Let's estimate the droplet size from a number of clues. First, we note that the cloud floats. If the terminal velocity of the droplets was much larger than random motions of the air, it would have a hard time doing this. We can also recall instances of seeing fog up close (even from the shower), and noting that individual particles do not descend noticeably fast—except some fogs have large enough particles that there will be a discernible downward drift, and we may even upgrade it to a mist. Let's say that the downward drift rate is less than 1 cm/s. What must the particle size be? For viscous (Stokes) drag, we know from previous lectures that $F_\mathrm{d} = 6\pi\rho_\mathrm{a}\nu r v = mg$, where the parameters are air density, air viscosity, radius and velocity of the particle, and the final equality is to establish terminal velocity. For a spherical geometry, we find that

$$r = \sqrt{\frac{9\rho_\mathrm{a}\nu v}{2\rho_\mathrm{w} g}} \approx 90\,\mu\mathrm{m}\cdot v^{\frac{1}{2}} \approx 10\,\mu\mathrm{m}.$$

We should check that viscous drag is the right regime: $\mathrm{Re} = rv/\nu \approx 10^{-5}\cdot 10^{-2}/1.5\times 10^{-5} \sim 10^{-2}$ is most firmly in the viscous/laminar regime. The water droplets see the air much like we would see peanut butter.

From another angle, we might explore how far we can see in a cloud (a.k.a. fog). We can use our mean free path machinery to ask how far a photon will travel without being disrupted by a water droplet. In the familiar symbols, $\lambda = 1/n\sigma$, where n is the number density of droplets, and σ is the cross-sectional area of the drop. For spherical geometry, we have $\sigma = \pi r^2$, and note that the density of water in the cloud, $\rho_\mathrm{c} = \frac{4}{3}\pi n \rho_\mathrm{w} r^3$. Putting these pieces together, we have that $\lambda = 4\rho_\mathrm{w} r/3\rho_\mathrm{c}$, or $r = 3\lambda\rho_\mathrm{c}/4\rho_\mathrm{w}$. In an airplane, sometimes the wingtip is obscured within a cloud. So let's say $\lambda \sim 5$ m. We will also go with a cloud density that corresponds to a cloud whose base is not on the ground, so $\rho_\mathrm{c} \sim \rho_\mathrm{air}/200$, or ~ 5 g/m^3. These numbers produce an estimate of $r \sim 20\,\mu$m.

A final line of attack: we see rainbows when sunlight interacts with the spherical geometry of water droplets, producing a pile-up of light at a 42° angle from the anti-solar direction. Yet we do not see the same phenomenon when sunlight strikes a cloud: there must be rain. But the cloud droplets must be spherical, with the same refractive index as larger droplets—so the same optical trick should apply. Why would smaller drops fail to produce the result? The answer is diffraction. Each droplet acts like an aperture, and this leads to an spread in the angle of emerging light. The rainbow feature is narrow, and once diffraction is about 0.5° across, or about 0.01 radians, the pattern will wash out. Setting $\lambda/D > 0.01$ means $D < 100\lambda \approx 50\,\mu$m. This translates to a drop radius of $r < 25\,\mu$m, in line with many of our previous estimates.

Raindrops

While we're talking about droplet size, let's figure out what limits the size of raindrops. As soon as the terminal velocity exceeds vertical airflow rates, water will fall. So this lets raindrops get arbitrarily small for still conditions (until we have a mist). But what about the upper end? What is it that holds a water droplet together? It's *surface tension*. For water, surface tension is $\gamma \sim 0.070$ N/m. It looks like a spring constant. One might expect it to carry units of area instead of length in the denominator. But think of it this way: the surface energy associated with holding a drop together should scale with the area. Indeed, $E \approx \gamma A_{\rm surf}$ has units of N·m. A pressure scale can be established by evaluating the energy density associated with surface tension, or $P \sim E/V \sim \gamma A_{\rm surf}/V$. Associating this with drag pressure sets the scale at which ram pressure can tear the drop apart: $P_{\rm drag} = F_{\rm drag}/A_{\rm proj}$, where $F_{\rm drag} = mg$ at terminal velocity and $A_{\rm proj}$ is the projected area seen by the onrushing air. One gets reasonable values via this approach. This technique can also be used to understand droplet sizes that survive on a waxy surface, the over-full phenomenon of a water glass, and sets limits on bug sizes that can walk on water. Fancy approaches may deviate from spherical assumptions.

Lightning and Arcing

When a large enough charge separation builds up in clouds, the air may break down and start to conduct electricity in a phenomenon we know as lightning. The critical phenomenon is that an electron liberated from an air molecule will be driven by the electric field until it encounters an air molecule. When this happens, the collision resets the energy of the electron, and the accelerated drift starts again. If, in the intervening drift, the electron accumulates enough kinetic energy, it will have enough oomph to liberate another electron. In this case, we get a cascade of charge flow.

What is the critical electric field? First we need to know how much distance the electron will go before finding a molecule to hit. This is the familiar mean free path: $\lambda \approx 1/n\sigma$. In air, $n = 2.5 \times 10^{25}\,\mathrm{m}^{-3}$, and we will take $\sigma \approx (2 \times 10^{-10}\,\mathrm{m})^2$. This puts the mean free path at $\lambda \approx 1\,\mu$m.

Let's say the electron needs to have 3 eV of kinetic energy to liberate an electron from the air molecule. This means that the electron must experience a potential difference of 3 V across its $1\,\mu$m joy-ride. That would make the critical electric field, $E = 3 \times 10^6$ V/m, which is indeed the breakdown field for dry air.