26 June 2007

Part II: What Ångström didn't know

Filed under: **Climate Science** — group @ 7:23 AM

By raypierre , with the gratefully acknowledged assistance of Spencer Weart

In **Part I** the long struggle to get beyond the fallacious saturation argument was recounted in historical terms. In Part II, I will provide a more detailed analysis for the reader interested in the technical nitty-gritty of how the absorption of infrared really depends on CO_2 concentration. At the end, I will discuss Herr Koch's experiment in the light of modern observations.

The discussion here is based on CO_2 absorption data found in the **HITRAN** spectroscopic archive. This is the main infrared database used by atmospheric radiation modellers. This database is a legacy of the military work on infrared described in **Part I** , and descends from a spectroscopic archive compiled by the Air Force Geophysics Laboratory at Hanscom Field, MA (referred to in some early editions of radiative transfer textbooks as the "AFGL Tape").

Suppose we were to sit at sea level and shine an infrared flashlight with an output of one Watt upward into the sky. If all the light from the beam were then collected by an orbiting astronaut with a sufficiently large lens, what fraction of a Watt would that be? The question of saturation amounts to the following question: How would that fraction change if we increased the amount of CO_2 in the atmosphere? Saturation refers to the condition where increasing the amount of CO_2 fails to increase the absorption, because the CO_2 was already absorbing essentially all there is to absorb at the wavelengths where it absorbs at all. Think of a conveyor belt with red, blue and green M&M candies going past. You have one fussy child sitting at the belt who only eats red M&M's, and he can eat them fast enough to eat half of the M&M's going past him. Thus, he reduces the M&M flux by half. If you put another equally fussy kid next to him who can eat at the same rate, she'll eat all the remaining red M&M's. Then, if you put a third kid in the line, it won't result in any further decrease in the M&M flux, because all the M&M's that they like to eat are already gone. (It will probably result in howls of disappointment, though!) You'd need an eater of green or blue M&M's to make further reductions in the flux.

Ångström and his followers believed that the situation with CO_2 and infrared was like the situation with the red M&M's. To understand how wrong they were, we need to look at modern measurements of the rate of absorption of infrared light by CO_2 . The rate of absorption is a very intricately varying function of the wavelength of the light. At any given wavelength, the amount of light surviving goes down like the exponential of the number of molecules of CO_2 encountered by the beam of light. The rate of exponential decay is the absorption factor.

When the product of the absorption factor times the amount of CO_2 encountered equals one, then the amount of light is reduced by a factor of $1/e$, i.e. $1/2.71282\ldots$. For this, or larger, amounts of CO_2, the atmosphere is *optically thick* at the corresponding wavelength. If you

double the amount of CO_2, you reduce the proportion of surviving light by an additional factor of 1/e, reducing the proportion surviving to about a tenth; if you instead halve the amount of CO_2, the proportion surviving is the reciprocal of the square root of e , or about 60% , and the atmosphere is *optically thin*. Precisely where we draw the line between "thick" and "thin" is somewhat arbitrary, given that the absorption shades smoothly from small values to large values as the product of absorption factor with amount of CO_2 increases.

The units of absorption factor depend on the units we use to measure the amount of CO_2 in the column of the atmosphere encountered by the beam of light. Let's measure our units relative to the amount of CO_2 in an atmospheric column of base one square meter, present when the concentration of CO_2 is 300 parts per million (about the pre-industrial value). In such units, an atmosphere with the present amount of CO_2 is optically thick where the absorption coefficient is one or greater, and optically thin where the absorption coefficient is less than one. If we double the amount of CO_2 in the atmosphere, then the absorption coefficient only needs to be 1/2 or greater in order to make the atmosphere optically thick.

The absorption factor, so defined, is given in the following figure, based on the thousands of measurements in the HITRAN spectroscopic archive. The "fuzz" on this graph is because the absorption actually takes the form of thousands of closely spaced partially overlapping spikes. If one were to zoom in on a very small portion of the wavelength axis, one would see the fuzz resolve into discrete spikes, like the pickets on a fence. At the coarse resolution of the graph, one only sees a dark band marking out the maximum and minimum values swept out by the spike. These absorption results were computed for typical laboratory conditions, at sea level pressure and a temperature of 20 Celsius. At lower pressures, the peaks of the spikes get higher and the valleys between them get deeper, leading to a broader "fuzzy band" on absorption curves like that shown below.
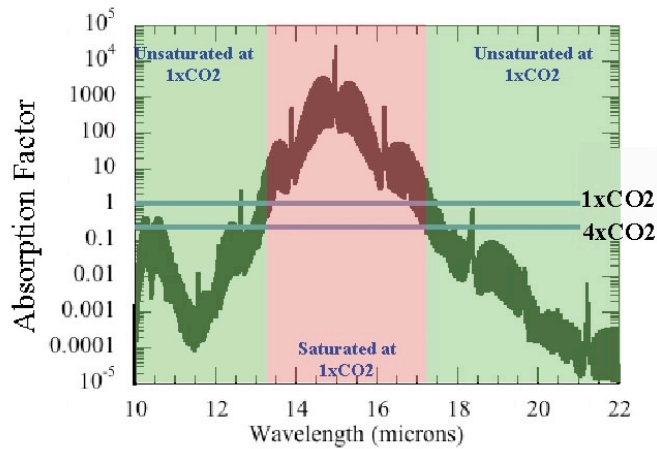

We see that for the pre-industrial CO_2 concentration, it is only the wavelength range between about 13.5 and 17 microns (millionths of a meter) that can be considered to be saturated. Within this range, it is indeed true that adding more CO_2 would not significantly increase the amount of absorption. All the red M&M's are already eaten. But waiting in the wings, outside this wavelength region, there's more goodies to be had. In fact, noting that the graph is on a logarithmic axis, the atmosphere *still* wouldn't be saturated even if we increased the CO_2 to *ten thousand* times the present level. What happens to the absorption if we quadruple the amount of CO_2? That story is told in the next graph:
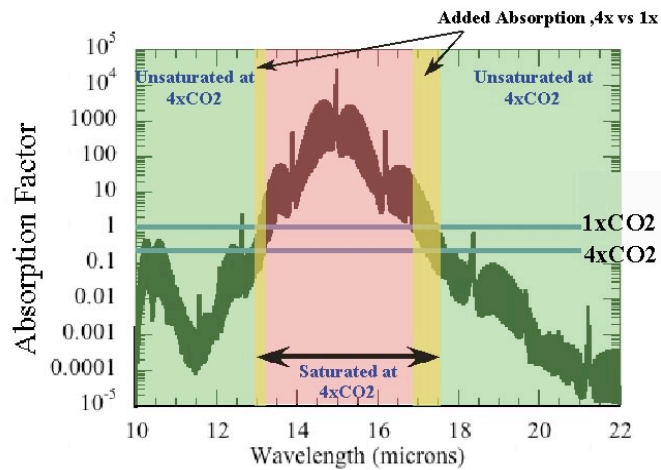
The horizontal blue lines give the threshold CO_2 needed to make the atmosphere optically thick at 1x the preindustrial CO_2 level and 4x that level. Quadrupling the CO_2 makes the portions of the spectrum in the yellow bands optically thick, essentially adding new absorption there and reducing the transmission of infrared through the layer. One can relate this increase in the width of the optically thick region to the "thinning and cooling" argument determining infrared loss to space as follows. Roughly speaking, in the part of the spectrum where the atmosphere is optically thick, the radiation to space occurs at the temperature of the high, cold parts of the atmosphere. That's practically zero compared to the radiation flux at temperatures comparable to the surface temperature; in the part of the spectrum which is optically thin, the planet radiates at near the surface temperature. Increasing CO_2 then increases the width of the spectral region where the atmosphere is optically thick, which replaces more of the high-intensity surface radiation with low-intensity upper-atmosphere radiation, and thus reduces the rate of radiation loss to space.

Now let's use the absorption properties described above to determine what we'd see in a typical laboratory experiment. Imagine that our experimenter fills a tube with pure CO_2 at a pressure of one atmosphere and a temperature of 20C. She then shines a beam of infrared light in one end of the tube. To keep things simple, let's assume that the beam of light has uniform intensity at all wavelengths shown in the absorption graph. She then measures the amount of light coming out the other end of the tube, and divides it by the amount of light being shone in. The ratio is the *transmission*. How does the transmission change as we make the tube longer?

To put the results in perspective, it is useful to keep in mind that at a CO_2 concentration of 300ppm, the amount of CO_2 in a column of the Earth's atmosphere having cross section area equal to that of the tube is equal to the amount of CO_2 in a tube of pure CO_2 of length 2.5 meters, if the tube is at sea level pressure and a temperature of 20C. Thus a two and a half meter tube of pure CO_2 in lab conditions is, loosely speaking, like "one atmosphere" of greenhouse effect. The following graph shows how the proportion of light transmitted through the tube goes down as the tube is made longer.

Transmission Through a Tube of CO_2
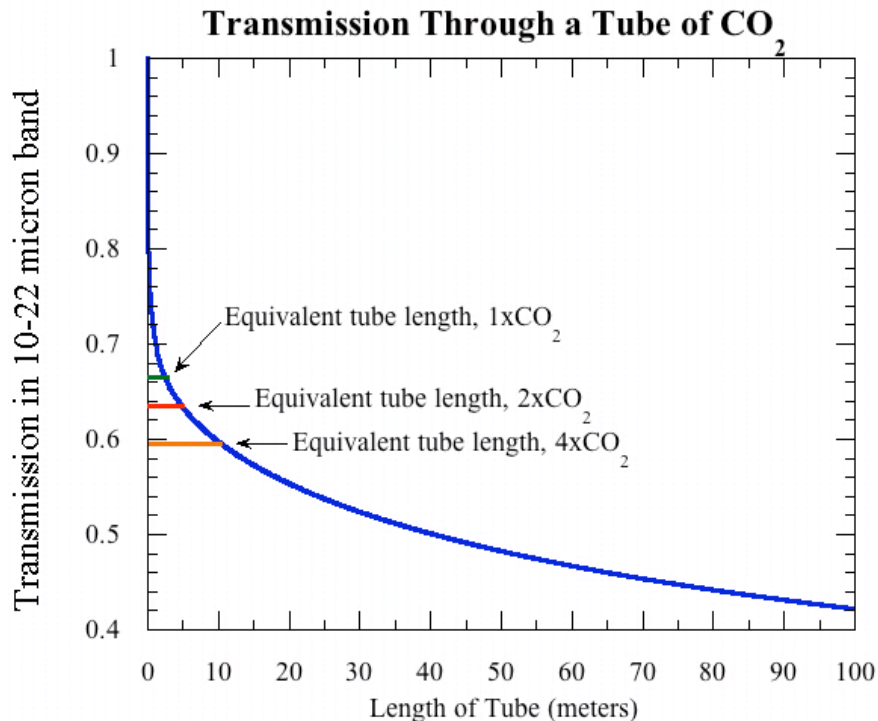
The transmission decays extremely rapidly for short tubes (under a centimeter or so), because when light first encounters CO_2, it's the easy pickings near the peak of the absorption spectrum that are eaten up first. At larger tube lengths, because of shape of the curve of absorption vs. wavelength, the transmission decreases rather slowly with the amount of CO_2. And it's a good thing it does. You can show that if the transmission decayed exponentially, as it would if the absorption factor were independent of wavelength, then doubling CO_2 would warm the Earth by about 50 degrees C instead of 2 to 4 degrees (which is plenty bad enough, once you factor in that warming is greater over land vs. ocean and at high Northern latitudes).

There are a few finer points we need to take into account in order to relate this experiment to the absorption by CO_2 in the actual atmosphere. The first is the effect of *pressure broadening*. Because absorption lines become narrower as pressure goes down, and because more of the spectrum is "between" lines rather than "on" line centers, the absorption coefficient on the whole tends to go down linearly with pressure. Therefore, by computing (or measuring) the absorption at sea level pressure, we are overestimating the absorption of the CO_2 actually in place in the higher, lower-pressure parts of the atmosphere. It turns out that when this is properly taken into account, you have to reduce the column length at sea level pressure by a factor of 2 to have the equivalent absorption effect of the same amount of CO_2 in the real atmosphere. Thus, you'd measure absorption in a 1.25 meter column in the laboratory to get something more representative of the real atmosphere. The second effect comes from the fact that CO_2 colliding with itself in a tube of pure CO_2 broadens the lines about 30% more than does CO_2 colliding with N_2 or O_2 in air, which results in an additional slight overestimate of the absorption in the laboratory experiment. Neither of these effects would significantly affect the impression of saturation obtained in a laboratory experiment, though. CO_2 is not much less saturated for a 1 meter column than it is for a 2.5 meter column.

So what went wrong in the experiment of poor Herr Koch? There are two changes that need to be made in order to bring our calculations in line with Herr Koch's experimental setup. First, he used a blackbody at 100C (basically, a pot of boiling water) as the source for his infrared radiation, and measured the transmission relative to the full blackbody emission of the source. By suitably weighting the incoming radiation, it is a simple matter to recompute the transmission through a tube in a way compatible to Koch's definition. The second difference is that Herr Koch didn't actually perform his experiment by varying the length of

the tube. He did the control case at a pressure of 1 atmosphere in a tube of length 30cm. His reduced-CO_2 case was not done with a shorter tube, but rather by keeping the same tube and reducing the pressure to 2/3 atmosphere (666mb, or 520 mm of Mercury in his units). Rather than displaying the absorption as a function of pressure, we have used modern results on pressure scaling to rephrase Herr Koch's measurement in terms of what he would have seen if he had done the experiment with a shortened tube instead. This allows us to plot his experiment on a graph of transmission vs. tube length similar to what was shown above. The result is shown here:
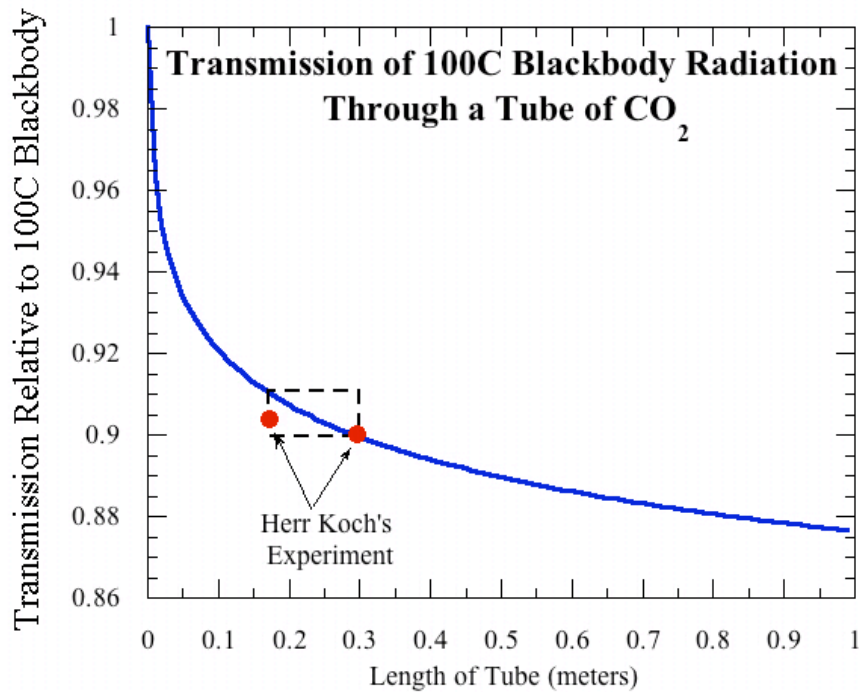
Over the range of CO_2 amounts covered in the experiment, one doesn't actually expect much variation in the absorption — only about a percent. Herr Koch's measurements are very close to the correct absorption for the 30cm control case, but he told his boss that the radiation that got through at lower pressure increased by no more than 0.4%. Well, he wouldn't be the only lab assistant who was over-optimistic in reporting his accuracy. Even if the experiment had been done accurately, it's unclear whether the investigators would have considered the one percent change in transmission "significant," since they already regarded their measured half percent change as "insignificant."

It seems that Ångström was all too eager to conclude that CO_2 absorption was saturated based on the "insignificance" of the change, whereas the real problem was that they were looking at changes over a far too small range of CO_2 amounts. If Koch and Ångström had examined the changes over the range between a 10cm and 1 meter tube, they probably would have been able to determine the correct law for increase of absorption with amount, despite the primitive instruments available at the time.

It's worth noting that Ångström's erroneous conclusion regarding saturation did not arise from his failure to understand how pressure affects absorption lines. That would at least have been forgivable, since the phenomenon of pressure broadening was not to be discovered for many years to come. In reality, though Ångström would have come to the same erroneous conclusion even if the experiment had been done with the same amounts of CO_2 at low pressure rather than at near-sea-level pressures. A calculation like that done above shows that, using the same amounts of CO_2 in the high vs. low CO_2 cases as in the original experiment, the magnitude of the absorption change the investigators were trying to measure is almost exactly the same — about 1 percent — regardless of whether the experiment is carried out at near 1000mb (sea level pressure) or near 100mb (the pressure about 16 km up

in the atmosphere).

Powered by **WordPress**