Lab 15B • 11/29/11

Chemical shift

The influence of a neighboring atom can cause a change in chemical shift. Specifically, if we had some kind of really electronegative substituent that was located next to a nucleus we're observing, then that electronegative substituent is going to pull electron density away from that nucleus being observed. Because of that decrease in the electron density surrounding it, the nucleus is going to be more exposed to the machine's magnetic field. If it's more exposed to that magnetic field, that means there's going to be a larger energy gap that that nucleus is going to experience. With a larger energy gap, that means you need a photon of higher energy in order to cross that gap. Photons of higher energy correspond to light of higher frequency. Since chemical shift is a machine-independent quantity that is directly derived from frequency, that means you're going to have an increase in the chemical shift. Electron-withdrawing groups that cause deshielding will cause a nucleus to have a higher energy gap, which corresponds to higher energy per photon, higher frequency of radiation, and higher chemical shift.

Chemical shift is only one piece of information that we get from NMR spectra. There are two other important pieces of information. The first of those pieces of information is integration. Even more importantly, we need to talk about splitting, which is a phenomenon that occurs between protons that are chemically inequivalent to each other. Let's address integration first, because that's by far the simpler topic.

Integration is useful in proton NMR because in proton NMR, there is nearly a linear relationship between the strength of the signal generated during the NMR experiment and the number of hydrogens present. Where the integrals show up is in something like this. There are two groups of peaks. Our discussion of splitting will address why do we have a group of four somewhat symmetric peaks and another separate group of three. Integrals [are often] displayed above [the peaks]. The reason why [it's often done this way], it draws what's called a trace, is because the computer is simply adding the area underneath these curves as it scans across the spectrum from left to right.

In older machines, it was not possible to do Fourier transform NMR, where you might do one frequency of observation at a time, or you might change the magnetic fields strength and scan at one frequency. The results from those scans had to be plotted in some way. Of course, back 40 years ago, there weren't inkjet printers like we have today. Instead, they had what are called strip charts. What the would do is the would have, for example, a constant change in, for example, observation frequency. That change would be coupled to the track of a piece of paper coming off of this strip recorder. There would be a plotter pen, literally a little pen, that would move back and forth across an arm as that paper's continually pulled out. Again, that paper would progress at a constant pace, in the same way that the frequency that you might be observing at would change at a constant pace. So as the frequency would change, the signal is generated, that plotter would respond, and eventually, by the end of the experiment, you would have a plot that's been generated. Why am I mentioning this? Integration, that's the area underneath a curve. What if this was 40 years ago where not every lab necessarily had a computer available to do that integration. What would be done is you'd take that plotter paper and you'd cut out the individual piece, because if the density of the paper was roughly a constant, then that means there was a linear relationship between the area of those peaks and the mass of them. If you had to try to determine the integrations, you would cut out those pieces of paper, weigh them, and use the masses to figure out the relative ratios of the area. We don't have to do that any more; we do have this computer technique.

If the paper was a constant thickness and the paper was a constant density, then there would be a linear relationship between mass and surface area. If you took the mass of the different peaks, you indirectly got at least a ratio of the different surface areas. But it's exactly that idea of ratio that comes up in this example, even with the assistance of a computer. I'm drawing up two different kinds of cyclic ethers, the second one of which is often called a crown ether. If you used your imagination and imagined that the oxygens are like the tops of the spikes or a fancy crown. The reason it's called a crown ether is those oxygens are able to circle and encompass different kinds of ions. Exactly which ions they prefer to bind to depends on how large the molecule is; whether it's oxygen versus nitrogen at those different positions on the interior of the ring; they're extraordinarily useful because it lets you take ions, complex them with this organic molecule, and then have those ions therefore be dissolved in organic solvents, even though they're ions.

But there's still one gotcha, even with the computerize technique, which is something that we can see in the following pair of examples.

If we were to ignore considerations of stereochemistry, how many unique hydrogens do we have on the lefthand compound, 1,4-dioxane? Another way to ask that is if I radically halogenated that molecule and talked about all of the possible monohalogenated products, and again ignored stereoisomers, how many unique products would I get from monohalogenation? The answer is one, because the molecule's entirely symmetric. Each carbon is chemically equivalent to the other, other than any discussion about diastereotopicism or enantiotopicism.

If we look at the second example, it's actually exactly the same situation: it's entirely symmetric and each one of those hydrogens, ignoring stereochemistry, is equivalent to the other. If we were to take a scan of either one of these compounds, they would, in fact, each generate just one signal. In fact, I would guess that these are close enough structurally to each other that that signal very likely would show up at, if not the same, then awfully close chemical shift for both of those compounds.

What about integration, though? What would the computer call either of these integrations? It wouldn't, because although the intensity of the signal is linearly related to the number of hydrogens, let's say that we prepared two different samples with two different concentrations of the these two different materials? What if each compound had a slightly different sensitivity to the NMR spectrometer itself? In other words, if you look at the first compound and say, "that's an integration of one", cause it has no other point of reference; it could look at the other compound and say, that's also an integration of one, because it has no other reference. When we do integration, you are having to tell the computer the number of hydrogens that do correspond to a signal, because although the signal strength is proportional, the computer doesn't know without measurement what that proportionality constant is. There are ways to determine that. Both of the ethers shown above would generate nearly identical proton NMR spectra, since, ignoring stereochemistry, each compound contains only one type of hydrogen. Again, although the number of hydrogens is linearly proportional to signal strength, that proportionality constant can only be determined by measurement. Without measurement, without that constant, the machine would be unable to distinguish between the two compounds.

[Will always be given integrals for problems]

You might remember that for NMR solvents, we often use not completely deuterated solvents, but mostly deuterated solvents, so that that little bit of residual solvent that has hydrogens could be used as a calibration. In NMR solvents, often you'll add TMS to have a calibration. What if you prepared an NMR sample where that TMS's concentration was known, let's say, to four decimal places. When you put that sample into the NMR spectrometer and acquire the signal with TMS in it, you could get that signal strength and, with the structure of TMS having twelve hydrogens per molecule, and knowing the concentration of TMS, you could figure out exactly the number of moles of hydrogens that are in that solution and how that corresponds, then, to the signal strength that's generated. If, at the same time, in that same solution, you also have a compound that you want to analyze, [and] you know its exact concentration, then you could take that same ratio and figure out: all right, if TMS has this response and it generated a signal this strong and it corresponds to 12 hydrogens at this concentration, then using the concentration of your compound and its signal strength, you could figure out how many hydrogens are on it. That's a quantitative NMR experiment; it is one way you could resolve the difference between these two compounds, but that does require known the concentration of the compounds.

Splitting (as in splitting headache)

[glossing over how/why splitting occurs] Hydrogens that are chemically inequivalent to each other can affect each other's magnetic fields. Hydrogens that are equivalent all have the same magnetic environment, and so they don't interfere with each other's magnetic fields. Hydrogens that are chemically inequivalent to each other can perturb each other's magnetic fields. Hydrogens that are chemically equivalent are part of the same magnetic environment, so one equivalent proton will not perturb the magnetic field of another. What exactly is this interaction that occurs? Let me create an example molecule so we can discuss a very concrete case. To simplify all of the arguments that I'm going to make today, I'm going to start with this disclaimer statement: in all of the examples shown, we will ignore the effects of stereochemistry.

The first example molecule is going to be the following: 1,1-dibromo-2,2-dichloroethane. The molecule only has two hydrogens on it, and those two hydrogens will definitely be chemically inequivalent to each other, because we have one hydrogens that's on a carbon that has two chlorines, one hydrogen is on a carbon that has two bromines. If we wanted to predict what we know so far about NMR theory what the spectrum of this compound might look like, we could guess the following: that each unique type of hydrogen generates its own NMR signal, so we should see two peaks in the spectrum. Hydrogen A versus hydrogen B, which of these should be more downfield? Which of these would be more to the left of the spectrum: Ha or Hb? Ha? Why? Chlorine's more electronegative than bromine, so there's going to be more electron withdrawing that occur around position A, so it's going to have the higher chemical shift. So, our guess, not knowing anything about splitting, might be the spectrum would look like this. But, one proton does affect the other. What will be seen in the spectrum is this: that each peak is, instead, present as a pair of peaks. I've drawn all of those peaks being equal height. It is possible that, in a real-life spectrum, one pair might be higher than the other, in terms of height, but because both of those positions contain exactly one hydrogen, the area of the curves should end up being the same.

Why is it that we have this doubling effect? Imagine the following. Let's say I'm proton A, and I'm being observed. I'm interacting with proton B next door. Proton B has a different magnetic environment than me, proton A. Proton B might have its spin aligned with the external magnetic field, the one that the entire machine is under. If so, then from my perspective, A's perspective, that neighboring spin effectively adds to the overall magnetic field. If it adds to the overall magnetic field, that's going to cause a wider energy gap, which, again, will cause the chemical shift to shift to the right.

But another possibility is proton B, my neighbor, instead has a spin that's down, and if it's down, oriented opposite of the main magnetic field, then that's going to decrease the magnetic field strength. If you decrease the magnetic field strength, that means the chemical shift is going to end up shifting to the right. Statistically, both possibilities occur. As the sample is being irradiated, the spins are going back and forth and effectively they're in roughly equal chance as far as what kind of interaction you have of a neighbor with a hydrogen being scanned. Throughout the whole solution, half of those spins will be [oriented] up, and so half of the signal will be shifted to the left, and the other half, the spins will be [oriented] down, and so the peak will shift to the right. Both occur equally, so one signal ends up splitting into to. That is the splitting effect. Let me write out this example. We're going to be talking about the effect of B on A. I could easily talk about the effect of A on B, the exact same thing's going to happen, but just so I don't tell the whole story twice, I'm going to focus just on B interacting with A.

A neighboring proton – in this case B – that is non-equivalent to a proton being observed – which is going to be A – can effectively add to or subtract from the magnetic field, as experienced by A. The neighbor can therefore cause an increase or decrease in the chemical shift for proton A. Statistically, both possibilities occur equally, so the signal for proton A appears to split. When we have this phenomenon, where one neighbor causes another hydrogen to have its signal split into two, that's called a doublet. There's a graphical way in which the split can be represented, which is if we imagine that we start out with the signal for A, and we first imagine that that signal starts out as a single signal, if we say that there's no interaction with neighbors, we can then consider that this whole thing is going on in the machine, and we can assign the direction of the machine's magnetic field. If we say that neighbor B is aligned with the magnetic field, then it adds to the magnetic field as experienced by A, so if we're looking at chemical shift, it's going to cause that chemical shift to move to the right. What I'm showing is the result of having that neighbor spin up, but the other possibility is that it is spin down. That will have a subtractive effect on the magnetic field to the same degree that it would have been additive if it had been flipped the other direction. This is what results: one peak splitting into two. Depending on exactly which hydrogens on the compound are involved, that will determine how much of a split occurs. The degree to which that split occurs is expressed by a quantity called J, known as the coupling constant – how much interaction is there between two different types of hydrogens. [We can use coupling constants] to try to identify, in an alkene, which pairs of protons might have a cis relationship to each other, a trans relationship, or a geminal relationship, geminal meaning the two hydrogens involved are on exactly the same carbon. This is the case of having just one inequivalent neighbor.

We're going to next do a case of two inequivalent neighbors. Second case is very similar. We'll still have one hydrogen, which I'll still label Ha, that will be on a position that has two chlorines here, but now I'll only put one bromine on the other carbon, which means that carbon will have two hydrogens that, if we ignore stereochemistry, are chemically equivalent to each other. Where position A is, that is not a stereocenter, so at most, these protons would be enantiotopic, not diastereotopic; as such, if all I do is say that all of this is being discussed in an achiral solvent, they would still be chemically equivalent to each other. To simply, that's why I made that catch-all statement at the beginning that we're going to ignore stereochemistry in our discussions.

Each of these Hb protons could have their own spins, independently being oriented up versus down, relative to the external magnetic field. Since each spin is independent of the other, that means we really have four unique spin combinations, where both spins are up, where the first spin is down the second spin is up; an equivalent energy case but still a distinct case where the first proton is up and the second is down; and then the last case where you have both of the spins pointed down. I say that these are two equivalent cases, one up/down versus one down/up, because since the spins are independent of each other, then both of those possibilities are distinct statistically. I say they're equal energy cases because you have one spin up and one spin down, those two would effectively cancel each other's effect, so whatever peaks would result would be in exactly the same place that the original peak would have been in, but because you have two unique versions of that same result, the peak will be double the height in comparison to the other peaks that show up. Let me show you this graphically so you can see what I'm talking about.

I'll do this in two different ways. Let's first talk about the different possible energy states. I said one energy state is where the two neighbors are both spin up; another energy state is where both neighbors are spin down; then we have the two unique cases of one up/one down, one down/one up – that would be the same in energy as the original peak to begin with, and because there's twice as many possibilities of it occurring, then that middle signal would be, if not exactly double in height, double in area, compared to the two other signals that would be generated. This kind of formation, because it consists now of three subpeaks, is known as a triplet. Where one spin would add to the magnetic field, the other would subtract again. If we did a splitting chart, we'd see the same thing, or what's called a splitting tree.

The way a splitting tree works is much like what we did up above. We start with one signal, and you say, what would happen if this neighbor interacts with it. But then you could keep on adding additional neighbors and seeing how that would further split the signals that are generated. In the first case, we started with a signal for A; we said allow one neighbor to interact with it. That caused the generation of two signals that were separated by a certain distance expressed by this thing called the coupling constant. The top would be no neighbors that are interacting; the next is one neighbor. But then, each of those would be split would be split by a second neighbor, because its spin is independent of the first one, so each of these two signals that I've just drawn, you would need to split further.

But because it's the same neighbor interacting, the split will have exactly the same coupling constant. As such, the signal on the left, the greater chemical shift, when it gets further split, it will generate one peak that's in the same position that will result when the other signal is split, again because they have the same J values, the same extent of splitting is occurring. That means, coincidentally, those energies would line up with each other – exactly what we predicted just on the basis of the spin states themselves. Really, these are the same picture; I've just developed it in two different ways: one is just to look at all of the spin possibilities; the other one is two tackle the splitting one proton at a time. They both give you the same results that you'll have overall integration of one to two to one. We're just arbitrarily assigning the magnetic field of the machine to be pointed in that same direction.

Notice the pattern's that developed here. With one neighbor, we got a split into two peaks; with two neighbors, we got a split into three peaks; what would you imagine would happen if we had a methyl group? If you had three neighbors (that's a methyl group), you might think that we should get four peaks – that is exactly what happens. If you have three spins, three independent spins, you've got four unique energy possibilities: all of them up; two of them up; one of them up; none of them up. But those aren't all of the possibilities period, because you could have different spins up or down. In fact, you can have the spins up/up/up, up/up/down, up/down/up, up/down/down, down/up/up, down/up/down, down/down/up, and down/down/down – eight unique possibilities. All I did right then was say the numbers zero through seven in binary; it's another way of thinking about the possible spin states. We're developing simple splitting trees at the moment where we're only worried about one type neighbor. What we're going to come up against next is the situation where you'll have three different neighbors; they'll each have three different coupling constants, and so you won't have the different spin states aligning on top of each other nearly so cleanly as you have in this case. We're tackling these simple cases first because we need to use the same technique to analyze more difficult situations in the future.

Let's go ahead and see this splitting tree and see how these eight different energy states would align. Example molecule would be one now in which we have a methyl group; we still have the same Ha on a carbon with just two chlorines attached. There are the eight different energy states, but some of them are degenerate, which means equal in energy. The eight possibilities were [up, up, up], [up, up, down], [up, down,up], [up, down, down], [down, up, up], [down, up, down], [down, down, up], and [down, down, down]. Notice that we've generated those four unique types of energy that I mentioned – all up; two up; one up; none up. Notice that two of those energy states have degeneracies; the second and third, there are three of that same energy statistically, which means the peaks generated will be three times the height of the ones on the outside. You would see represented on a spectrum something like this, with relative heights of 1:3:3:1. We could do what we did just previously and draw a splitting tree. Each time we split, we're still being split by the same Hb, the same neighbor; that's why these energy states are going to coalesce and overlap on top of each other. We'll start out with just the one peak where we had no neighbors; we'll have the first split that makes a doublet; when that doublet splits, you make a triplet; and when that triplet splits, you'll make a quartet. This is known as a quartet, because you have four peaks within that sub-peak.

There is a generalized rule. First, let's go back and look at this pattern. Notice that the first case we had peak intensities of 1:1; the second case, the triplet, we had peak intensities of 1:2:1; the quartet, 1:3:3:1; if somehow we could end up with four neighbors – which we couldn't have that on one individual carbon, because if we had four hydrogens on a carbon, that's the molecule methane, which doesn't have any neighbors. But what if you had two neighbors on one side, two neighbors on the other side, and the four of them were chemically equivalent to each other? You could imagine that situation in a molecule like this. Those four hydrogens that I've indicated, ignoring stereochemistry, would all be equivalent to each other. If you had four neighbors, you might guess that you'd end up making a quintet, a set of five peaks. What do you predict that the heights of those peaks wil be? 1:4:6:4:1. How [do] you get that? What is this really that we're developing here? Pascal's triangle. If you don't remember what Pascal's triangle is, you might have ended up discussing that in your math classes when discussing polynomial expansion. If you had, for example, (x + y) raised to the first power, that is 1x + 1y. If you have (x + y) squared, that's 1x^2 + 2 xy + 1 y^2. If you had (x + y) cubed, that's 1 x^3 + 3 x^2y + 3 xy^2 + 1 y^3. If you just let x = spin up and y be spin down, then the patterns generated by Pascal's triangle, the patterns generated by polynomial expansion are exactly the same as the patterns generated by splitting. It's just one of those interesting mathematical coincidences. If you don't remember Pascal's triangle, the way that you generate each succeeding row of the triangle is to look up above, and where you have the two numbers, add them together and put the new result between the two numbers. Notice that's what's effectively happening where you have splitting where the coupling constant is the same for each interaction. One plus two ends up making three. In the case of generating Pascal's triangle itself, we add ones to either side of the rows after we've done all these different additions.

Let me summarize these effects in what is called that (n + 1) rule: that when observing a proton with n equivalent neighbors, a multiplet – that's a generalized word to express all of these different spin possibilities – a multiplet with (n + 1) peaks will be generated.

Let me briefly show you how this information's eventually going to be used. [way information will be provided in non-spectral form normally] What information do we have from this? We see that there's a group of signals that seems to be centered at chemical shift of 1.5, and there's a group of peaks that's centered at chemical shift 3.2, so we'd be given two listings. Delta, that's the symbol for chemical shift. Delta 1.5, delta 3.2. Next to each of those numbers would be listed the integration.

I'll just tell you for the peak on the left, the integration is two, and the peak on the right, the integration is three. You'd then also be given the splitting information. Can you tell me what kind of signal I have for this lefthand peak? It looks kinda like a quartet. It's not perfect, but that may just be due to the phasing of whoever originally took this spectrum; with a more careful experiment, you might be able to get a more perfectly one-third the height for each of these shorter peaks compared to the taller peaks. We have absorbance at delta 3.2 with an integration of 2 split into a quartet. Then we have a peak at delta 1.5 that has an integration of 3 split into what appears to be a triplet, because you have 3 peaks and, very roughly, the heights appear to be 1:2:1. Again, the heights themselves technically don't have to be 1:2:1, the areas under those curves have to be 1:2:1. What if I tell you that this chemical formula is C2H5Br. Well, with just that piece of information, you can instantly figure out the chemical formula, cause there's only one formula possible from C2H5Br; let me show it to you – it's bromoethane.

But, how does this spectrum match the fact this is bromoethane? I told you earlier that the peak on the left integrated to two and the peak on the right integrated to three. The peak on the left is further downfield, which would seem to mean that it's closer to the bromine. We have the bromine that is next to only two hydrogens as its immediate neighbor; that would match with the fact that this peak has an integration of two. If we look at the carbon that's further away, the methyl group, it should have a lower chemical shift because it's further away from the bromine, which is does. That's the peak that I indicate had an integration of three. As far as splitting, the methylene – remember that's our common name for a CH2 group – the methylene next to the bromine has three neighbors, which means it should be split into a quartet. That's again exactly what we see. Then the methyl group has two neighbors, which it's split into a triplet; yet again, that's exactly what we see. This is the way that I'll the molecule together. For example, whenever we see a quartet, our first instinct will be to say: oh, that means there's a methyl group right next door. If you can demonstrate that, you've found the end of a molecule, and from there you could rationalize your way into creating the rest of the structure of molecule by comparing how many hydrogens are located at that position, how many hydrogens are located next to that position. Integration will tell you how many hydrogens you have in one place, and splitting will tell you how many hydrogens you have next door. If you tried to conceptualize that, imagine that you had tiles, where one number on the tile is how many hydrogens there are, then next number's how many are next to it, and then you assemble all of that conceptually together. You make the structure of the molecule. [plan for next quarter]

In a full problem, I'll give you an NMR spectrum; I might even give you an IR spectrum, because it's often easier, quicker to determine certain functional groups' presence through IR instead of NMR, and any bit of hint that we get about the structure, it makes structural elucidation even easier. Besides those two things, I'll also almost always give you the chemical formula as well. From the chemical formula itself, we can get some structural information, because there's something known as the degree of unsaturation. The term unsaturation has come up before when we were talking about compounds having only single bonds versus double or triple bonds. Technically, though, even compounds with just single bonds might not be fully saturated; we're following the more rigorous definition of what unsaturation is. Saturated refers to a compound that has the maximum number of hydrogens possible given the number of carbons present. I'll just state that, aside from other factors that we'll discuss in just a moment, that the number of hydrogens that could be present is equal to 2C + 2.

Let's see a simple example that we could use to rationalize this formula – that simple example would be something like hexane. Hexane has two methyl groups and four methylene groups. Each of those middle methylene groups has two hydrogens each. That's where the 2C portion of this formula comes from. Even on the methyl groups, we of course have two hydrogens, but on each methyl group, we also have a third hydrogens. On a straight-chain molecule like this, we only have two ends, so two extra hydrogens, that's the source of the plus two in the equation. So, 2C + 2. Even if we didn't have a straight-chain compound, even if we had something that had branches in it, we're still gonna find the formula will apply. Let's say we had something like isopentane. We could look at it and say: there are three methyl groups, so 3 times 3 hydrogens; there's a tertiary carbon that has one hydrogen; there's a secondary carbon that has two hydrogens; all together, that would be 12 hydrogens. If C = 5, 2C + 2 = 12; that matches our prediction.

But what would happen if we tried to make a cyclic compound, if we had something like cyclohexane? Cyclohexane only has the formula C6H12, even though 2 x 6 + 2, that should give us 14, if we have six carbons. But we have two less because, in order to make the ring, we would have effectively have to have removed a hydrogen from either end of plain old hexane so that a new carbon-carbon bond could form. In fact, every time that we make some extra connection – whether it'd be to form a ring like this, or whether it is to make a double bond like hexene – both of those, in fact, have the same chemical formula with six carbons, because in both cases we would have had to pull off a pair of hydrogens to make either the closure of the ring or this double bond. One pair of hydrogens removed, that's referred to as a degree of unsaturation. Both a ring and an alkene – the present not even of an alkene, any kind of double bond – those would each count for one degree of unsaturation. Again, by degree of unsaturation, we effectively mean two hydrogens are missing. The only other structural feature that will cause unsaturation is a triple bond. If we had something like pentyne, that has five carbons, but only eight hydrogens. C5, that should give us 2 x 5 + 2, which is 12 hydrogens; we only have 8, so two pairs of hydrogens are missing, each one corresponding to one of the pi bonds now present. Triple bonds end up counting as two degrees of unsaturation; that corresponds to four hydrogens effectively being missing. We could come up with a generalized formula for degree of unsaturation, because since the number of degrees of unsaturation depends on the number of pairs of hydrogens missing, we compare how many hydrogens we expect to have, subtract out how many hydrogens we really do have, and divide that number by two.

So, our initial formula for degree of unsaturation will be: number of hydrogens based on the number of carbons minus number of hydrogens present, and that whole thing divided by two.

I say our initial version of this formula, because we do have to take a couple of things into consideration. So far, we've only seen hydrocarbons, but what if we had an oxygen, or a halogen, or a nitrogen, very common atoms found on an organic compound. Let's see how those different types of atoms have, or don't have, an effect on how many hydrogens we think should be present. I'm going to divide this up into three cases, the three cases that we'll normally encounter in regular organic compounds. I'm going to use ethane as my comparison compound, and I'm only going to write out formally the hydrogens on one of the two carbons of ethane. My first comparison is going to be between ethane and ethanol, where we've added an oxygen. Notice, though, that that oxygen doesn't seem to have had any effect on the number of hydrogens that would need to be present in order to generate a structure that doesn't have charge and has the correct number of connections to oxygen. Oxygen, in other words, likes being divalent, which means it can just be slipped in between a carbon-hydrogen bond. Chalcogens, such as oxygen and sulfur, have no effect on the hydrogen count, because they're divalent.

Another case, then, would be to compare ethane and ethanamine, where we have a nitrogen present. Notice that, in order for the amine to be neutral, there is one additional hydrogen. That's because pnictogens like nitrogen prefer to be trivalent. So, nitrogen and phosphorus add to the hydrogen count because [they] are trivalent. That's nitrogen being left of oxygen on the periodic table. What if we go to the right and look at halogens; we end up with the opposite situation. Although, halogens are able to form polyvalent compounds: think of the ion iodate, for example, iodine having connections to three oxygens. In organic compounds, the far more common case is for a halogen to be monovalent, meaning only having one connection. If that's true, then if it's going to take a position on a carbon-based molecule, you have to get rid of a hydrogen, wherever that halogen would be substituted. Effectively, we are missing one hydrogen. Halogens effectively subtract from the hydrogen count, since they are normally monovalent.

If we wanted to take these considerations an incorporate them into our original formula, we could come up with the overall degree of saturation formula, which is that you still start with that base number of 2C + 2, you then add one for each nitrogen, phosphorus, or other pnictogen, you then subtract for each halogen; oxygen and the other chalcogens we saw don't affect the hydrogen count, so we don't even need to mention them in this formula. This, so far, would be the adjusted number of hydrogens that should be on a fully-saturated compound. We're still then going to compare that to how many hydrogens are really present, so we'll subtract that number of hydrogens, and then that whole darn thing we divide by two.

Why is this at all useful? Let's say that we have one of these problems where you have both IR and NMR available. Let's say that you have a chemical formula that you're given also that indicates you only have one degree of unsaturation. Let's say in your IR, you see a carbonyl, a carbon-oxygen double bond, because those are very readily apparent by IR. A carbon-oxygen double bond, that is the degree of unsaturation itself. If you see in your IR spectr[um] that's there's carbonyl, you know there are no carbon-carbon double bonds; you also know that, since you've already accounted for the one degree of unsaturation, there would be no rings on that compound. Sometimes, just those little pieces of information can be enough that it allows you to much more easily determine the structure of a compound from a NMR spectrum.

You can, in some circumstances, still determine the structure without a formula. But, if you think of those two different ethers I showed you earlier, that would be a situation where you would need to have a formula, or some other quantitative information. Sometimes you can do it without, be there are also cases where you do need it.

Imagine that you're sticking nitrogen between a carbon-nitrogen bond. Then one connection from the nitrogen would be to the carbon, a second connection would be to the hydrogen that's already there, and then you need one more hydrogen to make that third connection, but only one more.

Gas chromatography

Let me show you a very simplified diagram of what physically a gas chromatograph looks like. It's usually got four major components to it. This is gas chromatography, which means the mobile phase is a gas itself. If we're analyzing samples in the gas phase, but most of our organic samples are going to be solids or liquids, we have to vaporize that sample, so the first step in a GC experiment is, generally, to inject a very tiny portion of that sample into a heated injection port. That sample, then, passes into what's referred to as a column. The word column would indicate some tall, cylindrical object, but if we imagine that we took a cylinder and just continued to stretch it and stretch it and stretch it and make something that loops around, that's still, by extension, referred to, in chromatography, as a column. That column is generally placed inside an oven, so that the sample can be maintained in the gas phase as it passes through the column. Once it gets through the column, it will then hit some form of detector, which yet again is usually heated to maintain everything in the gas phase until the sample's completely out of the machine.

This column, in turn, is usually like the following, where the outside of that column is metal, usually stainless steel or something else that is likely to be unreactive; that's to provide physical support to the column, since this is something that's going to get heated. On the inside of that column is the stationary phase, the separatory material that can distinguish, separate different kinds of compounds. Your sample is going to pass through the column by the action of what's known as a carrier gas; there's one more thing in this system, the carrier gas itself. The carrier gas would be much like the solvent in a TLC experiment. In the TLC experiment, we spotted compounds on a plate; those compounds did move, just like the compounds will move in this GC. But something moved those compounds in TLC, which was the development solvent. Similarly, something will move the compounds in GC; that is the carrier gas. [The] carrier gas is floating through, therefore, at the same time that the sample is floating through. Based on what kind of sample we're trying to analyze, you might actually have different temperatures at your injection port, your oven, and your detector. On even fancier chromatography devices, you might have temperature profiles or temperature programs, where maybe you start everything off at a lower temperature to encourage only the lighter molecular weight, lower boiling point materials to come through quickly, but then if you have larger or, to anthropomorphize it, more recalcitrant compounds, you could heat the whole device up to cause those compounds to come through more quickly.

In gas chromatography, a liquid sample is injected into a port, which is heated, so that the sample can be vaporized or volatilized upon injection. The sample is then passed through a column that is also heated, so as to maintain the samples in the gas phase. A carrier gas, which is generally some form of inert gas, which usually means nitrogen, helium, or argon, is used to help push the sample through the column. The sample is passed through a column that is the stationary phase, and then a carrier gas, which would qualify then as the mobile phase – it's usually an inert gas, such as nitrogen, helium, or argon – that gas is used to help push the sample through the column.

The behavior of gasses is obviously different than that of liquids. There's going to be two major effects on why would compounds separate from each other as they pass through this column. One of the reasons is polarity – exactly the same reason we had separation occur in TLC. If you imagine that we have a very polar coating on the inside of this column, and we pass a polar compound through it, the compound will interact heavily with that compound and won't make it through as quickly to the end of the machine. Same scenario, polar coating, a non-polar compound passes through; it won't have much interaction with that coating, so it'll be more easily pushed through by the carrier gas. So, there's going to be a time factor involved. But besides just plain polarity, there is the fact that different mass molecules have a different velocity to them, even at the same temperature. So there's some influence of the size of the molecule or the boiling point on the molecule on how rapidly it'll pass through the machine. On the basis of polarity, and I'll simplify by expressing it as boiling point, the compounds will be separated as they pass through the column. As the sample passes through the column, the components of the sample are separated, based mainly on the basis of polarity and boiling point.

As the compounds successfully pass completely through the column, the quantity of a compound that's exiting the machine at any one time is measured by a detector. There are different kinds of detectors that can be used for different sorts of circumstances. The type of detector that we have is a flame ionization detector. To simply how that works, imagine that you have a really hot, thin piece of wire, a filament it's called. Imagine that we pass a current through that wire. That wire is particularly thin and of the right material, it's conductivity can be influenced by different compounds that pass over that wire. As these different gas sample come through the column and come across the detector, that change in conductivity corresponds to the quantity of compound that come through. There are also spectroscopic detectors; you might have a general, broad UV detector that could be used to determine the quantity of compound that comes through. Exactly which detector you use depends on your application.

[The results look like the following:] the graph generated – I say graph, not spectrum, because you're not passing different frequencies of light through this, you're really just measuring how long does it take for a compound to come through; that is the x-axis, the retention time. That retention time is the time it takes for a compound to pass from the injection port to the detector. The graph may look something like this. For a well-established system, or if you have a set of reference compounds, you can determine what the retention time should be for the various individual components and then, much like the fact that in TLC if you have a reproducible protocol Rf values can be used to identify compounds, in the same way, if you have well-established protocol for a gas chromatograph, retention time is used to determine the identity of compounds. Since, in most applications, the response by the detector is linearly proportional to how much of a compound is passing through it at any one point, we can use the integral, the area of these curves, to determine the relative quantity of each one of the components present. If you know the initial volume of the sample, you can figure out the volume of the products coming off and therefore analyze the composition of a mixture. Retention time is used to determine identity; the quantity of a compound present is determined by the area of the peak observed.

[lab exercise]

_____

Integration – For 1H-NMR, there is a linear relationship between the strength of an NMR signal and the number of hydrogens that correspond to that signal.

Both of the ethers shown above would generate nearly identical 1H-NMR spectra, since (ignoring stereochemistry(each compound contains only one type of hydrogen. Although the number of hydrogens is linearly proportional to signal strength, that proportionality constant can only be determined by measurement. Without that constant, the machine would be unable to differentiate between the two compounds.

Splitting – Hydrogens that are chemically inequivalent to each other can perturb each other's magnetic fields. Hydrogens that are chemically equivalent are part of the same magnetic environment so one equivalent proton will not perturb the magnetic field of another.

In all examples below, the effects of stereochemistry are ignored.

A neighboring proton (B) that is non-equivalent to a proton being observer (A) can effectively add to or subtract from the magnetic field (as experienced by A). This neighbor (B) can therefore cause an increase or decrease in the chemical shift observed for A. Since both possibilities occur with equal likelihood, the signal for A will appear to be split [in two]. –> doublet

n + 1 rule – when observing a proton with 'n' equivalent neighbors, a multiplet with n + 1 peaks will be generated.

Degree of unsaturation

Saturated – a compound that contains the maximum number of hydrogens possible, given the number of carbons present.

#H = 2C + 2; 3 x 3 + 1 + 2 = 12; C = 5; 2C + 2 = 12

A ring an a double bond each count as one degree of unsaturation.

A triple bond counts as two degrees of unsaturation.

D.O.U. = (# of H based on # C) - (# of hydrogens present) / 2

Oxygen (and sulfur) does not affect the hydrogen count as it is divalent.

Nitrogen (and phosphorus) add to the hydrogen count because it is trivalent.

Halogens subtract from the hydrogen count since they are (normally) monovalent.

D.O.U. = ([2C + 2] + N - X - H) / 2

Gas chromatography

In GC, a liquid sample is injected into a port which is heated, so that the sample can be vaporized upon injection. The sample is then passed through a column that is also heated so as to maintain the samples in the gas phase. A carrier gas (mobile phase) (an inert gas, usually N2, He, or Ar) is used to help push the sample through the column.

As the sample passes through the column, the components of the sample will be separated based mainly on the basis of polarity and boiling point. As the compounds successfully transit through the column, the quantity of the compound exiting the machine is registered by a detector.

retention time is used to determine the identity of a compound. the quantity of the compound is determined by the area of the peak observe.
retention time –> time it takes for a compound to pass from the injection port to the detector.

Structures *(remaining structures identical to lab 15A)*

11/29/11B lab • 1