

Lab 11A • 11/02/11

The next experiment is dehydration.

The dehydration experiment, the mechanism I'm going to show you for it, portions of it are very similar to what is known as an elimination reaction. So far, we've talked about nucleophilic substitution, we've talked about S_N1 and S_N2 , but we will be learning about these $E1$ and $E2$ mechanisms, these elimination mechanisms. Let's see the mechanism to see what I'm talking about.

What is the starting material being used in today's lab? 4-methyl-2-pentanol. What's the reagent you're going to react with it? Sulfuric acid, and what else, effectively? Heat. You're going to be heating the crap out of this to accomplish a dehydration. The first step of the mechanism is much like you saw for the previous experiment. In that previous experiment, we also had an acid around. If we had an $-OH$ group, an alcohol, and that's the only functional group on the molecule, that's really the only thing that's going to be able to react with an H^+ . So the $-OH$ group gets protonated; we end up making water. This water can leave as a way of relieving the positive charge on oxygen. It's at this point where something different can occur, and high temperatures help influence this outcome: instead of something coming in and substituting, what if, somehow, instead a hydrogen next door somehow reacts. Some kind of base and come along; water's not a great base, but compared to sulfuric acid, acid is a base. Some base can come along and grab that proton. If that occurs, the carbon-hydrogen bond will break, and that lone pair that would exist in that bond can collapse back into the carbocation, which is going to end up forming an alkene. This mechanism is very related to elimination, as it does produce an alkene.

Let's revisit why this mechanism would be possible; why would a hydrogen want to come off? In order to do that, I want to draw a SMOG. I don't want to draw the whole SMOG for the whole molecule; I just want to write a SMOG for this portion of it – not even including water, just this portion of the organic part. The carbocation itself in that position, what hybridization is going to be? How many hydrogens are on that place where the carbocation is? One. Pluses and minuses seem to still confuse a lot of people. If you see a positive charge on a carbon, it's because there's a hydrogen missing. The hydrogen went missing with its pair of electrons. If you see a negative charge, again that means a hydrogen went missing, leaving behind the lone pair. If you see a single dot, it means you're missing a hydrogen, and it only took one electron away with it. If I have a plus charge written there, and carbon has to have four bonds, and the plus charge is effectively counting for one of them, and then I see two carbon-carbon bonds that are formally written in, then that leaves one hydrogen on that position. If I have the one hydrogen and the two carbon-carbon bonds, that means it's going to be sp^2 -hybridized. What about either of the carbons on either side of the carbocation? sp^3 . [admonition to study SMOGs] So we draw that SMOG. sp^3 means we have four orbitals on that atom, four hybrids, that is. sp^2 means we have three. The rest of the molecule I'll just label as R, which stands for radical. sp^3 -hybrids have nothing else on them; sp^2 -hybrid means there's going to be a p orbital there. Of course, we do have the hydrogens running around. [review of color scheme]

What is hyperconjugation? Partial donation of electrons from one orbital to another in a situation that would not be classified as bonding. I have to be careful; I cannot use the word 'non-bonding' any more, because you might think I'm talking about a non-bonding orbital. It's a partial donation in situations in which it doesn't involve the formation of a bond. Hyperconjugation is different from conjugation. Conjugation is like the allyl system – either the allyl cation or the allyl anion, where you have a p orbital that is full-on interacting with another p orbital and another p orbital; in future molecules, we learn even more p orbitals. Full overlapping which causes the formation of new molecular orbitals. Hyperconjugation is where just a little bit of density is transferred. Like in this case here: we have a p orbital, but we don't have p orbitals on either side of it to make a pi bond, so we're not going to make a new bond by this interaction. But, the sigma bond that's here has some electron density that can be donated in. That bit of electron donation helps make a more stabilized carbocation.

I want to draw SMOGs for the following four molecules. The central atom on all four of these is going to be sp^2 -hybridized; wherever else we have carbons, those are going to be sp^3 -hybridized. I'll label these 3, 2, 1, 0, because they involve a tertiary, a secondary, a primary, and a methyl carbocation. Let's write those SMOGs out. The central atom in each one is sp^3 -hybridized. Notice that each time I write my sp^2 s, I write it in such a way that none of the three orbitals is directly oriented pointing up and down. That leaves the room so you can go back and put a p orbital in later. On the methyl, all we have left are hydrogens. On the ethyl, we have two hydrogens. On the isopropyl, we only have one hydrogen, and there is no hydrogen on the tertiary system. Where we don't have hydrogens, we have sp^3 -hybridized carbons. Of course, on all of those sp^3 -hybridized carbons, we've got more hydrogens.

In the methyl case, there's no interaction that can occur by the bonds with that p orbital. Why? Because they're 90° angles to each other. So HC is going to mean hyperconjugation. In the primary carbocation, there's one neighbor that can undergo hyperconjugation. In the secondary case, there are two neighbors that can undergo hyperconjugation; and in the tertiary case, there are three possible hyperconjugation interactions.

Realize that those single bonds are rotating continuously (we can assume we're at room temperature), so it's not that all these interactions are happening at the same time, necessarily; it's just that, in the tertiary case, you have three neighbors that can help out, secondary, two neighbors, primary, only one, methyl group, nothing. Although to use the word stable to compare energy between molecules we really need to be talking about isomers, it is a general statement that generally is true that, for a methyl group, because you have to form that positive charge and it doesn't have any supposed, that is more difficult to do. The more that you can spread that positive charge out, even just a little bit – this is not true delocalization because it's not resonance and it's not conjugation – but if you can spread that positive charge out at all, or bring electrons towards the positive charge, that's going to make it an easier system to form. So, tertiary carbocations form more easily than secondary than primary than methyl, because of the greater extent of hyperconjugation.

Why are we talking about hyperconjugation at all in this case? Let me go back and remind ourselves what the real story is today: that we're taking this alcohol and trying to dehydrate it. We end up at this point in the mechanism where I show hydrogen coming off. How can I justify that? Imagine this: what if I had a hyperconjugation interaction that turned into conjugation? Resonance is where you have full interaction of whatever orbitals are involved, where the electrons are completely delocalized; they can move around between any of those atoms. In this case, you have a sigma bond, and the bond is really going to be located between these two atoms, as opposed to p orbitals or pi bonds, where they're fully spread out. You notice it's only this part of the orbital that's interacting with that part of the orbital, so this is not going to form a double bond, this interaction. But, since these are close enough together in space, a little bit of electron density kinda slides over and covers a little bit of that p orbital. It's not that that lone pair moves over, you're not making a bond; it's just, you've got plus here and minus here, that minus is going to interact a little bit with the plus. That's hyperconjugation. But if, when that hyperconjugation begins, that small interaction begins, then that hydrogen comes off. That interaction could just become stronger, since there's nothing else keeping that lone pair any more. When that happens, this will turn into a sp²-hybridized carbon, which means you will form a double bond. What happens here is that hyperconjugation sets up a pathway that lets those electrons start interacting, and then when the hydrogen gets ripped off, that interaction becomes complete.

Conjugation means, if I have electrons here on this atom, some of the time it'll be on this atom, some of the time it'll be on that atom. All of the atoms are fully interacting. They make new bond, antibonding, and nonbonding molecular orbitals; that's conjugation. Hyperconjugation does not form new bonds, it does not form new molecular orbitals. It causes a small shift in the energy of the orbitals that are already there. What effectively happens is you lower the energy of the p orbital; it helps stabilize by a fraction of the electron density moving over. For conjugation, the whole electron moves, across the molecule. Hyperconjugation is more localized, and it's only a partial interaction.

If that partial interaction is allowed to be completed, pull the hydrogen away, let the lone pair move into being with the p orbital, you'll make a pi bond. But now imagine this situation: what if, because of this interaction, the hydrogen just slips over? Because look, I have a secondary carbocation that's formed, but tertiary carbocations would be more favorable. Recall that I had said there's a regiochemical problem with cationic mechanisms, where you can have this thing called a hydride shift. That's what happens. If you've got this partial interaction already going on, the what if the hydrogen just moves over? How would that look like mechanistically? Let's start with the carbocation already formed, and instead of the hydrogen disappearing, it just moves over. That is not resonance; resonance does not involve breaking or forming single bonds. You've made a totally different molecule. What if it happens a second time? Now, I have a made a tertiary carbocation. The secondary and tertiary carbocations, those are isomers of each other, so we could make a formal comparison of energy. We'll find that the one with the tertiary carbocation is lower in energy than the secondary. If this is allowed to happen, it will happen. Because you're heating the crap out of this, you're providing plenty of energy for this to be able to happen, so it does happen, which means you're going to get five different products from this reaction. This mechanistic step is called a hydride shift (H⁻ is hydride); it's called a hydride shift because the hydrogen, along with its bond, moves from position to position. What's the rationalization of why this might be possible? Again we look at the SMOG and say, ok, if we had hyperconjugation, then what if it just goes into hyper mode? Then the hydrogen slips on over. It will do so if you can end up with a lower-energy molecule, which is exactly what it's going to do in this case. Besides being called a hydride shift, this is also called carbocation rearrangement. Yes, you have to know it happens, because the whole long-term goal in all of this is that we're going to learn how to synthesize molecules. If we're trying to synthesize something specific, but the mechanism that we use makes a mess, makes multiple products, then that's a real-life consideration we've got to take into account, so you have to know that this happens. Carbocation rearrangement happens nearly any time carbocations are involved. This just happens to be the first place we're really going to see it.

Let's see the full mechanism for this reaction. The only logical thing that can happen [at the start] is the oxygen gets protonated. Water leaves, forming a carbocation. Multiple things now happen. I'm going to write several of these things on the same structure here, just to reduce the number of times I'm writing these molecules. Recognize that it's different possibilities, and not one molecule is doing these different things simultaneously. One possibility is you take a hydrogen from the end of the molecule, eliminate it, and make a double bond. The one I showed you earlier is possibility number 2, that the middle hydrogen gets removed and makes a double bond. Except, it would just make the product that I showed you, because that double bond could form with either the cis or trans configurations. Automatically, even if that one hydrogen's removed, you have the possibility of making two different products. Rewriting the carbocation, so I can show you the third possibility: hydride shift.

From this point, two things can now happen: one, we have elimination; one, we have hydride shift. Elimination is a possibility; hydride shift is again a possibility. This last hydride shift gets us to what should be the most stable of the carbocations. From here, our last possible product can be shown. It doesn't matter in this case which methyl group I eliminated from, you'd get exactly the same molecule, so there are five unique products. The proportion of the products formed is going to depend on how quickly elimination occurs, versus how quickly the hydride shift occurs.

Since this is the first time seeing an alkene being formed, why don't we continue with our nomenclature of alkenes. We have talked about cis versus trans and E versus Z. I'll remind us that E is somewhat interchangeable with the term trans, and Z is somewhat interchangeable with the term cis. To determine if something is cis or trans or E or Z, remember that I dissected a double bond in this manner: if we want to talk about this part of the double bond versus that, the left and the righthand parts of it, I could use the word 'side' there, but I could also use the word 'side' to talk about the top and the bottom. To not get confused with terminology, and time I use the word side, it's going to mean to refer to the two different carbons in that double bond. If I want to talk about the other division, then those will be the different faces of the double bond. Cis and trans are terms that you can use as long as you only have one substituent on either side of the double bond. If we have more than one substituent, we have to use E and Z. E and Z, you use exactly the same priority rules as you do for R and S – higher atomic number takes precedence; higher mass number (after that) takes precedence; and if you can't make your decision, you go to the next position out.

Let's go back and look at all of these cases. The one that's in red, the one I called possibility #1, is that cis or trans? The right answer is, no. The question wasn't where it was cis 'or' trans, the question was, is it 'cis or trans', and the answer's no. If you swapped and you moved it up here, you'd get the same molecule back again, because you have two hydrogens on one side, so it doesn't matter which way you put the two hydrogens, you're going to generate exactly the same molecule; it does not have two different forms, so it doesn't get the label cis or trans, it doesn't get the label E or Z. For numbering this compound, the alkene is more important than alkyl substituents, so we're going to number from the righthand side. Double bonds always have two numbers that, in theory, could be associated with them if the double bond lies in the main chain. So we number it to give it the lower of the two numbers. We're going to call this a 1-pentene instead of a 2-pentene. We have a methyl group at the 4 position, and so this is 4-methylpen-1-ene. What about possibility 2? We do have cis and trans you have to worry about. I'm going to always use E and Z instead, because you can always use E and Z, but you can't always use cis and trans. If you use E and Z, you do it exactly like stereocenters: if you only have one double bond that has a configuration, you don't need numbers with it. If you have multiple double bonds that have Es and Zs, then you put numbers with it. The top of these is going to be (E)-4-methylpen-2-ene, and the bottom one will be (Z)-4-methylpent-2-ene. What about possibility number 3? Is that molecule cis 'or' trans, or is that molecule 'cis or trans'? In other words, do we need E or Z at all? The answer is no, because you have two of the same substituent on one side; flip the substituents, you get the same molecule again, which means you don't have the case of either E or Z, neither of them applies here, so we don't need a label. This would simply be 2-methylpent-2-ene. Similarly, the last molecule, because you only have substituents on one side, or, another way of seeing it, since you have two hydrogens, two of the same substituent on one side, it doesn't need E or Z, so this is going to be 2-methylpent-1-ene.

We're going to take this alcohol, we're going to cook it up in sulfuric acids, we're going to produce this mixture of products. We're going to do two forms of characterization: IR and gas chromatography. It turns out that, whether you have one substituent or two, three, four, or none, you get different vibrations of that double bond. Gas chromatography we're going to use to figure out how much of each of the compounds we produced.

[physical set-up for lab]

Heating at reflux. A very typical setup for heating at reflux would look like this: a round bottom flask attached to a water condenser. Why is it called a water condenser? If you look at it carefully, the inside channel of that condenser is not connected to the outside. Stuff can flow through the middle of it, but it has nothing to do with the glass that's surrounding it. The opposite's true as well: you can flow water through here, and it won't leak into your reaction flask. The idea is this: many reactions occur more quickly with heat – more heat, more energy, more ability to get over the rate-limiting step, so faster reaction. You want to heat it up as much as you can sometimes, but of course, at some point, you heat a material enough, your solvent's going to boil. What you do is this: you do go ahead and let it heat up, and slowly over time, the vapor is going to start flowing up into that condenser. Once it flows up into that condenser, it cools off, it recondenses and falls down again. That process of heating, rising, condensing, falling, that's called reflux. Heating at reflux means, heat it so it's boiling but it's not just boiling away, it's being recollected and condensing back into the flask again. You can watch that process happen. When you first turn the heat on, you won't see anything, then you wait, and you wait, and you wait. Eventually, you start to see what looks like liquid slowly rising up. It's not that the liquid is rising up; vapor has started to condense. Condensation's an exothermic process, so that heat gets released to the glassware, which means the glassware is getting heated up, which means the vapor makes it a little further the next time before it condenses again. When it condenses, it releases more heat, heat's the glassware up, more vapor rises. That keeps going and going and going. Theoretically, you would just keep heating it and it would disappear, unless you somehow tried to contain that vapor by forcing it to condense. That's why you use water – or, some reactions with lower boiling points, you could just get away with using the air condenser.

In the case where we use a water condenser, the water goes in the bottom, out the top. Why? Because the water's going to be cold when it goes in, potentially warm when it comes out. To make sure that we trap as much of the compound as possible, you want to cool it as close to the reflux as possible, so that by the time that the water has warmed up, there's not as much vapor up at the top for it to be trying to trap. When you heat at reflux, you through the solvent and reagents in, turn the heat on, allow the boiling to happen, but very little of the solvent disappears because it recondenses. Can this be done in a closed system? No, not unless you want to blow your glassware up, cause you heat it up, you're going to build up pressure; if you don't give that pressure somewhere to escape, well

That's reflux; what's distillation? Where instead of having it condense back into the same pot, you have it drop into another reception flask. So you have your starting flask with your mixture of materials. The lower boiling point material is going to rise first. Low boiling point means high vapor pressure, means easier to evaporate, so it will climb up in reflux first. But if we give it another way to move – in other words, if we do something like this – this is an example of a distillation setup. Same type of thing, you heat the stuff in the flask, it rises up the tube, it reaches towards the top, it starts to condense. When it condenses, some of it may fall straight back down into the reaction flask again, but some of it's going to fall sideways again. If we put a condenser here, that encourages the product to condense, because it's going to condense most easily wherever it's coldest. In this setup, water again goes in the port closer to the point of reflux, to encourage the condensation. Your products are going to be collected in this side flask. In general, distillation separates compounds on the basis of boiling point. The lower boiling point material comes off first because it boils at a lower point. You could separate that out, then we could swap our collection flask, pull one material off, then pull the second material off, then pull the next one off.

Why are we doing it for this particular reaction? Why is this a distillation combined with reflux? Because alkenes have lower boiling points than the alcohols that you can make them from, in this dehydration reaction. Since the alcohol has a boiling point, it won't rise up first; the alkene that forms will rise up, and then it can condense in your product flask, while your alcohol sits there and continues to react. Why is the drying tube there? What does drying reagent do? Get rid of water. What is one of the products of this reaction? Water. If you keep removing water as you have a reaction go on, if you remove a product, what does that do to a reaction? Pulls it forward; Le Châtelier's principle. Your removing water as it forms to take the reaction and force it to move forward.

In general, it's a very bad to heat to dryness, to get rid of all of the liquid. In this case it doesn't mean water, it means getting rid of all of the liquid. Liquid, the solvent helps dissipate heat. What can happen is that you could have a reaction that leaves just a film of by-products around. Those by-products might have very different reactivity than what you're really try to react. So if you have no solvent to dissipate heat and you have these random by-products, add heat to that, add a lot of heat to that, and boom, a bad reaction can happen. It's bad practice, in general, to ever heat completely without liquid – unless you know for sure it's a safe reaction.

In this dehydration reaction, you're going to get a lot of by-products out. Besides the mechanism that I showed you, there are even other products that can form. The longer you heat it, the more those products build up, and eventually you just get this black goo, this gel that can be difficult to remove from your reaction flask.

If we wanted to carefully separate compounds from each other, it's good to have a lot of distance from where you're heating it to where the separation's occurring, but it take forever. All of this glassware has to reach the same temperature of the boiling point of the liquid, and it only gets to that temperature as the liquid evaporates and condenses. If you have a large glass setup, it might take you an hour before even the first drop comes over. We aren't going to care about carefully separating the compounds. I want to see what is the product distribution, what is the mixture of products that you're going to get. We can use spectroscopy to check whether we at least got it away from the alcohol or not. To make this happen within our lifetime, we're going to insulate this whole top area of the flask. It does defeat part of the purpose of distillation – if you're making it easy to evaporate, you're making it harder to separate.

Don't forget a stirbar, because you've got to constantly agitate the liquid in order to get it to boil. If you don't have it stirring, one thing that can happen is that it can bump, which means vapor builds up and all of the sudden the liquid just ... all at once. That's bad. Or, it just takes forever for the evaporation to occur, so you going to want it constantly stirring. You want this flask buried as much as possible into the sand bath to get as much heat transfer as possible.

[lab directions]

The starting point for a discussion of NMR is the fact that some nuclei have spin, in the same way that electrons have spin. Unlike electrons, not all nuclei have spin; some have spin 0. Some nuclei have spin other than $1/2$. You can analyze them using NMR, but it ends up being a bit more complicated. Complicated, there are two nuclei – ^1H and ^{13}C – that both have $1/2$ spin, and since carbon and hydrogen are part of organic compounds, that's why we're able to analyze organic compounds using this technique. ^1H and ^{13}C are atoms with nuclei of spin $1/2$. If we have carbons and hydrogens that are floating out in the middle space, with no electric fields around them, no magnetic fields around them, then these individual spins will be randomly oriented.

The fact that a nucleus has spin does not mean that it is sitting there physically spinning – it's a really bizarre quantum-mechanical property – but it has some of the same physical effects of something that is spinning. If you have a charge that spins, what does that generate? A magnetic field. Rotating charge generates a magnetic field; current going around in a ring generates a magnetic field. So each of these little nuclei, in a way, can act like a little bitty tiny magnet.

If you then put this sample inside a very strong magnetic field, then imagine the following: imagine that you have a bunch of weathervanes on the top of a building. [description of what a weathervane is] Imagine that there's no wind; then those weathervanes can be oriented however they want to be. If all of a sudden you get a strong wind gust, then [they're] going to turn one way or the other, so that it's in line with that wind. If you take a bunch of little magnets – imagine compasses – if we had a whole bunch of compasses and turned the Earth off, turn the magnetic field of the Earth off, those compasses would point in any random direction, but near a magnetic field they're all going to line up. That's what going to go on in this case, except it's closer to the wind analogy. If you have a magnetic field, and you have one of these other magnetic fields that's not aligned with it, there's going to be torque, which means it's going to twist, and it's going to twist one way or the other until it's aligned with the magnetic field. But it turns out it can be either totally parallel, or it could be antiparallel; you can have a field that oriented N/S in one direction, and you could have the spin oriented S/N. That's fighting against the magnetic field, but as long as it's in that direction, it's not going to be twisted any more, unless it loses energy somehow and then flips around the other way. What I'm trying to say is, you start out with all of these randomly oriented spins, but then put them into a magnetic field and they line up, either with the magnetic field or against the magnetic field (B).

You start with all of these spins randomly oriented if it's not in a magnetic field, you put them in a magnetic field, and they become oriented. There's an energetic consequence of that. The spins that are aligned with the magnetic field are now going to end up being lower in energy, cause it's like it's flowing down the river. If you're fighting against the river, that takes energy, and so the spins that are antialigned will be higher in energy. This is for the case of something that's spin 1/2. There's an energy gap that is associated with it, which means it corresponds to a particular frequency of light, and [therefore] a certain energy per photon. There are three factors that can affect this energy gap. The largest factor is the strength of the magnetic field. Beyond that, different nuclei have different extents to which they will interact with a magnetic field; some really are sensitive to magnetic fields, others hardly bothered by them. Hydrogen, for example, is four times as sensitive as carbon is. That sensitivity is then going to change that energy gap, and so carbon and hydrogen, you have to look for at different frequency ranges. This is represented by something known as the magnetogyric constant. Gyr, gyrate, to turn; magneto means related to magnetism; so magnetogyric means how much does it get tweaked by being interacted by a magnetic field. The last one, which is the most important one for us, is what we'll call chemical environment. If I have a nucleus that's interacting with a magnetic field, what if it's got a lot of electrons around it? Those negative electrons might mask, somewhat, the interaction of that positive nucleus with that external magnetic field. If I through some fluorines right next to that position, though, they would draw electrons away, then I could expose the nucleus to have a greater interaction with that magnetic field. A greater interaction means a greater energy gap, which means you'll see it at a different frequency. Even though that's the most important of these effects, it's going to be the smallest of them, so we need a pretty sensitive machine in order to see these differences.

Let's talk about in general what a spectrometer setup looks like. Generally, there's some kind of chamber in which you'll suspend your sample dissolved in some kind of solvent; we'll call that the NMR tube. That chamber, depending on your machine, is inside some form of magnet. [There are] what are called fixed magnets, meaning they are like that all the time – big, hulking, iron-core magnets. As big as they are, they aren't the strongest magnetic fields in the world. This last effect of chemical environment is a very subtle effect, so the stronger the magnetic field we can get in the machine, the more we can differentiate between these different energy levels, the better analysis we can perform. What 'real' spectrometers are are superconducting magnets, which means they have to be cooled way down to liquid helium temperatures. In fact, the way that these are constructed is that you have your sample chamber that is thermally isolated from the magnet – the magnet's got like this big thermos around it. Inside of that is a liquid nitrogen tank, and inside of that is a liquid helium tank, and inside of that are the coils, which would put an enormous charge in. You send a heck of a lot of current going around and around and around in a circle. Current going around like that generates the magnetic field. Some of these machines, the magnetic fields are so strong that if they're put into buildings that are just being built, where possible, these machines are put in the basement, then there is no first floor above those locations; that much room, a whole two-stories worth, is put in between these machines and anything else, because that's how strong the magnetic fields are, and that's how sensitive the machine is. These magnetic fields are so strong, in some cases, that if you've got the wrong kind of jewelry on, and if you get close enough to these machines, they're just going to yank you towards the machine because they're that strong. People with pacemakers can't walk up to these machines, because the magnetic fields are so strong they'll interfere with the action of a pacemaker. Strong stuff. In general, you put the sample suspended in some kind of magnetic field. [contrast in cost of the types of machines] That's what goes on inside the machine.

How is the signal actually acquired? There's the old way and the new way to do it. Older technology, when they started figuring out that this effect occurred, one way to do it was to vary the magnetic field. If you have an iron-core magnet, you can change the size of that magnetic to change to suddenly change the magnetic field. You could change the field strength and do a scan; you'd be scanning at one frequency and looking for a particular energy gap. As you change the machine's magnetic field, different kinds of nuclei would then have that right kind of gap to be scanned.

You could scan from low magnetic field to high magnetic field and see where interaction occur. But that can be a problem. Even if you had a fixed magnetic field but your scanning individually at different frequencies, if each scan took five seconds, let's say, and let's say that you needed 1000 scans to make a decent enough spectrum, that's 5000 seconds, which is a little over an hour and a half. That's in a good case; what if it takes a minute a scan? Then if you needed 1000 scans, that's more time than we can possibly afford.

A more modern approach is to do this: you take a whole bunch of range of frequencies and zap the sample all at one, with all of those different frequencies. Only those frequencies that match the different energy gaps you have in that sample, those are the only ones that are going to be absorbed. After that radiation's been added, you then 'listen' to what comes back off of it, because once the radiation's been absorbed, it has to get released. So the signal that you measure is all these different frequencies of light that get released by the nuclei after they've been excited. In most modern NMR spectrometers, a sample suspended in a magnetic field is irradiated with a range of electromagnetic frequencies – you could call it, if you wanted, the shotgun approach, so you just blast it with all of these different forms of radiation. Only those frequencies that correspond to the energy gaps generated by the nuclei being in the magnetic field will be absorbed. Once you've added the light in, you turn the switch off, and then again that energy is released. So, after irradiation, the various frequencies of light coming from the sample are observed. That's how a modern NMR experiment occurs – put the sample in, zap it with a bunch of radiation, turn the radiation off, and listed [to] what comes back off again.

But what does come back off again? It's all of these different frequencies coming at you at once. The signal that you get might look something like this. You might see some higher-order pattern on top of all of this other wiggling that's going on. Realize that these are frequencies stacked on top of frequencies stacked on top of frequencies. As that energy is being released by the sample, the amount of energy released at any one time's going to decrease, because it's releasing the energy. This kind of pattern is called the free induction decay (FID). I'm going to use a couple of music examples here because they're [best] thing comparable I can think of to describe. Ever live in houses where you have a piano just sitting around in one of the houses? Any of you ever play music in that same room? What happens if you play really loud music and you suddenly turn the music off? Different notes on the piano would continue to resonate (if you haven't done this and you have a piano sitting around, try it). What's happening is that you listen to this really loud music, and all of these different frequencies are coming out – they're turned to the right frequencies to match some of these strings on the piano, the piano's strings start to resonate in the response. Turn the music off, and each one of these different frequencies still continues. That's kind the same thing that I've describe with this NMR experiment: you blast it with a range of frequencies and then listen for the ones that come back again.

But here's another problem. This signal that you get is a mess, because it is all of these different frequencies that come out. Anybody know what an equalizer is? Simplistically, for sound, we usually talk about treble, which is high notes, base which is low notes. [On some equalizers] you can actually access different frequency ranges. It used to be, to show off a little bit, you'd have these flashing light bars, and you could listen to the music and see what different frequency ranges show up. The same thing goes on in a modern NMR spectrometer, because you get all of these different frequencies out. We want to know how much response is happening at each one of these frequencies. There's a mathematical operation we can perform – a Fourier transform. You have one function that you compare to another function, and if the two match, you get a non-zero value by doing this operation. If they don't match, you get nothing out. What you do is you individually overlay a particular frequency on top of this signal and see how much does this signal match, then you do it with another frequency, and you do it with another frequency. Our cell phone have enough [computing] power to do this. Realize that back in the 1940s and 1950s when this technology was originally being developed there wasn't that computing power, which is why back they they might have had to scan at individual frequencies or individual magnetic field strengths. In our modern approach, we blast this with all these different frequencies, we get this signal out, and we do what's called a Fourier transform, which in this case is going to convert from the time domain – meaning time being our x axis – to $1/t$ being our x axis, frequency instead. When we take this signal and deconvolute it, we'll end up with a much simpler looking spectrum. Realize what we start with is not a spectrum; it's a signal that decays over time. This is a spectrum where we're able to analyze individual frequencies. It's only because we have the computing power to do this transformation that modern NMR is possible.

We don't really use frequency as our x axis. Why? Going back to those different things that affect the energy gap. All of this is about these poor little nuclei suspended in a magnetic field and they've got energy gaps that form because of the magnetic field. The machine magnetic field strength itself has a huge impact on the energy gap. Well, it's virtually impossible to construct two machines that complicated to have exactly the same magnetic field, down to the fiftieth decimal place. That means, whatever frequency we're measuring on one machine is not going to be the same as the next machine is not going to be the same as the next machine. We have something that's called chemical shift. Chemical shift is both a unit of measure and the change in this unit of measure that's caused by chemical environment. Chemical shift, which is usually assigned lower case delta, is the observed frequency of some particular nucleus minus what's called the machine frequency, divided by that same machine frequency, multiplied by a million. Because you multiplied by a million, the units for this are actually parts per million. Why are we multiplying by a million? Because that's how subtle these chemical environment effects are. Out of 90 MHz, then maybe you'll only vary by 100 – 200 HZ within that 90 MHz. Very small effects.

What's this machine frequency defined as? It is the frequency measured on a particular spectrometer for a standard reference compound. For organic compounds, the reference compound almost always used is tetramethylsilane (TMS). Silane, that has a similar sounding name to alkane, this -ane ending. Silanes are compounds where you have silicon as the backbone instead of carbon. Tetramethylsilane would be to take silane, SiH_4 , and substitute each of the hydrogens with a methyl group. Why is this particular compound used as the reference standard?

Let's say we have this particular reference compound. I'm just going to tell you that this happens to generate a frequency that is lower than almost any other organic compound that you would measure. Because of that, if you subtract it from any other compound you're measuring, you're going to get a positive number. Why are you subtracting that? So we can have a zero reference scale; we can have some point on this scale where we say, that's zero, that's this reference compound. Why are we dividing by the machine frequency? This is an equation that shows that energy gap is linearly related to field strength. Frequency is linearly related to the energy gap. That means frequency is linearly related to the machine strength. If we have frequency in the top of the fraction, frequency in the bottom of the fraction, we have that magnetic field strength, effectively, on the top of the fraction and the bottom of the fraction, which means it cancels out. It means that the machine's magnetic field strength, it could be anything it wants to be, and it doesn't appear in this expression. This is a machine-independent quantity. The importance of that, that it's a machine-independent quantity, the upshot is that you'll get the same value for chemical shift regardless of the spectrometer that you use. Which is great: we want to have some unit of measure where you don't have to constantly quote what machine that you took it on.

Now let's get back to the point: why is TMS chosen as the standard compound. Let me tie that into a discussion of what an NMR spectrum looks like. A typical NMR spectrum looks like this: chemical shift, your x-axis; intensity, your y-axis, and 0 on the righthand side of the spectrum. Why is zero on the right? There's a historical reason. Chemists, instead of saying low chemical shift, high chemical shift, sometimes will use the phrase[s] shielded and deshielded. Let me tell you a story about something being deshielded. Let me give you a molecule to look at. Let's say we had fluoroethane, and let's say we were observing the carbon right next to that fluorine. Fluorine is going to pull electron density away from the carbon. When it pulls that electron density away, you're effectively exposing that nucleus to the magnetic field. If there's that greater interaction between the nucleus and the magnetic field, that means you're going to have a greater split in energy levels. That greater split in energy levels corresponds to a greater frequency of observation, which therefore gives you a larger chemical shift. When a nucleus being observed is close to an atom that is electron-withdrawing, electron density will be pulled away from the nucleus being observed. This means that nucleus is more exposed to the magnetic fields, which means it will experience a larger energy gap. If you pull the electrons, you expose this nucleus, it's going to have a larger energy gap, which results in a larger observed frequency, which results in a larger chemical shift. Notice on this spectrum, I have deshielded on the left, which frequency increases to the left, so those two things match. The more deshielded something is, the more you're pulling electrons away, the bigger the energy gap, the bigger the frequency, the bigger the chemical shift.

Now I can backtrack and explain why TMS is used as the reference compound for organic molecules. Think about the periodic table. Where is silicon in relation to carbon? Right underneath, so what's more electronegative: carbon or silicon? Carbon. Carbon, relative to oxygen and nitrogen, is where? To the left, so that means nitrogen and oxygen are more electronegative than carbon, which are more electronegative than silicon. Where are sulfur and phosphorus relative to silicon? On the right, which means they're more electronegative than silicon. So, carbon, nitrogen, oxygen, hydrogen, fluorine, chlorine ... all of these different common elements found in organic compounds are all more electronegative than silicon. Which means, if you have something connected to silicon, it is more shielded than other types of atoms might be, because silicon is not pulling those electrons. If it's not pulling those electrons away, then the nucleus is more masked against that external magnetic field. If it's not experiencing that external magnetic field as much, the energy gap is not as great, which means the frequency, and the chemical shift observed, are lower. Silicon is less electronegative than most elements typically found in organic molecules. Because of this, atoms next to silicon are, by comparison, shielded, which means the nuclei are not as exposed to the magnetic field, which means the energy gap is not as large, which results in a lower observed frequency and chemical shift.

Look at TMS. How many different kinds of hydrogens would you say exist on this compound? Chemical equivalency – how can you say this hydrogen's the same as that one? In this case, all the methyl groups, they're symmetric, and the three hydrogens, if you switched the positions of the three hydrogens on the methyl group, it doesn't change what molecule you've got. You could talk about bonds rotating and making rotomers, but the hydrogens themselves are equivalent. So there's only one type of signal that would be generated by these hydrogens; there's only one type of signal that would be generated by these carbons, and since they're hooked up to silicon, they're going to be at lower frequency than almost anything else for any organic compound that you're going to observe, which is why TMS is used as the machine frequency, because you're not going to get a number lower than it, normally. Compared to silicon, even carbon itself is more electronegative. So make, instead of tetramethylsilane, tetramethylmethane, put a carbon in the middle, and have these methyl groups around it. They'll be at higher frequency than TMS, because carbon is more electronegative than silicon is.

Let's tackle a little trickier problem. If I look one more time at a typical spectrum. Here's our zero point for chemical shift, chemical shift increases to the left. We've already assigned these terms shielded and deshielded, but it turns out that there's also this terminology used: upfield and downfield. Where does this come from?

Imagine the following: imagine we had three different types of nuclei, hanging around in the magnetic field. Let's say there's one nucleus that's of average type, and it's got a companion that's deshielded – more exposed to that field, which means it's going to have a larger energy gap – and let's say the third nucleus is more shielded – which means it's going to have a smaller energy gap because it's not as exposed to that nucleus. Let's say that we have an old school machine that only scans at one frequency. So if you can only scan at one frequency at a time, you have to have exactly the right energy gap in order for that gap to be seen. For example, the line right here represents me tuning the machine so that it matches the energy gap for this middle case nucleus. In order to get a shielded nucleus to have the right kind of energy gap to be registered at this frequency, I'm going to have to increase the strength of that magnetic field. Conversely, if I have something that's deshielded, which means it's more exposed to the magnetic field, then to get it to respond at the correct frequency, I'm going to have to decrease the field strength. That's where these terms upfield and downfield come from. Upfield means more field, downfield means less – a magnetic field, therefore, used to be the x axis, and lower magnetic field was naturally on the left. For a system that's at a lower magnetic field, that would generate a frequency equivalent to some other situation that might have been more shielded – in other words, if we had something that, at the same magnetic field was deshielded, it would create too high a frequency. I'd have to turn down the field strength to get the same frequency. Conversely, if I wanted to have something that was shielded to show up at that correct frequency, I'm going to have to turn up the field strength. Since modern spectrometers tend to have fixed magnetic fields and scan at a bunch of frequencies, that's why we instead use the terminology shielded or deshielded, because at a constant magnetic field, the shielded nucleus is going to have the smaller energy gap, which is the smaller frequency, which is the smaller chemical shift. The one that is deshielded is going to have the higher energy gap, the higher frequency, and the higher chemical shift.

If a spectrometer can only scan at a fixed frequency, then the magnetic field used must be varied in order to make the energy gap for different nuclei match the observation frequency. If an atom is shielded, a stronger magnetic field will be needed to create the equivalent energy gap – that's a situation we refer to as upfield. Conversely, if an atom is deshielded, a weaker magnetic field must be used to generate the appropriate energy gap – which means that it will be downfield.

[next topics: NMR solvents; chemical equivalency; splitting]

The greater the extent of hyperconjugation, the easier it is to form a + charge, since that charge is spread out over the molecule.

hydride shift = carbocation rearrangement; H:- ≡ hydride

E ⇌ trans; Z ⇌ cis

NMR – ¹H and ¹³C are atoms with nuclei of spin 1/2

Factors that affect this energy gap: strength of the magnetic field (B); magnetogyric constant (how sensitive a nucleus is to magnetic fields); chemical environment

In most modern NMR spectrometers, a sample suspended in a magnetic field is irradiated with a range of EM frequencies. Only those frequencies that correspond to the energy gaps generated by the sample being in the magnetic field will be absorbed. After irradiation, the various frequencies of light emitted (and originally absorbed by the sample) are observed.

Chemical shift (ppm) : $\delta = \frac{\text{observed frequency} - \text{machine frequency}}{\text{machine frequency}} \times 1,000,000$

Machine frequency – the frequency measured on a particular spectrometer for a particular reference compound. For organic compounds, the reference is almost always tetramethylsilane (TMS).

Chemical shift is a machine-independent quantity, meaning the same value will be obtained regardless of the spectrometer used.

Deshielded: When a nucleus being observed is close to an atom that is electron-withdrawing, electron density will be pulled away from the nucleus being observed. This means the nucleus will be more exposed to the magnetic field, so it will experience a large energy gap, which results in a larger observed frequency, which results in a larger chemical shift.

Silicon is less electronegative than most elements typically found in organic molecules. Atoms attached to silicon are therefore shielded by comparison, which means their nuclei are not exposed to the external magnetic field, which means the energy gap is not a large, which results in a lower observed frequency and chemical shift. ∴ TMS is used as a reference compound since the frequencies generated by it will be lower than those of most organic compounds.


If a spectrometer can only scan at a fixed frequency, then the magnetic field used must be varied in order to make the energy gap for different nuclei match the observation frequency.

If an atom is shielded, a larger magnetic field will be required to generate the appropriate energy gap → upfield.


If an atom is deshielded, a weaker magnetic field must be used to generate the appropriate energy gap → downfield.

Structures


11/02/11 lab • 1




11/02/11 lab • 2




11/02/11 lab • 3




11/02/11 lab • 4



11/02/11 lab • 5



11/02/11 lab • 6



11/02/11 lab • 7

