A Colorful Introduction to Complex Numbers and Fourier Analysis

Zack Booth Simpson, May 2014
zack.simpson@gmail.com

Introduction

People rarely set out to make a subject opaque. They don't say: "I think this subject is too clear, I better make up some bizarre words to obfuscate it." While this seems true enough it is very hard to believe that in the case of our current topic someone didn't deliberately obfuscate it.

A common explanation for obscurity is history. When a subject is new, it is like a dark forest of opaque ideas. As the trailblazers hack through the shadows they lay a trail of vocabulary exploited by followers who return and venture further. Technical jargon developed in such adventures is often a kind of *fossilized obscurity* – a record of the tortuous path of broken metaphors followed by the first explorers.

One could hardly find a subject which suffers from the opacity of its creators more than the field of "complex numbers." The title alone is so off-putting that it practically slams the door in your face. The key innovation of the field is the concept of an "imaginary number" – a term which implies the whole thing is the silly whim of a mad mathematician locked in an insane asylum. And terms like "quaternion," "conjugate," and "modulus" don't help.

Motivation

Complex numbers is the math of oscillators: processes that repeat themselves.

The subject of complex numbers has many applications, especially in the field of *signal processing*. Our modern world of computers and cell phones is built on a foundation of complex numbers and its most popular derivative: Fourier Analysis. Complex numbers are also at the heart of quantum mechanics and as such are part of the mysterious fabric of reality. Learn complex numbers for the signal processing; respect them for the quantum mechanics!

Reinventing the wheel

A natural way to represent a cycle is with a wheel. Imagine a bicycle wheel where one spoke is painted green. As the wheel spins we keep track of the progression of a cycle by watching where that spoke is along the perimeter of the wheel. Eventually it comes back around to where it started and the cycle repeats itself.
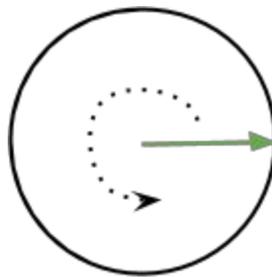
In order to track a cycle with a wheel we need to pick some starting direction. Let's pick right-facing as drawn above. We'll call that direction "unrotated." Why pick that direction and not some other? Because that's the one Descartes picked in the seventeenth century and mathematicians love precedent. (They love it even more than lawyers because it's almost never overturned!)

Just as arbitrarily, we need to pick a direction of spin and call that "positive rotation." We'll use counterclockwise (also as drawn) because, again, that's the way Descartes did it.

The resulting wave is called a "sine wave" named after the trigonometric function. Folloowing is a picture of one. Visit http://happyfuncoding.com/run/1388 to see this animation in action.
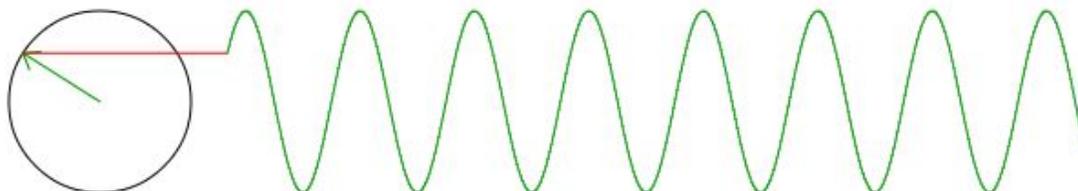

The etymology of "sine" is even more bizarre than most mathematical words. The word comes from Sanskrit via Arabic. The original term was Sanskrit *jya* meaning "bowstring." That's a nice

graphical term since it imagines the curve of an arc in the form of the bow and a radius in the form of the string. In Arabic it was translated to *jiba*. Arabic has a nasty habit of omitting vowels when written so there are words that are spelled the same but are pronounced differently. A medieval translator mistook the word *jiba* for *jaib* which means "fold" and he translated it to Latin *sinus* which means "inlet" like the sinuses in your nose. And from there we got *sine* not to be confused with *sign* which comes from Latin *signare* meaning "to mark". What a mess![1]
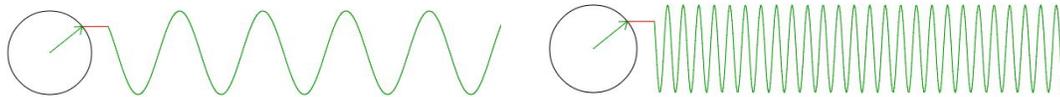
When a wheel turns faster it produces a sine wave with a higher frequency (more cycles per second). Frequency comes from *frequens* meaning "crowded."[2]
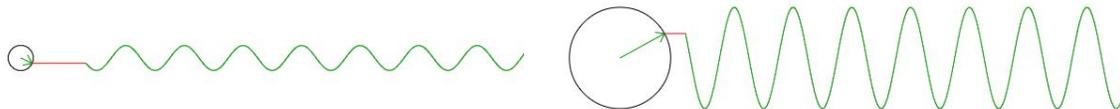
A bigger wheel makes a larger amplitude than a small wheel. Amplitude comes from Latin *amplus* meaning "large," same root as ample.[3]
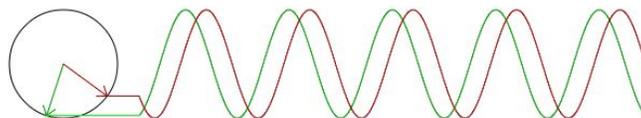
A wheel that starts turning at a different time has a different phase. Phase comes from Greek *phasis* meaning "appearance" as in the phase of the moon.[4]

Try adjusting the above three parameters in the interactive app located here: http://happyfuncoding.com/run/1398

When you add waves together you get more complicated signals. In fact, if you allow yourself to draw from any frequency, any amplitude, and any phase then you can recreate any signal you wish. This non-obvious fact is what Fourier discovered.

Try it with the interactive signal generator I made here: http://happyfuncoding.com/run/1389 This interactive app lets you add 5 different frequency components and see the resulting function.

Following is a sample function from the signal generator. I encourage you to use the interactive one and get a feel for how adding different frequencies makes up a complicated waveform.
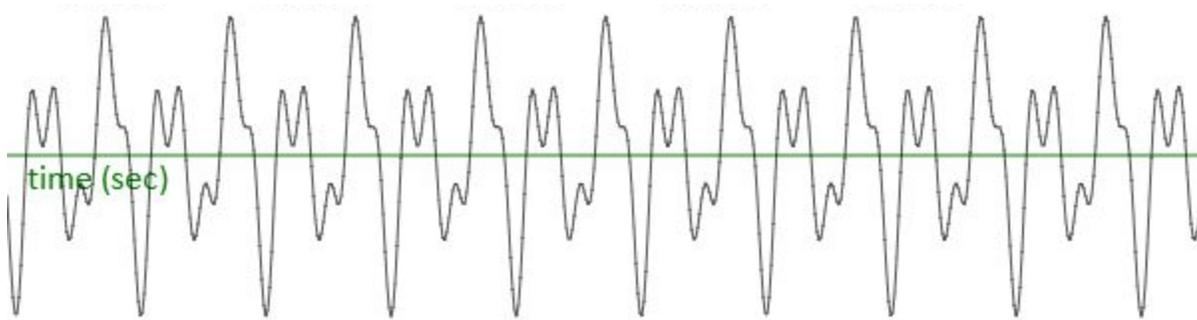

The more interesting thing about thinking of signals as sums of sine waves is that you can go backwards. Given an arbitrary signal you can decompose that signal into distinct waves of different frequencies. This is extremely valuable because it means that you can extract different components of a signal for separate processing. For example, in the radio spectrum there's cell phone traffic, talk radio, and emergency bands all sharing the spectrum. With a device that implements a Fourier decomposition you can extract out just the bands of interest without being confused by the others (for example, ignoring talk radio!) Such devices are called *filters*. The word *filter* comes from Latin *filtrum* meaning "felt" which the Romans used for filtration[5]. Filters are ubiquitous from cell phone chips to old fashioned radios. Even passive devices like optical lenses and your ear can be thought of as Fourier filtering devices. We'll look at various examples as we proceed.

The simplest filter that extracts a specific frequency out of a signal is accomplished by the process of "multiply and accumulate." Depicted below is a signal in black with red reference that we happen to know is well represented in the signal -- you can see that it aligns well by eye.
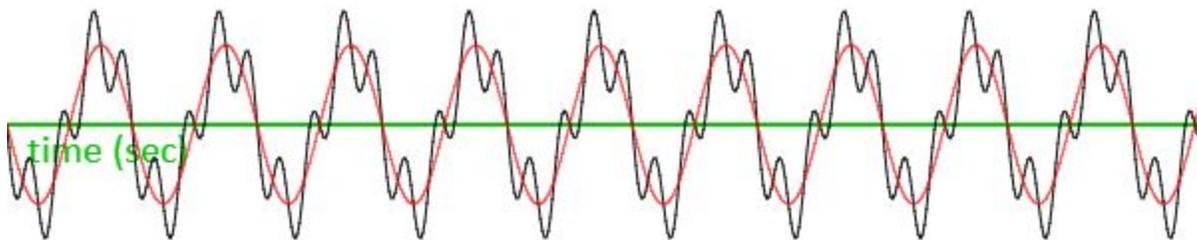
To find out how much of the red reference there is in the black signal we walk along from left to right and multiply the reference at that position times the signal at that position, and we sum up the result. That's the "multiply and accumulate."

In the following diagram you can see a zoomed in view. At each horizontal tick you multiply the length of the red arrow by the length of **black arrow** and sum the result as you walk from left to right.
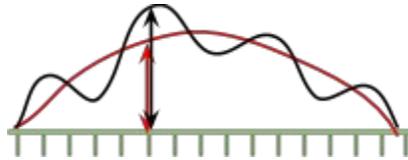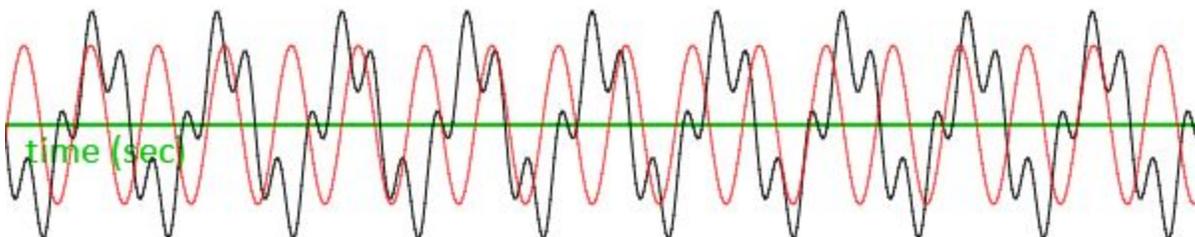

To generate the red reference we use a sine function. It is traditional to use the Greek letter omega "ω" to represent the angular velocity of the sine function in radians per second. So our reference function is sin(ω*t) where t is time. The frequency component, c, at frequency omega is the sum over the product of the reference and signal as described. Written in standard notation where we have N samples and s(t) is the signal:

$$c(\omega) = \frac{1}{N} * \sum_{t=0}^{N} sin(\omega * t) * s(t)$$

The 1/N in front is so that you get the average value instead of the total sum. Incidentally, if you're familiar with linear algebra you might recognize the "multiply and accumulate" as a "dot product" between the reference sin(ω*t) and the signal s(t).

In the example above we see that the red reference aligns well with the signal. Each time the red reference goes below the x-axis (green) we see that the signal also tends to do the same. And this is true for when it goes above the axis too. When it's below it will tend to be a negative number times a negative number which will be positive when it's above the axis it will be positive times positive which will also be positive. Therefore the sum across the whole signal will tend to be a large positive number -- that large number tells us that the reference is indeed well represented in the signal. Contrast that with the next illustration.

In the above example the red reference does not align well with the signal. As a result the reference is just as likely to be above the axis while the signal is below as vice versa. In this situation we find that sometimes we have a positive number times a negative number resulting

in a negative value while sometimes we have a negative times a negative or a positive times a positive and resulting in a positive value. The result is that the moments when the reference aligns well will cancel out with when they don't. Therefore the sum across the whole signal will tend to be a small number close to zero.

Phase

Filtering as illustrated allows us to measure the quantity of the reference in the signal; but, there's a huge problem. We have to know the phase of the reference! To illustrate the problem, consider the following example where the red reference has simply been shifted over a bit so it no longer aligns well with the signal. Now we have regions where the product will be positive and other regions where the product is negative. They will cancel each other out and we will conclude that there is none of the reference in the signal which is wrong; there is a lot of reference in the signal, it's just at a different phase.

You could scan along the phase repeating the filter over and over until you found the maximal fit. But that is silly, there's a much better way. We just need a mathematical tool that keeps track of the phase for us. And fortunately such a tool exists, it is called "complex numbers." Time to learn about them.

Complex Numbers One Step at a Time

Let's enroll in a simple dance class. We're going to imagine that there's four steps we can take.

You can combine these steps to create a little jig. For example:

The black steps take us on a path, adding one step after another. At the end of the dance we end up at a destination which we represent with a blue arrow. We're joining a very strange dance school that only cares about the destination, not the path -- any two dance paths that lead to the same place will be called *equal*. In the below example, the orange and the black paths are equal because they go to the same place.

The destination path of the dance is called a "vector." The word *vector* comes from latin *vehere* meaning "to carry;" it has the same root as *vehicle*.[6] The word makes a little bit of sense: vectors carry us to a destination.

Rotation

Now that we know how to take a dance step, we need to learn to turn. The following is an equation of rotation. It says if you are facing east and then you turn 90 degrees counter-clockwise then you end up facing north. The black circular arrow says "do turn" on left arrow using right arrow.

The orange arrow is telling us how much to turn the green arrow. Orange shows a rotation of 90 degrees relative to the unrotated direction which we previously decided was east-facing so the result is facing north. In the following we see another example: if we start facing north then turn by 90 degrees we end up facing west.

The following are two more examples rotating by 180 and 270 degrees respectively

The orange rotation arrow doesn't really need that little blue indicator that shows it is being turned, I just put that there to make it clear. We know that east-facing is unrotated so we can just leave that off. For example, the following is a rotation shown without that blue indicator:


Order Doesn't Matter

It's not obvious, but the order in which you do the rotation doesn't matter. For example, the following demonstrates that the result is south facing either way you choose to do the rotation:
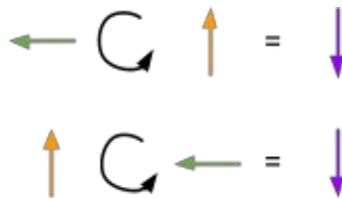
This is true for all rotations. If you don't believe it, try filling in the following table. You'll see that it is symmetric just like a multiplication table (foreshadowing!).

The rotation operator allows you to rotate more than one arrow at a time. You simply rotate each one and stack them all up. For example:
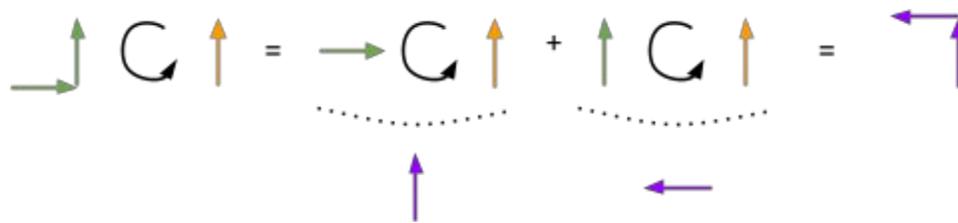
The above says that you can rotate the two green arrows one at a time and then add up the result. Notice that the result is the same thing as taking the combined green arrows and rotating by 90 degrees. Don't overlook the "+" sign in the above equation.

Remember that in our strange dance we're only interested in the destinations. So the above is really a statement about the destination vectors, not the steps that we take to get there.


Because the order of the operation doesn't matter, the second operator can have more than one arrow too as in the following. You can end up with complicated rotations where you simply rotate by each combination of arrows. Recall that the order in which you take the steps doesn't matter, we only care about the destination.
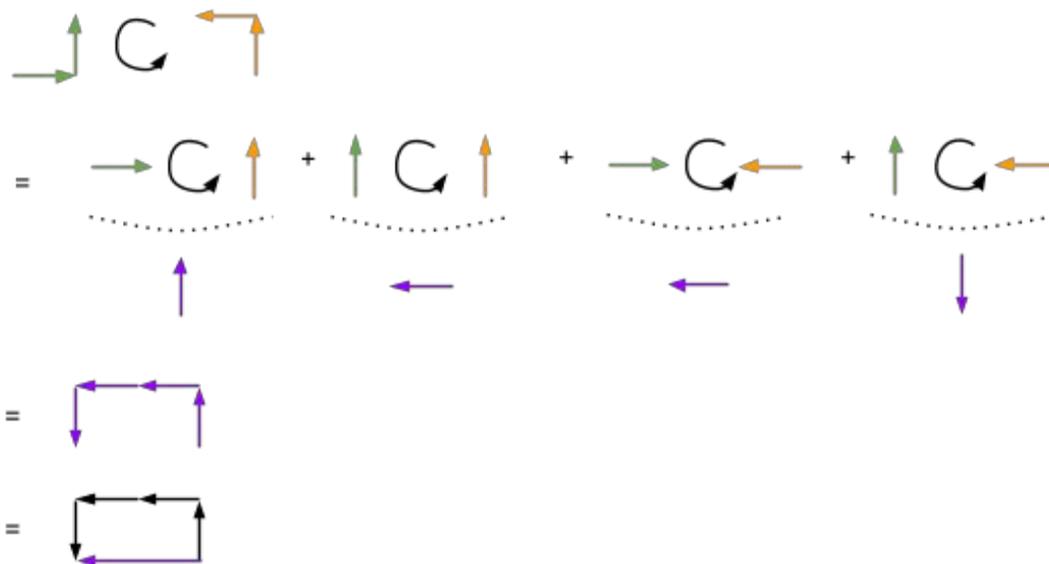
In the following example, it is shown that we can visually group the rotations together. The entire right side is rotated by each step of the green arrows and sums to get the result.
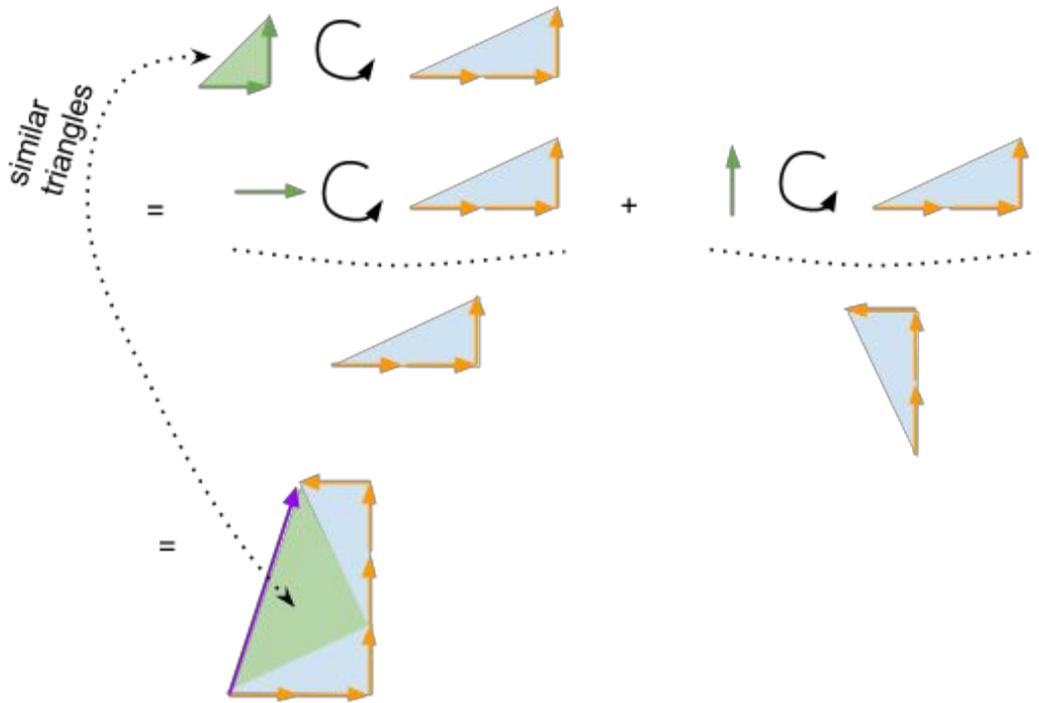
Notice something in the above illustration. We see the blue triangle twice in the result -- once for each of the green arrows in the input. But we also see a new green triangle in the result that is the same shape as the green input triangle but it is rotated and somewhat larger but nevertheless it's the same shape. It turns out this is always true whenever we rotate one complicated set of arrows by another. It means that the angle of the resulting direction is the sum of the angles of the two starting vectors. Following is the same example but I've enlarged the image and denoted the angles with dark triangles. Notice how the resulting angle is the sum of the two starting angles.
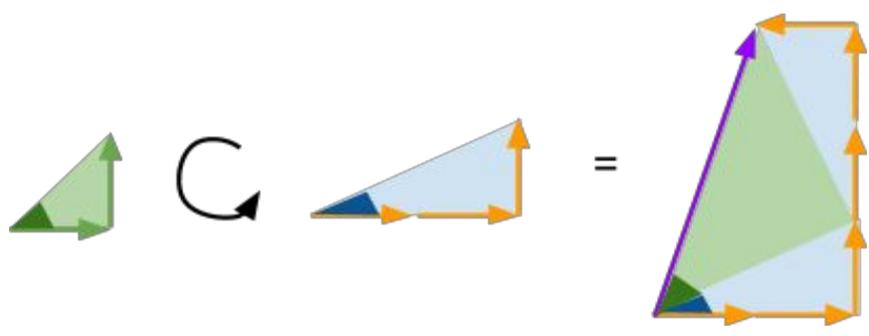
This is a quite surprising fact. When we rotate one vector by another the angle of the result is the sum of the angles of the inputs.

What about the length of the result? The two green triangles are "similar" meaning they have different sizes but the same angles. So we can use the similar triangle rule to get the length.
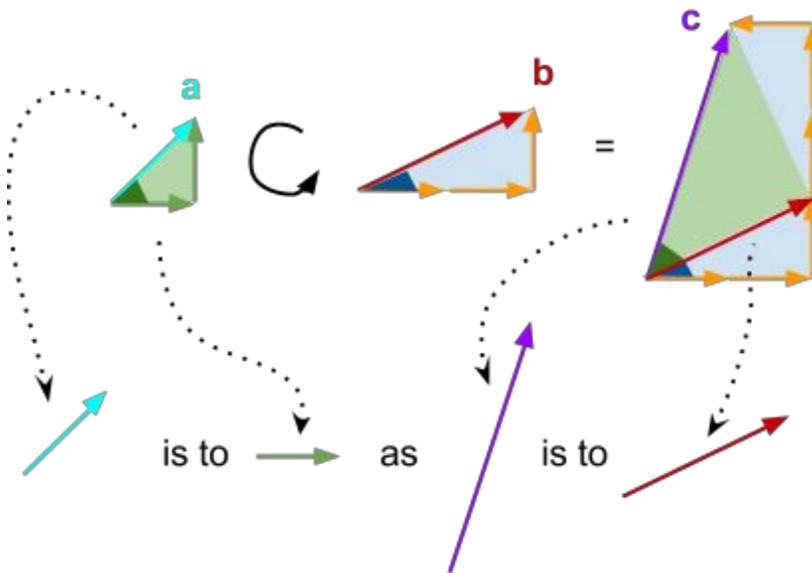
length(a) is to 1 as length(c) is to length(b).   *Stated in words*
length(a) / 1 = length(c) / length(b)   *As an equation*
length(a) * length(b) = length(c)   *Cross multiply*

So we've just discovered by the rule of similar triangles that the length of the resulting vector is equal to the product of the lengths of the inputs.

To see this more clearly, consider the following unrotated problem:

2 unrotated steps rotated by 3 unrotated steps gives us 6 unrotated steps. In other words, when the vectors are unrotated our rotation operator is the same thing as plain-old elementary school multiplication!

I have a confession. There is no such thing as the rotation operator, I made it up. Standard math notation confusingly uses the multiply symbol for rotation. I made up the rotation operator because it helps me remember that a multiply is the same thing as a rotation.

Even more confusingly, I like other computer-nerds use star * to represent a multiply whereas in elementary school you learn the symbol *x*. The elementary-school symbol is very confusing

because it looks like the letter x for the unknown and is therefore dropped in advanced standard notation.

Standard mathematical notation doesn't use the little arrows I use either. The have different names for the four arrows. Here's the translation from my symbology to the standard:

My symbol	$\rightarrow$	$\uparrow$	$\leftarrow$	$\downarrow$	$\circlearrowright$
Standard symbol(s)	1	i	-1	-i	*
		$\sqrt{-1}$		$-\sqrt{-1}$	x

So what's up with the crazy symbols mathematicians use? The "1" makes sense, after all it is unrotated and one step long. And "-1" makes sense since it is pointing in the opposite direction of 1. But what's up with *i* and $\sqrt{-1}$? It is an abbreviation for the *imaginary number.* What kind of stupid name is that?! Why not call it "perpendicular" or "orthogonal" or "right angle" or something meaningful. Calling it "imaginary" makes it sound scary and difficult when it is really just a ninety degree rotation.

And now a parable about how we got a name as stupid as "imaginary":

Once upon a time there were positive numbers. People intuitively grasp the concept of a positive number. One finger, two fingers, three fingers. Then one day an accountant ended up owing someone money. To balance the books the accountant had to write an equation like: "one debit plus x credits equals zero." Or, written as an equation "1 + x = 0".

"So what's x?" he asked. The clever accountant then said: "I'll just make up a little dash symbol and call it 'negative' from the Latin 'negare' meaning 'to deny' and that will stand for the situation when I owe someone else money."

This did not go over well. "Negative numbers aren't real things!" people cried, "You just made them up! That's absurd"

"Yes," replied our mythical bookkeeper, "but debt is a real thing and that symbol just means debt. It's just a bookkeeping convention, get over it."

Most people were not satisfied with negative numbers and kept called them derogatory names for a long time. After a long time, people got comfortable with the concept and stopped making fun of them.

Then one day two hypothetical land-owning accountants had the following conversation:

"You have one square kilometer of land in a perfect square; the length of its side is obviously one kilometer. I'm going to borrow that land from you to graze my cattle, so now I have negative one square kilometer of land in debt. No problem. By the way, what's the length of the side of that land?"

12

And the other accountant said, "Still one kilometer of course!"

And the first accountant replied, "Ah then, I'll just write in my books that I have land that's one kilometer on a side. And since one times one is *positive* one square kilometer I no longer owe you the land."

"Wait! That can't be right!" exclaimed the lender. "You can't change a liability to an asset by simply changing which way you measure the land. That's crazy."

"Ok then, we've got this equation: '$x^2 = -1$' so what is x? It can't be one and it can't be minus one because minus one times minus one is still positive one! So what is it?"

"I know, we'll imagine that there's a solution to this equation just like the accountant of long ago who made up negative numbers. We'll say that there exists a solution to $x^2 = -1$ and be done with it."

"That's a terrible idea. Those are just imaginary numbers!"

"It's just a bookkeeping convention, get over it."

Although a parable, it is true that negative number were called "absurd" and that the term "imaginary" was a derogatory term used first by Descarte.[7] Also true is that the imaginary number was thought of as a solution to an algebraic equation before the geometric interpretation that I've introduced was later described by Wessel.[8]

To see the solution to the parable's equation in geometric notation, remember that:

which is the same thing as:

$$i * i = -1 \qquad i^2 = -1 \qquad i = \sqrt{-1}$$

Because the vertical numbers were derogatorily coined "imaginary" the horizontal numbers were given a name in contrast: the "real" numbers. Any number which is a composite of both the horizontal numbers (reals) and the vertical numbers (imaginary) are called "complex numbers." "Complex" as in a *group,* not as in *complicated*.

Recall that the reason we went down this road to understanding complex numbers was that we are looking for a way to do filtering of signals without having to know the phase of the reference ahead of time. You can now see that complex numbers allow the representation of a rotation. And a rotation represents where you are in a cycle -- also known as phase. We're going to use complex numbers the way we used the special spoke on the bicycle wheel -- it will spin around telling us the phase. But, before we can show that, we need to learn one piece of relatively easy calculus.

Euler's equation

Imagine you have a balloon filled with air. You open the end and at first there's a lot of pressure to so it comes out quickly. Because you're venting the air after a while there's less pressure so the air comes out less quickly. Eventually there's no pressure and no air comes out at all.

Although not perfect, this is roughly an exponential. The word exponential comes from exponere "to put forth."[9] (Why that term was chosen is beyond me.) Exponentials occur in any situation where the rate at which quantity is changing (e.g., the volume of air in the balloon) is proportional to the amount of that quantity there is.

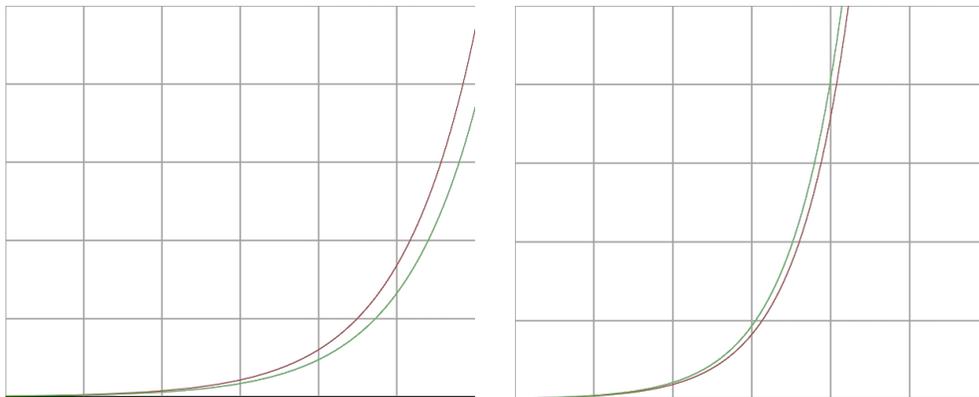As an equation this is written as:

$$\frac{d}{dt}P(t) = k * P(t)$$

This is in Leibniz's notation. It says: "The change in the pressure P with respect to time is proportional by the constant k to the pressure at that moment t."

There's an elegant solution to this equation...

$$P(t) = P(0) * e^{k*t}$$

… elegant but tricky. P(0) is the value at time zero and e is a magic number as magic as pi is. Where does it come from?


Above left is a graph of a^x where a is about 2.2; the plot on the right has a set to about 3.0. The value of a^x is plotted in red and the rate of change of a^x is plotted in green. When the constant *a* is 2.7 or lower the rate is slower than the value (green below red) is shown left. When a is 2.7 or higher the rate is faster than the value (green above red) as shown above right. There's a

magic sweet spot at 2.718… where the rate of change and the value are exactly equal. That sweet spot is called "e" for the "exponential."

You can play with in interactive app of the above http://happyfuncoding.com/run/1393. It allows you to adjust the exponential constant. Note what happens at 2.71. @ADD labels to know what green and red are in this interactive

Now we return to the spinning bicycle wheel. Suppose that the wheel with its special painted green spoke is turning at some speed measured in radians per second. Why radians? Because radians are the natural way to measure a circle. 360 degrees -- that's an invention of the Babylonians. That's not natural nor is any other whole number. Using the radius as a measuring tape makes a lot more sense. The circumference of a circle is 2*r*pi so there are 2*pi radians in a circle.

The instantaneous rate at which the radius (green spoke) is changing is perpendicular to the radius itself. (If you were spinning a rock tied to the end of a string and you let go, the rock would go in the direction perpendicular to the string at that moment.)
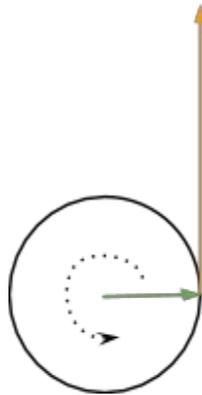
In other words, the velocity vector is the radius vector turned 90 degrees and scaled proportionally to how fast the wheel is turning. Remember how to write a 90 rotation in standard notation? Multiplying by i.

So if r(t) is the radius vector (our special spoke) at time t then the velocity is i * k * r(t) where k is the angular velocity in radians per second.

$$\frac{d}{dt}r(t) = i * k * r(t)$$

And we know that the solution to this is exponential.

$$r(t) = r(0) * e^{i*k*t}$$

The above says in English: "The radius vector, r, starts at r(0) and is rotated by an angular velocity k for t seconds." It is more traditional to use the Greek letter ω (omega) instead of "k" in this equation measured in radians per second and to assume that r(0)=1 meaning it always starts east-facing and has length one. So we get:

$$r(t) = e^{i*\omega*t}$$

This is a very important equation that we will be using a lot. It describes our bicycle wheel with its special green spoke spinning at a constant speed ω (omega).

Let's plug in some specific round numbers into this equation. We'll assume that the radius vector starts off unrotated so r(0)=1. We'll let ω be a nice round pi radians per second and we'll run the experiment for 1 second. Which way will the radius vector point after 1 second?

$$r(1) = 1 * e^{i*\pi*1}$$

If it is going pi radians per second for one second it must be rotated by pi radians. Which way is that? It's half way around. That is, r(1) is pointing in the -1 direction.

$$-1 = 1 * e^{i*\pi*1}$$

Rearrange and simplify that and you get:

$$e^{i*\pi} + 1 = 0$$

That's an amazing little equation as it uses all five of the primitive algebraic constants (e, pi, i, 0, and 1), and it uses all three of the primitive operations (addition, multiplication, and exponentiation). It is in effect the summary of much of algebra in one tidy equation. And all it says is: "When a wheel is turning at half a circumference per second for one second it is turned halfway around" which, stated in English, is pretty obvious. This tidy little equation is called "Euler's equation" -- pronounced "Oiler." However, Euler was such a prolific mathematicians that there are a lot of Euler's equations!

A visualization is available at http://happyfuncoding.com/run/118.

Complex Conjugate

Often we wish to know about the magnitude of a complex-valued vector without worrying about the phase. Remember that the phase is the angle that is spanned by the vector from the right-pointing horizontal axis, and the magnitude is the length of the vector.
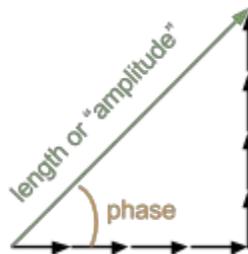

There's a simple way to eliminate the phase: rotate by the negative of the phase angle. That negative angle is below the horizontal axis and therefore is the same thing as simply flipping the sign of the vertical part as demonstrated below.
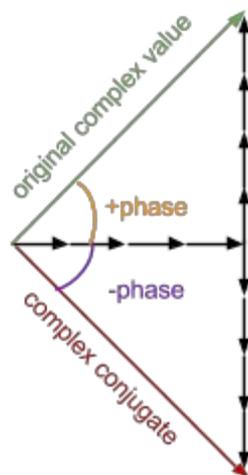
The vertically flipped vector is now called the "complex conjugate" which, as we've come to expect, is a strange name for it.

The product of a complex number and its conjugate will be unrotated (positive angle + negative angle = 0) and its length will be the product of the vector lengths. Because a times its own conjugate both terms are the same length so the result is the square of the vector length. In the below example the length of the original green vector is √2 so when multiplied by its complex conjugate the result is two units long.

original vector   complex conjugate
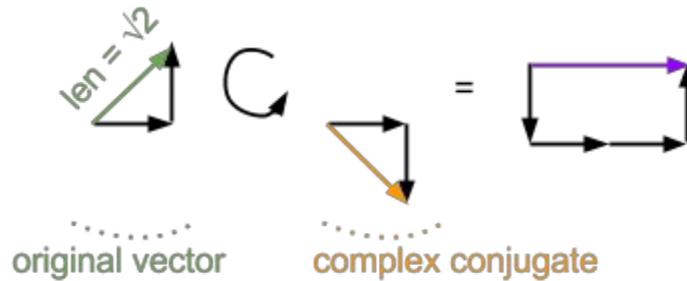
The conjugate is denoted in standard notation with a star like: a*. As we've come to expect, this is a weird convention -- it looks like multiply. In standard notation the multiply symbol is often left out so that *a* times *b* is "ab" and thus the star usually means conjugate for example aa* means to multiply a by its conjugate. But what about a*a is that a "multiply a times a" or is it "a conjugate times a"? In standard notation it would typically be "a conjugate times a." The disadvantage of standard notation where you leave out the multiply symbol is that you can't name any variable with more than one letter (or you have to resort to using subscript) and functions like "sin" look like a triple product: "s*i*n" -- one is supposed to "just know" about the inconsistency and deal with it.

Filtering with Complex Numbers

Now that we have complex numbers at our command we are ready to return to the problem of filtering where we don't know the phase of the reference relative to the signal.

The trick is that we make the reference be not a real-valued sine wave as it was before but rather a complex-valued oscillator. We do this using our little mathematical spinning bicycle wheel of radius one with its special green spoke spinning at omega radians per second for t seconds that we learned above:

$$e^{i*\omega*t}$$

Engineers and physicists have a special name for the above term and unlike most of the other vocabulary from the domain of complex numbers it is kind of cool sounding: "phasor." (Note that Star Trek ripped off that name for their weapons because it sounds neat.) Whenever you see the above form think "phasor" and visualize a wheel with its special painted spoke spinning at ω radians per second. Specifically the term is saying "give me a vector that is the result of rotating omega radians per second for t seconds."

As an example, if ω is π radians per second and t is 1 second then the above form is saying that the phasor is west pointing as we learned above or, in standard notation, its value is -1.

$$e^{i*\pi*1} = -1$$

Another example: pick ω as 2 * π radians per second for 2 seconds. In that situation the phasor will have wrapped all the way around twice and be back to pointing in the unrotated (east facing) direction, or "1" in standard notation.

$$e^{i*2*\pi*2} = 1$$

To filter using the phasor form we substitute the phasor in replacement of the sine wave we used previously. Recall the filtering equation we used before where we had a sine wave as a reference.

$$c(\omega) = \frac{1}{N} * \sum_{t=0}^{N} sin(\omega * t) * s(t)$$

Now we replace that sine wave with the phasor. The following equation F(ω) is now properly called the "Fourier Transform".

$$F(\omega) = \frac{1}{N} * \sum_{t=0}^{N} e^{i*\omega*t} * s(t)$$

$F(w$

This is a really neat equation. The s(t) is real-valued but the phasor is complex-valued. When you multiply a complex-valued term by a real-valued term you are scaling the length of the complex-valued vector. So, when s(t) is big and positive you're making the phasor long so it contributes a lot to the sum. When s(t) is small you are making it short so it contributes very little. When s(t) is negative you are just reversing the direction of the phasor.

The sum is now a sum of complex-valued vectors. Recall that the sum of a vector is just stacking the vectors end to end.

I've written an interactive viewer to help visualize these summations.
http://happyfuncoding.com/run/1395

In the interactive app watch as the red reference moves across the signal when you press "step". Note that it is rotating as it goes. A each step, the red arrow is being multiplied by the signal (orange) and accumulated into the blue series which is complex-valued. The contribution of each individual blue vector starts to really add up so you can zoom out to see the big picture of the sum as it grows.

There are draggable bars in the middle of the screen. Drag the top of those bars to change the speed at which the reference is rotating. Note that as you change the reference frequency the

19

shape of the blue summation changes. When the reference frequency is not well represented in the signal the blue summation curls up and stays near zero. When you hit one of the two well represented frequencies you get a long blue result vector. Drag the lower slider to adjust the phase of the reference. Notice that the blue summation simply rotates around.

The Fourier Transform

We've seen how to use complex-valued filtering to extract out of a signal the amplitude component at a given frequency ω. If we scan over a range of frequencies then we have implemented the discrete Fourier transform and we now have a spectrum of results representing the component of each frequency. Recall that the Fourier component of ω is *complex-valued* -- it is a vector not a scalar and therefore it can't be easily plotted. Fortunately it is often the case that you do not care about the phase but only the magnitude of each component in the spectrum. Therefore it is typical to plot the square magnitude of the component known as its "power" which is scalar.

Recall from the discussion of complex conjugates that you can get the square magnitude of a complex number by multiplying by its own complex conjugate. Multiplying by its own conjugate simply removes the phase and leaves a real-valued product. That product is plotted for each frequency ω to create the "power spectrum."
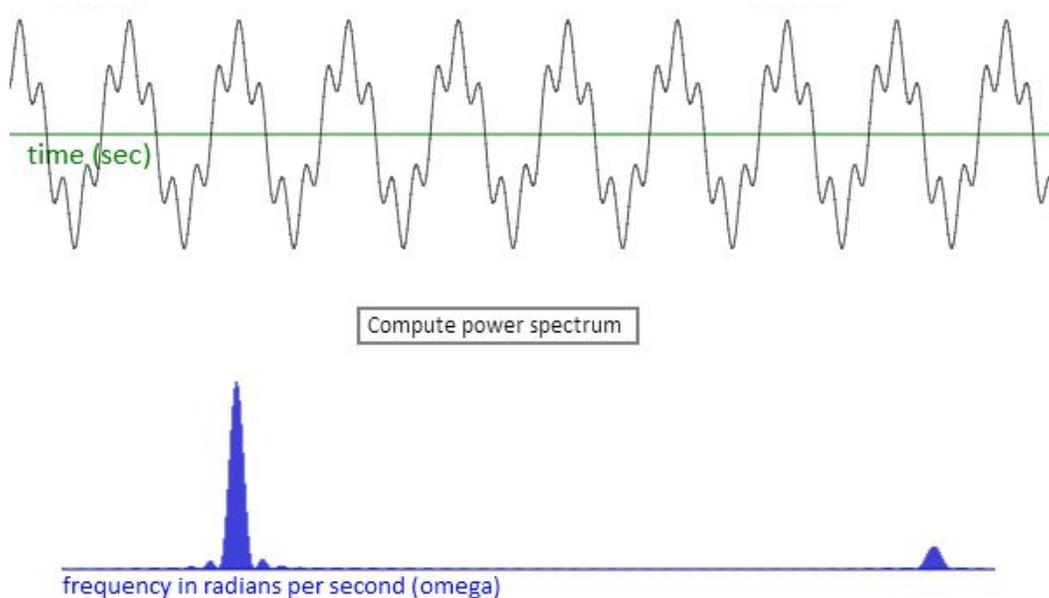
The above top plot is called the "time domain" -- it is our signal as we usually think about it. The bottom in blue is called the "power spectrum." The horizontal axis of the blue graph is frequency in radians per second (ω) and the vertical axis is power (square of the Fourier component length at that frequency). In the blue chart we see two peaks for the two components that are in our signal. The tall one is the dominant low frequency component and

the second smaller one is the little ripple that riding along the main surf. This is a screen-grab from http://happyfuncoding.com/run/1389 after you press the "compute power spectrum" button

Note that the two peaks have thickness -- they are not infinitely narrow even though we put in pure sine waves into the signal. Why is that? This is because we have a finite sample length. Imagine that the signal went on and on and on. In that case the omegas that are close to, but not quite equal to, the true represented frequency would have time to loop around and cancel out. But because we have a short finite sequence the close-but-not-quite-right frequencies return a smaller but still positive sum. You can see this when you are adjusting the frequency bar in the interactive demo from above.

The fact that a pure tone in a finite sample has width is related to the Heisenberg Uncertainty Principle from quantum mechanics @TALK TO PRASHANT

Fast Fourier Transform

The way that I've taught you to compute the Fourier transform via filtering is a computationally expensive but accurate way. Computationally expensive because the number of operators needed to do the work grows to the square of the sample length, N. That is, for each frequency omega (you will have N such frequencies) you have to compute a filter which takes N steps. Thus it takes N^2 operations.

There's an alternative way to do it called the "Fast Fourier Transform" or FFT that does the work in N * log(N) number of operations which is way faster. However, the technique imposes the constraint that N has to be a power of two which can be difficult to deal with in practice.

Like so many other mathematical concepts, the FFT was first discovered by Gauss in the early nineteenth century but he failed to publish it. Apparently Gauss even invented the Fourier Transform before Fourier, but again, didn't publish it. His FFT algorithm had to wait another 160 years before Cooley and Tukey published it in 1965.[10]

References

[1] http://www.etymonline.com/index.php?term=sine
[2] http://www.etymonline.com/index.php?term=frequency
[3] http://www.etymonline.com/index.php?term=amplitude
[4] http://www.etymonline.com/index.php?term=phase
[5] http://www.etymonline.com/index.php?term=filter
[6] http://www.etymonline.com/index.php?term=vector
[7] via Wikipedia Martinez, Albert A. (2006), Negative Math: How Mathematical Rules Can Be Positively Bent, Princeton: Princeton University Press, ISBN 0-691-12309-8, discusses ambiguities of meaning in imaginary expressions in historical context.

[8] via Wikipedia Rozenfeld, Boris Abramovich (1988). A history of non-euclidean geometry: evolution of the concept of a geometric space. Springer. p. 382. ISBN 0-387-96458-4., Chapter 10, page 382

[9] http://www.etymonline.com/index.php?term=exponential

[10] http://www.math.ethz.ch/education/bachelor/seminars/fs2008/nas/woerner.pdf