Here’s what synthesis is all about:
Is it a short sound, or long one?
That’s the envelope. The envelope generator makes an up-then-down set of values. Hit a note, and the volume will go from nothing to high then back down again. Patch the envelope generator into an attenuator (a VCA is a voltage controlled attenuator, although we don’t have to be dealing with voltages these days, we originally did, so you could think of it as being a Value-Controlled-Attenuator (or amplifier if you want to inverse the thinking) and it’ll still mean the same). Short sounds: low values for attack, decay, sustain, release. Long sounds, low value for attack (if it starts suddenly, or higher value if it smoothly creeps in and takes you by surprise when you realise it’s standing beside you); long sounds: high value for decay, long sounds maybe high value of sustain (which is not time, like the others, it’s an actual level), and long sounds, high value of release. Piano type of long sounds, make the sustain lower. The decay is still long but not quite as much, but the release can be longer. Okay, is it to be a harsh sound or soft mellow sound? Harsh sounds (a dog barking, or sawing wood, or frogs) have lots of harmonics and possibly noise. Smoother mellow soporific sounds have not as much harmonic complexity. A soft gentle sound is more like a sine wave (like whistling that doesn’t want to draw attention, or an owl). How do we get harmonics? Start with a ridiculous wave shape, like square or sawtooth. How do we get not as much harmonics? Start with a sine wave or the nearest thing to it, like a sawtooth wave. What about pulse wave? They sound nice, hollow, and lacking in energy, yet complex in a small way (especially if you can vary the width of the pulse from tiny to not quite as tiny but still tiny). Now that we know how to make long vs short, and raucous vs mellow, how do we change the nature of the sound? In subtractive synthesis we hack some of the sound out. We chop out some aspects of it but leave other parts behind. Like yeast turns sugar to alcohol but leaves behind a certain profile of sugar (or the beer would have no taste), we use a filter to carve bits out. Not only that, a VCF - voltage (or value) controlled filter allows us to vary the nature of our filtering over time, if you let something control it. If you started with a harsh sound, you can filter out the harshness (the harmonic complexity) and smooth it out. The filter is usually a low-pass filter. This means it lets the low frequencies pass unhindered, as if some kind of border guard or night club bouncer. Your papers are in order, you may pass. The higher frequencies, on the other hand, get stopped and turned away (attenuated). The filter – a low pass filter (LPF) has a couple of controls: One is to decide what that cutoff frequency

is (or maybe labeled centre frequency, as that's the centre of the decision below which fundamental frequencies pass and above which harmonics will be prevented from passing). The other control, resonance, is a funny one. That parameter is often called Q, which stands for quality, yet it means how 'peaky' the centre cutoff point is. The slope can either be a gentle cutoff, where you hardly notice it's doing any cutting off, but it is. Or it can be a sudden slope where everything above the cutoff point is definitely being lost. Or you can go further into ridiculousness and make the peak resonant (i.e., go up - yes, actually up in value compared to all below or above frequencies neighbouring it, and this gives a pronounced effect at that frequency). Go even further and the filter misbehaves and starts eating its own output and feeding it back in again at that one centre frequency – i.e., it starts oscillating like any badly designed amplifier will. On purpose. Young people like that effect. Don't use it. Get to the point where it almost oscillates but not quite, and be satisfied with that. You've got responsibilities, you know. What else is there? Nothing really. Well, maybe more in future episodes. I'll tell you what, I'll have to put some videos together instead. That'll be better. There is some more to explain - LFOs, Sample and hold, the crazy notion that you can squeeze the envelope values into something other than just amplitude, maybe the filter or – crazy talk – the oscillators themselves at the beginning of the chain – to make their frequencies go up and then down for each note pressed? No way! Yep, seen it done.

I might as well explain a few other models. Additive synthesis. It's like how an organ works. You add partials together to make a more complex sound. There have been a few additive synths on the market, but lacking a cohesive model of how they work, each one tends to stand alone somewhat. The essential idea is that you start with many oscillators or wave generators, and add them together. Each oscillator can be pitched differently yet sound at the same time, and set at a different level relative to the others, and perhaps even given a separate amplitude level or envelope over time of the note. All per one press of the key. If that's the case then there's no need for a filter (VCF) or even the final VCA as the same effect is achieved earlier up line. Imagine back in the above previous example of subtractive synthesis, you have a complex-ish wave (that in reality if you were to break it down, consists of a fundamental (sine wave- ish) and the next harmonic up and the next harmonic up from that), and you hack the upper harmonics out with the filter, over the duration of the note thanks to an envelope generator tweaking the filter, to get that “oww” effect. Well, if you started with a single oscillator doing the fundamental alone, and added another oscillator doing the next harmonic up, and yet another doing the next harmonic up from that, and gave each of those extra oscillators a varying amplitude envelope such that they’re staggered in duration (the fundamental lasts the longest, then the next one dies out before that, and the higher one dies out the soonest), you get pretty much the identical result. Except without an actual filter, and associated things – just a bunch of oscillators whose amplitudes are under control by individually set (but related) envelope shapes.

Okay, another one: FM. If you have a waveform, let’s say, a sine wave, it’ll be a relatively soporific uninteresting featureless tone. Let’s say you increase the frequency to another higher tone. Same thing, only higher pitched. Let’s decrease the frequency to a lower tone than we started with. Pretty much the same thing except lower. Waste of time? No. Let’s put a bit of lag in the way we go to higher then lower frequencies (i.e., portamento – it’s a thing that ’slurs’ or lags step changes into having sloping curves along the way (also called differentiation)). As the frequency shifts higher then lower, you’re going to change the shape of the sine wave a bit. This cycle we’re on now and the next cycle to occur will be a slightly different frequency – if we’re going up, then the wave shape at the start of the waveform will be as if it is a static frequency, but hang on, we’re going up, so by the time the waveform is finished, it’s going to be all crushed up a bit. Therefore it’s no longer a sine wave. Therefore there’s additional harmonics in there. Therefore we’ve given it additional partials, at least, partially. If we do this very quickly (for example, instead of prodding a keyboard with portamento on it, we use an actual sine wave oscillator, and crazy as it may seem, one that is running at an audio frequency) then we’re going to introduce wave shape twisting bending crushing distorting or otherwise altering at a very fast rate. If we’re doing that at an audio frequency, we’re introducing harmonic presence at a rate of another harmonic presence. (Then we get into that daft stuff we didn’t learn at school, like quadratic equations, which it turns out do have a use - carrier freq plus modulation freq plus carrier+modulation freq plus carrier-modulation freq. Or something. Who cares.) So, FM is distorting a wave shape to make harmonics occur. If you FM an oscillator with another oscillator, there’ll be a complex set of additional harmonics. Now wait – if you do that, but in-between the two oscillators there’s an enveloped VCA, you now have the possibility of a wide variance of harmonics across the time duration from the start of the note to the end.

For example, harmonically rich content at the beginning, fading away to lower harmonic content by the end of the note (because the modulating oscillator starts all heavily modulating, then its modulation effect fades away thanks to the envelope). You know, that “oww” effect again. By the way, phase distortion does it the same way, but legally is different.

Hands up who wants another? In the meantime, something to think about: are all the parts of the synthesiser, as we know it, really that different? Perhaps not. What is an oscillator? Most would say it is a unit that outputs a continuous waveform. However, arguably it’s a badly designed amplifier! We’ve all experienced feedback – an open mic in a venue that can hear its own amplification and speakers. A badly designed amplifier might reach oscillation if it is so good at amplifying that it becomes unstable and steals part of the output to leak back into the input. On the other hand, I previously referred to a filter as having a Q parameter, which in some cases can be overdriven into oscillation – most of your synths can do that if you push the Q up into oscillation. If you do that, isn’t it now an oscillator? Yes it is. Is a filter an amplifier (even if a unity gain amplifier, i.e. amplifies input level 1 all the way up to output level 1, but no bigger)? Yes, but a selective one – it is fussy about what it lets through. Therefore a filter is sometimes an amplifier and sometimes an oscillator and an amplifier is sometimes an oscillator and an oscillator is sometimes a… ? Well, sometimes it is a free-running transient generator (i.e., transient generator another name for an envelope generator) except that it is going to continuously run cycles of transients (i.e. the wave form) – you can’t stop it from doing that. Or can you?

Alright, one more and that’s the lot for now. Physical Modelling synthesis (Karplus-Strong, waveguide, etc). Go and get a bottle (plastic for safety, glass at your own risk) and thwack the opening with your palm. Or even, go and get a kitchen paper roll (even with the paper still on) and thwack one open end with your palm. Or get a glass and carefully hold it without muffling it while gently pinging the glass with your fingernail. Then hold the glass fairly securely in your hand to muffle it and ping the glass with your fingernail. As I was saying above, a lot of the parts of a synth are not as separate as we’d think. Filters that drive themselves into oscillation, or maybe almost oscillation. Why do they oscillate at a certain frequency? It is because you set the cutoff or centre frequency at a specific setting, then pushed the Q up high. Why does the open microphone feed-back at a certain frequency? Why does a glass ping at a certain frequency? Why does the kitchen roll cardboard tube make a dull thud at a certain frequency? It is not the frequency of your palm hitting it. It is also, by the way, a longer note than your palm hitting it was – quite a lot longer. Pinging the glass lasts quite a long time. The initial attack is short, but the decay is very long. Yet your fingertip initiating this event was short, and had no discernible tuning. If you've got a kind of synth that allows this, get your VCF and have nothing going into it (no oscillators etc) at all, and push the Q up to almost oscillation but not quite there. If it starts to, take the Q back down a pinch. Now you need a change or transient. If you can (again, depends on the synth), send a pulse – a single pulse, not a continuous one, perhaps even a keyboard on/off gate transition – into the filter audio input (yep, an on-off or off-on switch going into your filters audio input!). It will likely cause the filter to "ring". This will have its own natural decay (maybe), and it will have a specific frequency. It is much like hitting the bottle (not hitting the bottle, but actually hitting the, oh never mind) It is much like thwacking the cardboard tube. You've sent an impulse into a highly resonant and almost chaotic unstable situation (that has nevertheless got some amount of equilibrium (i.e. you get a particular note of oscillation, not some nonsensical noise)). Moving on, let's imagine a delay line. Like a reverb spring, or an echo unit. Imagine a short delay. Most of you have had experience with delay units, driving the feedback up a bit too much, and taking the delay length down a bit too short. What happens? The feedback does what feedback does – it self oscillates. Your clothes are red. I mean, your delay unit is now an oscillator. Pinch it back to slightly pre-oscillation and it is still very audibly peaky and resonant. Sweep the delay time a bit back and forth but still very short and you get a peculiar effect on the signal going through it. Your clothes are black. Your delay unit is now a filter. Not necessarily a low-pass filter, but perhaps a more complex one. As a quick aside, let's go through a few filters. Low pass, as we saw earlier lets the low frequencies pass unhindered. It is only the high frequencies that it diminishes or removes. High pass, guess what, lets the high frequencies pass and increasingly prevents the low frequencies. Put the two together and what have you got? Highlowpass. Lowhighpass? Hlowigh? Lihogh? Nope, we call it a band pass. If the high pass passes only the high, cuts the low, and if the low pass passes only the low, cuts the high, then if you set the high pass to pass 700Hz and above, and the low pass to pass 1.3KHz and below, then you'll basically be left with a little slab between 700Hz and 1.3Khz (as you can see, the band width – the width of that band – is 600Hz). Bring them closer together (hpf is 900Hz and lpf is 1.1KHz) and the bandwidth that gets passed is 200Hz. Move them both up at the same time and you move the band pass centre frequency up. Move the ganged pair back down, you shift the centre frequency back down. What if you had a filter that instead of passing a band, swapped the positions (e.g., hp cutoff is 1.1KHz and lp cutoff is 900Hz)? Aha, it notches. It passes everything except a small band which it tries to remove. Funnily, human hearing is not sensitive to this effect. Until you sweep the notch. Then we sometimes can detect it. What is a phaser (or flanger)? Actually it's a bunch of notch filters. We call this a 'comb' filter, due to its resemblance to a hairbrush toothbrush comb. Again, we can kind of hear this, not very well when left alone, but particularly well when swept through the spectrum. You can also kind of hear it if a friend of yours has just set up a magnificent new stereo system and you stand in the centre and you move your head side to side a bit. Then you look at how they've wired it up and quite astoundingly they've managed to cross the phase of one speaker's wiring with respect to the other. So, a delay line can be an oscillator and can also be a notch filter or swept comb filter. But it is also a delay line as well, so stuff comes in and stops coming in, yet stuff comes out and keeps coming out. Your thwack on the cardboard tube, or ping on the glass, was a short impulse, yet it resulted in a resonant longer lasting output. The tube became a delay line with resonance and an element of feedback, and the feedback passes through physical characteristics to modify the new regeneration of signal. In the case of the cardboard, not much regeneration and less 'tuned', in the case of a glass, more specific and also more of it. A string pluck on a stringed instrument is an impulse, and the string against the body cavity forms a tuned delay line, with feedback (the 'sustain' guitarists like). A piano is similar. A blown instrument is a train of impulses (blowing what is essentially a raspberry into a shiny assemblage of expensive plumbing (or in the case of a didjeridoo, not)). A drum is, well, a drum. A violin is a series of impulses (the untuned bow) across the same tuned string and body cavity as a guitar pluck. The Godley-Creme Gismo supplied a series of impulses to an electric guitar. A triangle is an impulse and a resonant delay line (the, er, triangle part). Tubular bells – impulse, resonant feedback-prone delay. It goes without saying, but I'll say it, that we can do all of that in a computer (it can even be, and has been, done using purely analogue delays). We can initiate an impulse, and we can set up a state where that impulse drives a resonant delay line into (or not into) oscillation and give parameters to the feedback loop, and tap off outputs. The characteristics of the delay and feedback govern the envelope and development of harmonics and what happens to them. Given several delay lines in parallel (initiated from the same driver or impulse), we can see how different tonal characteristics can be had by different development or dissolution of harmonic content as the note dies away. And (stand back, or sit down) – their interaction with each other! (Yep, that happens in real instruments). The driver can be a simple impulse like a tap or click or it can be more complex. There's a recently introduced app (Mersenne – though I haven't got it) that even does a quick bit of FM synthesis to the driver impulse to make it more fruity to start with (that's your "hit") and adds some filtered noise (which also contributes to a rich impulse), and then leaves you to specify parameters of the resonant chamber (which is effectively a delay line filter-y thing fed-back in almost oscillation) to do the "body" of the sound. There's other ways of doing physical modelling, such as imagining they are physical connected bodies with stickiness or resistance, versus slipperiness, such as the Spring Sound app. But essentially, they all involve a driver impulse and then a resonant or delayed feedback thing.